

CYBERSECURITY LAW & STRATEGY

AI Transparency: Clear Explanations Matter More Than Disclosure

As AI becomes embedded in everyday business and legal operations, organizations are confronting a new expectation: simply disclosing AI use is no longer enough. A critical shift is taking place in the legal industry: transparency is no longer just about disclosure; it's about comprehension.

BY CHRISTOPHER WALL

APRIL 30, 2026

As AI becomes embedded in everyday business and legal operations, organizations are confronting a new expectation: simply disclosing AI use is no longer enough. I had the opportunity to moderate a recent HaystackID® webcast, *"Meaningful Transparency in AI: What Privacy Laws Actually Require."* Our guests — Aleida Gonzalez, Ken Suh and Patrick Zeller — emphasized a critical shift taking place in the legal industry: transparency is no longer just about disclosure; it's about comprehension.


This shift reflects growing pressure from all sides — judiciary, regulators, legal counsel and users. People want to know not just *that* AI is being used, but how it works, what data it's using, and what questions or decisions it's trying to support. Vague, generic statements no longer meet that standard.

The Risk of Vague Messaging

Many organizations rely on broad phrases like "we use AI for internal business purposes." As Ken Suh pointed out, that language immediately raises a critical question: *What does that actually mean?*

Without clarification, such statements fail to provide context and may expose companies to legal risk. Aleida Gonzalez emphasized that disclosures are often too vague or overly technical, leaving people confused rather than informed.

Inconsistency can compound the problem. When messaging differs across privacy notices, websites



and internal documents, it signals deeper governance issues. Regulators increasingly look for alignment between what organizations say and what they actually do, and discrepancies can trigger scrutiny.

Transparency Starts Internally

A key takeaway from the discussion is that meaningful transparency begins inside an organization. Before companies can disclose and accurately explain how they are using AI for external audiences, they must understand it internally.

Patrick Zeller noted that many organizations lack a clear inventory of the AI applications used across the organization and how they are actually used. Without that visibility, accurate disclosure is nearly impossible. Foundational steps must be taken, including mapping data flows, documenting AI use cases, and maintaining governance records.

This internal governance is essential because regulators often compare public disclosures against actual practices. If behavior does not match policy, the risk of errors and associated fines increases significantly.

Transparency Should be Simple

Across global frameworks, a consistent set of expectations is emerging. Organizations must clearly explain whether AI is used, what it does, what data the model relies on and how privacy has been protected. Disclosures must be understandable to a reasonable person, not just legally compliant.

In the U.S., enforcement tends to rely on consumer protection laws. International regulations reinforce similar principles but often with greater regulations around user approvals and personal data protections. The direction is clear regardless – transparency must be practical, not theoretical.

Disclosures should be provided in clear, plain language and prioritize user understanding over legal precision alone. For example, instead of saying AI is used “to improve services,” organizations should explain how it is being applied. Whether AI is being used to screen job applications, support customer service or monitor employee performance, the details directly affect individuals and shape their expectations.

The panel emphasized that disclosures do not need to exist in a single place. Layered approaches, such as brief summaries supported by more detailed explanations, can make complex information more accessible.

Learning from Mistakes

Recent rulings and enforcement actions are demonstrating the consequences of poor transparency, particularly in cases involving biometric data and facial recognition. Few illustrate the stakes of AI transparency failures more starkly than *In Re: Clearview AI, Inc. Consumer Privacy Litigation, No. 1:21-cv-00135 (N.D. Ill.)*. The company built one of the world’s most powerful facial recognition databases by scraping billions of publicly available images – photos pulled from social media, news sites, and across the open web – without notifying or obtaining consent from the people in them. The legal reckoning that followed was swift and costly: Clearview ultimately agreed to a \$51.75 million settlement tied to violations of

biometric privacy laws, including Illinois’ Biometric Information Privacy Act (BIPA). Clearview’s case is just one example. A lawsuit against JCPenney has raised parallel concerns, alleging that its AI-powered virtual try-on tools quietly captured and analyzed customers’ facial geometry without proper disclosure or consent – a far less dramatic use case, but rooted in the same underlying failure.

These examples highlight a core principle: if data is collected without proper notice or consent, its use in AI systems remains problematic. Gonzalez compared this to the “fruit of the poisonous tree” – if the initial data collection is flawed, everything built on it is at risk.

A New Standard for AI Trust Is Required

Meaningful transparency is about trust. As AI plays a larger role in decision-making, users expect clear, honest communication. Organizations that rely on vague or inconsistent disclosures risk regulatory action and reputational harm. A new standard is needed for the legal process. To improve AI transparency, organizations must:

- Build and maintain an AI inventory and data flow maps
- Align disclosures with actual practices
- Ensure proper notice and consent are built into the process *before* using the AI
- Use clear, audience-focused language
- Test disclosures with non-experts
- Provide information at key interaction points

The takeaway is straightforward: transparency is no longer about saying you use AI – it’s about explaining it in ways people can understand.

We invite you to listen to the full webcast here.

Christopher Wall is the Data Protection Officer and Special Counsel of Global Privacy and Forensics at Haystack ID. In his role as Special Counsel for Global Privacy and Forensics, Chris helps HaystackID clients navigate the cross-border privacy and data protection landscape and advises clients on technical privacy and data protection issues associated with cyber investigations, data analytics, and discovery. Prior to joining HaystackID, Chris worked at Ernst & Young, where he led cross-border cybersecurity, forensic, structured data, and traditional discovery investigations.