Shadow AI is outpacing healthcare email security

What happens when AI adoption moves faster than healthcare security can respond

PAUBOX

PAUBOX

Table of contents

PAUBOX

Rapid AI surge without a safety net

Healthcare organizations are rushing to use AI at breakneck speed, but security and compliance teams are already stretched thin and can't keep up.

Shadow AI, the use of AI tools without IT or compliance approval, has become one of the most dangerous blind spots in healthcare security.

To measure the scale of this risk, Paubox conducted a survey of healthcare IT and compliance leaders about the hidden risks and organizational pressures driving shadow AI in email security. This report examines how rapid-paced AI adoption is colliding with under-resourced compliance, creating risky gaps.

Why it matters: Shadow AI usage in healthcare introduces a new layer of security and compliance risk. Employees leveraging AI without proper guardrails can unintentionally expose sensitive patient data. Research shows that cybersecurity risks are highest in healthcare among industries where shadow AI is present, largely because unauthorized usage bypasses oversight.[1]

BY THE NUMBERS

95%

of organizations report staff are already using AI tools, yet **25%** have not formally approved any staff AI email use

62%

have observed staff experimenting with ChatGPT or similar tools even though they're unsanctioned

41%

of leaders feel confident they could detect improper AI use before a HIPAA violation occurs, yet **16%** admit compliance was never consulted before AI email tools were enabled

75%

believe employees assume tools like Microsoft Copilot are automatically HIPAA compliant

The push for fast AI adoption without oversight

Executive enthusiasm is high. Executive leadership is enthusiastic about AI usage, seeing it as a game-changer for efficiency and patient care. They're pushing rapid rollout of tools like Microsoft 365 Copilot or Google's Gemini to gain an edge. Many applications that staff use daily now come with embedded AI.

Compliance and IT are struggling to keep up. IT and compliance teams are under-resourced and many weren't fully prepared for AI's sudden surge. 83% of healthcare IT and compliance leaders have raised concerns about AI security. Policies and training haven't caught up to the new tech. 75% of respondents say AI has added confusion, not clarity, to email compliance at their organization.

> "Shadow AI develops when speed and departmental innovation are rewarded, often bypassing IT and compliance oversight."[3]
>
> **A. Omar & H.R. Weistroffer**
> "From shadow IT to shadow AI – threats, risks, and governance."

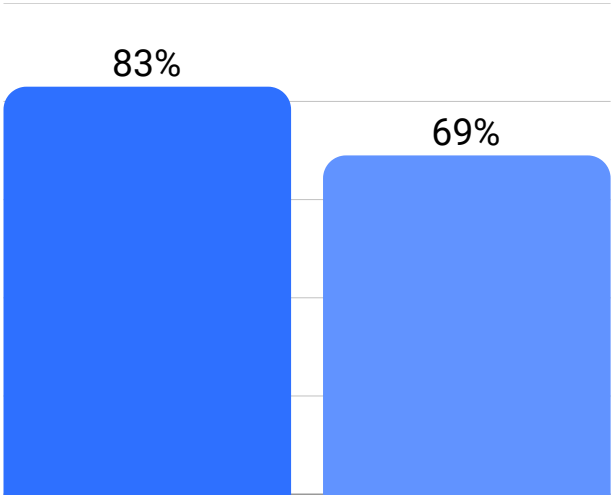

83% of clinical leaders have raised concerns about the use of AI in healthcare.

69% of IT leaders feel pressured to adopt AI faster than they can secure it.

Adoption is moving faster than security teams. Two thirds of IT leaders also feel pressured to adopt AI faster than their organization can secure it, creating a widening gap between enthusiasm and readiness.

Oversight is often an afterthought. Compliance officers are finding themselves in reactive mode, trying to understand tools that may already be in use by staff. Even worse, 16% of compliance leaders were not even consulted before AI features were activated in Gmail or Outlook at their organizations. It's impossible for compliance to secure what they don't even know about.

PAUBOX

75%

of IT leaders say AI has added confusion to email compliance at their organization

21%

of teams believe a BAA isn't required for an AI email assistant

69%

of IT and compliance leaders feel pressured to adopt AI faster than their organization can secure it

16%

of compliance leaders were not consulted before AI features were activated in Gmail or Outlook at their organization

BOTTOM LINE

Leadership's drive to keep up with the AI Joneses has outpaced the usual checks and balances. Compliance and IT teams are now playing catch-up, updating policies and guidelines on the fly. Cue the growth of shadow AI.

Knowledge gaps: Some teams misunderstand compliance requirements for AI. For example, 21% believe a Business Associate Agreement (BAA) isn't required for an AI email assistant. Most don't realize that any tool that touches Protected Health Information (PHI) must be covered by a BAA or risk noncompliance.

Overconfidence: An overwhelming 94% are confident their organization could detect AI misuse before a HIPAA violation happens, but 62% have seen staff experimenting with ChatGPT even though it isn't approved in their org.

Policy updates in progress: On the positive side, most organizations are attempting to respond – nearly all (94%) have begun updating security policies to address generative AI threats in email. However, a policy on paper is only as good as its enforcement and understanding among staff. The current state leaves a gap between rapid adoption and slower-moving compliance safeguards.

Frontline staff using AI under the radar

While leadership pushes AI from the top, frontline staff aren't waiting for permission. This "shadow AI" by employees is widespread in healthcare settings:

Shadow usage is pervasive: Nearly 95% of organizations suspect their staff are already using generative AI for work email or content. In our survey, 62% of leaders have directly seen employees experimenting with ChatGPT (or similar AI) even though it wasn't officially approved. Whether it's drafting patient emails, summarizing records, or brainstorming responses, staff are jumping into using AI to speed up their work.

Often unsanctioned: Much of the AI use is happening without formal approval or guidance. 25% of organizations have not formally approved any staff use of AI in email, yet the majority of those organizations' employees are forging ahead anyway. Employees might install unvetted AI writing assistants or use free online AI services on their own.

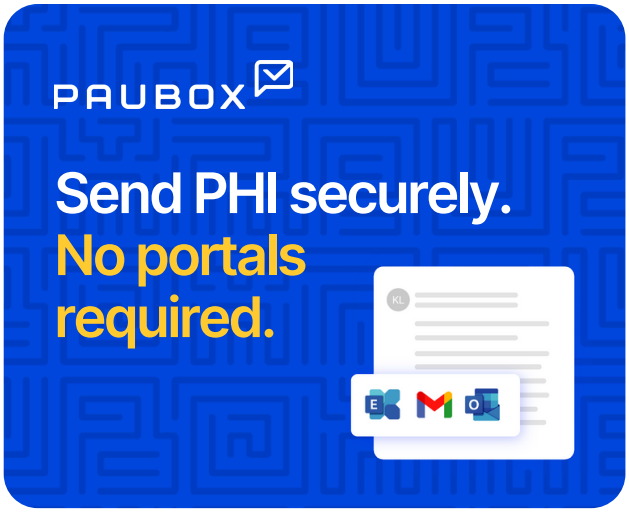

> "People tend to do it without thinking, just wanting to speed up their work … you just uploaded a bunch of company data…and your security team does not know about this."[3]
>
> **Limor Kessem**
> "Shadow AI a growing risk for hospitals"

Assumptions of safety: Many employees assume “if the tool exists, it must be fine to use.” In fact, 75% of IT leaders believe their staff assume tools like Microsoft Copilot are automatically HIPAA compliant. There’s a general trust in technology to “do the right thing”.

Unfortunately, this trust can be misplaced if the tools haven’t been vetted. For example, an employee might use generative AI to compose an email containing patient info, assuming the AI or email platform will handle that data securely (which might not be true without proper agreements and settings).

> **75%** of IT leaders believe their staff assume tools like Microsoft Copilot are automatically HIPAA compliant

Why staff go rogue: Clinicians and administrators face heavy workloads, and see AI as a quick helper to draft messages, analyze data, or streamline tasks. If organizations haven’t provided an approved solution or if approval processes are slow, staff take matters into their own hands.

90% of leaders believe staff are using AI email tools without formal approval. In many cases they may not even realize they’re breaking any rules; they view it as just another productivity tool.

The surge of shadow AI means sensitive patient data is likely to be shared outside the view of IT. According to research by cybersecurity company Netskope, “Healthcare workers routinely expose sensitive data such as PHI by using generative AI tools such as ChatGPT and Google Gemini” without oversight.[6]

62% of IT leaders report seeing staff experimenting with ChatGPT even though it’s unsanctioned.

25% of organizations have not formally approved any staff use of AI in email.

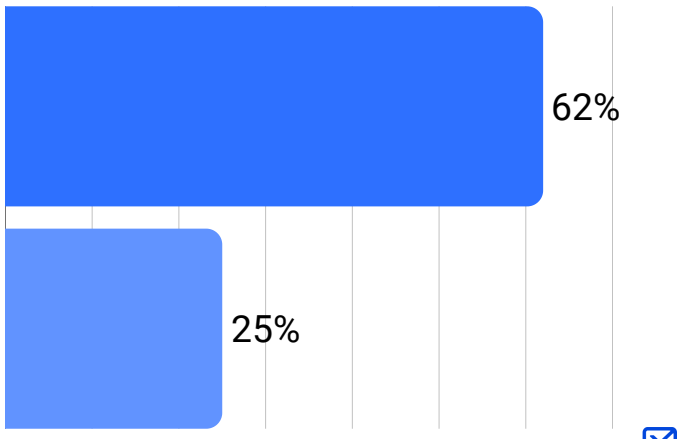

PAUBOX

Security and compliance dangers of shadow AI

When AI tools are adopted without policies in place, organizations face a new wave of security and compliance challenges on top of the threats they're already facing.

Data leakage and HIPAA violations: The most immediate risk is exposure of PHI. If an employee pastes a patient's medical summary into an AI chatbot, that data may be stored or even used to train the AI model, potentially violating HIPAA privacy rules. Without a BAA in place, sharing PHI with an AI service is equivalent to a breach.

In a recent IBM poll, over one-third of employees (38%) admitted to sharing sensitive work information with AI tools without employer approval.5 Shadow AI can easily undermine carefully built compliance protections by letting data slip out through a channel that's difficult to monitor.

> **94%** of IT leaders feel confident in their ability to detect AI misuse within their organizations

Compliance blind spots: Many leaders feel confident in detecting AI misuse (94%), yet that confidence may be unwarranted if they lack tools to actually monitor AI activities.

Traditional security controls (like email DLP or encryption gateways) may not catch someone using a web-based AI service. If staff assume that AI is compliant, they may not take manual precautions they normally would.

Shadow AI is a vector for breaches and compliance failures if left unaddressed. The combination of high employee trust in AI and low visibility by IT is a dangerous mix.

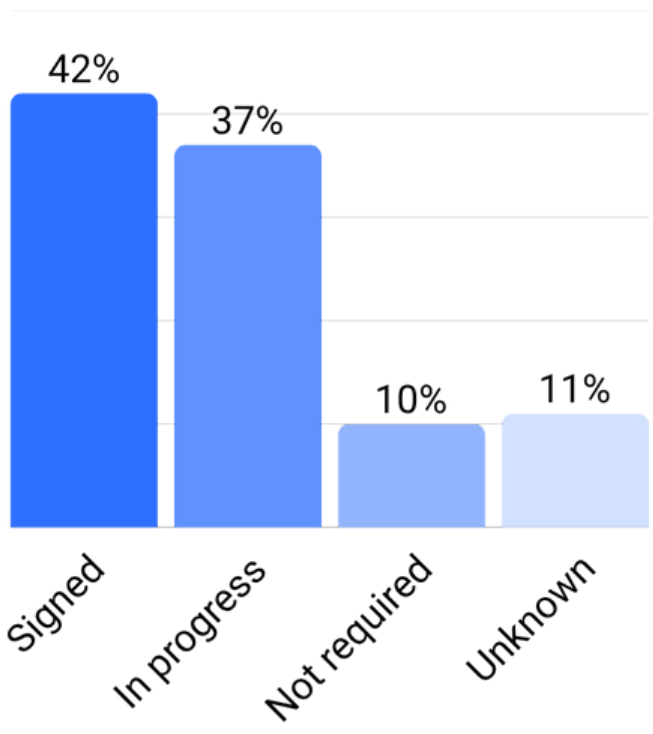

Only **42%** of organizations have signed a BAA covering any AI assistant used in email.

PAUBOX

AI email security safeguards

Consider solutions that add an extra layer of protection for email and AI. Email security and compliance platforms like **Paubox Email Suite** ensure every message is encrypted by default, keeping communications HIPAA compliant without extra steps for staff or patients.

Beyond encryption, technology can help prevent and detect unsanctioned AI usage. Updating Data Loss Prevention (DLP) rules to flag large text blocks being copied to external websites or unusual patterns that might indicate AI misuse. Some email security platforms can also identify sensitive content in outgoing emails and issue warnings or block messages that risk compliance violations.

Products such as **Paubox's Inbound Email Security** can be leveraged to detect AI-generated phishing attempts or email behavior that is outside the usual communication patterns for your organization. It can also help organizations mitigate newer risks like prompt injection attacks, where attackers try to manipulate an AI system into leaking sensitive data or taking harmful actions through crafted email content. By understanding what "normal" communication looks like for your organization, Inbound Email Security can flag anomalies before they reach an inbox.

These tools can form a safety net for AI-driven communication.They allow healthcare organizations to embrace AI benefits while ensuring that sensitive data stays protected and that compliance is maintained. These solutions are enablers of safe innovation, giving staff the freedom to explore AI with the guardrails needed to prevent costly mistakes or malicious attacks.

How compliant is healthcare AI adoption?

Every IT and compliance team in healthcare should urgently work to close this oversight gap before it leads to an incident. The goal is to enable the benefits of AI while maintaining security and compliance. Here are key strategies and best practices to consider:

AI usage policies: 94% of orgs have updated or are updating security policies to address AI in email.

Vendor due diligence (BAAs and beyond): AI providers should be treated like any other vendor handling sensitive data. Only 42% of organizations have signed a BAA for an AI email tool so far (with ~37% in progress, and 21% thinking it's "not required").

> "Traditional security philosophies, such as validating and sanitizing both input and output to the models, can still apply in the AI space."[5]
>
> **Royal Hansen**
> VP of Privacy, Safety & Security Engineering, Google Cloud

Education and training: Simply having a policy isn't enough; AI needs to be added as a new layer of cybersecurity training. Only 16% of orgs have trained most of their staff (75-100%) who have access to PHI on AI usage in email.

By combining policy, education, oversight, and technology, healthcare organizations can reap the benefits of AI safely. The aim is to integrate AI into workflows under a watchful eye, rather than in the shadows.

94%

of healthcare orgs have updated or are updating security policies to address AI in email.

58%

have not signed a BAA for an AI email tool so far. 21% think it's not required.

84%

have not trained most of their staff who have access to PHI on AI usage in email.

PAUBOX

PAUBOX

Stop email phishing attacks with AI security

Surface hidden threats with generative AI email security that learns and evolves.

Talk to an expert

Sources

1. Balogun, A., et al. “The ethical and legal implications of shadow AI in sensitive industries.” ResearchGate, 2025. https://www.researchgate.net/publication/388960052_The_Ethical_and_Legal_Implications_of_Shadow_AI_in_Sensitive_Industries_A_Focus_on_Healthcare_Finance_and_Education
2. Kessem, Limor. “Shadow AI a growing risk for hospitals.” Chief Healthcare Executive, 2025. https://www.chiefhealthcareexecutive.com/view/shadow-ai-a-growing-risk-for-hospitals
3. Omar, A., & Weistroffer, H. R. “From shadow IT to shadow AI – threats, risks and governance.” Journal of Strategic Contracting and Negotiation, 2025. https://onlinelibrary.wiley.com/doi/abs/10.1002/jsc.2682
4. IBM. “Shadow AI.” IBM Think Blog, 2025. https://www.ibm.com/think/topics/shadow-ai
5. Hansen, Royal. “Spotlighting shadow AI: how to protect against risky AI practices.” Google Cloud Blog, 2025. https://cloud.google.com/transform/spotlighting-shadow-ai-how-to-protect-against-risky-ai-practices
6. HIPAA Journal / Netskope. “Healthcare workers violating patient privacy by uploading PHI to AI tools and cloud accounts.” HIPAA Journal, 2025. https://www.hipaajournal.com/healthcare-workers-privacy-violations-ai-tools-cloud-accounts