

Horizon Search Institute

HORIZON SCAN · 001

PILLAR I · RESPONSIBLE AI

AI Governance in *Regulated Industries*

*Lifecycle accountability and the institutional gap in
financial services and healthcare.*

AUTHORS

Cynthia Chen and Ashwin Telang

With contributions from Hernando Liu and Gloria Chen. Edited by David Lovejoy.

READING TIME

~35 minutes · 9,100 words

STREAM

Quarterly Horizon Scan

Contents

00	Executive Summary	3
I	Financial Services	6
	Actionable Strategies for Financial Institutions	12
II	Healthcare Industry	17
	Three Converging Pressures	21
	Actionable Strategies for Healthcare Institutions.....	23
III	Cross-Sector Synthesis	26
IV	Implications for Institutional Leaders	30
§	Bibliography	33

00 • EXECUTIVE SUMMARY

Deployment is outpacing accountability.

Financial services and healthcare are converging on the same AI governance expectation, continuous lifecycle oversight of deployed systems, through opposite institutional paths. Finance spent fifteen years building model risk infrastructure around SR 11-7. In April 2026 the banking agencies retired it, replacing it with SR 26-2 — and placed generative and agentic AI expressly outside the new guidance's scope. Institutions now hold a modernized set of governance principles and no directly applicable supervisory framework for the systems already entering production. Healthcare has no equivalent scaffold and is improvising one in real time, largely by inserting human verification back into workflows that automation had stripped out: patient disclosure at the point of care, physician sign-off on AI-influenced coverage denials, mandatory human utilization review.

Both sectors share a common problem that runs underneath their different regulatory architectures: the gap between any framework and the institutional capacity to operate inside it. Goldman Sachs has embedded Anthropic engineers to co-develop autonomous compliance agents, inside a sector whose validation apparatus is periodic by design. Thirty-seven percent of insurers report using or planning to use AI for prior authorization within a year, and fifty-six percent for utilization management more broadly — a determination process that was already producing high overturn rates among the denials that get appealed, and appeal rates low enough that most are never challenged at all. McKinsey's 2026 survey finds only about one-third of organizations report governance maturity at level three or higher across agentic AI controls. Deployment is outpacing accountability, and the corrective signals that should pull the system back into balance are not reaching the institutions that need to act on them.

37% / 56%

Insurers using or planning to use AI for prior authorization within a year, and for utilization management broadly.

~82%

Overturn rate on appeal found in a Medicare Advantage study, documented before AI entered utilization review.

~33%

Organizations reporting governance maturity at level three or higher across agentic AI controls.

Healthcare carries a second driver that finance does not. Federal AI legislation has not passed, state legislatures filled the vacuum with more than 250 health AI bills introduced in 2025, and a federal preemption agenda is now actively challenging those state structures. A health system operating across Colorado, Texas, California, and Maryland faces four distinct compliance obligations for the same clinical workflow, on a regulatory perimeter that may shift mid-construction. Healthcare institutions are managing two governance problems simultaneously: building accountability they did not previously need, and tracking a perimeter that is fragmented and politically contested.

For institutional leaders, the practical question over the next 18 to 24 months has shifted from which rule applies to how the organization is built to answer any rule, in any jurisdiction, at the moment a supervisor asks. Lifecycle governance requires reporting lines, monitoring infrastructure, cross-functional authority, and documentation that operate continuously, which point-in-time legal compliance does not produce.

The institutions treating AI governance as organizational design will set the template. Those treating it as legal compliance will become the case studies.

I

Financial Services

The governance infrastructure that U.S. and global financial institutions rely on to manage model risk was engineered for a category of system that is rapidly ceasing to exist.

For fifteen years, SR 11-7 — the Federal Reserve's 2011 guidance on model risk management — assumed that models are discrete, inventoriable, periodically validated, and stable between review cycles. Agentic AI systems — systems that autonomously parse trade data, resolve compliance exceptions, and initiate multi-step workflows with limited human direction — violate each of those assumptions simultaneously.¹ In April 2026 the agencies replaced SR 11-7 with SR 26-2.² The revision modernizes the framework in ways the industry had asked for, and then removes generative and agentic AI from its scope altogether.

Regulators in the United States, the European Union, and Singapore are converging on the same structural prescription: lifecycle governance, meaning continuous post-deployment monitoring and intervention rather than point-in-time clearance. The consequential gap in financial services AI governance is therefore organizational. Compliance functions inside banks remain organized around quarterly review cadences and committee-driven validation cycles designed for the Basel II era. The regulatory expectation is moving toward continuous, risk-proportionate oversight. The institution has yet to rebuild itself to receive it.

Understanding why this gap has opened requires examining the framework the new guidance replaced. SR 11-7, issued jointly by the Office of the Comptroller of the Currency and the Federal Reserve Board in 2011 and adopted by the FDIC in 2017, governed model risk in American banking until April 2026. It established the supervisory expectation that banks maintain strong model governance, conduct independent validation, and document model behavior throughout the lifecycle. Its foundational principles remain, as the Global Association of Risk Professionals noted in February 2026, "conceptually robust." The stress lies in the operational tools SR 11-7 prescribed to implement those principles. Validation approaches such as conceptual soundness assessments, outcomes analysis, and benchmarking were designed for models whose structure and behavior remain stable between review cycles. For models that recalibrate autonomously or adapt based on ongoing interaction, as GARP analyst Krishan Sharma observed two months before the framework was replaced, "these tools may lose effectiveness, as material changes in behavior can occur without a formal redevelopment event."³ The recurring debate over whether advanced AI systems should be governed as software assets or as models is a symptom of this definitional strain.

1. Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, *SR 11-7: Supervisory Guidance on Model Risk Management* (April 4, 2011). Adopted by the FDIC in 2017. Superseded April 17, 2026.

2. Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation and Office of the Comptroller of the Currency, *SR 26-2: Revised Guidance on Model Risk Management* (April 17, 2026). Generative and agentic AI are placed outside scope at note 3 of the attached guidance.

3. Krishan Sharma, "SR 11-7 in the Age of Agentic AI," Global Association of Risk Professionals, February 27, 2026. Published before the April 2026 revision.

The deployment pipeline that is generating this strain is already operational. In February 2026, Goldman Sachs disclosed that it had spent six months working with embedded Anthropic engineers to co-develop autonomous AI agents in two operational areas: accounting for trades and transactions, and client vetting and onboarding. Marco Argenti, Goldman's chief information officer, told CNBC the bank had been "surprised" at how capable Claude proved to be at tasks beyond coding, especially in accounting and compliance functions that combine parsing large volumes of documents and data with applying rules and judgment. In a subsequent interview with *American Banker*, Argenti described the resulting agents in operational detail: "reviewing documents, extracting entities, determining, for example, whether you need to ask for another document or determining if you have an ownership on a certain company and your spouse also has an ownership, then you need to do a separate KYC for that." The work, he added, involves "a lot of micro decisions within boundaries, and those micro decisions are not rules, but based on reasoning, on a chain of thought."⁴ This is precisely the category of work where prior rules-based automation has plateaued: high volume, structured at the top of the workflow, exception-laden at the bottom, and dependent on interpretive judgment to resolve the cases that fall outside the rules.

Goldman's experience is representative. Accenture's *Top Banking Trends for 2026* report found that banks are turning to cloud provider platforms to build agentic AI agents aligned with compliance and service standards, and that AI agents are already improving performance in software engineering, risk management, and customer service. BCG and OpenAI jointly projected that agentic AI has the potential to increase bank profitability by thirty percent and reduce costs by thirty to forty percent by 2030. McKinsey's 2026 AI Trust Maturity Survey, conducted across approximately five hundred organizations, found that financial services firms lead in responsible AI maturity, yet only about one-third of organizations across sectors report maturity levels of three or higher in governance and agentic AI controls.

4. Hugh Son, "Goldman Sachs taps Anthropic's Claude to automate accounting, compliance roles," CNBC, February 6, 2026; Penny Crosman, "Goldman equips AI agents to do trade accounting, onboarding," *American Banker*, February 13, 2026.




FIGURE 1 AI risks: considered relevant vs. actively mitigated, 2025. Across every risk category, awareness exceeds active mitigation. The widest gap — fifteen percentage points — is for autonomous and unintended system actions, the category most directly implicated by agentic AI deployment. Source: McKinsey & Company / Stanford AI Index Survey 2025, n ≈ 500. Reproduced from Stanford AI Index 2026, Figure 3.3.5.

The framework was modernized in April 2026. The systems this report is about were written out of it.

This last finding captures the structural problem. Financial institutions are deploying agentic systems at speed while the internal governance apparatus designed to oversee them operates on a fundamentally different clock. Validation cycles at most large banks are periodic rather than continuous. The SR 11-7 framework contemplated a world in which a model is developed, validated before deployment, and then re-examined at scheduled intervals. An agentic system that autonomously adapts its behavior in response to incoming data can undergo material behavioral drift between those intervals.

That gap was visible to regulators. The Government Accountability Office found in May 2025 that SR 11-7 may not explicitly address AI models, and both the Consumer Bankers

Association and the American Bankers Association raised the point in formal submissions, the ABA describing a "complex overlay" of supervisory expectations designed for an earlier generation of technology.

On 17 April 2026 the agencies responded. SR 26-2 supersedes both SR 11-7 and the 2021 interagency statement on BSA/AML model risk, replacing them with a risk-based framework scaled to model materiality and to the size and complexity of the institution, and expected to be most relevant to banking organizations above \$30 billion in total assets. It retains the disciplines that made SR 11-7 durable: conceptual soundness, outcomes analysis, ongoing monitoring, effective challenge by objective and independent reviewers, model inventory, documentation. It also tightens the definition of a model, excluding spreadsheet arithmetic and deterministic rule-based software from scope — resolving, in one direction, the long-running argument over whether advanced systems should be governed as models or as software assets.

And then, in a footnote, it removes the systems this report is about. Generative and agentic AI models are described as novel and rapidly evolving, and placed outside the guidance's scope. Institutions are directed to apply their own risk management and governance practices to anything the document does not cover. The guidance is also explicitly non-binding: non-compliance will not itself draw supervisory criticism, though supervisory action may still follow from unsafe or unsound practices stemming from insufficient management of model risk.

This is the sharpest version of the institutional gap in financial services, and it is not the one the sector expected. The problem is no longer that the governing framework is fifteen years old. The framework has been modernized. The problem is that the modernization was drawn to exclude precisely the category of system that Goldman Sachs, and others, have already put into production. Banks now hold sound principles, a supervisory expectation that they will exercise judgment, and no directly applicable framework — for agentic systems making consequential decisions today.

Singapore's approach offers the clearest picture of where lifecycle governance is heading and what it demands from institutions. The Monetary Authority of Singapore issued its consultation paper on Guidelines for AI Risk Management in November 2025, proposing requirements that apply to all financial institutions and cover oversight, AI risk management systems, lifecycle controls, and organizational capabilities. The guidelines explicitly address generative AI and AI agents. They require board-level accountability for AI risk, mandate the establishment of cross-functional oversight committees for institutions with material AI risk exposure, and prescribe continuous monitoring of fairness across credit and insurance decision systems. The toolkit includes an operationalization handbook or-

ganized around four sections — scope and oversight, risk management, lifecycle management, and enablers — each aligned with the proposed MAS guidelines. This iterative, sector-specific trajectory distinguishes Singapore from the EU's horizontal risk classification approach and from the United States' continued reliance on retrofitting pre-existing supervisory guidance.

KPMG's analysis of the MAS consultation emphasized that the guidelines introduce formal requirements focused on strengthening AI risk management across the entire AI lifecycle. The EU follows a different route to the same destination: rather than issuing new sector-specific AI guidelines, the EBA's position is that the AI Act's high-risk obligations — deferred in July 2026 to 2 December 2027 for standalone Annex III systems, a category that includes creditworthiness assessment and credit scoring — can be implemented through existing financial regulatory structures and supervisory cooperation. The mechanism differs, but the expectation of continuous, lifecycle-oriented oversight does not.

K&L Gates' analysis of the AI Act's financial services implications noted that the regulation's global reach means AI providers and financial institutions interacting with EU users must comply regardless of where they are established. For U.S. institutions operating internationally, these obligations create a compliance surface that extends well beyond domestic supervisory expectations. The EU's Digital Omnibus proposal, introduced in November 2025, further consolidates rules on AI, data access, privacy, and cybersecurity for the financial sector. A parallel development is unfolding in Asia: Vietnam's Law on Artificial Intelligence, effective March 2026, is the first standalone binding AI law in Southeast Asia, with an eighteen-month transition period for the financial sector. The ASEAN Digital Economy Framework Agreement, expected by the end of 2026, will create legally binding AI governance interoperability rules across the region.

The directional convergence across these jurisdictions has a clear implication for executives: the institution's operating model, not its compliance documentation, is what determines whether it can meet the standard. MAS requires continuous monitoring and lifecycle controls. U.S. banking regulators, while operating through guidance rather than statute, have signaled through the GAO report, NIST's December 2025 Cyber AI Profile, and Congressional testimony that expectations around AI accountability are tightening. The Financial Stability Board's November 2024 report on the financial stability implications of AI identified model risk, data quality, governance, and third-party concentration as key vulnerabilities requiring regulatory attention. Georgetown CSET's October 2025 mapping of over 950 AI governance documents found that governance failure and AI system security vulnerabilities were the most heavily covered risk subdomains across the global gov-

ernance landscape, a finding that confirms regulators worldwide have identified the same structural weakness.

For financial institutions, the operational implication is concrete. The compliance organizations that currently manage model risk were built for a regulatory environment that assumed stable models and bounded use cases. Agentic AI systems require monitoring infrastructure that can detect behavioral drift in production, escalate anomalies in real time, and maintain audit trails that satisfy regulators who increasingly expect lifecycle documentation. McKinsey's 2026 survey found that security and risk concerns are the top barrier to scaling agentic AI, that inaccuracy and cybersecurity remain the most frequently cited AI risks, and that active mitigation lags behind risk awareness across nearly every AI risk category. Eighty percent of AI projects fail to deliver intended business value, according to RAND Corporation data synthesized by Pertama Partners, with governance failure a recurring cause.

The institutions that will navigate this transition successfully are those treating AI governance as organizational design. They will restructure reporting lines, build continuous monitoring capabilities, and invest in the cross-functional expertise that MAS, the EBA, and the Federal Reserve's supervisory posture all increasingly demand. The calendar has loosened; the work has not.

The EU AI Act's high-risk obligations for standalone Annex III systems, including credit scoring, moved from August 2026 to 2 December 2027 under the Digital Omnibus that entered into force on 27 July 2026. Obligations for AI embedded in regulated products moved to 2 August 2028. The Article 50 transparency obligations still apply from 2 August 2026. MAS closed its consultation on 31 January 2026 and has proposed a twelve-month transition following finalization, which had not occurred as of this edition's cutoff. U.S. banking regulators have issued their revised model risk framework and left agentic systems outside it.

The delay is not a reprieve on the substance. The European Commission's own reasoning was that the harmonized standards and the designated national authorities were not ready — an admission that the supervisory apparatus is behind, not that the requirements are softening. Institutions reading the new dates as permission to defer will arrive at the 2027 examination cycle with a governance program in mid-build. The institutions that use the sixteen months will arrive with one that has been stress-tested.

Actionable Strategies: What Financial Institutions Should Do Now

The strategies that follow apply to financial institutions specifically. Healthcare-specific actions appear in the next chapter; cross-sector postures appear in Implications for Institutional Leaders.

01 Build a Living AI Inventory That Catches Shadow Deployments

The Treasury's Financial Services AI Risk Management Framework, released February 2026, opens with the AI Adoption Stage Questionnaire for a reason: institutions cannot govern what they cannot see. Automated discovery is now the baseline. Scanning tools maintained by the second line of defense should sweep API endpoints, embedded LLM prompts, vendor integrations, end-user computing files, and notebook environments on a scheduled cadence, populating a central repository with risk tier, owner, intended use, data lineage, and change history. The complete model inventory expectation, carried from SR 11-7 into SR 26-2, has been reinforced by the GAO's May 2025 findings on AI oversight gaps: incomplete or fragmented inventories remain a recurring examination concern, compounded by shadow deployments that accumulate outside the second-line view. A monthly reconciliation between procurement records, cloud spend, and the central inventory catches the drift between what the enterprise has authorized and what is actually running in production.

02 Operate a Tiered Validation Cadence Calibrated to Agentic Behavior

SR 26-2 calibrates validation to model purpose, methodology, materiality, and the frequency and scope of model changes. Agentic systems sit outside its scope entirely, which makes cadence an institutional judgment rather than a supervisory instruction. Building on the principles SR 26-2 carries forward and MAS's proposed lifecycle controls, institutions should assign every agentic system a validation frequency tied to its autonomy level, the reversibility of its actions, and its exposure to non-stationary data, with shorter cycles for systems that recalibrate autonomously between reviews. Credit decisioning agents that adjust weighting in response to new application data warrant quarterly validation. Trade reconciliation agents operating on structured ledger data may sustain annual cycles. Champion-challenger testing, where a candidate model runs in parallel with the production model on live data, is a benchmarking technique named in both the old framework and the new, and should be standard practice for any system that recalibrates autonomously, and is increasingly expected under continuous-monitoring supervisory regimes.

03 Engineer Continuous Monitoring Into the Production Stack

Continuous monitoring is the lifecycle governance expectation that MAS, the EU AI Act, and the Treasury's FS AI RMF all converge on. Implementation requires three layers. The first is statistical: automated drift detection on input distributions, output distributions, and population stability indices, with thresholds that trigger automatic alerts to the second line. The second is behavioral: telemetry that records every agent action, tool call, memory state, and decision rationale, stored in tamper-evident audit logs. The third is fairness: ongoing disparate-impact testing across protected classes for any model that influences credit access, pricing, or claims adjudication, consistent with the Consumer Financial Protection Bureau's dynamic monitoring expectations under the Equal Credit Opportunity Act. Integrated governance platforms that consolidate these three layers produce the examination-ready documentation supervisors increasingly expect on short notice, in line with the MAS, Treasury FS AI RMF, and SR 26-2 documentation standards.

04 Install Kill Switches, Circuit Breakers, and Graduated Autonomy Ramps

Agentic systems fail faster than human operators can respond. Infrastructure-level safety controls become essential. Building on NIST AI RMF GOVERN function guidance and the Treasury FS AI RMF's incident response expectations, institutions should design a global hard stop that revokes tool permissions and halts queues in seconds, paired with per-agent circuit breakers for high-cost actions and objective-based breakers for repetitive patterns. Spend and rate governors cap tokens, API calls, and per-task budgets. The AWS Prescriptive Guidance on incident response for agentic AI adds a critical operational layer: business continuity plans that assume the agent is offline, with staff trained to handle the workload manually during the degradation window. New agents should return to production through a graduated autonomy ramp: full human-in-the-loop mode at launch, graduating to human-on-the-loop after a defined confidence period with zero anomalies.

05 Extend Third-Party Oversight to Foundation Model Providers

Goldman Sachs's Anthropic partnership illustrates a new category of vendor risk: institutions depending on foundation models whose training, fine-tuning, and update cadence they do not control. The Treasury FS AI RMF includes explicit control objectives for vendor evaluation, service-level agreements, and incident response coordination with AI providers. Contracts should include right-to-audit clauses, change notification requirements, performance guarantees tied to the institution's actual token volume, and exit strategies that address data portability. MAS's consultation paper requires compensatory testing of third-party AI to verify robustness and detect bias, plus contingency plans for vendor discontinuation. A standing quarterly review of every third-party AI contract against the FS AI RMF vendor control objectives prevents the accumulation of unpriced risk.

06 Rebuild the Three Lines of Defense Around Agentic Accountability

The structural problem in agentic AI accountability is diffuse ownership: product owns the model, engineering owns the infrastructure, compliance owns the policy, and when a decision is challenged each team points elsewhere. The remedy is explicit accountability mapped to specific action categories and documented in writing before the system enters production. A named senior manager should own each material agentic system with authority to pause, modify, or decommission it. The first line deploys and operates. The second line validates independently, monitors in production, and challenges. The third line provides assurance to the board risk committee. The BIS Consultative Group on Risk Management's 2025 report on AI governance recommends adapting this existing model to cover AI-specific threats and vulnerabilities, which preserves governance coherence as the institution scales.

07 Elevate AI Risk to the Board Risk Committee

GARP's February 2026 analysis argues that AI oversight belongs with the board risk committee because AI risk now cuts across credit, market, operational, and compliance domains. The board should approve the institution's AI risk appetite, set boundaries on automated decisioning, review material deployments before launch, receive quarterly AI risk reporting with defined KRIs, and review model incidents and failures. The CFA Institute's 2026 bank risk framework recommends board-level engagement specifically to reduce ambiguity in decision rights and signal to regulators that AI governance is treated as a core risk discipline. Board materials should include the institution's current FS AI RMF adoption stage, the gap to the next stage, incident counts and remediation status, and cross-jurisdictional compliance posture for MAS, EU AI Act, and U.S. supervisory expectations.

08 Red-Team Before Deployment and Continuously After

The Cloud Security Alliance's Agentic AI Red Teaming Guide details specific test categories: permission escalation, hallucination, orchestration flaws, memory manipulation, and supply chain risks. Gray Swan AI's 2025 benchmark pitted 22 frontier AI agents against 44 realistic deployment scenarios and observed nearly two million prompt injection attacks, with over 60,000 successful. Institutions should conduct structured red-team exercises before every material deployment, mapping findings to the OWASP Top 10 for Agentic Applications. The red team should sit in the second line of defense, report independently to the chief risk officer, and rerun test suites after every material model update. HackerOne's 2026 agentic taxonomy provides a mapped weakness framework that translates red-team findings into compliance reporting, closing the loop between testing and remediation.

09 Document Everything for Examination Readiness

Regulators across jurisdictions are converging on the same documentation expectation: institutions should be able to produce, on short notice, the full lifecycle record of any material AI system. This includes the intended use statement, training data lineage and quality assessment, architecture description, validation evidence, production monitoring logs, incident history, change management records, and third-party documentation. Following the model documentation expectations carried from SR 11-7 into SR 26-2, on which MAS, the EU AI Act, and Treasury FS AI RMF all build, effective documentation demonstrates rationale (why this model over alternatives), assumptions (underlying principles, data sources, methodologies), testing performed (pre-implementation and ongoing), and decision-making (why validation and deployment choices were made). Examination-ready documentation is the tangible output of every other strategy on this list, and it is the artifact regulators evaluate first.

II

Healthcare Industry

The U.S. healthcare AI regulatory environment presents a structurally different version of the same problem.

The institutions deploying AI, including health systems, insurance payers, and clinical software vendors, are outpacing both the regulators supposed to oversee them and their own internal governance capacity. By pulling back at the center, the federal government has pushed compliance accountability down to the state level and, ultimately, inside the organizations themselves. The following sections map the redistribution across four state legislative lanes and argue that the central risk in 2026 is not regulatory uncertainty but the gap between what these organizations are deploying and what they have built the internal infrastructure to govern.

The clearest signal of the federal government's current posture came on January 6, 2026, when the FDA issued revised final guidance on Clinical Decision Support Software, superseding its 2022 version. The agency re-issued the guidance on January 29, superseding its own January 6 text; both dates appear in the citation below. The guidance draws a line between clinical decision support tools that are regulated as medical devices and those that are not. Under the updated guidance, FDA will now exercise enforcement discretion for CDS tools that provide a singular output, including, for example, software that recommends a specific FDA-approved drug based on a patient's symptoms and medical history, as long as the tool meets the other non-device CDS criteria.⁵ This was a meaningful shift from the 2022 guidance, which had excluded software providing a "specific preventative, diagnostic, or treatment output or directive" from the non-device category. FDA Commissioner Marty Makary framed the changes at CES as intended to "cut unnecessary regulation and promote innovation" and said the agency needed to move "at Silicon Valley speed." Legal analysts noted that while the guidance provides helpful clarifications, it does not represent a fundamental deregulatory shift. FDA oversight still applies fully to tools that analyze medical images or make recommendations that are not independently reviewable by a clinician.

The CDS guidance is the most substantive federal healthcare AI action of 2026, and its practical effect is to expand the category of AI-enabled clinical software that can reach the market without premarket FDA review. For health system procurement teams, this creates a gap: software that previously required FDA clearance as a condition of vendor due diligence may no longer carry that signal. The accountability structure changes while clinical risk stays the same. No federal legislation directly governing health AI has passed, and the Trump administration has signaled it may challenge state-level AI laws through DOJ litigation as incompatible with its "minimally burdensome" national policy framework. The result is a federal posture that is simultaneously deregulatory on AI tools and adversarial toward the state-level structures attempting to govern the consequences.

5. U.S. Food and Drug Administration, "Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff." Issued January 6, 2026; re-issued January 29, 2026, superseding the January 6 version.

As a response, states have moved with unusual speed and specificity. In 2025, 47 states introduced more than 250 health AI bills, with 34 passed and enacted into law across 21 states, and over 240 bills across 43 states have already been filed in 2026. The result is a fragmented national framework: disclosure requirements vary in threshold and timing, mental health chatbot rules differ in scope, and payer mandates impose different standards of human review depending on the state. A health system operating across Colorado, Texas, California, and Maryland faces four distinct compliance obligations for the same AI-enabled workflow, with no federal floor to unify them.

47

States introducing health AI legislation in 2025.

250+

Bills introduced. 34 enacted into law across 21 states.

240+

Additional bills filed across 43 states in 2026 as the legislative pace continues.

Health system leaders have publicly described the operational difficulty of deploying AI tools when state requirements diverge, citing the resources required to maintain different workflows for the same clinical use case across multiple jurisdictions. State legislative activity in 2026 is concentrating in four areas: mental health chatbots, patient disclosure and consent, preventing AI tools from presenting as clinical providers, and payer use of AI in coverage decisions. Each area reflects a different institutional failure that legislators are trying to correct from the outside, without coordination across jurisdictions.

The mental health chatbot lane is driven by documented harm: California's law, effective January 1, 2026, bans chatbot deployment without protocols to prevent the production of suicidal ideation content and imposes stricter requirements when the user is a minor. Illinois, effective August 2025, prohibits AI from making independent therapeutic decisions or generating treatment plans without licensed professional review. These statutes are not precautionary but responses to specific documented incidents of consumer-facing AI products causing therapeutic harm to vulnerable users, including minors.

The patient disclosure and consent lane is the most organizationally demanding because it requires operational changes at the point of care. Texas, effective January 1, 2026, requires written disclosure to patients before any AI system is used in their treatment or services. Ohio has introduced legislation requiring written informed consent. The practical implication for health systems, particularly those that have embedded AI scribes, diagnostic aids, or decision support into clinical workflows, is that disclosure obligations now

sit with the provider at the point of service, not with the software vendor. This creates a governance and training problem, on top of the legal aspect.

The most consequential regulatory action in 2026 is playing out in the payer lane, and it is driven by documented evidence of harm at scale. Research published in *Health Affairs* in January drew on a 2024 survey of ninety-three large insurers across sixteen states to show how far AI has moved into utilization review: 84% of insurers reported using AI for some operational purpose, 37% for prior authorization now or within a year, 44% for claims adjudication, and 56% for utilization management broadly. In the large-employer group market, 70% are using or exploring AI for prior authorization⁶ — the process that determines whether patients receive covered care before it is delivered.

The governance concern is not theoretical, but it is easily misstated, and the distinction matters. The evidence base concerns the process AI is entering, not AI's performance within it. Investigations going back more than a decade have documented high denial rates and high overturn rates on appeal in prior authorization: one Medicare Advantage study found an overturn rate of nearly 82%. In a separate study of Affordable Care Act Marketplace plans, denial rates reached 20%, fewer than 1% of denials were appealed, and nearly half of the appeals that were filed resulted in reversal. Both findings predate AI's arrival in utilization review, and Medicare Advantage plans approved more than 93% of prior authorization requests between 2019 and 2023.

That is the correct way to read the risk. AI is being introduced into a determination process already characterized by high overturn rates among the decisions that get challenged, appeal rates low enough that most are not challenged at all, and weak accountability for the difference. The Stanford researchers are explicit that evidence on AI's effectiveness in insurance settings remains scarce, and that erroneous decisions may equally be made by human reviewers or by older non-AI algorithms. Without governance, AI does not have to be worse than what it replaces to do damage. It only has to scale what is already there.

The researchers identified structural reasons why internal governance has not corrected this: human reviewers at insurers may lack the clinical expertise to meaningfully evaluate AI-generated recommendations; the opacity of algorithmic determinations makes errors hard to detect and harder to challenge; and training data drawn from insurers' historical decisions repeats the biases and flaws of those decisions. The technology deployed at scale inside these institutions is making consequential decisions that the organizations themselves cannot adequately audit or explain, and the incentive structure, where incor-

6. Michelle M. Mello, Artem A. Trotsyuk, Abdoul Jalil Djiberou Mahamadou and Danton Char, "The AI Arms Race in Health Insurance Utilization Review," *Health Affairs* 45(1), 2026, reporting a 2024 National Association of Insurance Commissioners survey of 93 insurers in 16 states.

rect denials are financially beneficial and rarely challenged, explains why there is no internal pressure to fix it.

States have responded by trying to impose accountability structures from outside. California's SB 1120, effective January 2025, prohibits health plans from basing medical necessity determinations solely on AI and requires human clinical review. Maryland, Arizona, Connecticut, Nebraska, and Texas have enacted comparable requirements. The legislative frontier is now moving further. Louisiana SB 246, introduced by Senator Jay Luneau, goes further than any currently enacted law. Under the bill, AI could not be used to delay, deny, or modify healthcare services on its own; any adverse determination would require independent judgment from a human utilization reviewer, and a physician who personally reviewed the medical record would be required to sign off on each denial. The bill's most structurally significant provision is its burden-shifting rule: any adverse determination in which AI materially contributed would be presumed invalid unless the insurer could demonstrate the decision was independently reached through documented clinical judgment without reliance on algorithmic output.

Additionally, if a denial is appealed on AI grounds, the insurer could not use AI in any subsequent review of that claim. This inverts the current default. Rather than patients having to prove algorithmic error, insurers would have to prove human independence. Louisiana SB 246 represents the direction the legislation is moving, even if this specific bill does not pass in its current form.

What connects these four legislative lanes is a common problem: the institutions deploying AI in healthcare have not built the internal governance infrastructure proportionate to the risk of what they are deploying. The FDA's CDS guidance expansion means more software will reach clinical environments without premarket review, shifting validation responsibility to health systems and payers at the point of procurement and deployment. State disclosure mandates mean frontline clinical workflows now carry legal obligations that require training, documentation, and oversight infrastructure that most organizations have not yet built. Payer AI mandates mean that prior authorization workflows, which have been heavily automated with AI for efficiency gains, now require documented physician-level review at scale, a requirement that is operationally expensive and organizationally complex precisely because automation was adopted to avoid it.

The corrective is not new technical oversight infrastructure but the legislated reinsertion of human judgment, supported by training, documentation, and workflow redesign.

There is a common mechanism underneath these four lanes, and it is worth naming directly. Each is, in effect, legislating humans back into workflows that automation had stripped out. Patient disclosure mandates require a clinician to name the AI's role at the point of care. Mental health chatbot rules require licensed professional review of any therapeutic decision. Payer AI mandates require physician sign-off on adverse determinations, and Louisiana SB 246 goes further by presuming any AI-influenced denial invalid unless human independence can be documented. The corrective is not new technical oversight infrastructure but the legislated reinsertion of human judgment, supported by training, documentation, and workflow redesign. This matters for how institutions should plan. The compliance investment is not principally a model risk platform or a monitoring stack. It is the operational capacity to route AI-influenced decisions through qualified humans at scale, in workflows that were redesigned over the past decade to avoid exactly that step.

Tim Hwang, General Counsel at AI scribe company Abridge, has argued that regulatory clarity would actually accelerate AI adoption in healthcare. His argument understates the near-term organizational challenge. Clarity about what the rules are is not the same as organizational capacity to comply with them. Health systems operating across multiple states now face disclosure requirements that vary by jurisdiction, payer AI mandates that differ in threshold and scope, and a federal regulatory boundary for clinical software that has just been redrawn. The compliance question is whether these organizations have the governance structures, the trained personnel, and the documented processes to demonstrate accountability to multiple regulators simultaneously.

Three Converging Pressures and the Institutional Gap They Expose

The regulatory trajectory in healthcare AI is determined less by any single law or court ruling than by the interaction of three converging pressures: state mandates tightening around payer accountability, a federal preemption agenda with specific enforcement deadlines, and an FDA whose published governance ambitions for AI-enabled devices now exceed its demonstrated institutional capacity to enforce them. Each puts pressure on the same institutional weakness, which the chapter closes on.

P1 January 1, 2027: Colorado Rewrites Its Own Framework

Colorado did not wait to be preempted. On 14 May 2026, weeks before its own AI Act was due to take effect, the state repealed and reenacted it. SB 26-189 replaces the 2024 statute's duty of care, risk management program and impact assessment requirements with a narrower disclosure regime governing automated decision-making technology in con-

sequential decisions, health care among them. The new obligations take effect 1 January 2027, with Attorney General rulemaking required before that date. Colorado also enacted a Chatbot Safety Act and a separate health care AI statute on the same timeline. The DOJ AI Litigation Task Force, established in January 2026 with the sole mission of challenging state AI laws, remains active, and legal analysts at White & Case advise continued compliance with state AI laws until there is greater clarity.

The retreat is the finding. The first and most comprehensive state AI framework in the country was dismantled by its own legislature, under federal preemption pressure, before it ever operated. For health systems that had begun building to it, the sunk work is real and the target has moved. For every other state legislature watching, the signal is that the most demanding compliance architecture is also the most politically fragile. The perimeter is not merely fragmented. It is retractable.

P2 June 9, 2026: FCC Rulemaking Deadline

The executive order's June 9 deadline directs the FCC to initiate rulemaking on federal AI reporting and disclosure standards specifically designed to preempt conflicting state laws. Unlike DOJ litigation, which challenges statutes case by case, an FCC standard would establish affirmative federal requirements, a structurally more durable preemption mechanism that poses a direct threat to patient disclosure mandates now embedded in clinical workflows across Texas, California, and Maryland.

P3 The FDA's Two-Track Posture and Capacity Gap

The January 2026 CDS guidance pulling back on clinical decision support oversight should be read alongside the FDA's January 2025 draft guidance proposing a Total Product Lifecycle framework for AI-enabled medical devices. The agency is deregulating one category while building more structured controls for another, and the boundary between them is where the most contested healthcare AI products sit. That boundary question is complicated by institutional capacity: CDRH approved only nine new or substantially changed high-risk medical devices in the quarter following 2025 DOGE staffing cuts, down from 13 the prior year, even as pending applications increased. An agency publishing more sophisticated governance standards while losing specialized reviewers creates a gap that health system procurement teams should factor into vendor due diligence.

THE INSTITUTIONAL GAP THESE PRESSURES EXPOSE

Actionable Strategies: What Healthcare Institutions Should Do Now

The strategies that follow apply to health systems and payer organizations operating in the U.S. regulatory environment. They translate the four legislative lanes and the three converging pressures into operational priorities for the next twelve to eighteen months. Cross-sector postures appear in Implications for Institutional Leaders.

01 Build Multi-State Disclosure Infrastructure at the Point of Care

Texas's January 2026 disclosure requirement, California's mental health chatbot law, and pending Ohio consent legislation share a common operational implication: the disclosure obligation lives with the clinician or provider, not the software vendor. Health systems operating across state lines need workflow infrastructure that scales with the strictest applicable standard rather than maintaining different processes for the same clinical use case across jurisdictions. EHR-integrated disclosure prompts, captured patient acknowledgment, and clinician training programs are the operational components. Building once to the highest common denominator is more durable than rebuilding per state as the legislative map shifts under the federal preemption agenda.

02 Restructure Prior Authorization for Documented Physician-Level Review at Scale

California SB 1120 already prohibits AI-only medical necessity determinations, with comparable requirements in Maryland, Arizona, Connecticut, Nebraska, and Texas. Louisiana SB 246 signals the direction of travel: burden-shifting that presumes AI-influenced denials invalid absent documented human independence. Payer organizations should redesign utilization review workflows so every adverse determination flows through, and is documented as flowing through, qualified clinical review. The infrastructure question is staffing, audit trails, and case management, not algorithmic improvement. The overturn rates documented in Medicare Advantage before AI entered the process indicate the magnitude of the corrective work required, and the standard any AI-assisted process will be measured against.

03 Establish a Governed AI Inventory Across Clinical and Operational Environments

The FDA's narrowed CDS guidance shifts validation responsibility from premarket review to procurement and operations teams at the point of deployment. Health systems need a central registry of every AI-enabled tool in clinical use, including EHR scribes, decision support, imaging, prior authorization tools, and scheduling algorithms, with risk tier, validation evidence, vendor documentation, and post-deployment monitoring linked to each. Without this inventory, an institution cannot answer the basic supervisory question of where AI is operating inside its walls. With it, the institution has the foundation for the post-market obligations that the FDA's Total Product Lifecycle framework will eventually impose.

04 Prepare for FDA Total Product Lifecycle Obligations Before They Finalize

The FDA's January 2025 draft guidance on AI-enabled device lifecycle management proposes post-market monitoring obligations, including model drift detection, performance documentation, and predetermined change control plans. Most health system IT and procurement functions are not currently structured to handle these. Building drift monitoring, performance dashboards, and incident reporting infrastructure now, even before finalization, produces examination-ready capability when the guidance lands. CDRH's reduced reviewer capacity following 2025 staffing cuts means that, when finalization comes, supervisory expectations will likely outpace the agency's bandwidth to provide guidance during transition.

05 Stand Up Cross-Functional AI Governance with Clinical Authority

Most health systems hold IT, clinical leadership, legal, compliance, and quality in separate reporting lines. Lifecycle AI governance routes through all five. The structure that works is a named senior owner per material AI deployment, a cross-functional oversight body with authority to pause or modify deployments on demand, and clinical leadership with explicit decision rights, not a committee that only reviews. Payer organizations need the equivalent structure for utilization review AI, with accountability anchored at the level of the chief medical officer rather than distributed across operations and information technology.

III

Cross-Sector Synthesis

Financial services and healthcare operate under fundamentally different regulatory traditions, yet they are converging on a common governance expectation from opposite directions.

The dominant supervisory posture in both sectors is shifting from point-in-time review to lifecycle accountability. The Monetary Authority of Singapore's November 2025 consultation paper, the FDA's January 2025 draft guidance on AI-enabled medical device lifecycle management, the EU AI Act's high-risk obligations now applying from December 2027, and the April 2026 replacement of SR 11-7 by SR 26-2 differ substantially in form, but each imposes the same underlying requirement: continuous monitoring, post-deployment intervention, and documented behavioral oversight across the operational life of a deployed system.

What has not shifted at comparable speed is the organizational architecture of the institutions these instruments govern. Compliance functions at most large banks still operate on quarterly or semi-annual validation cycles. Most health systems and payers have not yet built the monitoring infrastructure or cross-functional accountability structures their own AI deployments now require. Regulators are writing rules for 2026. Most institutions are still running governance on a 2018 operating model.

The asymmetry has two layers, and they need to be held separately. The first is organizational: in both sectors, deployment is moving faster than governance capacity, and the resulting gap is structurally identical regardless of which industry the institution sits in. The second is regulatory. Finance operates inside a coherent supervisory perimeter. SR 26-2, the EU AI Act, the MAS guidelines, and BIS expectations are different instruments, but they are issued by recognizable bodies and apply across the institution's footprint with reasonable predictability. Healthcare does not have that. With federal AI legislation absent and a federal preemption agenda actively challenging state structures, the regulatory perimeter itself is fragmented and contested. The organizational driver is shared across sectors. The regulatory driver is not.

Even with the organizational driver shared, the institutional starting point differs. Financial services carries the weight of fifteen years of institutional model risk infrastructure. SR 11-7 established in 2011 a jointly supervised expectation that banks document model lineage, conduct independent validation, and maintain governance artifacts regulators can examine. The question confronting the sector now is whether that infrastructure, engineered for stable and periodically validated models, can be adapted to oversee agentic systems that recalibrate autonomously between review cycles.

Healthcare arrives at the same point with no comparable scaffold, and what it is being asked to build looks structurally different from what finance is adapting. The state-level mandates already in force or pending (patient disclosure at the point of care, physician sign-off on AI-influenced coverage denials, mandatory human utilization review) bolt human verification back into workflows that automation had largely displaced. Louisiana SB

246 makes the pattern explicit: any AI-influenced adverse determination is presumed invalid unless the insurer can document independent human clinical judgment. This is less the construction of novel technical infrastructure than the legislated reinsertion of human review into automated processes, supported by training, documentation, and workflow re-design that most organizations have not yet built at scale. Finance is adapting model risk infrastructure to a new class of system. Healthcare is operationalizing human-in-the-loop accountability that the deployment curve had bypassed.

What connects the two sectors is institutional deployment outpacing institutional accountability; the regulatory architecture sits downstream of that fact rather than driving it. Goldman Sachs has embedded Anthropic engineers inside its technology organization to co-develop agentic systems capable of extracting entities from complex documents, determining ownership structures, and triggering downstream compliance checks. The bank's model validation functions continue to operate on quarterly cycles designed for an earlier paradigm, even as production behavior adapts continuously between those cycles.

The same pattern, reframed, describes the payer lane in healthcare. Thirty-seven percent of insurers use or plan to use AI in prior authorization, and fifty-six percent in utilization management broadly — a process in which pre-AI studies already found overturn rates near 82% among appealed Medicare Advantage denials, and appeal rates below 1% in the Marketplace population. Deployment has scaled faster than the organizational structures required to audit, challenge, or explain algorithmic determinations, which means most incorrect denials are never contested and the denial rate itself carries no corrective signal back into the system.

McKinsey's 2026 AI Trust Maturity Survey, conducted across approximately 500 organizations, found that only about one-third report governance maturity levels of three or higher across agentic AI controls. The finding generalizes: any regulated industry deploying AI at operational speed while governance capacity lags its deployment curve is exposed to the same structural risk, and the one-third data point is the best current proxy for how unevenly that risk is distributed across the sectoral landscape.

Over the next 18 to 24 months, the meaningful split is between two kinds of institution: those treating AI governance as organizational design, and those treating it as legal compliance. The operational consequences of that split are large. Legal compliance functions are oriented toward demonstrating adherence to a specified rule at a specified moment. Organizational design produces something different in kind: the reporting lines, monitoring infrastructure, cross-functional authority, and documented lifecycle posture that allow an institution to answer any rule, in any jurisdiction, at the moment it is asked.

The gap between regulatory expectation and organizational reality will close in one of two ways: because institutions rebuild themselves, or because supervisors force them to. The first path is governance. The second is enforcement.

IV

Implications for Institutional Leaders

The synthesis names the pattern. The institutions best positioned to navigate it are acting on four postures now, before the regulatory expectations fully arrive.

All four are achievable inside the existing supervisory environment without waiting for federal clarity. Each reflects a decision about what AI governance is inside the institution, and which functions are accountable for it.

01 AI governance is organizational design, not legal compliance.

The institutions that will succeed in this transition treat AI governance as a question about how the organization is structured rather than about which rule applies. Legal compliance answers whether the institution is in adherence at a specified moment; organizational design answers whether the institution is built to demonstrate adherence to expectations that have not yet been finalized, across jurisdictions that do not yet agree on the standard.

The difference matters for two reasons. First, lifecycle expectations cannot be satisfied through point-in-time attestations. MAS's lifecycle controls, the EU AI Act's post-deployment monitoring obligations, the FDA's Total Product Lifecycle framework, and SR 26-2's ongoing monitoring expectations all presume an institution with continuous visibility into its AI systems, continuous authority to intervene, and continuous documentation of both. Those are organizational capabilities built into reporting lines and monitoring infrastructure, not policies the legal function can draft in isolation.

Second, the fragmentation of state-level healthcare regulation and the cross-jurisdictional scope of financial regulation make any single-rule compliance posture structurally fragile. A health system operating across Colorado, Texas, California, and Maryland faces four distinct AI compliance obligations for the same clinical workflow, and Colorado rewrote its own obligation mid-year. A bank operating across the United States, the EU, and Singapore faces three converging but non-identical lifecycle frameworks. Building the organization to meet the highest common expectation is more durable than building the legal function to parse each jurisdiction in sequence, and the institutions best positioned are already treating governance architecture, not rule interpretation, as the primary question.

02 Build continuous monitoring before the regulation demands it.

MAS's proposed twelve-month transition period following finalization, the EU AI Act's high-risk obligations now applying from December 2027, and state-level payer mandates already in force converge on the same operational expectation: continuous post-deployment monitoring. What that means in practice varies by system type. An agentic trading system needs drift detection and behavioral telemetry. A clinical decision support tool needs post-market performance documentation. A prior authorization model needs fairness testing across protected populations and audit trails that connect denials to documented clinical review.

Retrofitting monitoring after the deadline is meaningfully harder than building it before. The infrastructure operates in three layers, and none can be purchased as a product. Statistical monitoring requires integration with production systems. Behavioral telemetry requires instrumentation at the agent-action level. Fairness monitoring requires longitudinal data pipelines that most organizations have not built. An institution that starts the work before the deadline has time to iterate on false positives, calibrate thresholds, and produce the documentation artifacts regulators will request; one that starts after faces supervisory scrutiny on infrastructure that has not been stress-tested, with no grace period to tune it.

The question inside the institution is no longer whether a monitoring capability exists, but on which systems, at what granularity, with what escalation paths, and reviewed by whom. These are twelve-month build decisions, not twelve-week ones, and institutions that treat them as a procurement exercise typically discover mid-build that the supporting data infrastructure is the binding constraint.

03 **Invest in cross-functional expertise now.**

Most institutions currently hold model risk, compliance, legal, information security, and technology in separate reporting lines, each with distinct leadership, accountability, and examination posture. That structure works for a regulatory environment organized around separable compliance domains, but lifecycle AI governance routes through all five functions simultaneously and requires coordinated decisions that the existing committee architecture was never set up to produce.

The financial services chapter's treatment of the three lines of defense and the board risk committee, and the healthcare chapter's treatment of physician-level review obligations and state disclosure mandates, describe the same organizational fact in different domains. Agentic systems and clinical AI both produce consequential outputs on a continuous cadence, and the authority to pause a deployment, investigate an anomaly, or document an adverse decision therefore cannot be distributed across five committees on quarterly meeting schedules. It has to sit in a named, empowered, cross-functional structure with decision rights, which is an organizational design most large institutions have not yet built even where the component functions all exist in mature form.

The institutions best positioned are building this structure now: a named senior owner for each material AI system, a cross-functional oversight body with authority to convene on demand, and model risk, compliance, legal, and technology represented in a single governance forum rather than in parallel ones. These are not exotic organizational forms; the same patterns already exist in cybersecurity and market risk at most large institutions,

and the task is transposing them onto a category the organization has not yet reorganized around.

04 Build documentation infrastructure, not documentation artifacts.

Regulators across both sectors are converging on the same documentation expectation: institutions should be able to produce, on short notice, the full lifecycle record of any material AI system. That record includes the intended-use statement, training data lineage, validation evidence, production monitoring logs, incident history, change management records, and third-party vendor documentation. The first artifact a supervisor now requests is no longer the model itself but the record of how it was developed, validated, deployed, monitored, and updated.

An institution that treats documentation as a byproduct of its validation cycle, generated at the end of each review, cannot produce it on examination timelines; an institution that treats documentation as real-time infrastructure, generated automatically as systems run, can. The practical gap between the two is both a platform choice and an accountability choice: which governance technology sits between the AI system and the supervisor, and who inside the institution owns the freshness, completeness, and auditability of the record.

Neither choice can be delegated to a vendor, and both have roughly eighteen-month consequences. An institution making those choices in 2026 will enter the 2027 examination cycle with evidence infrastructure already stress-tested; one still in mid-build by then becomes visible to supervisors at the worst possible moment, with a documentation project in motion rather than complete. In a regime that evaluates AI governance maturity through examination-ready artifacts rather than policy statements, that timing gap is what distinguishes demonstrated governance from described governance.

None of these postures is novel in isolation. Each has an analog in an existing institutional discipline: model risk in finance, post-market surveillance in medical devices, three-lines-of-defense in risk management, examination-ready documentation across every regulated sector. What is new is the requirement that they operate together, continuously, across a category of system that is still in active deployment. Institutions that rebuild around that requirement will set the template for the next decade of regulated AI. The institutions that delay will be the case studies that motivate the next round of supervisory action.

METHODOLOGY, SOURCES AND VERSION HISTORY

This second edition incorporates regulatory developments through 5 August 2026 and revises several passages concerning U.S. model risk guidance, EU and state implementation timelines, and the evidence base on AI in prior authorization. Three instruments material to the report's analysis — the interagency revision of U.S. model risk guidance, the EU's deferral of high-risk obligations, and Colorado's replacement of its AI Act — were issued or agreed in April and May 2026 and were not reflected in the first edition. The report's central argument and comparative framework are unchanged, and in the case of U.S. model risk guidance the revision strengthens them.

This Horizon Scan draws on primary regulatory instruments, peer-reviewed and policy literature, industry survey data, and public statements by named officials and executives. Where a quantitative claim appears in the body, it is tied to the specific source rather than to the bibliography alone.

Regulatory research for the first edition closed on 30 April 2026. Research for this second edition closed on 5 August 2026. AI regulation in both sectors is moving faster than a quarterly publication cycle; readers should treat every date in this report as current to the second cutoff and verify before relying on it. Where an instrument was in consultation or pending adoption at the cutoff, that status is stated in the text.

First edition published 13 May 2026.

Second edition published August 2026. Revised the U.S. model risk analysis to reflect SR 26-2, which superseded SR 11-7 on 17 April 2026 and placed generative and agentic AI outside its scope; updated EU AI Act implementation dates following Regulation (EU) 2026/1744; rewrote the Colorado section following the repeal and reenactment of the state's AI Act; corrected prior authorization adoption figures and the attribution of Medicare Advantage overturn data; clarified the issuance and re-issuance dates of the FDA clinical decision support guidance; corrected a bibliography link; and added this methodology and source-cutoff statement.

Corrections may be sent to corrections@horizonsearch.org. Material corrections are recorded here.

§ • REFERENCES

Bibliography

Sources are organized by chapter. URLs shortened for print.

FINANCIAL SERVICES

34 ENTRIES

Accenture. (2026, January 14). *Top banking trends for 2026: Unconstrained banking — A new age of possibility*. Accenture. [accenture.com/us-en/insights/banking/accenture-banking-trends-2026](https://www.accenture.com/us-en/insights/banking/accenture-banking-trends-2026)

American Bankers Association. (2025, December 10). Statement for the record: From principles to policy — Enabling 21st century AI innovation in financial services [Submission to House Financial Services Committee]. [aba.com/advocacy/policy-analysis/sfr-on-ai-innovation](https://www.aba.com/advocacy/policy-analysis/sfr-on-ai-innovation)

- Amazon Web Services.** (n.d.). Best practices for incident response in agentic AI [Prescriptive guidance]. AWS Documentation. docs.aws.amazon.com
- Bank for International Settlements, Consultative Group on Risk Management.** (2025, January). Governance of AI adoption in central banks. BIS Representative Office for the Americas. bis.org/publ/othp90.pdf
- Riemer, S., Pauly, M., Poddar, B., Sack, D., Elkin, K., & Alwell, K.** (2026, March 9). How retail banks can put AI agents to work. Boston Consulting Group. [Joint BCG/OpenAI initiative.] bcg.com/publications/2026/how-retail-banks-can-put-agentic-ai-to-work
- CFA Institute.** (2026). AI is reshaping bank risk. *Enterprising Investor*. rpc.cfainstitute.org
- Cloud Security Alliance.** (n.d.). Agentic AI red teaming guide. cloudsecurityalliance.org/artifacts/agentic-ai-red-teaming-guide
- Consumer Bankers Association.** (2025, October 27). Comment letter to the White House Office of Science and Technology Policy (Docket OSTP-TECH-2025-0067). consumerbankers.com
- Crosman, P.** (2026, February 13). Goldman equips AI agents to do trade accounting, onboarding. *American Banker*. americanbanker.com
- European Banking Authority.** (2025, November 21). AI Act: Implications for the EU banking and payments sector [Factsheet]. EBA. eba.europa.eu
- European Commission.** (n.d.). Regulatory framework for artificial intelligence [EU AI Act]. digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
- Federal Reserve Board & Office of the Comptroller of the Currency.** (2011; FDIC adopted 2017; superseded 2026). SR 11-7: Guidance on model risk management. Board of Governors of the Federal Reserve System. federalreserve.gov/supervisionreg/srletters/sr1107.htm
- Federal Reserve Board, Federal Deposit Insurance Corporation & Office of the Comptroller of the Currency.** (2026, April 17). SR 26-2: Revised guidance on model risk management. Supersedes SR 11-7 and SR 21-8. federalreserve.gov/supervisionreg/srletters/SR2602.htm · Attachment (PDF)
- European Parliament & Council of the European Union.** (2026, July 24). Regulation (EU) 2026/1744 (Digital Omnibus on AI), amending Regulation (EU) 2024/1689. *Official Journal of the European Union*. Entered into force 27 July 2026; defers Annex III high-risk obligations to 2 December 2027 and Annex I to 2 August 2028.
- Financial Stability Board.** (2024, November). The financial stability implications of artificial intelligence. FSB (hosted at BIS). bis.org/fsi/fsisummaries/exsum_23904.htm
- Global Association of Risk Professionals.** (2026, February). Black boxes in boardrooms: AI oversight and board governance. GARP. gapr.org/risk-intelligence/culture-governance/black-boxes-boardrooms-260220
- Georgetown Center for Security and Emerging Technology.** (2025, October). Narayanan, M., et al. Mapping the AI governance landscape: Pilot test and update. CSET. cset.georgetown.edu
- Zou, A., Lin, M., Jones, E., Nowak, M., Dziemian, M., Winter, N., et al.** (2025, July 28). *Security challenges in AI agent deployment: Insights from a large-scale public competition* (arXiv:2507.20526) [Preprint]. arXiv. arxiv.org/abs/2507.20526
- Gautam, S.** (2026, March 24). Advancing AI red teaming: Agentic taxonomies and governance-ready reporting. HackerOne. hackerone.com/blog/ai-red-teaming-agentic-taxonomies-governance-ready-reporting
- K&L Gates.** (2026, January). EU and Luxembourg update on the European harmonised rules on artificial intelligence. klgates.com
- KPMG Singapore.** (2025, November). MAS guidelines: Artificial intelligence risk management (AIRG) 2025. kpmg.com/sg

- McKinsey & Company.** (2026, March). State of AI trust in 2026: Shifting to the agentic era. McKinsey Tech Forward. mckinsey.com/capabilities/tech-and-ai
- Monetary Authority of Singapore.** (2025, November 13). Consultation paper on proposed guidelines on AI risk management for financial institutions. MAS. mas.gov.sg/publications/consultations/2025
- Monetary Authority of Singapore.** (2026, March 20). MAS partners industry to develop AI risk management toolkit for the financial sector. MAS. mas.gov.sg/news/media-releases/2026
- National Institute of Standards and Technology.** (n.d.). AI risk management framework. U.S. Department of Commerce. nist.gov/itl/ai-risk-management-framework
- National Institute of Standards and Technology.** (2025, December 16). Cybersecurity framework profile for artificial intelligence (NIST IR 8596 iprd) [Preliminary draft]. U.S. Department of Commerce. csrc.nist.gov/pubs/ir/8596/iprd
- Open Worldwide Application Security Project.** (n.d.). OWASP Top 10 for Agentic Applications. owasp.org
- Pertama Partners.** (n.d.). ASEAN AI governance and ethics guide for Southeast Asia. pertamapartners.com
- Pertama Partners.** (n.d.). Vietnam AI Law 134/2025 compliance guide. pertamapartners.com/insights/vietnam-ai-law-134-2025-compliance-guide
- Ryseff, J., De Bruhl, B. F., & Newberry, S. J.** (2024). The root causes of failure for artificial intelligence projects and how they can succeed: Avoiding the anti-patterns of AI (RR-A2680-1). RAND Corporation. rand.org/pubs/research_reports/RRA2680-1.html
- Sharma, K.** (2026, February 27). SR 11-7 in the age of agentic AI: Where the framework holds and where it strains. Global Association of Risk Professionals. garp.org/risk-intelligence/operational/sr-11-7-age-agentic-ai-260227
- Son, H.** (2026, February 6). Goldman Sachs taps Anthropic's Claude to automate accounting, compliance roles. CNBC. cnbc.com/2026/02/06/anthropic-goldman-sachs-ai-model-accounting.html
- U.S. Department of the Treasury.** (2026, February). Financial services sector AI risk management framework. home.treasury.gov/news/press-releases/sb0401
- U.S. Government Accountability Office.** (2025, May). Artificial intelligence: Use and oversight in financial services (GAO-25-107197). gao.gov/assets/gao-25-107197.pdf

HEALTHCARE

20 ENTRIES

- Anderson, M.** (2026, March 25). What's the state of healthcare AI regulation? *Healthcare Brew*. healthcare-brew.com
- Arnall Golden Gregory LLP.** (2026, January 23). Texas requires providers to make disclosures to patients related to the use of AI in healthcare services. agg.com
- Baker Botts LLP.** (2026, January). U.S. artificial intelligence law update: Navigating the evolving state and federal regulatory landscape. bakerbotts.com/thought-leadership/publications/2026/january/us-ai-law-update
- California Legislature.** (2024). SB 1120: Health care coverage: Utilization review: Artificial intelligence (Physicians Make Decisions Act). California Legislative Information. leginfo.legislature.ca.gov
- Covington & Burling LLP.** (2026, January 6). 5 key takeaways from FDA's revised clinical decision support (CDS) software guidance. cov.com
- Emerson, J.** (2026, March 19). Louisiana bill would require human review of AI-driven coverage denials. *Becker's Payer Issues*. beckerspayer.com

- Food and Drug Administration.** (2025, January 6). Artificial intelligence-enabled device software functions: Lifecycle management and marketing submission recommendations [Draft guidance]. U.S. Department of Health and Human Services. [fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-software-medical-device](https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-software-medical-device)
- Food and Drug Administration.** (2026, January 6; re-issued January 29). Clinical decision support software: Guidance for industry and Food and Drug Administration staff. The January 29 document supersedes the January 6 issuance. U.S. Department of Health and Human Services. [fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software](https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software)
- Illinois Department of Financial and Professional Regulation.** (2025, August). Gov. Pritzker signs legislation prohibiting AI therapy in Illinois. State of Illinois. idfpr.illinois.gov
- Colorado General Assembly.** (2026). SB 26-189: Automated decision-making technology in consequential decisions. Signed 14 May 2026; repeals and reenacts SB 24-205 (Colorado AI Act); effective 1 January 2027.
- National Association of Insurance Commissioners.** (2025, May). *Health insurance artificial intelligence/machine learning survey report*. NAIC. Survey of 93 insurers across 16 states; underlying source for the adoption figures reported in Mello et al. (2026).
- Louisiana Legislature.** (2026). SB 246: Establishes requirements for health insurance issuers using artificial intelligence or automated decision systems, 2026 Regular Session. Louisiana Senate. Accessed May 13, 2026. legis.la.gov/legis/BillInfo.aspx?s=26RS&b=SB246
- Manatt Health.** (2025, December 16). 2025 AI policy and health care priorities in review: What was hot; what was not. Manatt, Phelps & Phillips, LLP.
- Manatt Health.** (2026, April 9). *Health AI Policy Tracker*. Manatt, Phelps & Phillips, LLP.
- Mello, M. M., Trotsyuk, A. A., Djiberou Mahamadou, A. J., & Char, D.** (2026). The AI arms race in health insurance utilization review: Promises of efficiency and risks of supercharged flaws. *Health Affairs*, 45(1). doi.org/10.1377/hlthaff.2025.00897
- Newsom, G.** (2025, December 31). New in 2026: California laws taking effect in the new year. Office of the Governor of California. gov.ca.gov/2025/12/31/new-in-2026-california-laws-taking-effect-in-the-new-year
- Ohio House of Representatives.** (2026). Lawmakers advance new limits on mental health chatbots [News coverage via Rep. Christine Cockley]. ohiohouse.gov
- Stanford University.** (2026, January 6). AI-driven insurance decisions raise concerns about human oversight. *Stanford Report*. news.stanford.edu/stories/2026/01/ai-algorithms-health-insurance-care-risks-research
- Stefanescu, V.** (2025, April 8). As FDA slashes workforce, number of new medical devices reaching the public has fallen. *The Minnesota Star Tribune*. startribune.com/fda-layoffs/601322062
- White & Case LLP.** (2026, January 15). State AI laws under federal scrutiny: Key takeaways from the executive order establishing federal AI policy framework. whitecase.com

§ · COLOPHON

This report was set in Source Serif 4 and Inter, both designed under Adobe's open-source typography programme. Body text is set 11 on 16.5 point, justified, with hyphenation enabled. Headings are set in Source Serif 4 at multiple weights.

AUTHORS AND AFFILIATIONS

Cynthia Chen, Georgetown University

Ashwin Telang, Northwestern University

Hernando Liu, Northwestern Medill School of Journalism

Gloria Chen, New York University

David Lovejoy, Founding Executive Director, Horizon Search Institute (Editor)

PUBLICATION

Published by Horizon Search Institute, a Delaware nonprofit corporation (EIN 42-1954110). The Institute publishes the Horizon Scan series quarterly as part of Pillar I: Responsible AI. First edition published 13 May 2026; second edition published August 2026.

HOW TO CITE

Chen, C., Telang, A., Liu, H., & Chen, G. (2026). *AI governance in regulated industries* (2nd ed.; D. Lovejoy, Ed.; Horizon Scan No. 001). Horizon Search Institute.

CONTACT

Horizon Search Institute — horizonsearch.org