Lowenstein Sandler

Client Alert

Data Privacy, Security, Safety & Risk Management

January 6, 2026

AI Platform Risk Assessments: Why 2026 Is the Year for Action

Preparing for High-Risk AI Uses, California Deadlines, and Federal Interest with Executable AI Governance Under NIST

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

What You Need To Know:

- **Act Now:** Boards and regulators expect visible progress on AI governance.
- **California Deadline:** Risk Assessments including certain high-risk AI usage due by December 31, 2027.
- **Use NIST AI RMF:** Adopt a defensible, sector-agnostic framework.
- **Start with Infrastructure Mapping:** Define ownership and accountability early.

Don't Wait for Regulatory Perfection on the AI Executive Order - California Risk Assessments Are Coming for Certain Use Cases

Leverage the NIST AI Risk Management Framework for Immediate, Defensible Action

Companies spent 2025 racing to adopt artificial intelligence (AI). As 2026 begins, the data shows that AI didn't just create new risks; it also acted as a high-speed searchlight, exposing the infrastructure gaps many organizations have carried since the late '90s. We aren't just closing a year; we are also closing an era of deferred maintenance. (See our alert Entering 2026 and the Age of AI, highlighting "the 1999 Problem" of information governance tech debt.)
Why Act Now? Many of you are facing risk now. Your board is asking about AI risk. Your engineers are deploying models faster than Legal can review them. Your vendor contracts don't address who owns training data. And regulators are watching. The recent executive order establishing a national AI policy framework signals (see our prior alert Executive Order Establishes National Policy Framework for Artificial Intelligence; Sets Up New Federal-State Flashpoints) that heightened regulatory and enforcement may heat up, even if a preemption battle ensues.

Stakeholders, regulators, and boards now expect visible, defensible action. Building a robust governance framework takes time, so organizations that begin now will be better positioned to meet future requirements. Notably, under California's new mandatory risk framework, AI risk assessment is a required component of enterprise risk assessment, with a compliance deadline of **December 31, 2027**. For more details, see our recent update on California's regulatory developments.

Mitigate Risk and Use the NIST AI RMF as Your Operating Spine. The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF) provides a sector-agnostic, defensible structure for AI governance and is quickly becoming the industry standard. It offers practical tools, including an implementation playbook and crosswalks to other governance frameworks. These tools enable organizations to align legal, risk, and engineering teams while maintaining traceability from policy to practice. Click the link for the playbook and the crosswalks.

Function	What to implement now	Evidence to retain
Govern	Charter an AI risk committee; define roles and responsibilities, thresholds, and escalation	Charter; RACI (Responsible, Accountable, Consulted, and Informed matrixes), meeting minutes, risk appetite statements
Map	Inventory AI use cases and systems architecture, data flows, stakeholders, and potential harms	Inventory, system ownership (individual and joint), data lineage diagrams, impact assessments, risk register prioritized, legal, operational, and engineering alignment
Measure	Define metrics for performance, robustness, bias, privacy, and security	Test plans (pre- and post-deployment), datasets, results, acceptance criteria, sign-offs
Manage	Monitor in-production; implement rollback, retraining, and incident handling	Monitoring dashboards, drift alerts, incident logs, change approvals

This mapping enables organizations to speak a common language across legal, risk, and engineering teams and to demonstrate continuous improvement even as regulatory requirements evolve.

Make It Actionable: Three Critical Foundations

1. Infrastructure Reality: Define Accountability at Every Handoff in Policies, Procedures, and Data Maps

Who owns AI outputs when customer data is trained by engineering, deployed by product, and used for decisions Legal is liable for? Map data flows, model lineage, and system ownership. Identify who owns data at each stage, from training and fine-tuning through deployment and decision-making, and who has authority to pause or override systems when risks emerge. Policies that cannot be executed in production are not governance; they create risk without a roadmap for execution.

2. Legal-Engineering Alignment: Test Policies Against System Reality

Can you honor all your data subject access requests in a trained model? Ensure that privacy, deletion, access, and transparency commitments are technically feasible. Can you explain decisions your algorithm makes? Legal teams must understand how systems operate in practice; engineering teams must understand the legal consequences of design choices.

3. Board-Ready Oversight: Ground Reporting in Infrastructure Reality

Document AI risk appetite, unacceptable uses, and testing standards. Provide quarterly dashboards on high-risk systems, incidents, and regulatory milestones. Board reporting should reflect system operations and risk reality, not just compliance status.

Then Operationalize Across Your Organization:

- Incident Response: Update playbooks for AI-specific issues–bias, drift, adversarial events. Define escalation paths and document decisions.
- Contracts and Third Parties: Update templates for training data rights, safety/bias/privacy warranties, model change disclosures, audit rights, and
- State Law Compliance: Maintain a register of obligations by jurisdiction. Adopt the strictest common denominator for enterprise standards.
- Tabletop Exercises: Conduct realistic scenarios that mimic real incidents (e.g., technical partial information, time

pressure, and competing priorities). Include Legal, Engineering, Product, Communications, and the leadership team. Pull actual logging interfaces in the tabletop so you are aware of what logging is available for your most critical AI platforms.
- Regulatory Monitoring: Assign responsibility for tracking Department of Justice, Commerce Department, agency rulemaking, and state updates.

Long-Term Planning: Phased Approach with Time Frames

Phase	Time Frame	Focus
Phase 1	0-3 months	Mapping, governance, ownership, documentation
Phase 2	3-9 months	Testing, contracts, technical controls
Phase 3	9-18 months+	Monitoring, reporting, transparency

Phase 1: Governance and Documentation (0-3 months)

- Map AI usage
- Assign accountable owners for AI risk and compliance
- Form cross-functional and diverse review groups (Legal, risk, IT, business)
- Create a system of record for all AI systems in use
- Update incident response plans for AI-specific risks
- Assess and update policies

Rationale: These foundational steps establish oversight and visibility. They can be launched immediately and should be completed quickly to demonstrate good faith to regulators and stakeholders.

Phase 2: Strengthen Testing and Controls (3-9 months)

- Broaden testing protocols (e.g., for subgroup fairness, privacy, security)
- Revise contracts and agreements for AI-specific obligations (training data rights, audit rights, model change disclosures)
- Implement technical controls for monitoring, rollback, and retraining
- Schedule the first tabletop for AI response

Rationale: This phase builds on the governance foundation. It requires coordination across teams and may involve vendor negotiations and technical upgrades. Regulators increasingly expect demonstrable progress within the first year, with an improving compliance narrative over time.

Phase 3: Continuous Monitoring and Reporting (9-18 months and ongoing)

- Shift to ongoing monitoring (alerts, dashboards, drift detection)
- Implement quarterly reporting to boards and leadership team
- Prepare public summaries or model cards as needed for transparency

Rationale: Continuous monitoring is an ongoing commitment. Initial systems should be in place within 12 to 18 months, with regular updates and improvements as the regulatory landscape evolves.

What Regulators Will Ask for by Priority

Regulators do not expect perfection; they expect visible progress and a credible improvement narrative. Here is what to have ready on a prioritized basis because full compliance is not feasible immediately.

Have Now (Foundation):

- AI system inventory with risk tiers and ownership
- Updated incident response plans for AI-specific risks
- Charter for AI governance committee

Build in Year 1 (Demonstrating Progress):

- AI policy and standards; iterate–it will not be comprehensive initially.
- Written protocols for testing and validation, but this may need to be done sooner rather than later, particularly where systems affect employment, housing, or vulnerable populations such as children or seniors.
- Vendor diligence questionnaires and updated contracts.
- Impact sector-specific assessment templates for the Health Insurance Portability and Accountability Act ("HIPAA"), the Gramm-Leach-Bliley Act ("GLBA") Family Educational Rights and Privacy Act ("FERPA"), Customer Proprietary Network Information ("CPNI"), NIST/Cybersecurity Maturity Model Certification
- Minutes from risk governance committees and training records.

Maintain Ongoing (Operational Maturity):

- Model cards and data sheets (examples of model cards in action from Hugging Face; see a helpful article on model card standardization available at Cornell University, Model Cards for Model Reporting)
- Change logs and approval records
- Compliance mapping for state and federal laws (living document)
- Monitoring dashboards and drift alerts

Key Takeaway

Regulatory uncertainty is real, but defensible steps exist. Use the NIST AI RMF as your foundation, stay compliant with state laws, monitor federal updates, and implement ongoing oversight. Acting now reduces enforcement risk, demonstrates leadership in responsible AI practices, and enables prepared, measured judgment if—and when—an AI incident occurs.

We want your team to avoid a scenario such as

> *discovering your customer service AI was making eligibility decisions it wasn't authorized to make. Legal thought they'd prohibited automated decisioning. Engineering thought the model was advisory-only. Product thought they'd disclosed it. Nobody had mapped who owned the output or who could stop the model. As a result, you fumble around for hours trying to find out who has access to shut down the model.*

AI is moving quickly, and operational documentation will ensure you have sufficient knowledge to act when necessary. Organizations building AI governance programs in 2026 should begin with infrastructure mapping and governance chartering. Early action positions you ahead of evolving requirements and ensures your AI tools are reliable and compliant.

For guidance tailored to your organization's structure—including assessment scope, privileged pre-assessment strategy, and implementation timelines—contact Amy Mushahwar, Trish Wagner, or Kai Knight.

Lowenstein's Data360 approach brings integrated legal and technical expertise to data, infrastructure, and AI governance implementation and is grounded in pressure testing in more than 20 years of incident response. Our team includes practicing attorneys who are also former chief information security officers, certified information systems security professionals, developers, and ethical hackers, as well as veterans of the New York Attorney General's Office, the FBI Cyber Division, and other enforcement agencies. This depth enables us to deliver technical fluency, operational reality, and stellar legal strategy.

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
amushahwar@lowenstein.com

P. KAI KNIGHT
Counsel
T: 202.753.3828
kknight@lowenstein.com

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

This Alert has been prepared by Lowenstein Sandler LLP to provide information on recent legal developments of interest to our readers. It is not intended to provide legal advice for a specific situation or to create an attorney-client relationship. Lowenstein Sandler assumes no responsibility to update the Alert based upon events subsequent to the date of its publication, such as new legislation, regulations and judicial decisions. You should consult with counsel to determine applicable legal requirements in a specific fact situation. Attorney Advertising.

© 2026 Lowenstein Sandler LLP | One Lowenstein Drive, Roseland, NJ 07068 | +1 973.597.2500