CYERA

Protecting Patients and PHI with Cyera AI Guardian

Artificial intelligence is rapidly transforming healthcare, from predictive diagnostics to operational efficiency. Clinicians, researchers, and administrators are embracing AI to improve patient outcomes, reduce costs, and unlock insights from vast datasets. Yet as AI becomes more deeply embedded in clinical workflows, the risks associated with misuse of Protected Health Information (PHI), model bias, and shadow AI adoption are mounting. In healthcare, these risks are not abstract—they affect patient privacy, regulatory compliance, and trust.

This makes healthcare one of the most high-risk sectors for AI governance. Regulators are watching closely, with HIPAA, GDPR, and emerging AI-specific regulations creating stringent expectations. To thrive in this environment, healthcare organizations must embed privacy-by-design and security-by-design into every stage of their AI adoption. Cyera's **AI Guardian** delivers the controls required to achieve this balance.

The Special Risks of AI in Healthcare

Healthcare data is among the most sensitive information an organization can process. With the rise of AI, the stakes are higher than ever:

Exposure of PHI

Ransomware attackers are increasingly targeting hospitals and clinics with AI-powered phishing campaigns. At the same time, 71 percent of healthcare workers still use personal AI accounts, though these may not be HIPAA compliant.

Diagnostic Bias

Training models on incomplete or skewed datasets risks unequal outcomes across demographics.

Shadow AI Tools

According to IBM, one fifth of organizations have suffered a data breach due to unapproved AI applications, and the cost of a breach for organizations that make extensive use of Shadow AI is $670,000 higher.

Research Data Leakage

Sensitive clinical trial data could be copied into public LLMs, undermining intellectual property and patient privacy.

Unchecked, these risks could result in regulatory penalties, reputational damage, or direct harm to patients.

The Governance Controls Healthcare Requires

To address these risks, healthcare organizations must implement specialized governance controls:

Privacy-by-Design

Ensure PHI is minimized, anonymized, or pseudonymized before it reaches AI systems.

Strict Access Control

Enforce least-privilege access to patient data, with clear accountability for every access request.

Auditability

Maintain detailed logs to demonstrate compliance with HIPAA and GDPR.

Real-Time Monitoring

Real-Time Monitoring: Detect and block risky AI activity before data is exposed.

These requirements demand tooling that understands both sensitive healthcare data and the unique risks of AI systems.

CYERA

For more information visit cyera.io or scan→

CYERA

Protecting Patients and PHI with Cyera AI Guardian

How Cyera AI Guardian Secures Healthcare AI

Cyera's AI Guardian extends data-centric security and compliance capabilities directly into the AI ecosystem. It operationalizes privacy and security by design across healthcare AI pipelines.

AI Security Posture Management (AI-SPM)

AI Guardian discovers and classifies all AI systems in use - including shadow AI - while mapping PHI and other sensitive datasets. This visibility allows healthcare organizations to detect unapproved tools accessing patient data before violations occur.

AI Runtime Protection

AI Guardian continuously monitors prompts, responses, and agent actions in real time. If a clinician attempts to paste PHI into a SaaS LLM, AI Guardian instantly blocks the action and alerts security, preventing a HIPAA breach.

Identity and Access Context

By correlating data access with human and AI agents, AI Guardian enforces least-privilege policies. Healthcare organizations gain clarity into who is accessing patient data and why.

Audit-Ready Reporting

AI Guardian generates detailed logs that healthcare organizations can present to regulators or auditors, demonstrating continuous compliance with HIPAA, GDPR, and sector-specific AI governance requirements.

A Realistic Scenario

Imagine a nurse using a generative AI tool to help summarize complex patient histories. Without guardrails, PHI could be copied into a third-party system, creating a HIPAA violation. With AI Guardian:

- AI-SPM detects that the generative AI tool is not approved for PHI use.
- Runtime Protection blocks the attempt to paste PHI into the tool.
- An alert is sent to compliance teams, and the incident is logged for audit.
- The organization prevents a breach, avoids penalties, and maintains patient trust.

This proactive, real-time protection is what sets Cyera apart.

Looking Ahead: Preparing for Stricter AI Regulation

Global regulators are tightening oversight on AI in healthcare. The EU AI Act classifies medical AI systems as high-risk, while U.S. agencies are preparing sector-specific AI governance guidance. Healthcare organizations cannot wait for enforcement—they need controls now.

By embedding Cyera's AI Guardian, healthcare providers and researchers not only meet today's compliance requirements but also future-proof their AI governance programs against evolving regulations.

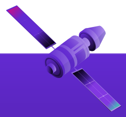

Conclusion

AI is revolutionizing healthcare, but without strong governance, it risks violating privacy, eroding trust, and harming patients. Privacy-by-design and security-by-design are regulatory and ethical imperatives. Cyera's AI Guardian delivers the visibility, enforcement, and auditability needed to secure AI in healthcare. With AI Guardian, healthcare organizations can confidently innovate with AI while protecting patients and their data.

CYERA

For more information visit cyera.io or scan→