

When AI Possibilities Create Cyber Risk: What Healthcare and Device Counsel Should Be Doing Now

Overview

As AI-related capabilities are increasingly integrated into medical devices and healthcare operations, they offer important opportunities to improve patient care and outcomes, ranging from AI-assisted diagnostics and imaging tools to AI-facilitated patient interactions and record management. However, with this new advancement comes increasing cybersecurity risk and complexity in maintaining compliance with existing regulatory frameworks. IBM's annual Cost of a Data Breach Report released on July 29, 2026, noted that for the fourteenth consecutive year, the healthcare industry experienced the highest average cost per data breach of any industry (\$6.64 million). Industry reporting has also uncovered increased use of AI by threat actors throughout the lifecycle of a cybersecurity incident, increasing the speed and impact of such incidents, with IBM finding that AI-driven incidents increased by 56% year over year and added an average cost of \$1 million per breach.

The nature of these attacks is also evolving. A 2026 report by RunSafe Security based on a survey of nearly 600 global healthcare professionals familiar with their organizations' cybersecurity practices found that while malware and network intrusions remain common, remote access exploitation has quickly emerged as a major threat, accounting for 38% of incidents. In response, 77% of healthcare organizations reported increasing their cybersecurity resources in 2026.

The healthcare industry is also mobilizing to collectively address these risks. In August 2026, the Coalition for Health AI (CHAI) announced the creation of a dedicated work group focused on countering cybersecurity risks posed by frontier AI models in healthcare. The creation of the work group signals growing recognition within the industry of the unique security challenges presented by advanced AI systems that require coordinated, sector-specific attention.

Entities integrating AI into healthcare and medical devices must not only contend with these risks, but also map AI's functionalities onto stringent security frameworks and regulatory expectations from a wide variety of sources, including FDA guidance, the Health Insurance Portability and Accountability Act ("HIPAA") requirements, and a growing number of state regulations, such as the Colorado AI Act and California's S.B. 53. In many instances, these frameworks predate the recent exponential expansion in AI capabilities, and regulatory efforts to update frameworks and guidance lag behind current

Authors



Vinita Kailasanath

Transactions
Partner, Palo Alto
+1 650 231 5875
VKailasanath@cov.com



Caleb Skeath

Cybersecurity
Partner, Washington
+1 202 662 5119
CSkeath@cov.com



Christina Kuhn

Life Sciences
Special Counsel, Washington
+1 202 662 5653
CKuhn@cov.com



Elizabeth Brim

Life Sciences
Associate, Washington
+1 202 662 5850
EBrim@cov.com



AI developments. However, many of these frameworks are built around existing risk-based approaches that provide both regulators and regulated businesses with tools to adapt to the evolving AI risk landscape. This article highlights key aspects of the current regulatory environment, discusses how AI capabilities are changing the contracting and compliance playbook, and offers proactive approaches that healthcare organizations and device manufacturers can take to stay ahead of cybersecurity risks.

The Current Regulatory Picture

Although AI's capabilities have created new possibilities (and risks) for healthcare and medical devices, they do so against an existing backdrop of stringent security-focused regulations applicable to both devices and data. Over the past several years, both FDA and HHS have continued to update their regulations and guidance regarding cybersecurity considerations, including by implementing changes in response to increased AI usage in this sector. These frameworks document important compliance and risk management considerations for entities integrating AI into healthcare and medical devices.

FDA released its draft guidance on “Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations” in early 2025. Notably, the draft guidance included an entire section on AI-related cybersecurity risks for medical devices, including data poisoning by injecting inauthentic or maliciously modified data, data leakage if sensitive training or inference data is exposed, and model evasion if attackers intentionally craft or modify inputs to deceive the model, leading to incorrect outputs. The draft AI guidance cross-references FDA's then-current 2023 premarket

cybersecurity guidance, encouraging organizations to integrate the “unique considerations related to AI cybersecurity” into the cybersecurity risk assessment and risk management processes outlined in the 2023 guidance. While FDA has not yet released an updated or finalized version of the 2025 draft AI guidance, it released an updated final version of its premarket cybersecurity guidance earlier in 2026, underscoring the priority FDA places on cybersecurity as a core component of medical device safety.

In addition to its draft AI guidance, FDA has separately finalized guidance on predetermined change control plans (PCCPs) for AI-enabled device software functions. PCCPs permit certain pre-specified post-market software modifications, including some model updates, to be implemented without a new marketing submission, provided that the changes fall within an FDA-authorized plan that includes defined validation and risk-control measures.

AI-enabled devices and related systems might also process electronic protected health information (“ePHI”), creating obligations under HIPAA for covered entities and business associates (“regulated entities”) to secure this data in addition to FDA requirements applicable to the device itself. The HIPAA Security Rule requires regulated entities to implement safeguards, based on a risk-based approach, to protect the confidentiality, integrity and availability of ePHI. The U.S. Department of Health and Human Services (“HHS”)’s proposed HIPAA Security Rule update, released in 2025 but not yet finalized, would strengthen the Security Rule to add additional and more specific requirements around risk assessments, asset inventories, patch management, and other controls. In its Notice of Proposed Rulemaking (“NPRM”), HHS explicitly noted that part of the rationale for proposing these changes to the Security Rule was to “clarify our expectations for when and how regulated



entities need to consider, prepare for, and address” the risks posed by increased AI adoption and capabilities. Notably, even in the absence of finalized updates to the Security Rule, HHS explicitly stated in its NPRM that “as technology such as AI evolves, the Department would expect a regulated entity to perform a risk analysis [under existing Security Rule requirements] to consider the effects of such changes on the confidentiality, integrity, and availability of ePHI.”

While several pieces of the regulatory puzzle remain in process, the existing guidance and regulations emphasize risk analysis, risk management and reasonable safeguards that provide a basis for regulators to scrutinize AI-related security gaps. Industry-led initiatives, such as CHAI’s frontier AI cybersecurity work group, reflect a parallel effort to develop consensus-based best practices that may inform future regulatory standards. Below, we highlight some practical contracting and compliance measures that healthcare organizations and device manufacturers can take to translate this regulatory backdrop into day-to-day risk management.

Why AI’s Capabilities Change the Compliance Playbook

While healthcare and medical devices have always operated within a highly regulated environment, integrating AI into that environment complicates the task of demonstrating that compliance has been maintained over the potentially long life of a device or system. In light of the rapidly evolving AI capabilities and their growing adoption in the healthcare and medical sectors, proactive, rather than solely reactive, compliance measures can reduce legal and business risk while easing adherence to evolving regulatory requirements and expectations. Organizations that are integrating or utilizing AI

should recognize that the risk landscape has shifted, adjust their risk-based decision-making accordingly, and document how they are addressing the practical complexities of AI-enabled technology.

As an initial matter, lifecycle-based risk assessment should replace point-in-time approaches to compliance and risk management. A model updated through an authorized change-control pathway, retrained on new data or otherwise modified to improve performance is a moving target. Similarly, HHS has made clear that assessing and managing AI-related risks to ePHI is an ongoing process as AI use cases and capabilities evolve. Risk assessments, validation records, performance monitoring and security reviews should therefore be tied to each meaningful model update as well as post-market experience for devices, or revisited as use cases and technological changes warrant for ePHI, and not necessarily limited to an annual compliance refresh.

Second, incident- response and crisis management processes should account for AI-related events or crises that can have a significant impact on patients, data, or devices but might not fit into traditional “data breach” or “cybersecurity incident” categories. Data poisoning, model drift attributable to adversarial input, training-data leakage, unexplained model performance degradation and integrity issues may not begin with a confirmed intrusion or even be linked to a malicious action. Existing policies and contracts that define a security incident only by reference to unauthorized access, or are framed to only respond to malicious actors, might be too narrow. Counsel should consider whether internal escalation triggers, vendor notification provisions and crisis management playbooks sufficiently capture anomalous model behavior and potential integrity impacts, among other scenarios that could give rise to events that create meaningful legal risk and require a significant response.



Third, device maintenance and patching may become more complex. AI can increase the speed at which vulnerabilities are discovered and exploited, but AI-enabled devices may also be difficult to update for a combination of clinical, technical, regulatory and operational reasons. Organizations should assess whether their patch management processes, vulnerability disclosure intake, software bill of materials practices and change-control procedures are equipped to address both conventional software vulnerabilities and model-related issues over the full device lifecycle.

Finally, AI risk should be addressed before rather than after contract signature or transaction closing. Health systems and device manufacturers may find themselves bearing operational, regulatory and patient-safety risk that is only partially controlled by vendors. Boilerplate cybersecurity language, standard SaaS liability caps and generic indemnities are unlikely to allocate AI-enabled device risk in a way that reflects clinical consequences or evolving regulatory expectations.

Practical Steps to Update Contracts and Compliance Programs for AI


For legal counsel, protecting against these evolving risks can be challenging, particularly when entering into longer-term agreements or transactions. We have listed below practical steps for organizations deploying, manufacturing or acquiring AI-enabled healthcare technologies to consider in addressing these risks.

- **Revisit risk assessments.** Update risk analyses, device cybersecurity assessments and enterprise risk reviews to account for AI-specific vectors, including training-data integrity, model performance monitoring, model inversion risk, adversarial input, data provenance and post-market update processes. For HIPAA-regulated entities, this work is valuable regardless of whether HHS finalizes the proposed Security Rule in its current form because existing regulations already require an accurate and thorough assessment of risks to ePHI.
- **Strengthen secure development and model governance.** Integrate security into the development and training of AI-enabled device functions, including controls around training data, validation data, documentation, model evaluation, vulnerability handling and post-market monitoring. Entities procuring these tools should interrogate how the potential or existing vendor manages model updates, validates changes, monitors performance and detects anomalous behavior. Organizations may also benefit from monitoring and participating in industry-led initiatives to develop shared frameworks for managing risks unique to advanced AI models in healthcare settings.
- **Build AI diligence into procurement and transactions.** Incorporate questions about adversarial testing, model provenance, training and validation data, bias and performance evaluations, software and machine-learning bills of materials, update governance and prior security or performance incidents into procurement and M&A diligence. If the device relies on a predetermined change



control plan or similar update process, ongoing delivery of updated documentation should be considered a contractual requirement.

- **Plan for increased vulnerability volume and faster response expectations.** While advancing AI capabilities have already increased the speed and volume of identified vulnerabilities, increased AI use by threat actors might also shorten the time to patch these vulnerabilities before they are exploited. To prepare for this new reality, organizations can evaluate whether they have sufficient triage capacity, documented patch timelines, escalation criteria, and governance if a fix requires device validation, regulatory analysis and/or coordination across clinical operations.
- **Update and test incident-response and resilience planning.** Test incident-response plans through tabletop exercises against AI-specific scenarios, including suspected data poisoning, unexplained model degradation, loss of availability of AI-enabled clinical tools, integrity concerns affecting device outputs, and uncertainty about whether patient-facing recommendations are reliable. Business continuity, disaster recovery and clinical downtime procedures should likewise account for dependencies on AI-enabled systems. Organizations should consider having outside counsel with regulatory expertise involved to structure tailored, realistic tabletop exercises and incident simulations in a manner that preserves attorney-client privilege over resulting findings and recommendations.
- **Revise contract provisions.** Define AI-related security incidents in vendor agreements broadly enough to include anomalous model behavior, data poisoning, training-data leakage, adversarial manipulation, model drift attributable to malicious or improper input, and integrity issues affecting outputs. Notification triggers are particularly helpful when they extend beyond simply confirmed unauthorized access. Agreements should also consider addressing model update processes, documentation delivery, audit rights, vulnerability disclosure, cooperation during incident response, regulatory support and preservation of evidence.
- **Reassess liability and insurance.** Standard liability caps calibrated to license fees may not reflect the clinical, operational and regulatory consequences of AI-enabled device failures or security incidents. Organizations should consider whether the combination of caps, exclusions, indemnities and cyber insurance provisions (of both the vendor and the customer) appropriately address algorithmic error, patient harm, ePHI exposure, business interruption, regulatory investigations and the full range of potential third-party claims. Organizations should also consider periodically reviewing vendor certificates of insurance and coverage scope rather than assuming they are adequate.



AI-enabled healthcare and medical technologies offer substantial promise, but they also require legal and compliance teams to rethink familiar assumptions about cybersecurity, device maintenance and vendor risk mitigation. The organizations best positioned for the next phase of AI adoption will be those that treat AI cybersecurity as an ongoing governance and contracting discipline now, rather than waiting for regulators, litigants or threat actors to define the standard after an incident occurs.

Covington summer associate Kimberly Collins contributed to the article.