AI Governance: The Internal Framework Every Organization Needs

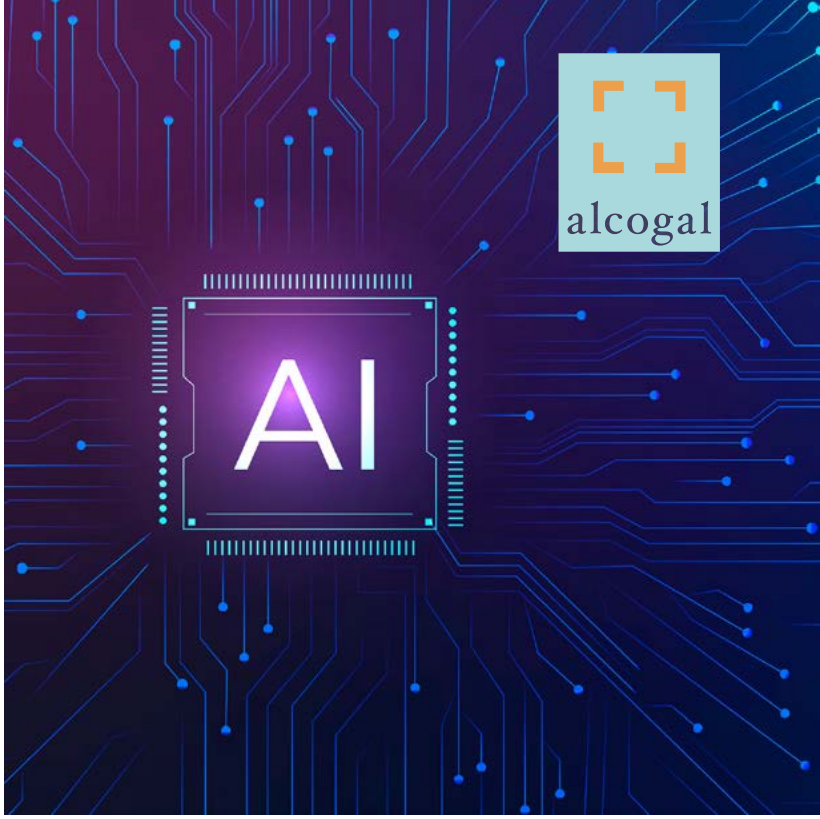

There is one figure worth keeping in mind: while 63% of CEOs worldwide consider artificial intelligence a strategic priority, 91% acknowledge they do not feel adequately prepared to implement it responsibly.

Governing artificial intelligence is not the same as securing systems, nor is it synonymous with compliance. Information security and compliance each have their own objectives and structures. Organizations have accelerated AI adoption in pursuit of efficiency and competitiveness, yet few have moved with the same urgency to define who is accountable for the systems they deploy, under what conditions those systems operate, and what happens when they fail.

Across Latin America, AI-specific regulation remains largely in development. More than 109 legislative initiatives are currently under discussion across eight countries, yet only three have reached the approval stage. Brazil marked the most significant exception in December 2024, when its National Congress approved Bill No. 2,338/2023, establishing a risk-based framework that classifies AI systems by their potential impact and assigns obligations in proportion. Chile is advancing a similar approach structured around four risk categories. In Panama, several legislative proposals are under discussion before the National Assembly, although none has yet been enacted.

Interpreting this regulatory gap as a zone free from responsibility would be a mistake. Whenever an organization operates systems capable of making decisions that produce legal effects on third parties, exposure exists regardless of whether a specific AI regulation has been enacted. What does not yet exist is a clearly defined AI-specific standard of care that organizations may rely upon in their defense. Building that framework internally—before legislators impose it—is the central premise of this article.

An effective AI governance framework rests on four essential pillars. The first is a systems inventory: a centralized register that identifies each AI tool in use, the decisions it automates, the individuals or groups it affects, and the level of risk it poses. The second is an internal policy. In this context, a policy is not merely a statement of principles; it is an operational instrument that assigns specific responsibilities, defines authorized uses, establishes which decisions require human validation, and sets out the consequences of non-compliance.

Today, sectors such as financial services, insurance, and talent acquisition face the greatest exposure, precisely because their AI systems can directly affect individual rights and interests.

The third pillar is accountability. Effective oversight cannot be diluted across departments or delegated indefinitely down the organizational hierarchy. A governing body that approves the deployment of an AI system without first reviewing the controls applicable to that system may ultimately bear responsibility for the consequences of its operation. Every level of the organization must have—not merely in form, but in practice the information, authority, and resources necessary to fulfill its role.

The fourth pillar is meaningful human oversight. Simply assigning a responsible individual to a process is not enough. Oversight is only genuine when that individual simultaneously has sufficient information to understand the system's decision, adequate time to review it, and the authority to modify or reverse the outcome. If any of these elements is missing, the control is only apparent.

To these considerations must be added the issue of explainability, which has moved beyond academic debate to become a practical business requirement. When an AI model generates a decision affecting an individual and the organization cannot explain the basis for that outcome, the challenge is no longer technical—it becomes an inability to demonstrate that the system operated according to legitimate and defensible criteria.
Algorithmic opacity is not a mitigating factor. In practice, it functions as a legal risk attributable to those who deploy and operate the system. Organizations should be able to explain, for every material decision, which variables determined the outcome, the relative weight of those variables, and under what circumstances a different result would have been produced. In several jurisdictions, the right to challenge decisions based exclusively on automated processing is already legally enforceable.

Finally, it is worth considering what formal regulation will ultimately bring when it arrives: a minimum standard and a compliance timeline. Organizations that have already implemented robust governance frameworks by that point will be in a position to demonstrate that their actions were driven by sound management and risk oversight rather than by regulatory pressure.

That distinction carries evidentiary value in the event of disputes and reputational value before clients, investors, , and business partners alike.
The starting point is not particularly complex: understanding which AI systems the organization operates, identifying who is accountable for each of them, and documenting that reality in a verifiable manner. In Panama and throughout Latin America, the opportunity to act proactively remains open. The question is no longer whether regulation will arrive, but rather in what position organizations will find themselves when it does.

Miguel Sáenz de Pipaon
International Attorney - Visiting