LS Lowenstein Sandler

March 18, 2026

AI Governance Is Not Policy. It Is Infrastructure.

By Amy S. Mushahwar and Tricia Y. Wagner CIPP/US, CISSP, CISA

Field Notes: From the Legalweek Vendor Floor and on Our Way to the RSA Conference. Productivity Is the Killer App; True Governance Legal Architecture Is Catching Up

If a regulator or litigant asked your organization tomorrow to explain how a specific AI system produced a decision, could you?

Most organizations would struggle to answer. Not because they lack policies. Because they lack reliable evidence.

AI governance that exists only in documents cannot scale. Real governance lives in systems, pipelines, identity controls, telemetry, and monitoring infrastructure that produce verifiable proof that governance is working.

Right now, most organizations have the first. Very few can prove they have the second. That gap matters. The evidence your systems generate will determine whether your governance holds or becomes a liability. Every organization has an AI “Pilot’s Manual.” Almost none have a “Black Box.” If your AI makes a billion-dollar error tomorrow, can you reconstruct the flight path, or are you just going to show the regulator your Ethics Charter? Policies promise governance. Pipelines prove it.

Like cybersecurity before it, AI governance is forcing a new wave of legal learning. Understanding AI risk now requires understanding the systems that produce it.

In practice, this gap becomes visible very quickly when you ask a small set of operational questions, the same questions a regulator, board, or incident response team will eventually ask. We outline five of those questions at the end of this document, and most organizations cannot answer them today.

Move From Governance Policy to Engineering Pipelines

As organizations deploy increasingly complex AI systems, governance is shifting from a policy exercise to an operational discipline. That shift is shaping how we approach AI governance conversations, and it is the reason we have spent recent months doing something lawyers rarely do deliberately: taking pre-conference demos, walking vendor floors, and asking infrastructure questions.

Both authors have spent more than two decades working in privacy and cybersecurity, including the period between roughly 2005 and 2015, when modern data protection and incident response frameworks were first operationalized. The law firms that provided the most effective advice during that period (and beyond) were not those writing the longest policies but those that understood the technical evidence produced by systems: trouble tickets, epics, logs, telemetry, and identity data. The same shift from policy to operational proof is now emerging in AI governance but in a more expedited time frame.

What follows is not a finished framework of technical operationalization mapped to AI compliance domains. It is working research from practitioners who are still learning. We attended Legalweek to observe the policy and buyer behavior side. We are heading to the RSA Conference to examine the infrastructure side. This document reflects where our thinking stands between those two data points, with the explicit expectation that it will evolve (in the next few weeks and years to come).

A Pattern Repeating Across Every Industry

Walk any major technology conference floor right now (e.g., legal, health care, financial services, manufacturing) and you find a version of the same exhibit hall. The dominant categories are productivity tools: workflow automation, AI-assisted research and documentation, intelligent analytics, decision-support platforms. These are valuable products that make organizations faster and more efficient.

What is largely absent, across nearly every sector, is the technical infrastructure that would make AI governance real: model inventory systems, runtime monitoring, data lineage, identity governance for AI agents, AI attack surface visibility.

Legalweek is instructive precisely because the legal profession takes governance seriously. The CLEs are thoughtful. Sessions on responsible AI deployment are substantive and well attended. The governance vocabulary is present and growing. It also happens to be an industry that we've each individually lived in for 20 years and felt comfortable fully evaluating.

At Legalweek, you get excited about the robust governance CLEs, then you walk the exhibit hall. The dominant categories are:

- Contract review and intelligence platforms
- AI-powered legal research
- E-discovery and litigation analytics
- Document automation
- Workflow tools for legal operations

The same split appears at HIMSS in health care: strong clinical AI adoption, governance infrastructure still catching up. At financial services innovation conferences: sophisticated AI deployment in credit, fraud, and trading, and model risk management frameworks still being written after the fact. At manufacturing and supply chain events: AI-driven demand forecasting and quality control, with minimal visibility into what those systems are doing to operational data.

The pattern is not a failure of any profession. It is a structural reality of how technology adoption works. Productivity is visible and measurable (and admittedly, we both love Harvey and other legal productivity tools). This is not an indictment of the future; we're excited for it and can't wait to see the next evolution of the legal profession. But as incident response and data attorneys, we know that governance is invisible until it fails and becomes very visible.

On the Legalweek floor: At one booth draped in "Responsible AI" branding, one of us asked how a platform enforced the data-handling policies it touted. The representative pulled up a screen showing a read-only PDF of an AI ethics charter and a static log of user logins — a digital filing cabinet, not a governance system. In some additional questions about where the enforcement controls lived, there was a long pause before the response that such functionality would be an enterprise environment configuration.

The CLEs promise governance. The floor delivers productivity. This is not a legal profession problem. It is an industrywide condition, because this is the way that technology is deployed. Productivity is the killer app; governance takes time.

Consider Legal Engineering Infrastructure to Prove Governance

Legal and compliance teams have spent years developing governance frameworks for emerging technologies. Those frameworks are essential. They establish principles, accountability structures, and risk management expectations.

But governance frameworks alone do not control systems.

AI systems are built through pipelines that include data ingestion, model training, model deployment, integration with applications, and continuous monitoring. Each stage of that pipeline creates potential governance risks. Policies describe how those risks should be managed. Infrastructure determines whether they are. The difference between the two is evidence.

Regulators, boards, and auditors are increasingly asking not simply whether governance frameworks exist, but whether organizations can demonstrate how AI systems are built, monitored, and controlled. AI governance is gradually becoming an evidentiary discipline.

Policies promise governance. Pipelines prove it. Understanding how those pipelines generate evidence is now core legal work for defensibility.

A Working Research Framework: The AI Governance Infrastructure Stack

As part of this fieldwork, we have begun mapping the operational layers that may ultimately support AI governance across the security, identity, data, and AI infrastructure ecosystem — the ecosystem you find at RSA, not at Legalweek.

Both of us want to be direct about the status of what follows. This taxonomy is not a finished architecture. The ecosystem is evolving quickly, many capabilities are still emerging, and we have not completed a full evaluation of the vendors under active review. We are presenting this as version 1.0 of a working framework, to share where our thinking currently stands and to invite scrutiny from practitioners, security engineers, and compliance professionals working through the same questions.

We expect to revise this materially after RSA and in the years to come as the industry matures. Tools are necessary to generate evidence, but tools alone are not enough; they still require governance, human oversight, and ongoing monitoring.

A note on cost and scale: The eight-domain framework below represents the full governance infrastructure landscape on the horizon, not at the starting line. For large enterprises and regulated institutions, building toward it comprehensively is the right goal. For emerging organizations, mid-market firms, and legal departments without dedicated security engineering, the full stack is prohibitively expensive if approached all at once. The question is not "How do we buy all of this?" It is "Where do we start, given what we have?"

A practical tiered approach:

Tier 1 — Any organization: Before purchasing anything, conduct a manual AI inventory: what tools are in use, by whom, and what data do they touch. This costs time, not budget, and is the prerequisite for everything else.
Tier 2 — Growing organizations: Many data security posture management (DSPM) and identity governance capabilities exist in tools already licensed — Microsoft Purview, Google Workspace DLP, Okta, Azure AD. Audit existing infrastructure before purchasing dedicated AI governance platforms.
Tier 3 — Organizations with compliance exposure: Purpose-built tooling becomes justifiable here. Open-source

options probe for security and safety failures, and those for governance — Garak, Giskard, MLflow — provide meaningful capability without enterprise licensing costs.
Tier 4 — Regulated industries and enterprises: For organizations subject to Securities and Exchange Commission (SEC), Office of the Comptroller of the Currency (OCC), HIPAA, or EU AI Act requirements (in addition to state law requirements, where applicable), the full eight-domain infrastructure is the standard you can be measured against. The question is sequencing and vendor selection, like how security stacks evolved after years of implementation.

1. Governance and Risk Orchestration

This is the policy coordination layer: platforms that organize governance processes, risk documentation, and compliance workflows. Among the tools we are reviewing in this category are Credo AI, IBM, Collibra, and Monitaur each with somewhat different cost structures and integration assumptions.

The core limitation we are evaluating across this category is consistent: These tools primarily document governance decisions rather than enforce technical controls. In highly regulated sectors, regulators want evidence as well as a well-organized memo. When the SEC, OCC, or Federal Reserve asks for proof of model integrity or nondiscrimination, they want outputs of a functioning risk management framework. Whether any platform in this category bridges the documentation-to-enforcement gap in practice is part of what we will evaluate at RSA and after. Some of the vendors named here and in the other seven domains are emerging companies that may not have the funding to exhibit at RSA, but they will grow with time.

**Key question: How are the governance decisions recorded in these platforms connected to the systems actually running AI?*

2. AI Discovery and Security Posture Management

One of the most persistent governance challenges organizations face is visibility into where AI is being used. Models may be embedded in applications, integrated through application programming interfaces (APIs), or introduced via third-party services, frequently without any centralized inventory or oversight. Shadow AI — undocumented or unknown AI systems and toolsets, is not a theoretical concern. It directly parallels the shadow IT and infrastructure risks that IT and security teams have spent decades trying to control.

Among the platforms we are actively reviewing are Pillar Security, Cranium, Wiz AI-SPM, Orca Security, and Noma Security. Our evaluation questions center on how each handles discovery across hybrid environments and what deployment requires, because cost, complexity, depth, and management approach vary significantly from the marketing materials. For organizations not yet ready for enterprise tooling, a structured manual inventory (a spreadsheet with defined fields for system, owner, data touched, and deployment environment) is the zero-cost starting point that makes everything else possible. As with early security, governance does indeed usually start with a spreadsheet!

**Key question: Can your organization identify every AI system touching your data right now?*

3. Agent Orchestration and Workflow Control

AI agents introduce a fundamentally different governance challenge. These systems do not simply answer questions; they take action. They call APIs, retrieve internal data, initiate transactions, and interact with external systems, often autonomously and at speed. As a result, IT leaders are understandably cautious about connecting these tools directly to production environments.

This is the domain where our research is most actively in progress, and where our thinking is least settled. The core governance question is straightforward but unresolved: What actions are agents authorized to take, under what credentials, and with what audit trail? Most platforms have not yet answered this cleanly.

We are bringing specific evaluation criteria to RSA focused on authorization controls and action logging, and we expect this portion of the framework to evolve substantially as a result. At a minimum, any organization deploying AI agents today should maintain a written authorization register identifying each agent, its permitted actions, and the accountable human owner before making any platform decision.

We are currently evaluating tools that address agent orchestration and workflow control, including LangGraph, Microsoft Semantic Kernel, Microsoft AutoGen, Amazon Bedrock Agents, Apache Airflow, Prefect, and Dagster, as well as low-code agent workflow tools such as Microsoft Copilot Studio.

**Key question: What actions are your AI agents authorized to take, and how are those actions recorded?*

4. DSPM for AI

Data governance becomes more urgent once AI systems are introduced. DSPM platforms are designed to identify, classify, and control sensitive data as it moves across environments and into AI systems. We are actively reviewing Netskope, Nightfall AI, Cyberhaven, and Bedrock Data, with particular attention to how each handles AI-specific data flows, distinct from traditional data loss prevention (DLP) use cases, because not all DSPM platforms have kept pace with real-world AI deployment patterns.

Many organizations already have partial capability embedded in tools they own today, including Microsoft Purview, Google Workspace DLP, and existing DLP configurations within their broader security stack. Before investing in a dedicated AI DSPM platform, organizations should first audit what their existing infrastructure already provides. In practice, the configuration gap is often smaller than the budget conversation suggests.

**Key question: What sensitive data has already reached your models, and did you know before today?*

5. Data Lineage and Pipeline Visibility

Regulators are increasingly asking organizations to demonstrate where training data originated, how it was transformed, and how model outputs flow through downstream systems. Data lineage tools provide the evidentiary record for these questions. We are reviewing platforms such as Collibra and Monte Carlo, among others, with a specific focus on how their lineage capabilities extend into AI pipeline visibility, a newer use case that not all lineage platforms support natively.

Without some level of pipeline visibility, governance frameworks describe processes that cannot be independently verified. Open-source tools, including Apache Atlas and OpenLineage, provide meaningful capability for organizations not yet operating purpose-built platforms. In many cases, data provenance can begin as structured documentation before it is automated into telemetry and continuous monitoring.

**Key question: Can you trace, step-by-step, how a specific model output was produced?*

6. Identity and Access Governance for AI Systems

AI systems interact with multiple types of identities, including human users, service accounts, APIs, automated workflows, and increasingly autonomous agents. Many governance questions ultimately reduce come down to identity questions: Who initiated a model action, which service account accessed training data, and what system authorized a model to access a particular dataset?

We are reviewing ConductorOne and other platforms in this space, with specific focus on how each addresses nonhuman identity life cycle management. Most identity governance solutions were originally designed for human users and are now being adapted to AI agents in real time. For organizations earlier in maturity, existing identity and access management platforms — such as Okta, Azure AD, or similar — can often be extended to cover AI system identities if deliberately configured, though meaningful gaps may remain.

**Key question: What systems can your AI agents access, and who last reviewed those authorizations?*

7. Runtime Protection and Behavioral Monitoring

Governance does not end at deployment. AI systems require continuous monitoring for prompt injection attempts, data leakage, model drift, and anomalous behavior. We are actively reviewing Lakera Guard, Lasso Security, Protect AI, Arthur AI, and Robust Intelligence (acquired by Cisco and integrated within Cisco AI Defense and Cisco Foundation AI) with one specific evaluation question driving the analysis: Does this platform detect risk or does it enforce controls?

That distinction is the most consequential one in this category and is not always clear from marketing materials. Dashboards describe problems; controls prevent them. For organizations not ready for enterprise runtime platforms, open-source options including Garak and Giskard provide real capability at no licensing cost, and drift telemetry can be instrumented directly in MLflow or similar tools many organizations already use.

**Key question: If a deployed AI system began behaving unexpectedly tomorrow, how quickly would you know, and what, if anything, would automatically stop it?*

8. AI Supply Chain and Model Integrity

When AI vendors update models, system behavior can change instantly across every application built on those models without notification to the organizations depending on them. Supply chain governance for AI addresses this: monitoring model updates, verifying third-party dependencies, detecting behavioral drift introduced by upstream changes.

This is the least mature category we are evaluating and the one where we have the most open questions going into RSA. Purpose-built supply chain governance tooling for AI is nascent (vendors are early-stage, capabilities vary widely, and the cost-to-risk calculus is not yet well established). At minimum, organizations should implement a vendor notification protocol — a defined process for tracking model versioning in third-party platforms and testing for behavioral changes after updates — before committing to any platform in this space.

**Key question: If your AI vendor pushed a model update tonight, how would you know whether system behavior changed?*

A note on application-embedded governance: *The eight-domain framework above describes enterprise infrastructure. But governance controls can also live at the application layer. Platforms like Harvey and Legora are building audit functionality, access controls, and data handling governance directly into legal AI tools (compensating, in part, for the enterprise governance infrastructure many organizations have not yet deployed). For legal departments evaluating AI tools, vendor selection is itself a governance decision: Choosing platforms with embedded compliance controls reduces the infrastructure burden on your security team.*

What We Are Taking to RSA

The Legalweek floor illuminated the policy and buyer behavior side of this problem. RSA is where we expect to examine the other half: how the security infrastructure ecosystem is approaching the operational requirements that AI governance actually demands.

We are going there with specific questions we do not yet have good answers to:

- Which platforms in the runtime monitoring category enforce controls versus merely surfacing signals, and what does enforcement cost to deploy at scale?
- How are agent orchestration vendors approaching authorization controls and audit trails, and which are embedding governance?
- What does the AI supply chain tooling landscape look like in practice? Is it mature enough to recommend to clients today?
- Where does the governance stack become affordable for mid-market organizations, and what do those organizations have to give up (or add capacity to) to get there?

We expect the framework above to change now and in the years to come. Some domains will consolidate. Some vendors will differentiate themselves in ways their marketing does not currently convey. Some categories may prove less developed than they appear, and we will recalibrate our research.

The Evidence Architecture

Whatever the final architecture looks like, the evidentiary requirement driving it will not change. These eight domains exist to produce one thing: proof that governance is working.

The artifacts that answer regulatory and board questions are not policy documents. They are audit logs, pipeline telemetry, model monitoring records, access control documentation, and supply chain verification records.

At any organizational tier, the goal is movement toward these five artifacts, through documentation first, then automation:

- Automated Model Inventory: A real-time, queryable list of every internal and third-party model, version number, and nested AI dependency. Start with a structured spreadsheet if purpose-built tooling is not yet in scope.
- Quantitative Validation Logs: Independent test results (i.e., backtesting, stress testing, proving the model performs within predefined risk parameters).
- Data Provenance Records: Documentation showing where training data originated, how it was transformed, and the legal basis for its use.
- Drift and Performance Telemetry: Continuous monitoring logs that flag when a model's output distribution shifts from its validated baseline.
- Machine Identity Audit Trails: A log of every action taken by an AI agent, tied to a specific authorized credential and a named human supervisor.

What This Means for Lawyers — and Anyone Advising on AI

The legal profession is increasingly being asked to advise on AI systems that most of us were not trained to evaluate technically. That is not a failure of effort or intelligence. It is a structural gap: a translation tax paid every time governance frameworks are written by professionals who cannot see the systems those frameworks are meant to govern. We've written extensively about legal-to-tech translation taxes in other contexts outside AI.

The same translation tax is being paid in hospital systems where clinical AI is outpacing compliance frameworks, in financial institutions where model deployment is outrunning model risk management, and in manufacturing environments where operational AI is being treated as just another enterprise software purchase.

This is why we are doing this fieldwork. Not to become a vendor evaluator, but because practitioners who understand the infrastructure and can ask informed questions about model inventory, runtime monitoring, identity controls, and AI attack surfaces will be the advisors who can help organizations govern AI systems at scale.

At RSA, the infrastructure is being built. At Legalweek, and at every sector conference where governance conversations are happening without the engineers in the room, accountability is being defined and debated. Those conversations need to find each other faster than they currently do. We are finally finding one another in general privacy and security after more than 20 years of implementation. Let's hope that AI doesn't take that long.

The organizations that navigate this well will not necessarily be the ones with the largest governance budgets. They will be the ones that understand where they are in the maturity curve, build deliberately toward verifiability, and learn to operate in both rooms at once.

That is what legal engineering actually means. And we are still learning what it requires.

Five Questions To Bring to Your CISO This Week

1. Can you show me our current AI model inventory — every system, every deployment environment, every vendor integration?
2. What does our runtime monitoring cover, and what does it not cover?
3. If a regulator asked us to trace how a specific AI output was produced, what documentation exists to support that?
4. What systems can our AI agents access, and who last reviewed those authorizations?
5. If one of our AI vendors pushed a model update tonight, how would we know if system behavior changed?

If any of those questions produces a long silence, you now know where to start. And if the silence is followed by "we have a policy for that," ask to see the pipeline.

*This is a living document. The authors will update the framework following the RSA Conference and welcome thoughts from practitioners, security engineers, AI platforms, and compliance professionals working through the same questions. As an industry, we will all work through this together and benchmark off of each other.

Contacts

Please contact the listed attorneys for further information on the matters discussed herein.

AMY S. MUSHAHWAR
Partner
Chair, Data Privacy, Security, Safety & Risk Management
T: 202.753.3825 / 703.283.3515
amushahwar@lowenstein.com

TRICIA Y. WAGNER CIPP/US, CISSP, CISA
Counsel
T: 202.753.3658 / 916.201.7657
twagner@lowenstein.com