Protecting PHI in the age of AI

How healthtech can innovate securely

WIPFLI

Overview

As AI drives both innovation and increasingly complex cyberthreats, healthtech organizations face growing risks to the security of electronic personal health information (ePHI). A single recent ransomware attack resulted in a PHI data breach that impacted nearly half of all Americans.

Regulations and certification measures are evolving in response. Proposed HIPAA updates introduce new requirements for handling ePHI, and HITRUST's new AI Assurance Program introduces specific controls to assess the security of AI platforms. However, until HIPAA updates pass and until organizations begin achieving standards that demonstrate AI security, the oversight landscape remains uncertain. How can healthtech organizations innovate with AI while protecting sensitive data?

In this e-book, we outline practical steps you can take now to safeguard ePHI, prepare for future regulations and continue driving innovation, including:

- What AI-driven cyberthreat trends you should watch out for.
- What potential federal regulatory changes to monitor and how they could impact healthtech.
- How to build AI security assurance, either through new or existing risk assessments.
- How to build an AI-resilient security culture, from people to best practice controls.

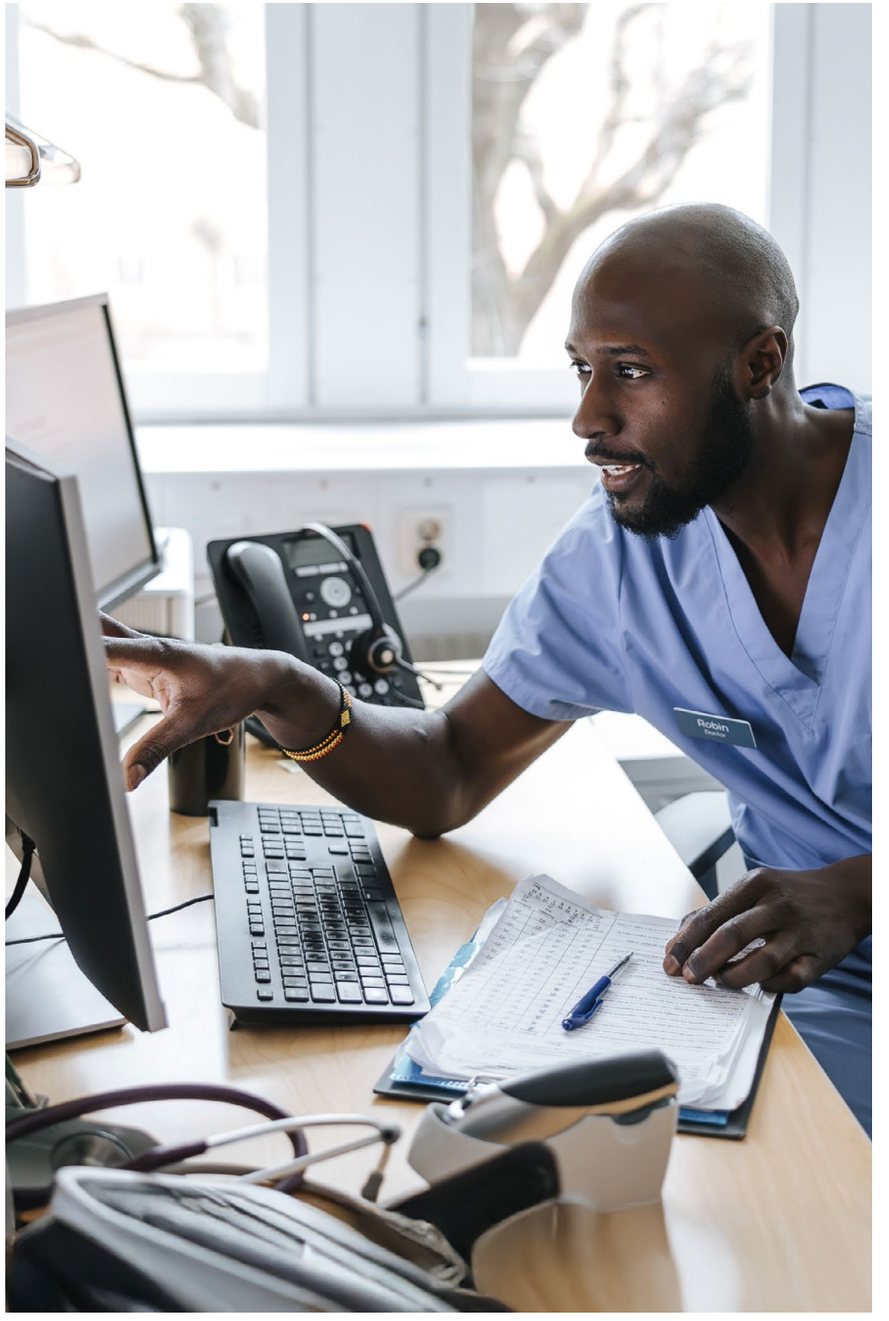

Understanding AI's impact on cybersecurity

In healthtech, protecting PHI is a business imperative for ensuring patient privacy and safeguarding reputational standing.

And the risks have never been higher. With the rise of AI-enhanced cyberattacks and increasing regulatory scrutiny, even a single breach can result in significant financial, legal and operational damage.

Just as AI can accelerate aspects of your business, AI makes it easier for hackers to dig and explore potential victims with significantly less legwork than before. And AI-driven attacks are more adaptive and more difficult for users to detect, making them especially dangerous for industries, like healthtech, that handle highly sensitive information.

High risk, high stakes

$9.77 million

The average cost of a data breach in healthcare in 2024 according to IBM's Cost of a Data Breach Report 2024

190 million

The number of Americans impacted by a PHI data breach from a single recent ransomware attack

588

The number of data breaches reported to the Department of Health and Human Services in 2024 alone

4 AI-advanced cyberthreat trends to watch out for

In Wipfli's work with clients across the healthtech, healthcare and technology sectors, we're seeing an increase in AI-advanced cyberthreats, especially across these four areas.

1. Social engineering

AI is making data mining easier. Threat actors are using highly personalized information to customize phishing attacks that are increasingly difficult to spot, especially in high-pressure, fast-paced work environments. For example, AI can scrape professional bios, social media posts and internal org charts to generate emails that reference real colleagues, job roles or upcoming events, making them appear legitimate and urgent.

2. Cloud security vulnerabilities

Thanks to its incredible processing power, AI can easily find misconfigurations in cloud environments and exploit them. For example, AI can quickly identify improperly configured storage buckets, exposed APIs or unpatched services, and use those weaknesses to gain unauthorized access to sensitive data.

3. Autonomous malware

AI-driven malware can learn from its environment in real time, adapting its tactics to avoid detection and persist longer in systems. For example, it might change file names, shift code behavior or mimic authorized applications.

4. Automated ransomware deployment

Threat actors use AI to scan networks, find vulnerabilities and launch attacks based on findings. For example, AI can automate lateral movement across systems, identify high-value assets like EHRs or billing systems, and trigger encryption once access is maximized.

Monitoring an uncertain regulatory landscape

In response to rapidly evolving cyberthreats, regulators are updating how they monitor risk and assign accountability.

Proposed rules — such as updates to the HIPAA Security Rule — signal growing federal oversight.

At the same time, proposed legislation like the state-level AI regulation moratorium in the 2025 federal budget reconciliation bill suggests that federal authority may override state-level efforts, despite recent AI laws passed in California, Colorado, Connecticut and New York.

While the timing and final form of these measures remain uncertain, two things are clear: AI is accelerating cyberthreats that demand stronger protection, and regulatory change is on the way. Here is a closer look at pending legislation and how it could impact your business.

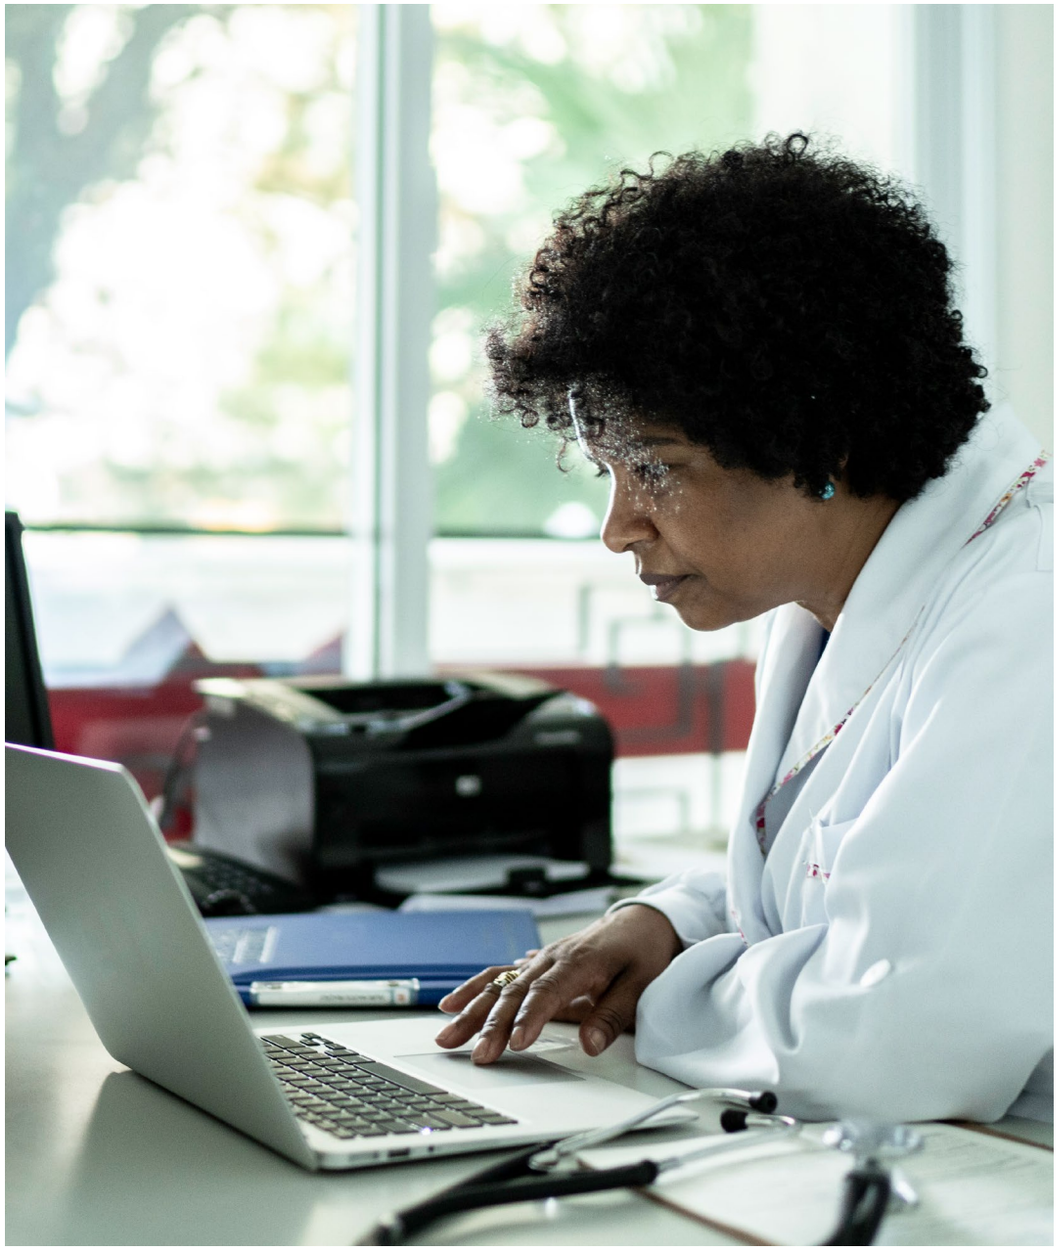

Proposed changes to HIPAA Security Rule

HIPAA is undergoing its biggest proposed update in a decade, aiming to address the rising use of AI in healthcare.

While the proposed update doesn't create AI-specific rules, it does aim to modernize the HIPAA Security Rule to better manage risks associated with AI. If passed, updates could go into effect in late 2025 or early 2026 and would apply to any organization that creates, receives, maintains or transmits ePHI.

Proposed changes include:

1. Mandatory risk analysis for AI systems

Organizations that handle ePHI would be required to conduct regular, written risk assessments that evaluate how AI models and tools might impact confidentiality, integrity and availability of ePHI, including risks related to training data, algorithm outputs and potential reidentification of deidentified data.

2. Stronger technical and administrative controls

These measures are designed to protect against unauthorized access, including vulnerabilities introduced by AI systems. The proposed rule would mandate encryption of ePHI both at rest and in transit, deployment of multifactor authentication, regular penetration testing and patch management, and network segmentation to isolate systems handling ePHI.

3. Clarification of AI's role under HIPAA

The update explicitly clarifies that AI tools must comply with existing HIPAA standards, including adhering to the "minimum necessary" standard for data use, ensuring proper de-identification of data and establishing business associate agreements (BAAs) with AI vendors handling PHI.

4. Comprehensive documentation requirements

Organizations would need to maintain detailed documentation of their security policies, procedures and risk analyses, including those related to AI systems. This ensures accountability and facilitates compliance reviews.

What to do now

Regardless of whether the proposed changes pass in their current state or evolve based on industry feedback, healthtech organizations should take an aggressive stance on safeguarding ePHI from AI-associated risks.

If you haven't already done so, you should:

- Apply HIPAA security controls to AI systems in scope.
- Spend time on due diligence vetting tools.

In Wipfli's experience with healthtech clients, vetting procedures often represent the most commonly glazed-over aspect of implementing AI into systems. You should spend time developing and following robust vetting procedures, which can provide a great idea of the security around the tools you want to implement. It's also essential to look for tools from vendors who have undergone third party assessments or audits such HITRUST and SOC.

What to look for next

Feedback delivered during the 60-day open comment period after the proposed update's release raised concerns about feasibility and affordability, especially for smaller organizations. You should watch for the HHS response and monitor any resulting changes to the final rule. Once finalized, organizations will have 180 days to comply.

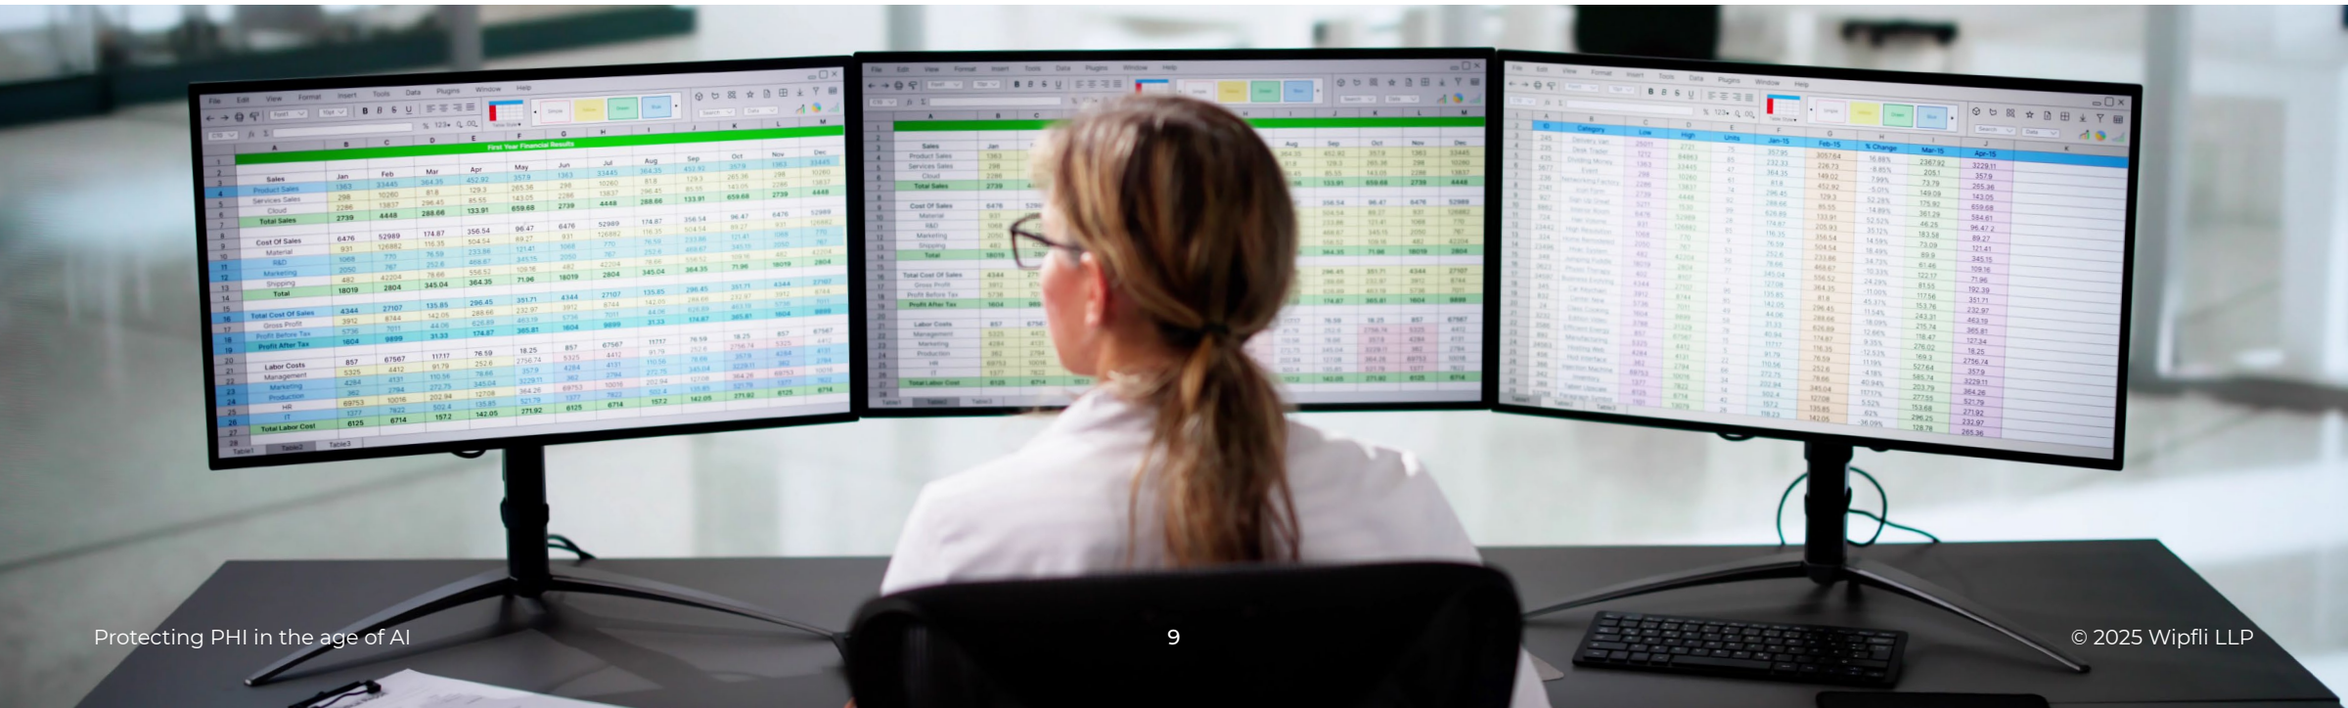

The evolution of AI risk assurance and compliance

Compliance authorities are also under pressure to adapt to new risks posed by AI-related cyberthreats.

While formal rulemaking is still in flux, the compliance ecosystem is actively evolving to help organizations assess and validate their AI-related security risks. Leading risk evaluators are adjusting their frameworks to account for emerging threats and the complexities introduced by AI systems.

Although no single standard is mandated, healthtech organizations should begin evaluating available options such as NIST AI RMF, ISO 42001 or OWASP AI Exchange, and certifications/attestations such as HITRUST, ISO or SOC 2.

It's also important to note that while most security frameworks can be adapted to AI tools and systems, HITRUST has developed AI-specific controls to help organizations integrate the appropriate security measures and address the added risk of AI tools.

Taking early steps to align with these evolving frameworks will better position organizations for both near-term risk management and long-term regulatory readiness.

HITRUST AI Assurance Program

Introduced in early 2025, the new HITRUST AI Security Assessment and Certification provides a structured framework to validate the security of AI systems. The program encompasses 44 controls specifically designed to address AI-related risks. It can be completed in conjunction with any e1, i1 or r2 assessment.

The program focuses heavily on data governance, risk management, technical security of the model(s) used and change management, including secure development processes for internally developed apps, vetting processes for externally developed apps, communication processes for changes and third-party due diligence — all areas where AI introduces unique risks.

While specific organizations have not publicly confirmed their certification efforts, HITRUST has signaled that early adoption is expected in 2025, particularly among cloud service vendors and infrastructure providers.

"As AI reshapes how healthcare operates and innovates, organizations need a clear, trusted path to manage risk," Jeremy Huval, Chief Innovation Officer at HITRUST, told Wipfli.

"The HITRUST AI Security Certification brings structure and assurance to a fast-moving landscape by aligning AI technologies with proven security and privacy controls — so organizations can adopt AI responsibly and confidently."

What to do now

As with other HITRUST assessments, full or partial inheritance will be available for certain requirements based on the shared responsibility model depending on what third-party providers you use — such as cloud platforms.

As mentioned above, the big cloud platforms have not provided the option for inheritance for the AI controls yet, but this doesn't mean you should wait to begin evaluating your environment. You should start discovery around the AI control set or work with a certified assessor to conduct a readiness assessment to clarify which requirements you'll inherit and what gaps you'll need to close to achieve certification.

What to look for next

Monitor HITRUST for framework updates and certification guidance. Insights from early adopters moving through the process may provide practical takeaways to guide your own efforts.

SOC 2

While there are not any pending AI-related changes to SOC 2 audits, their focus on established security-related criteria makes it a good vehicle to adapt for your AI-related tools and systems. A SOC 2 audit examines and reports on a service organization's internal controls relevant to the security, availability, confidentiality, processing integrity, and/or privacy of customer data.

If AI solutions are part of your cybersecurity program, then your SOC 2 audit could, and should, encompass AI-related risks.

7 strategies you can enact now to safeguard PHI

Healthtech organizations should proceed under the assumption that increased regulation is on the horizon — even if it ultimately looks different from today's proposals.

Core focus areas and methodologies like data governance, access control and model transparency are sure to remain in the regulations' final form. And while formal frameworks are still in the early stages of adoption, there are actionable steps you can take now to protect PHI while responsibly adopting AI.

One of the most common misconceptions we hear from clients is that AI is too complex to afford transparency. While it's true that AI systems can be intricate, they are not impenetrable. With the right development practices and thorough vetting, your organization can maintain visibility into how AI tools operate, ensure they perform as intended and build trust in their outputs. Here's how to get started.

1. Apply best practice controls to AI systems.

Best practice controls that are already implemented in your environment should be applied to AI systems. Focus on:

- **Encryption:** Encrypt data in transit and at rest.
- **De-identification:** Replace all unique identifiers with tokens.
- **Data minimization:** Retain only the minimum necessary data required, especially for AI models, to help ensure they aren't retaining ePHI.
- **Access controls:** Limit who has access to AI systems and focus on minimum necessary access.
- **Change management controls:** Ensure strong controls for making changes to your own developed AI systems and for monitoring changes made to third-party AI systems.

2. Test, test, test.

Testing is the best way to help ensure privacy and protect data, both for internally developed and third-party solutions. AI naturally looks for a solution, and this can be to your advantage when testing for holes.

- Code/train your internally developed tools to push back or create error alerts for spots where data could potentially be leaked.
- Spend time testing these features and pushing the limits of AI to see if you can identify areas for data leaks.
- Don't assume outputs are unbiased: Data sets can be trained and can naturally inherit bias.
- Consider using synthetic patient data (or dummy data) to mimic real patient data without risking PHI exposure when building or training an AI solution.

3. Foster a security culture by investing in teams and technology.

Having a knowledgeable team allows you to implement emerging technologies while balancing the risks. It also enables you to trust that everyone on your team, not just your security team, has the training and awareness needed to be human firewalls against AI-driven cyberattacks.

- Send security teams for AI-related training.
- Build security-related best practices into your program.
- Educate everyone about the high level of awareness needed to prevent sophisticated attacks that are increasingly easy to miss, especially in high-paced, high-pressure environments.
- Remind everyone not to put blind faith in AI tools because their outputs seem robust; tools miss threats and give off-base feedback. You still need extremely qualified human oversight and critical thinking to stay on top and stay ahead.

4. Evaluate third-party tools.

Before integrating any third-party AI solution into your environment, take time to vet its security posture and compliance readiness. With standards-based AI audits still emerging, this step is critical for minimizing risk exposure.

- Identify what compliance frameworks (HIPAA, HITRUST, SOC 2) the vendor claims to follow — and verify supporting documentation.
- Use a due diligence checklist to risk rate your vendors before procurement and implementation.
- Push for data compliance language to be included in contracts to define security obligations and ownership of liability in the event of a data breach.
- Require vendors to commit to security audits or certifications as part of the engagement process.

5. Leverage AI threat modeling systems.

While AI introduces new and dynamic attack surfaces, it also introduces new ways to protect against attacks. When paired with qualified security team oversight, AI-based threat modeling tools can help anticipate vulnerabilities and strengthen defenses.

- Use tools that constantly update and adapt based on the threat landscape.
- Validate model behavior under stress tests to check for bias, hallucinations or data leakage.
- Test how models respond to edge cases and adversarial prompts.
- Assign skilled team members to regularly review outputs and investigate anomalies.
- Incorporate threat modeling into your broader risk management and incident response processes.

$2.2 million

The average cost savings for organizations that used security AI and automation extensively in prevention versus those that didn't, according to IBM's Cost of a Data Breach Report 2024

6. Build strong governance.

As with any system that touches PHI, AI systems need to be governed by clear policies, executive oversight and regular reviews.

- Map how PHI enters, moves through or is generated by AI models.
- Define access controls, retention policies and audit trails specific to AI workflows.
- Ensure your governance structure includes executive leadership and board-level individuals who understand AI risks and hold their teams accountable for AI-related security.
- Develop internal guidelines for responsible AI use, especially if you provide clinical decision-making tools.
- Revisit existing PHI governance frameworks and adapt them to accommodate AI-specific challenges.

7. Begin preparing for AI-related compliance.

Your compliance efforts should begin now. Documentation and baseline control mapping will pay dividends when formal requirements arrive.

- Create an inventory of where and how AI is used across your organization.
- Document risk assessments for each AI system, including third-party tools.
- Align internal controls with existing frameworks like the HITRUST AI Assurance Program.
- Establish a plan for continuous monitoring and evidence collection to demonstrate accountability.
- Conduct a gap analysis with the help of an external assessor or qualified consultant.

Wipfli: Helping you grow securely

You work at the intersection of two complex industries: healthcare and technology. We have deep experience in both, and we use that industry-specific knowledge to help you grow, securely.

Whether you need help understanding the new HITRUST AI Assurance program or how to leverage SOC 2 for AI security, contact our risk advisory team. We provide a full range of support, from readiness to certification and compliance services.

Learn more at

wipfli.com/industries/tech/digital-health.

Perspective changes everything.

WIPFLI