Lowenstein Sandler

Client Alert

Data Privacy, Security, Safety & Risk Management

March 23, 2026

The Trump AI National Policy Framework: What Preemption Actually Means for Your Governance Infrastructure

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

Key Takeaways:

The Trump administration's AI National Policy Framework does not prevent accountability for AI systems; it redirects it. Preempting state AI laws removes one pathway for governance, but enforcement through the FTC, sector-specific regulators, and civil litigation remains. The gap between policy documents and verifiable technical infrastructure is now the primary legal risk. Organizations that cannot produce audit logs, model inventories, and pipeline telemetry may face enforcement risk.

What the Framework Does–and Does Not Do

On March 20, 2026, the Trump administration released its National AI Legislative Framework, the culmination of the December 2025 executive order (EO) directing a unified national approach to artificial intelligence (AI) regulation. The framework is the most significant federal AI governance action to date. It notably is also not a governance framework in any technical or operational sense.

The framework's primary mechanism is preemption. It calls on Congress to establish a national standard for AI development to displace the growing body of state-level AI regulation with what the administration describes as a "minimally burdensome" federal floor. The Attorney General's AI Litigation Task Force, established in January 2026, stands ready to challenge state laws that conflict with these priorities. Examples of state laws that could face potential legal challenges from the AI Litigation Task Force or express preemption under the framework include the Colorado AI Act, which intends to regulate the development and deployment of "high-risk" AI systems to prevent algorithmic discrimination in consequential decisions related to hiring, medicine, and other areas, and California's Transparency in Frontier Artificial Intelligence Act (SB 53), which requires annual comprehensive risk assessments, safety frameworks, and state reporting for security incidents.

What the framework preserves is significant. The carve-outs from preemption are narrow but telling, such as state laws relating to fraud and consumer protection, child protection, zoning, and state use of AI. All of the carve-outs may enable actions through which AI systems can face accountability in the absence of direct regulation.

A Potential Misreading: Preemption Is Not Protection

Organizations interpreting the framework as a reduction in regulatory burden are misreading the potential enforcement landscape. The preemption carve-outs are still powerful when enforced against AI companies and those improperly using AI. The following are potential examples of viable enforcement mechanisms under the framework.

Federal Trade Commission: Section 5–Unfair and Deceptive Acts and Practices (UDAP)

The December 2025 EO directed the Federal Trade Commission (FTC) to issue a policy statement on when Section 5 of the FTC Act, the prohibition on unfair and deceptive acts or practices, applies to AI models. The legislative framework did not disturb this policy statement. If enforcement occurs, an AI system that cannot document how it reached a decision is a plausible candidate for a deception claim. An AI system that promises governance it cannot demonstrate is a plausible candidate for an unfairness claim. Audit logs, pipeline telemetry, and data provenance records are not compliance overhead in this environment; they are the primary defense against FTC exposure.

Sector-Specific Regulation: The Full Stack Remains

Preemption of general state AI laws would not disturb sector-specific federal frameworks. The Securities and Exchange Commission's (SEC) model risk expectations, the Office of the Comptroller of the Currency's (OCC) guidance on algorithmic decision-making, the Health Insurance Portability and Accountability Act's (HIPAA) requirements for systems processing protected health information, and the EU Artificial Intelligence Act (EU AI Act) for any organization operating in Europe remain fully operative. The Treasury Financial Services AI Risk Management Framework, which we analyzed in our earlier alert, specifically addresses machine unlearning obligations, model drift, and what we termed "algorithmic disgorgement" risk, which remains intact.

Civil Litigation: A Potential Opening

Perhaps most consequentially, the federal framework preempts private civil litigation. The framework explicitly seeks to prevent states from "penalizing AI developers for a third party's unlawful conduct involving their models," which reads as a liability shield for AI platforms. But that shield may leave open AI platform liability where the model did not function as designed. AI platforms should carefully construct the evidence stack to prove it was the third-party conduct, not the model itself, that caused harm. To that end, AI platforms should consider model inventories, runtime behavioral logs, authorization records, and supply chain verification. Organizations will have difficulty asserting a third-party conduct defense without evidence that it was the third party, not the model, that caused harm.

Child Safety

Child safety is another area explicitly preserved from preemption, and the framework specifically mentions the Take It Down Act remaining at a federal level. The framework acknowledges that Congress should require AI companies to implement features that reduce risks of sexual exploitation and harm to minors, and it affirms that existing laws, including those prohibiting child sexual abuse material, remain in force. What the framework does not provide is any enforceable technical standard for how those obligations are met. The resultant vagueness can result in more litigation. When a child safety incident occurs and there is no federal enforcement mechanism specifying what "reasonable safeguards" required, the evidentiary record of what an organization's AI systems did becomes the entire case. For help starting a child safety platform review, see our recent video.

The Enforcement Shift: From Regulation to Evidence

Our March 18 article argued that AI governance is not a policy exercise; it is an infrastructure discipline. The Trump framework sharpens that argument.

In a regulatory environment defined by direct AI-specific rules (i.e., EU AI Act and state algorithmic accountability laws) governance proof may be satisfied by demonstrating compliance with defined requirements. In a preemption environment defined by a federal floor with no technical specifications, organizations face a more difficult task: They

must produce affirmative evidence of what their systems did, how they were authorized, and whether controls were operating in response to enforcement or litigation without a predefined checklist to point to.

> *THE GOVERNING PRINCIPLE*
> *Policies promise governance. Pipelines prove it. In a deregulatory environment, the pipeline is the policy. Systems evidence and telemetry provide the verifiable artifact that harm did or did not occur.*

The five questions we posed in our March 18 piece have become five enforcement questions.

1. Can you produce your current AI model inventory, systems, deployment environments, and vendor integrations in response to an FTC civil investigative demand?
2. What does your runtime monitoring cover, and what behavioral anomalies would it have detected before a harm event?
3. If a regulator or plaintiff's counsel asked you to trace how a specific AI output was produced, what documentation exists, and how long would it take to compile?
4. What systems can your AI agents access, and who last reviewed those authorizations? Can you demonstrate that review in a log?
5. If your AI vendor pushed a model update that changed system behavior and a harm followed, what would your supply chain verification record show?

For each, ask to see evidence or to understand how the company could gain visibility.

The Eight-Domain Infrastructure Stack: Now a Defense Posture

The eight governance infrastructure domains we outlined in our March 18 article were framed as a maturity framework, a map from where organizations are to where they should be. The Trump framework reframes them as a defense posture. Each domain now could correspond to a specific enforcement or litigation exposure.

Domain	Enforcement Exposure Without This Infrastructure
Governance and Risk Orchestration	Without documented risk decisions connected to actual system controls, there is no defense to "We didn't know" in FTC or SEC enforcement.
AI Discovery and Security Posture	Shadow AI is not a theoretical risk. Undocumented systems are the first things enforcement counsel finds. Inability to inventory = inability to defend.
Agent Orchestration and Workflow Control	Agentic systems that take actions without authorization records expose organizations to both liability and the inability to assert a design defense.
Data Security Posture Management (DSPM) for AI	Sensitive data reaching models without documented controls create direct UDAP and sector-regulator exposure. "We didn't know it was there" is not a defense.
Data Lineage and Pipeline Visibility	Regulators increasingly require organizations to show where training data originated and the legal basis for its use. No lineage = no provenance defense.
Identity and Access Governance for AI	Every material AI enforcement action eventually becomes an identity question: Who authorized what and when? Machine identity trails are the evidentiary record.
Runtime Protection and Behavioral Monitoring	Drift and anomaly detection are the record of what a system was doing when harm occurred–and what controls would have stopped it.
AI Supply Chain and Model Integrity	The liability shield for developer conduct requires proof that vendor model updates were tracked and tested.

Practical Guidance: What Organizations Should Do Now

The framework creates urgency without creating new requirements. That combination produces a difficult compliance environment, one where organizations assume they are safe because no specific law applies, while enforcement exposure builds through existing channels. The following guidance applies across the organizational maturity tiers we described in our March 18 article.

Immediate Actions for All Organizations

- Conduct a manual AI inventory of tools in use, by whom, what data AI touches, and what environment the AI stack operates in. This is the prerequisite for everything else.
- Map your existing UDAP exposure. Review AI-assisted customer-facing decisions for documentation–what did the system produce, on what basis, and what would you show a regulator tomorrow?
- Document child safety measures and compliance with existing laws now. If any AI system your organization operates or deploys touches users under 18, document the safeguards, the monitoring, and the authorization structure before a safety incident creates the documentation request.
- Review sector-specific obligations. Identify which federal sector frameworks apply to your AI deployments and assess whether your governance documentation satisfies their evidentiary standards–not their policy requirements.

Near-Term Actions (Organizations With Compliance Exposure)

- Audit existing infrastructure before purchasing dedicated AI governance platforms. Microsoft Purview, Okta, Azure AD, and similar tools already licensed may provide meaningful DSPM and identity governance capability if deliberately configured for AI use cases.
- Implement agent authorization registers. For any AI agent taking action in production environments, maintain a written register of authorized actions, permitted data access, and named human supervisors. This is the minimum viable supply chain document.
- Establish vendor notification protocols. Define a process for tracking model versioning in third-party AI platforms and testing for behavioral changes after updates–before a supply chain incident creates the need for retroactive documentation.

For Regulated Industries and Enterprises

- The eight-domain infrastructure framework is the standard against which you will be measured in SEC, OCC, HIPAA, and EU AI Act enforcement contexts. The question is not whether to build toward it; it is sequencing and vendor selection.
- Post-RSA update: We will revise our vendor assessments in the runtime monitoring, agent orchestration, and supply chain categories following the RSA Conference. We welcome engagement with organizations working through the same infrastructure questions.

Conclusion: The Framework Confirms the Thesis

The Trump AI National Policy Framework is, by design, a governance-light document. Its preemption architecture may remove certain sets of compliance obligations. It does not remove accountability.

What the framework confirms is the core argument of our March 18 analysis: In the absence of a comprehensive regulatory architecture, accountability arrives through enforcement, litigation, and sector-specific frameworks. The organizations that navigate this environment successfully will be those that can answer five operational questions

on demand with logs, telemetry, audit trails, and pipeline documentation.

The framework was released March 20, 2026. The RSA Conference begins this week. We will be there evaluating the infrastructure side of this equation–which platforms in runtime monitoring enforce controls versus merely surface signals, how agent orchestration vendors are approaching authorization and audit trails, and where the governance stack becomes genuinely affordable for mid-market organizations. We will update this analysis following the conference.

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

ERICH J. KALETKA
Associate
T: 862.926.2792
ekaletka@lowenstein.com