LAW.COM

Expert Opinion **Cybersecurity**

From Data Security to Decision Security: AI's New Legal Risk

The article explains how AI is transforming cyber incident investigations, shifting focus from just data security to "Decision Security"—ensuring AI-driven conclusions are accurate, justifiable, and legally defensible. AI helps handle massive and complex data quickly, but this increases legal risks because organizations must justify AI-based findings used for compliance, notifications, and investigations. To manage these risks, companies should implement transparent, validated, and accountable AI processes—such as explainable models, human oversight, and clear responsibility.

March 02, 2026 at 09:09 AM By **Jessica Talar** **Mercedes Kelley Tunstall** **Michael Sarlo** & **Anya Korolyov**

Cybersecurity

Managing legal and regulatory risk following a cybersecurity incident goes beyond the obvious. While it is tempting to focus on whether the data is appropriately secured or compromised, today's cyber incident questions are more complex and increasingly focus on the cyber investigation itself.

AI is being inserted into most cyber incident investigations—not only investigating the intrusion but the interpretation of its findings. Organizations are turning to AI to assist in determining what happened, what data was affected and what obligations may follow. The risk surface has shifted. The question is no longer just whether the data was secure; it is whether the AI-assisted determinations are reasonable, accurate and defensible. This evolution marks the emergence of

Decision Security: the discipline of reaching AI-assisted investigative conclusions that can withstand regulatory, judicial and commercial scrutiny long after systems are restored and headlines fade.

This move toward AI-enabled analysis has not risen from technological enthusiasm; it has emerged from necessity. Modern incident response operates within three structural pressures that stress traditional investigated methods following an incident:

1. Scale. Organizations are not only increasingly storing more information; they are starting to generate entirely new classes of data—collaboration platforms, ephemeral messaging, machine telemetry, and cloud-native artifacts—that did not exist when breach-notification frameworks were drafted. Historical archives are being digitized while real-time data multiplies. Investigations can now routinely span billions of data points across heterogeneous systems.

2. Speed. Regulatory clocks can begin running quickly. Notification statutes, contractual obligations and markets often demand answers promptly and typically well before the full scope of the incident is understood. This pressure mounts as victim organizations must make representations about incidents before being confirmed.

3. Knowledge. Enterprises may not always maintain accurate and comprehensive data maps to provide a full inventory of where sensitive data resides. Incident response can become an exercise in discovery under pressure—simultaneously mapping the organization's data environment and assessing compromise.

These constraints mean traditional, manual review and linear analysis can no longer scale effectively. Legacy methods that relied on lexical search often flagged documents with pre-defined terms that created vast false-positive populations requiring human review and validation. AI models can now evaluate semantic context, distinguishing between superficial references and substantive regulated content.

They can interpret meaning rather than simply detect keywords, which can be a significant benefit for cyber investigations (e.g., differentiating between a policy document describing medical benefits and a record documenting an actual clinical interaction). The distinction can dramatically reduce review burdens. The more AI interprets context, the more its reasoning must be auditable and subject to strict oversight by qualified professionals to ensure accuracy and completeness.

AI-driven data analytics is promising to compress time, contain costs, improve accuracy, and provide insights to stakeholders quickly. But efficiency gains must be matched by methodological transparency and intentional direction as part of the investigation. Acceleration has introduced a

new kind of exposure: if AI can help reach conclusions faster, organizations must be prepared to rigorously defend them.

Legal Consequences of AI Data Tooling in Incident Response

AI in incident response is often described as a tool for efficiency. That characterization understates its importance. In practice, AI systems reviewing data and leveraged by professionals in incident response are performing tasks that carry legal consequence:

• Identifying information that is subject to specific regulatory frameworks

• Assisting in collecting data to provide required notice of the incident

• Prioritizing evidence streams that shape investigative narratives

• Distinguishing between benign artifacts and information that requires action

These are not clerical activities; they are key inputs to inform interpretive judgments that shape compliance with statutory obligations, litigation exposure management, and public communications strategies. Organizations leveraging AI tools to review data during incident response are introducing a computational intermediary into the fact-finding process. Unlike a human analyst, that intermediary operates through probabilistic inference rather than explicit reasoning. This is where Decision Security becomes essential.

Cyber investigations are becoming subject to a new form of legal scrutiny. How will an organization's use of AI during a data investigation result in accurate and prompt notice to consumers, regulators and other stakeholders. Working with incomplete information and elusive facts, data reviews in cyber investigations tend to focus on volume. When in doubt, cyber investigators tend to review more. That approach, while expensive, can help demonstrate that investigative conclusions are reasonable, comprehensive and defensible.

AI has the potential to disrupt that model by enabling targeted reviews, reducing millions of records to those that truly require analysis. This precision will be welcomed news to cyber investigators looking to move quickly without sacrificing accuracy, but it also might create new concerns. Organizations relying on AI-enabled analytical narrowing must be prepared to justify why some data was excluded from the review and other data was emphasized.

This shift represents a potential fundamental change in how investigations are executed and evaluated. The focus is no longer on the sheer volume of material reviewed but instead on accuracy of finding the right information from the start and being able to reconstruct and justify the conclusions. Regulators and courts will likely probe on how AI tools are applied, including the AI methodologies used, how the models are trained and what validation steps are taken to ensure

that sensitive data remains protected and that exclusions or prioritizations are reasonable. They also might question whether alternative methods might have produced materially different outcomes.

Applying Decision Security in Practice

Many organizations are approaching AI adoption as a tooling decision—acquiring platforms without redesigning investigative governance. Decision Security demands that AI outputs be embedded within a structured validation framework and include three critical elements:

1. Explainable Analytical Pathways—Document how results were derived, not just results themselves. This includes model inputs, classification logic, sampling validation and escalation thresholds. The goal is reproducibility and accuracy, not technical disclosure.

2. Human-in-the-Loop Determinations—AI may triage and prioritize, but, ethically, attorneys and the experts who assist them are still responsible for those conclusions. Human validation is not a ceremonial step; it is what ensures the analysis is reasonable.

3. Context-Bound Learning—Adaptive models will improve through exposure to prior incidents. Without constraints, the risk of importing assumptions from unrelated matters increases. Each investigation must start anew and define the permissible scope of learning to avoid analytical drift.

Organizations must treat AI assistance less like software and more like a regulated process.

Elevating Professional Judgment

Public discourse often frames AI as replacing human expertise. Incident response demonstrates the opposite. AI removes the mechanical data burdens of sorting, clustering and pattern recognition while elevating the importance of professional judgment.

Practitioners must validate AI insights, translate probabilistic outputs into legal standards, and be prepared to defend the methodologies underlying those outputs months and years later. The professional role is shifting from a reviewer to an interpreter. Responsibility is not diminished; it is becoming more explicit and accountable.

We must move beyond viewing AI as merely a tacked-on tool for accelerating analysis to actually embedding it within a governance framework. This requires asking foundational questions at the time AI is integrated into investigative workflows rather than after they have influenced decisions, including:

1. Explain AI-assisted results to a regulator, court or expert in clear, documented terms. If an explanation relies on technical abstraction, vendor assurances or opaque models rather than traceable reasoning and preserved methodology, the process is open to attack.

2. Deliver reproduceable results if challenged. Reproducibility vs. model sophistication will serve as a true benchmark of reliability as decisions must withstand scrutiny months or years after initial analysis.

3. Explicitly assign responsibility for validating AI-generated insights. Accountability cannot emerge retroactively in response to scrutiny; it must be structured into the workflow through defined review roles, validation protocols and audit mechanisms established before deployment.

These considerations shift AI adoption from a narrow technology initiative into an enduring governance discipline, one that treats analytical outcomes as decisions requiring stewardship rather than outputs to be passively accepted.

Securing the AI Results

From the threats organizations face to the efforts needed to protect networks and data, the cybersecurity landscape is ever shifting. Considerable resources are being invested in areas like perimeter, endpoint and infrastructure security to prevent unauthorized intrusions. Organizations are also prioritizing data through encryption, governance frameworks, and regulatory compliance designed to safeguard it at rest and in transit. The challenge is no longer only preventing access or preserving confidentiality but ensuring that the interpretations drawn from interconnected data ecosystems are reliable, explainable and capable of withstanding legal and regulatory scrutiny.

This Decision Security evolution reflects that modern risk is tied not just to what data is exposed but to how organizations understand, act on and defend AI-assisted decisions. Organizations that recognize this shift will design investigative architectures where AI accelerates insights without obscuring accountability.

Jessica L. Talar *is a senior IP staff attorney in the Global Litigation Group at Cadwalader, Wickersham & Taft.* **Mercedes Kelley Tunstall** *is a partner in the firm's Financial Services group.* **Michael Sarlo** *is chief innovation officer and president of Global Investigations & Cyber Incident Response Services at HaystackID.* **Anya Korolyov** *is executive vice president of Cyber Incident Response and LDI Strategy at HaystackID.*