elastic

How to use Elastic to scale AI compliance efforts

Executive summary

As governments worldwide introduce laws and regulations governing the development and use of artificial intelligence (“AI”) within services and tools, organizations are adopting appropriate measures to ensure AI systems are transparent, risk-managed, and compliant with their respective legal requirements. The Elasticsearch Platform helps enterprise customers to implement comprehensive controls to monitor AI deployments, perform impact assessments, and maintain a more trustworthy ecosystem. This white paper explores key areas of AI governance, from understanding transparency mandates to ethically leveraging data for model training. We’ll provide a roadmap that empowers your enterprise to navigate regulatory expectations and to drive innovation with confidence. We’ll also share how Elastic’s powerful platform can help you track and manage compliance with your own requirements under applicable AI laws.

Please note: *This white paper is provided for informational purposes only and is not intended to constitute legal advice. Please consult your own legal counsel for legal advice.*

Background and primer on global AI laws

Over the past few years, the global regulatory landscape for AI has evolved significantly, with new laws around the world addressing the development, application, and oversight of AI. The rapid emergence of generative AI (GenAI) and large language models (LLMs) has enabled organizations and consumers to leverage data in entirely new and transformative ways. With this progress, the dynamic nature of these technologies has naturally raised questions about their legal, ethical, and practical implications.

While there are key distinctions between AI laws around the world, the laws share many overlapping principles. Notable examples include, among others, the EU AI Act, South Korea's AI Basic Act, Brazil's AI Act, the Colorado AI Act, the California AI Transparency Act, California's regulations related to automated decision-making technologies under its existing CCPA privacy law, and the Utah AI Policy Act. Beyond these current legislative initiatives, international and national voluntary frameworks such as the OECD's AI Principles, Australia's AI Safety Standard, Singapore's Model AI Governance Frameworks, or even contractual agreements with customers and end users can impose additional obligations regarding data usage and processing activities.

To navigate these requirements and expectations, organizations can strategically leverage Elastic to effectively track, manage, and improve compliance with AI laws. With Elastic as your partner, we'll help you build a future of responsible AI innovation.

Artificial intelligence and machine learning

AI has evolved dramatically over the past several decades. Early AI systems relied on rule-based programs designed to perform narrowly defined tasks following explicit instructions. Over time, with the advent of machine learning — a subset of AI where computer systems use statistical techniques to "learn" from data and improve their performance without being explicitly programmed — the field has progressed to sophisticated models that can perform complex tasks such as natural language processing, image recognition, and automated decision-making.

While the definitions for AI vary across laws and industry guidance, the EU AI Act provides a helpful starting point, describing an AI system as software that is developed using machine learning, logic-based, or statistical methods and which, for a given set of human-defined objectives, generates outputs — including predictions, recommendations, or decisions — that can influence real or virtual environments. A significant development within this landscape is GenAI, which refers to systems that interact with users through text, audio, or visual communications, and that process such communications to produce certain targeted outputs. This evolution from rigid, rule-bound processes to dynamic, data-driven learning systems has fundamentally revolutionized the possibilities for using data with AI.

Developers and deployers

Developers and deployers of AI systems fulfill distinct yet interconnected roles within AI's rapidly evolving ecosystem. Regulatory proposals and statutory frameworks, such as those found in the EU AI Act and several of the US State Acts, generally define "developers" as the persons or entities who design, create, train, and maintain AI systems. Their responsibilities typically include the technical and theoretical foundations of the systems, including algorithm design and model training.

By contrast, "deployers" generally refers to the individuals or organizations that determine the intended purposes of AI systems and then integrate them into products, services, or operational workflows. Deployers often bear the responsibility of ensuring that the AI systems they implement operate in compliance with the established standards for fairness, transparency, safety, and accountability.

Together, *developers and deployers* define the lifecycle of AI, from conceptual design and development through real-world application, and highlight the importance of clear accountability throughout AI system implementation.

Automated decision-making and profiling

In addition to laws that specifically regulate AI technology, there are a growing number of laws prohibiting the use of automated decision-making technologies, which would include AI, in ways that may result in illegal or unfair discrimination (even if such discrimination is inadvertent). For instance, Illinois law places limitations on AI usage that could lead to discrimination based on protected characteristics in personnel recruitment and retention. Similarly, New York City's Local Law 144 regulates certain "automated employment decision tools" that significantly influence employment decisions, mandating bias audits, among other requirements.

Further, various other proposed laws and regulations target AI systems that produce simplified outputs used to assist or replace discretionary human decision-making, such as scoring, classifications, or recommendations.

Additionally, to the extent that personal data is used in these systems, certain privacy laws such as, most notably, the EU GDPR impose restrictions and obligations on automated processing aimed at evaluating, analyzing, or predicting aspects of an individual's characteristics, behavior, economic situation, health, personal preferences, or interests.

Risk-based approach to AI legislation

Many of the new AI laws adopt a risk-based framework to classify AI applications according to their potential for harm. For example, the EU AI Act differentiates between Unacceptable, High, Limited, and Minimal risk applications, with systems like workplace sentiment and emotion analysis being prohibited. Similarly, in Colorado and under other proposed regulations in the US, there is an emphasis on assessing risk related to particular AI deployments. This indicates a growing trend toward regulating particular applications of AI, especially where the decisions can significantly impact groups or individuals.

Foundational AI principles

Prior to the enactment of AI laws, industry standards and best practices organically emerged to guide the responsible development and deployment of AI. These self-regulatory measures were introduced by industry actors, standard-setting organizations, and academic researchers, all seeking to address ethical and operational concerns associated with the rapid advancements in AI technology and deployment. Key principles emerged from these early efforts to ensure that AI systems operated in ways that were understandable, fair, and accountable.

1

Transparency

This principle is a commitment to openly share information about how AI systems are designed and operate, including disclosure of data sources, methods, and decision-making processes, so that users and stakeholders can understand and trust the systems' functioning.

2

Explainability

This refers to the capability of an AI system to provide clear, understandable, and interpretable justifications for its outputs or decisions, thereby enabling developers, regulators, and users to trace and assess the rationale behind the system's conclusions.

3

Protection from AI bias and algorithmic discrimination

This principle acknowledges the unfair outcomes that may arise from illegal or unfair bias in the data or design choices in AI systems. It emphasizes the importance of ensuring that the technology does not systematically disadvantage specific groups or individuals in an unfair or illegal manner.

These principles laid the foundation for later legal frameworks by highlighting the ethical imperatives for the responsible integration of AI.

The enterprise cost of non-compliance with AI legislation

Failing to comply with the growing body of AI legislation is not just a compliance oversight; it can be a real threat to an organization's financial stability, market position, and long-term viability. While there is a growing body of regulation that continues to evolve, existing penalties in current regulation are intentionally severe to reflect the significant societal and economic harms that unregulated AI systems can inflict. For example:

- The **EU AI Act** imposes fines of up to 7% of a company's global sales for breaches related to high-risk or unacceptable risk AI systems, or €35 million, whichever is higher. For a multinational corporation with billions in revenue, such a fine could amount to hundreds of millions, or even billions, of euros and pose a significant threat to profitability, investor confidence, and market capitalization. Other violations under the Act can cost up to 3%, and up to 1.5% for supplying incorrect information. The EU AI Act also has an extraterritorial effect, meaning that any provider offering an AI system in the EU market must comply, regardless of their physical location.
- **Brazil's draft AI Act** not only envisages financial penalties of up to R$50 million (US$9 million), but as proposed, it would also give regulators the power to order the suspension of non-compliant AI services and to mandate system adjustments.
- The **Colorado AI Act** considers failures to use "reasonable care" to avoid algorithmic discrimination in high-risk AI systems to be an unfair trade practice, allowing for penalties of up to $20,000 per violation and up to $50,000 per violation if committed against an elderly person.
- The **California AI Transparency Act** imposes fines of up to $5,000 per violation, per day, against covered providers and, in certain cases,

may impose injunctive relief. The “per day” imposition means delays in remediation or continuous non-compliance can quickly escalate into devastating financial burdens.

- The **Utah AI Policy Act** imposes fines of up to $2,500 per violation and allows for other relief, such as injunctions or disgorgement of monies made in violation of the law. Ongoing violations may result in a $5,000 penalty per violation. Companies are responsible for violations caused by their generative AI applications, even when the AI is directly responsible for the violative output. This shifts the burden of compliance entirely to the deploying organization.

Beyond quantifiable financial penalties, non-compliance with laws regulating AI carries intangible, yet equally impactful costs. Damaged brand reputation, loss of customer and stakeholder trust, and operational inefficiencies can lead to long-term market disadvantage and hinder growth and innovation.

Additionally, where monetary or injunctive penalties may not suffice, the Federal Trade Commission (FTC) in the US and other enforcement agencies can pursue “algorithmic disgorgement,” which would require an organization to delete not only the unlawfully obtained data, but also any algorithms or models that rely on that data. As AI becomes more integral to business operations, the financial and strategic implications of non-compliance continue to grow.

How Elastic can help enterprises streamline AI legal compliance

As a leader in innovative AI solutions committed to an open development process with transparent and direct engagement with our community, Elastic is also dedicated to building transparent, responsible, and explainable systems. This commitment directly empowers customers to manage their data confidently while ensuring robust compliance with evolving AI legal standards. Elastic offers a comprehensive suite of capabilities that directly address core compliance challenges posed by the new regulatory environment. With Elastic, you can transform complex compliance mandates into streamlined, automated processes.

Among the AI laws emerging over the last few years, there has been a trend toward safeguarding against potential AI harms, whether they relate to lack of transparency, concern around algorithmic illegal or unfair discrimination or bias, or automated decision-making generally. While many legal frameworks already regulated the underlying data being processed by AI solutions, few (if any) regulated the technology itself or the companies designing and/or using such solutions.

Thus, compliance with global AI laws requires an understanding of the entire ecosystem in which an organization's data resides and travels and how that data is otherwise processed. This is where Elastic can help our customers simplify and automate these processes, supporting your compliance frameworks.

The following table illustrates how Elastic can help organizations navigate various AI compliance use cases:

AI compliance challenge	Core regulatory need	Elastic capability	Key benefit
Transparency	Notice and disclosure	Centralized logging, metrics, audit trails	Demonstrate data flow and decision-making, simplify investigations
Documentation and data inventories	Data inventory	Data mapping and classification	Automate data governance, ensure accurate reporting
Identifying risks	Continuous monitoring	Real-time alerts and analytics	Proactive risk adjustment, dynamic control implementation
Conducting impact assessments	Algorithmic discrimination prevention	Search functionality, data lineage tracking	Streamline assessments, ensure foundational compliance
AI literacy and policies	Training	Comprehensive training platform	Operationalize AI knowledge, empower staff for oversight
Providing user choices	Individual requests	Data mapping and categorization	Respond to requests faster, streamline individual rights management

Transparency: Using Elastic to meet notice and disclosure obligations

AI systems are inherently complex, and ensuring transparency — whether in relation to a legal, regulatory, or contractual obligation — is fundamental to building trust with users, regulators, and stakeholders. Regulations, such as the EU AI Act, require organizations to provide insight into data usage and model decision-making. For example, while notice-related requirements under the EU AI Act and elsewhere vary based on the relevant industry or type of AI involved, most of these laws contain general obligations to inform end users when they are interacting with AI and, in certain circumstances, to present clear notice to users and maintain an inventory of data used to train models. In general, some of these emerging frameworks — such as in California and Colorado — may also require notices prior to use and, in some cases, before a consequential decision is made about the end user. Among the growing patchwork of laws, the obligation to understand and be able to convey the relevant data and processing involved remains consistent throughout US and EU laws.

The Elasticsearch Platform centralizes logs, metrics, and audit trails across environments, enabling real-time monitoring and historical traceability. This helps our customers demonstrate how data flows through their AI systems and how decisions are made based on such data. Specifically, our customers can leverage Elastic to implement measures that help facilitate compliance with these transparency obligations.

For example, Elastic customers can:

- Integrate diverse data sources across operations, AI applications, and user interactions to better understand their data inventory, so that users can identify, classify, and otherwise evaluate data used for training, testing, and validation (among other things)
- Conduct audit trails by maintaining logs that record data lineage and model activity for forensic analysis and compliance reporting
- Use tools, such as Kibana, to create dashboards that help users simplify data investigations by searching, aggregating, and visualizing how AI renders certain decisions

Elastic's customers can ingest and store detailed logs from their AI applications. This can include LLM prompts and responses and any errors or exceptions. This data is crucial for understanding AI system behavior.

Elastic's powerful search capabilities enable customers to index and make searchable vast amounts of structured and unstructured data, including technical documentation, training data details, and operational logs.

Elastic customers can access real-time monitoring with custom dashboards in Kibana to track the performance of their AI systems. Kibana features like log analytics, anomaly detection, and pattern analysis can help track AI system behavior. This helps identify anomalies or unexpected behavior that might necessitate disclosure.

Learn more about how Comcast uses Elastic to visualize data trends and anomalies and to share insights between teams.

Documentation and data inventories: Using Elastic to develop and promote appropriate use of AI systems

Related to transparency, the EU AI Act and certain US State laws, such as in California, require maintenance and publication of certain documentation concerning certain AI systems. For example, as of January 1, 2026, California's AB 2013 mandates that AI developers post documentation on their website prior to making generative AI systems available to consumers. A developer under California law refers to corporations that "design, code, produce, or substantially modify" AI systems. Among other things, the documentation requires a high-level summary of the datasets used to develop the generative AI system, including the sources of the datasets, a description of how the datasets further the purpose of the AI system, and whether the dataset includes aggregate or personal information.

As described above, Elastic enables effective data mapping to evaluate your data — including ways in which you may modify our search experience to better tailor the solution for your end users. Additionally, by enabling our customers to centralize, tag, and understand their data, we empower them to understand the obligations that apply to particular data, whether those obligations arise out of legal, contractual, fiduciary, or confidentiality obligations.

Learn more about how Sitecore uses Elastic Security to centralize data in one place and automate up to 96% of security workflows.

Identifying risks related to your data and potential AI use cases

As emerging AI laws prescribe different requirements based on data types and use cases, it is more important than ever to understand, manage, and protect your data.

With Elastic's continuous monitoring capabilities, customers are able to assess risk related to their data and potential uses thereof, so that they can more effectively adjust controls as risk levels change over time. For example, under applicable law, high-risk AI systems — such as those used to make medical or legal decisions — are subject to tighter controls. Elastic supports our customers' implementation of risk management frameworks by providing real-time alerts, customizable dashboards, and detailed analytics, so that users can set rules and parameters around how to address (and redress) potential harms arising from the use of AI systems, including our search functionality.

Learn more about how Ernst & Young uses the Elasticsearch Relevance Engine to improve accuracy and accelerate retrieval of key insights from unstructured data that are integral to compliance and innovation.

Conducting impact assessments

Similar to existing obligations under certain privacy laws to conduct data protection impact assessments, emerging AI laws — such as in the EU and in Colorado — require AI deployers to conduct impact assessments, which is especially salient for high-risk applications. These assessments generally require documenting key details about the AI use case, including various details about the system, its purpose, data used, intended benefits, risks of algorithmic discrimination, safeguards, and post-deployment monitoring.

Elastic enables customers to understand when and how to conduct these impact assessments. Particularly, knowing where data lives, how it is processed, and where it flows streamlines completion of impact assessments, which traditionally can require multifunctional support across business units to understand uses of personal data. These impact assessments in turn demonstrate foundational compliance while enabling organizations to limit processing of data to what is authorized under applicable law.

Learn more about how pharmaceutical companies are using Elastic to help researchers and compliance teams generate usage reports from ingestion through to search and to streamline reporting obligations to regulatory bodies.

Implementing AI literacy and risk management policies and procedures

Under the EU AI Act, AI providers and deployers should take measures to ensure that their staff involved in AI operation and use have a sufficient level of AI literacy (including, particularly, individuals who exercise human oversight functions). Moreover, the AI literacy objective expects that organizations develop and implement training programs tailored to ensuring that staff understand the opportunities, risks, and limitations associated with AI systems being deployed and, further, be able to recognize and mitigate potential harms. This expectation goes hand-in-hand with requirements elsewhere, like in Colorado, to implement risk management policies and programs to address potential algorithmic discrimination.

Elastic enables customers to scope, determine, and document what they consider to be sufficient AI literacy for the intended use case. Elastic can also support meeting these requirements by leveraging our comprehensive training platform, technical expertise, and integrated data solutions, particularly through our extensive library of on-demand training and virtual instructor-led courses including advanced concepts in machine learning and AI. The training subscription offers hands-on and practical exercises that reinforce theoretical concepts, making an abstract understanding of AI processing more concrete.

Providing users with choices

Many AI laws (and laws that impact certain AI deployments) compel organizations to provide their users with clear choices regarding their data and how decisions are made. For instance, regulations may require transparency around profiling and automated decision-making processes, and users may have the right to opt out or request human intervention.

Elastic's data mapping capabilities form the core foundation by which organizations can process data subject requests. Specifically, organizations can use Elastic's data mapping and data categorization capabilities to quickly ascertain how to determine the validity of such requests and respond to requests as appropriate or required, saving valuable time to enable compliance teams to respond within the short timeframes afforded by these laws.

Reducing algorithmic discrimination and conducting bias audits

AI systems rely on large volumes of training data, and the quality, diversity, and sourcing of that data directly impact the fairness and reliability of AI outcomes. Regulations increasingly focus on the provenance and bias of training data to ensure ethical AI deployments.

The Elastic platform can ingest and index data from diverse sources, including logs, training data, and outputs from machine learning models. Because Elastic enables search and analysis of all data types — without moving or rehydrating data — organizations can collect data from the full decision pipeline, from input data through to final outcomes, in one place. Elastic's platform allows customers to query, explore, and visualize datasets to assess their composition and identify potential biases or gaps that could affect the AI system's performance and fairness.

Additionally, by using Elastic's powerful query DSL, organizations can filter and drill down into data to compare outcomes across different demographic groups. For instance, customers can perform aggregation queries to detect whether an algorithm's decisions disproportionately affect certain populations.

By enabling customers to maintain detailed records on data and their sources, we provide them with the ability to have a 360-degree view of their data, so that decision-making can become less of a black-box.

Conclusion

Take control of the future of AI compliance with Elastic

Understanding your data and how your technology renders decisions is increasingly becoming as much an industry best practice as it is becoming a legal requirement. The ability to comply at scale with a growing body of AI-related requirements is expected to become a market differentiator and can support an organization's strategic success. Elastic streamlines critical steps of this process so that we put you in the driver's seat of compliance. By transforming regulatory challenges into strategic advantages, Elastic enables organizations to not only mitigate risks, but to innovate responsibly and confidently.

