From Flexibility to Operational Governance: The HIPAA Security Rule's Quiet Transformation

Lynn Sessions / Eric D. Morris

Lynn Sessions is a partner at BakerHostetler. Ms. Sessions leads the Healthcare Privacy and Compliance team and serves as the Houston office lead in the Digital Assets and Data Management Practice Group, as well as national co-lead of the Healthcare Industry team, demonstrating a career of advising healthcare industry clients across multiple areas of the law.

Eric D. Morris is an associate at BakerHostetler. Mr. Morris concentrates his growing practice on healthcare privacy and compliance. Focusing on incident response, he works with clients to develop a strategy once an incident has occurred, including formulating an initial response, conducting an investigation, and meeting all necessary state and federal regulatory requirements.

The Health Insurance Portability and Accountability Act (HIPAA) Security Rule (the Rule) was drafted to be indifferent to any particular technology. When the Department of Health and Human Services adopted the standards for protecting electronic protected health information (ePHI) in 2003, it deliberately declined to prescribe particular products, protocols, or controls.[1] Instead, it articulated objectives, namely the confidentiality, integrity, and availability of ePHI, and directed regulated entities to implement "reasonable and appropriate" administrative, physical, and technical safeguards calibrated to their size, complexity, capabilities, and risk environment. The Rule's implementation specifications were divided into those that were "required" and those that were merely "addressable,"[2] a structure intended to let a solo physician practice and a large hospital system comply with the same regulations without being held to identical technical mandates.

The decision to be indifferent to particular technical requirements has allowed the Rule to age gracefully. If the drafters had chosen a prescriptive, control-specific rule at the time, then the Rule would have been obsolete within a few years. Fortunately, the drafters chose a principles-based Rule that has remained viable as healthcare and its supporting technology have evolved. The industry has transitioned from on-premises servers to cloud servers, from desktops to mobile endpoints, from paper charts to electronic health records, and from local networks to API-connected ecosystems. The flexibility intended by the drafters allowed the Rule to survive changes in technology that would have made a more rigid rule obsolete.

But that flexibility is a double-edged sword. The discretion that let the Rule accommodate new technology also opened gaps between the Rule's minimum expectations and the security practices that the modern threat landscape demands. The healthcare industry is among the sectors most frequently targeted by threat actors,

and ransomware has evolved from a nuisance into an operational threat, while identity-based intrusions have replaced malware as the primary initial-access vector. Against that backdrop, "reasonable and appropriate" has evolved to mean more than it did at the time of the Rule's drafting, and the focus of regulators has shifted from whether an organization has policies to whether its safeguards meet the demands of the moment.

The Rule is evolving from a flexible, principles-based compliance framework into an operational model of cybersecurity governance. The proposed amendments, now working their way through the rulemaking process, most notably the elimination of the addressable/required distinction and the imposition of explicit technical mandates, are visible expressions of that shift.[3] For the most part, the amendments would codify expectations that the Office for Civil Rights (OCR) enforcement, contemporary cybersecurity standards, and operational necessity have established over the years. Organizations that understand the Rule as a governance obligation rather than a documentation exercise are, in effect, complying with the regime that is coming. The sections that follow trace how that transformation has unfolded and what it requires of covered entities in practice.

The Evolution of a Principles-Based Rule

The genius of the Rule drafters allows for a technology-neutral regulation that governs outcomes rather than specific mechanisms. The Rule never required a firewall or an intrusion-detection system by name; however, it did require that access to ePHI be controlled, that activity be recorded and examined, and that transmissions be protected. This abstraction is precisely what allowed the Rule to remain applicable as mechanisms changed and evolved over the years. It is difficult to overstate how much has changed. The healthcare industry has seen hyper-adoption of cloud computing and software as a service, and now artificial intelligence. The change in the technological landscape has transitioned ePHI from systems the healthcare entity owns to those it rents, and often without a clear understanding of the underlying infrastructure. Moreover, remote work has blurred traditional network boundaries, interoperability requirements have expanded the number of entities with legitimate access to patient data, and threat actors have become increasingly sophisticated, operating ransomware-as-a-service enterprises complete with affiliates, negotiators, and data leak sites.

Principles-based regulation depends on an evolving understanding of what the principles require, and for many years that understanding was provided by OCR informally, through enforcement, guidance, and prevailing industry standards rather than through the regulatory text. This enforcement is the manner by which a static rule became a changing standard. An implementation specification that was genuinely optional at adoption could become, through a decade of resolution agreements treating its absence as a violation, a *de facto* baseline. For example, encryption remains an "addressable" requirement under the Rule; however, its consistent role in OCR settlements involving lost or stolen unencrypted devices made it a practical expectation—*i.e.*, the standard of care, rather than a discretionary choice—highlighting how enforcement standards have advanced despite the regulation itself remaining unchanged.

The gradual drift between text and expectation is the pressure the proposed amendments are designed to relieve. By converting the implementation specifications from addressable to required, subject to specific and limited exceptions; mandating written documentation across the board; and naming specific controls—among them encryption of ePHI at rest and in transit, multifactor authentication

(MFA), network segmentation, technology asset inventories and network maps, vulnerability-scanning and penetration-testing cadences, and bounded restoration timelines—the amendments would close the gap between what the Rule says and what compliance looks like. The significance of this move is easy to misread. It is less a new set of obligations than a formalization of expectations that mature organizations already meet, or are expected to meet, and that enforcement has already imposed. The organizations most exposed by the change are those that have treated the flexibility of the current Rule as permission to do less, rather than as latitude to tailor a rigorous program to their environment.

Understanding the shift reframes the compliance question. The issue is no longer whether an organization can point to its policy addressing each standard but whether it can demonstrate that safeguards are implemented, governed, and effective. That reframing runs through every subject that follows, beginning with the discipline that OCR treats as the foundation of the entire regime: risk analysis.

Risk Analysis as Enterprise Cybersecurity Governance

If the Rule has a keystone, it is the requirement to conduct an accurate and thorough assessment of the risks to and vulnerabilities of the confidentiality, integrity, and availability of ePHI and to implement measures sufficient to reduce those exposures to a reasonable level.[4] OCR has said, in guidance and in enforcement, that the risk analysis is the foundation on which the rest of Security Rule compliance is built, and the agency's enforcement record bears this out with striking consistency. Deficient risk analysis is among the most frequently cited findings in OCR investigations, technical assistance, and resolution agreements, and it is the express focus of the agency's Risk Analysis Initiative, under which OCR has pursued a series of settlements based on the failure to conduct an adequate enterprise-wide analysis.[5]

A recent enforcement action illustrates both the substance and the timing that OCR expects. In that matter, OCR determined that a business associate had failed to conduct a sufficient risk analysis or develop a corresponding risk management plan before a ransomware incident affecting ePHI, resulting in a monetary settlement and a multiyear corrective action plan.[6] The recurring lesson is not merely that risk analysis is required but that it must be performed proactively and maintained continuously, not reconstructed after an incident to explain what went wrong. A risk analysis produced in the aftermath of a breach is evidence of the deficiency, not a cure for it.

The deeper point is that risk analysis is not a compliance artifact but instead the governance engine of the entire security program. It is the process through which an organization develops and maintains an accurate asset inventory; maps how ePHI is created, received, maintained, and transmitted across its systems and its vendors; models the threats applicable to its environment; and prioritizes remediation according to actual exposure rather than intuition. Every downstream safeguard decision—whether about what to encrypt, where to segment, which accounts warrant privileged-access controls, or how quickly systems must be restored—should trace back to findings that the risk analysis produced. An organization that cannot say with confidence where its ePHI resides cannot rationally decide how to protect it and cannot defend that decision when a regulator asks.

This is where the convergence between the legal standard and prevailing technical frameworks becomes concrete. NIST Special Publication 800-66, Revision 2, the successor to the long-standing HIPAA implementation guide, translates the Rule's risk-analysis obligation into a structured methodology and maps each

standard and implementation specification to the subcategories of the NIST Cybersecurity Framework and to the control catalog of NIST SP 800-53, Revision 5.[7] OCR references SP 800-66 in its risk-analysis guidance, and although the publication is not mandatory, aligning with it offers a defensible, recognized framework for demonstrating what "accurate and thorough" means. The 2024 release of Version 2.0 of the Cybersecurity Framework reinforces the same trajectory: its newly elevated "Govern" function makes explicit that cybersecurity risk is an enterprise governance responsibility to be integrated with organizational strategy and overseen by leadership, not a technical matter delegated wholesale to IT.[8] Read together, these frameworks describe the same thing the proposed amendments describe, and enforcement already rewards: a continuous, documented, enterprise-wide risk management process rather than a periodic questionnaire.[9]

Framed this way, risk analysis is the point at which the Rule stops being a checklist and becomes a management discipline. It is also the point from which every operational control derives its justification, which is where the analysis turns next.

Operationalizing Security: From Controls to Capacity

A persistent and costly misconception is that acquiring security technology is the same as achieving security. It is not, and the enforcement record is littered with organizations that owned capable tools that were misconfigured, inconsistently deployed, unmonitored, or disconnected from any risk-based rationale. The Rule has never rewarded the mere possession of controls; it requires that safeguards be implemented and, critically, that their effectiveness be maintained over time. The proposed amendments sharpen this expectation by naming specific technical controls, but the controls they name are meaningful only insofar as they function together as an operational capability. It is worth examining how the principal controls interlock, because coherence among them, not the presence of any one, is what identifies a defensible program.

Access control and identity governance have become the center of gravity of healthcare security, a shift driven by a corresponding shift in threat actor behavior. As threat actors increasingly authenticate with valid, compromised credentials rather than deploy detectable malware, the effectiveness of security programs depends less on the network perimeter than on the controls that govern and monitor access to sensitive systems and data. This is why MFA, which the proposed amendments would require for technology assets in a regulated entity's relevant electronic information systems, with limited exceptions, has moved from a recommended practice to a near-universal baseline because it is among the most effective controls against credential-based intrusion.

But MFA is necessary, not sufficient. The principle of least privilege, granting each account only the access its function requires, limits the blast radius when a credential is compromised, and privileged-access management, which brings administrative and service accounts under heightened control and monitoring, addresses the accounts an attacker most wants. Network segmentation, another control the amendments would require, provides the same containment at the architectural level, constraining an intruder's ability to move laterally from an initial foothold to the systems where ePHI lives. These controls express, in current technical vocabulary, the zero-trust premise that no user or device should be implicitly trusted by virtue of its location on the network, a premise that maps naturally onto the Rule's long-standing access-control and person-or-entity-authentication standards.[10]

Encryption belongs in the same integrated frame. Long classified as addressable and long treated in enforcement as effectively expected, encryption of ePHI at rest and in transit does more than protect confidentiality. Where it is implemented consistent with HHS guidance and the decryption key or process is not itself compromised, encrypted data may render the information secured so that its acquisition does not trigger the breach-notification obligation at all.[11] Encryption is therefore not only a protective control but a factor that shapes an organization's legal exposure after an event. The proposed amendments' move to mandate it, with narrow and documented exceptions, would formalize a standard that the interaction of enforcement and breach-notification incentives has already made close to universal in practice.

Detection and response have become increasingly critical as prevention is not guaranteed. The severity of an incident often depends more on how quickly an organization detects, investigates, and contains a threat through disciplined logging and monitoring, meaning the recording of system activity and, crucially, the examination of it. Although the Rule has always required organizations to record and examine such activity, they frequently satisfy only the recording part. Incident severity depends as well on vulnerability management and configuration management—knowing what weaknesses exist across the asset inventory and remediating them on a defined cadence, and maintaining systems in a hardened, known-good state rather than an accreted, undocumented one.

The proposed amendments' specific cadences, including vulnerability scanning at defined intervals, periodic penetration testing, and bounded restoration timelines, are best read not as arbitrary numbers but as an insistence that these capabilities be operational and measurable rather than aspirational. The relevant reference point for the response half of this equation, NIST SP 800-61, Revision 3 (which supersedes the withdrawn Revision 2), now frames incident handling through the functions of the Cybersecurity Framework rather than as a stand-alone linear lifecycle; an organization that has genuinely operationalized this model is positioned to satisfy both regulatory expectation and the practical demands of a live incident.[12]

Prevention still matters, but in the current environment the speed and competence of detection and response increasingly determine how damaging an incident becomes, and the reality that some intrusions will originate outside the organization's own walls is what makes third-party risk the next essential subject.

Third-Party Risk as Enterprise Risk

No development has reshaped the healthcare risk landscape more than the sector's growing dependence on third parties. Modern organizations rely on business associates and their subcontractors for technology hosting, claims processing, revenue-cycle management, analytics, clinical applications, and a widening array of managed services. Each relationship extends the organization's attack surface into infrastructure it does not control, and a covered entity cannot delegate its own responsibility by using a vendor. Since enactment of the 2013 Omnibus Rule, business associates are directly liable for their own Security Rule violations, and a covered entity is not automatically liable for every independent business associate failure. But it must obtain satisfactory assurances through a compliant business associate agreement and remains responsible for its own obligations.[13] Incident-response reports and enforcement matters continue to show that vendor access and vendor security controls can be material sources of healthcare cyber risk and that a third-party's controls do not always match those of the entity whose data it holds. Business associate risk

is not a procurement footnote; it is enterprise risk, and it warrants governance of the same rigor an organization applies to its own operations.

The volume and sensitivity of the data make the scale of the exposure difficult to dismiss. In one recent industry dataset, roughly a quarter of all incidents across sectors were attributable to a vendor, and for healthcare organizations specifically, the proportion was higher still, with more than one-third of incidents originating with a third party.[14] Regulators have responded in kind. OCR has signaled that business associates are an enforcement focus by entering multiple resolution agreements against them. Additionally, the proposed amendments would tighten the regime further by requiring business associates to provide annual written verification, prepared by a subject-matter expert, that they have deployed the required technical safeguards and by mandating that business associate agreements incorporate technical obligations. The trajectory is toward verified, ongoing accountability rather than the contractual assurances that have too often stood in the place of oversight.

The systemic nature of this risk deserves attention because it distinguishes third-party risk from other categories. A single compromised vendor can affect hundreds or thousands of downstream organizations simultaneously, as was seen in the February 2024 ransomware attack on Change Healthcare. In that incident, the concentration of risk was pervasive throughout the healthcare sector. It uniquely demonstrated how the compromise of a single clearinghouse could disrupt claims processing, eligibility verification, and pharmacy operations across the healthcare industry while ultimately implicating the ePHI of reportedly 192.7 million individuals, the largest healthcare data breach on record. It was the most visible of a series of major vendor incidents, not an aberration. The lesson is not that any particular vendor is dangerous but that interdependence is a risk that no single organization's internal controls can fully mitigate.

Managing that risk requires treating vendor oversight as a continuous governance function rather than a point-in-time contracting event. This is what NIST's cybersecurity supply-chain risk management guidance and the expanded supply-chain emphasis of CSF 2.0 address, and their logic aligns directly with risks faced by business associates.[15] Effective programs require diligence that is tailored to the sensitivity and volume of ePHI a vendor will handle, coupled with the establishment of clear security obligations and accountability through audit rights, breach-notification timelines, and subcontractor requirements, while maintaining ongoing oversight throughout the relationship rather than treating a signed contract as the end of the process. Contractual language allocates liability, but it does not reduce risk. As the proposed amendments' annual written verification suggests, the regulatory expectation is converging on exactly this model—third-party risk managed as an enterprise discipline, integrated with the organization's own risk analysis and governance rather than quarantined in a vendor-management silo.

Cybersecurity as Organizational Governance

The through-line connecting risk analysis, operational controls, and third-party oversight is governance—the structures, accountability, and discipline that determine whether safeguards are sustained or allowed to decay. This is the facet in which the Rule's administrative safeguards, long the least glamorous portion of the Rule, have proven the most consequential, and it is the dimension that the elevation of governance in contemporary frameworks has brought into sharp relief. A security program is not a collection of controls; it is a set of ongoing behaviors that require governance.

Workforce conduct remains among the most reliable predictors of security outcomes, a fact that decades of technical advancement have not altered. Phishing, credential misuse, misconfiguration, and misdirected disclosures account for a large share of incidents, and phishing in particular has consistently ranked among the leading root causes across years of incident response reporting, although its precise ranking varies by industry and dataset. The administrative safeguards' emphasis on training reflects this reality, but training that satisfies the current threat environment bears little resemblance to the annual module many entities believe passes for compliance.[16] Effective training is continuous, tailored to an individual's role, and designed to address evolving threats, such as credential harvesting and social engineering, particularly as threat actors increasingly target help desk and support personnel, who can be manipulated into resetting credentials or facilitating unauthorized access; accordingly, such training should be treated as a security control rather than viewed as a compliance exercise.

Documentation is the governance mechanism that makes every other element of a security program demonstrable and, far from being mere bureaucracy, provides the evidence through which an organization proves that its policies, procedures, risk analyses, training records, configuration baselines, and security decisions existed, operated effectively, and reflected reasoned judgment.[17] A well-designed control that cannot be evidenced is, from the regulator's vantage point, difficult to distinguish from a control that was never implemented, because regulators evaluate proof rather than intention. This is why the proposed amendments' move toward comprehensive written documentation is less of a new burden than a formalization of what defensibility has always required. The organizations that will adapt most easily are those that already treat documentation as the connective tissue of the program rather than as after-the-fact paperwork.

Accountability must extend to the top of the organization, and here the regulatory and framework trajectories are unmistakably aligned. The Rule places responsibility for reasonable and appropriate safeguards on the organization itself, and enforcement increasingly reflects an expectation of genuine executive and board-level engagement rather than nominal sign-off. The addition of the "Govern" function as a first-class element of CSF 2.0 codifies the same principle from the technical side. Cybersecurity risk is now understood as a category of enterprise risk that governing boards are expected to oversee with the seriousness they bring to financial risks. Programs that lack leadership buy-in often lack funding, authority, and durability, and the questions boards must be prepared to answer about cybersecurity are becoming more specific and more consequential. A governance structure that cannot answer them is itself a deficiency.

Incident response is where governance is tested under the most demanding conditions and provides the most appropriate illustration that a plan on paper is not a capability, as effective response requires roles, escalation paths, decision-making authority, and communication protocols that enable coordinated execution across legal, privacy, security, compliance, communications, and business leadership.[18] Response windows are often tight enough that improvisation is a liability, as incident-response studies report that the intervals from occurrence to discovery, discovery to containment, and discovery to notification frequently leave little room to establish roles and authority in the moment. Consequently, regular tabletop exercises are essential for identifying gaps such as undefined decision-making authority, untested backups, and unrealistic assumptions before a real incident occurs and for

transforming an incident response plan from an assertion into a true organizational capability. That distinction between improvisation and demonstrated capability is precisely the one toward which the regulatory regime is moving.

Preparing for the Next Phase of Security Rule Compliance

The trends traced through this article point to a conclusion about where Rule compliance is headed. The next phase will place weight on whether an organization's safeguards are operational, documented, and defensible in practice rather than merely in policy. This is not a prediction that depends on the fate of any particular rulemaking but rather is an observation about a direction that enforcement, contemporary standards, and operational reality have already established and that the proposed amendments would formalize rather than invent. Even if a given proposal were delayed, revised, or withdrawn, the underlying expectations would persist, because they are being driven by the threat environment and by OCR's enforcement priorities as much as by regulatory text.

For that reason, the prudent posture is preparation grounded in the program an organization should maintain regardless of the rulemaking's timeline. Much of what the proposed rule seeks to implement (*e.g.*, current and accurate risk analysis, encryption and MFA) constitutes defensible practice under the Rule and reflects the measures OCR rewards.[19] Organizations that build toward those standards convert a future scramble into current risk reduction, and they do so on their own timeline rather than under a compressed schedule that final rules typically impose. Waiting for regulatory certainty before acting inverts the risk calculus, because the exposure the program is meant to address does not wait.

Defensibility is the concept that ties preparation together. The flexibility the Rule still affords has always carried a corresponding obligation to make and document defensible decisions in order to be able to explain why a given safeguard was implemented, tailored, or, in the narrowing set of cases where discretion remains, not adopted. As the Rule moves toward mandatory specifications, the domain of discretion shrinks, but it does not vanish, and the decisions that remain within it become more important to document rather than less. The organizations best positioned for the next phase are those that can show their work, and that can produce, for any element of their program, the risk-based reasoning and the operational evidence behind it.

Conclusion

The HIPAA Security Rule has not been rewritten so much as re-understood. The principles-based architecture, once a grant of latitude, now functions as a demand for operational maturity, and the proposed amendments largely make explicit a standard that enforcement and prevailing practice had already established. The organizations that experience the coming changes as disruptive are, in the main, those that read the Rule's flexibility as license to do less. The organizations that read it as latitude to build rigorous, well-governed programs will find that the regime arriving is one they have substantially already met.

What has changed most fundamentally is the frame. Cybersecurity in healthcare is no longer a discrete privacy or IT compliance obligation to be managed at the edge of the enterprise. It is a governance responsibility that must reach the board; an enterprise risk spanning the organization and its entire vendor ecosystem; and, when ransomware halts scheduling, disables clinical systems, or corrupts a record, a patient safety issue measured in care delayed or degraded. Regulators and standards bodies have moved toward this understanding in parallel, and the direction of travel is not seriously in doubt.

The organizations that will meet the moment treat security not as a state to be certified but as a capability to be sustained. These organizations know where their ePHI lives, govern the vendors that touch it, operationalize the controls that protect it, exercise the plans that respond when those controls fail, and document the reasoning behind every material decision. That is the standard toward which the Rule has been evolving for two decades, and it is the standard the next phase will demand. Building it is no longer a matter of anticipating regulatory change; it is a present obligation of responsible governance, and, increasingly, of responsible patient care.

Endnotes

1. 45 C.F.R. §§ 164.302–164.318 (Security Standards for the Protection of Electronic Protected Health Information). The Rule requires covered entities and business associates to implement safeguards that are "reasonable and appropriate" in light of the entity's size, complexity, and capabilities, and the probability and criticality of potential risks to ePHI. *See* 45 C.F.R. § 164.306(b).
2. 45 C.F.R. § 164.306(d) (distinguishing "required" from "addressable" implementation specifications).
3. HIPAA Security Rule To Strengthen the Cybersecurity of Electronic Protected Health Information, 90 Fed. Reg. 898 (proposed Jan. 6, 2025) (to be codified at 45 C.F.R. pts. 160, 164). Available at *https://www.federalregister.gov/documents/2025/01/06/2024-30983/hipaa-security-rule-to-strengthen-the-cybersecurity-of-electronic-protected-health-information*. As of this writing the proposal remains pending and may be finalized, revised, or withdrawn.
4. 45 C.F.R. § 164.308(a)(1)(ii)(A) (risk analysis), (B) (risk management).
5. U.S. Dep't of Health & Human Servs., Off. for Civ. Rts., HIPAA Security Rule Guidance Material (describing the Risk Analysis Initiative, which focuses select investigations on compliance with the risk analysis provision). *See https://www.hhs.gov/hipaa/for-professionals/security/guidance/index.html*.
6. Press Release, U.S. Dep't of Health & Human Servs., "HHS' Office for Civil Rights Settles HIPAA Ransomware Security Rule Investigation with BST & Co. CPAs, LLP" (Aug. 18, 2025) ($175,000 settlement and two-year corrective action plan). Available at *https://www.hhs.gov/press-room/hhs-ocr-bst-hipaa-settlement.html*.
7. NIST, Special Publication 800-66 Rev. 2, Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule: A Cybersecurity Resource Guide (Feb. 2024). NIST hosts the mapping of HIPAA Security Rule standards and implementation specifications to NIST Cybersecurity Framework subcategories and SP 800-53 Rev. 5 controls in its Cybersecurity and Privacy Reference Tool (CPRT). Available at *https://csrc.nist.gov/pubs/sp/800/66/r2/final*.
8. NIST, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29 (Feb. 26, 2024). Version 2.0 adds "Govern" as a sixth core function alongside Identify, Protect, Detect, Respond, and Recover, and expands the framework's treatment of supply-chain risk. Available at *https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf*.
9. NIST, Special Publication 800-53 Rev. 5, Security and Privacy Controls for Information Systems and Organizations (Sept. 2020, updated Dec. 2020).
10. 45 C.F.R. § 164.312(a)(1) (access control), (b) (audit controls), (d) (person or entity authentication).
11. 45 C.F.R. § 164.312(a)(2)(iv) (encryption and decryption), (e)(2)(ii) (transmission security). Under the breach-notification framework, ePHI rendered unusable, unreadable, or indecipherable to unauthorized persons through encryption consistent with HHS guidance is not "unsecured" PHI. *See* 45 C.F.R. §§ 164.402, 164.404; 74 Fed. Reg. 42,740 (Aug. 24, 2009).
12. NIST, Special Publication 800-61 Rev. 2, Computer Security Incident Handling Guide (Aug. 2012) (describing the incident-response lifecycle of preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity). NIST has since advanced this guidance; *see* NIST SP 800-61 Rev. 3 (2025), which recasts incident response through the CSF 2.0 functions.
13. 45 C.F.R. §§ 164.308(b) (business associate contracts), 164.314(a) (business associate agreement requirements). A covered entity's use of a business associate does not relieve the covered entity of its own Rule obligations.
14. Incident figures in this article are drawn from aggregated healthcare incident-response data reflecting recent matters, in which vendors accounted for a substantial share of incidents and phishing has remained the leading root cause across successive reporting periods. *See https://www.bakerlaw.com/insights/bakerhostetler-releases-2026-data-security-incident-response-report-familiar-threats-new-pressures*.
15. *See* NIST, Special Publication 800-161 Rev. 1, Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations (May 2022); *see also* NIST CSF 2.0, *supra*, at the "Govern" function's Cybersecurity Supply Chain Risk Management category (GV.SC).

16. 45 C.F.R. § 164.308(a)(5) (security awareness and training).
17. 45 C.F.R. § 164.316 (policies, procedures, and documentation requirements; six-year retention).
18. 45 C.F.R. § 164.308(a)(6) (security incident procedures), (7) (contingency plan, including data backup, disaster recovery, and emergency-mode operation planning).
19. For additional practical guidance, *see* U.S. Dep't of Health & Human Servs., 405(d) Program, Health Industry Cybersecurity Practices (HICP), available at *https://405d.hhs.gov/*; and Cybersecurity & Infrastructure Sec. Agency, Cross-Sector Cybersecurity Performance Goals, available at *https://www.cisa.gov/cross-sector-cybersecurity-performance-goals*.

Reprinted from Journal of Health Care Compliance, Volume 28, Number 5, September–October 2026, pages 5–12, 46–47, with permission from CCH and Wolters Kluwer.
For permission to reprint, e-mail permissions@cch.com.