MENLO SECURITY

2026 HEALTHCARE SECURITY REPORT

Browser-First Attacks in Healthcare:

Why the Browser Is Healthcare's Biggest Security Blind Spot

REPORT

EXECUTIVE SUMMARY:

The Browser Security Gap in Healthcare

Healthcare providers and payers face a threat environment unlike any other sector. Patient data commands the highest price on the dark web, while operational downtime creates direct patient safety consequences that accelerate payment decisions for ransomware groups. And the browser — where clinicians access Electronic Health Records (EHRs) and communicate with patients, and where payers process claims, manage member portals, and interact with AI-assisted workflows — has become the primary access point for every major attack category.

This brief draws on Menlo Security's Q1 2026 platform telemetry and customer case studies from healthcare organizations currently protected by Menlo. The pattern is consistent: attacks that bypassed every existing layer of the security stack were caught at the browser session layer before they could execute.

1 in 5	276M+	58%
Phishing Links Go Completely Undetected	**Patient Records Compromised**	**of Healthcare Attacks**
Actively clicked by users, missed by legacy URL filtering (Menlo Threat Research)	In 2024 alone — 758,000 exposed every single day (Healthcare IT Today)	Are credential phishing — the browser is where they arrive (Menlo Threat Research, healthcare vertical)

The browser is not just where healthcare employees work. It has become the premier entry point for attackers because it is not visible to any of the tools your security stack was designed to protect.

Threats by Type

Last Quarter

Zero-Hour Phishing Site	44.9%
Zero-Hour Malicious Site	19.8%
Zero-Hour Fraudulent Site	16.0%
Phishing Site	12.7%
Malicious Site	5.1%
Fraudulent Site	1.4%

Threat Type Details

Malicious File Download: infected file downloads

Malicious Site Access: sites serving exploits / malicious scripts / malware

Zero-Hour Malicious Site Access: identified retroactively / detected by HEAT Shield

Phishing Site Access: sites that attempt to steal credentials and sensitive data

Zero-Hour Phishing Site Access: identified retroactively / detected by HEAT Shield

Fraudulent Site Access: sites that impersonate brands / entities with nefarious intent

Zero-Hour Fraudulent Site Access: identified retroactively / detected by HEAT Shield

The Browser Gap in Healthcare Environments

Healthcare security teams typically operate with strong perimeter controls, including network security, email gateways, endpoint protection, and data loss prevention. These tools are performing exactly as designed. The problem is that none of them were designed to operate at the browser session layer, and that is precisely where every major attack technique targeting healthcare organizations now lives. Closing that gap does not mean replacing what is working. Implementing proper browser security controls complements your existing stack by covering the one layer your other tools were never designed to see.

Three Structural Gaps

1. Network Tools Cannot See Inside the Browser Session

Your secure web gateway inspects connections, enforces URL policies, and blocks known-bad domains. What it cannot do is see what executes inside an encrypted HTTPS session after that connection is established. When a clinician navigates to what appears to be a legitimate patient portal or EHR login page and that page is actually a phishing page on a clean, newly registered domain, your SWG sees a valid HTTPS connection to an approved destination. Nothing fires.

In Q1 2026, 35% of highly evasive threats blocked by Menlo originated from domains already categorized as safe, including clinical collaboration platforms like Epic MyChart, Cerner PowerChart, Change Healthcare, DocuSign, Doximity, and Meditech portals, as well as general platforms such as SharePoint, Google Drive, and Dropbox that healthcare organizations use daily for document exchange. You cannot block these domains. Attackers know this and deliberately host attack infrastructure there.

Source: Menlo Threat Research, Q1 2026

2. EDR Only Sees What Happens on the Device - Not What Happens in the Browser

Endpoint detection and response tools monitor process execution, file writes, and registry changes on the device. In credential theft attacks, which are the dominant attack type in healthcare, the damage is done entirely inside the browser session, before anything touches the device. Session tokens are harvested. Credentials are entered on spoofed login pages. Neither event writes a file or triggers a process. From the EDR's perspective, a clinician simply logged in normally. There is nothing to flag because nothing suspicious happened on the endpoint — it all happened in the browser.

3. DLP Cannot See In-Session Data Movement

Traditional data loss prevention solutions monitor email attachments, file transfers, and API-level traffic. They cannot see data typed or pasted inside a live browser session because the content never becomes a file and never crosses an email gateway. As healthcare workflows increasingly rely on AI assistants, browser-based clinical tools, and web forms for patient data entry, this creates PHI / PII exposure events that generate no alert, no log entry, and no compliance signal.

The PHI Exposure Your DLP Cannot See

In a single healthcare customer environment, Menlo observed 276,000 file uploads to GenAI tools in a 30-day reporting period. Most organizations have no policy, no visibility, and no governance for these interactions. This is not a theoretical risk: the 2025 Verizon DBIR found that 15% of employees access generative AI platforms from corporate endpoints at least twice a week, with 72% signing in using personal email addresses — creating data flows that no enterprise DLP tool can see.

Sources: Menlo Security Threat Research, healthcare customer (2025); Verizon DBIR 2025

Q1 2026

The Healthcare Threat Landscape:

Menlo's Q1 2026 telemetry across healthcare customer environments documents three dominant attack patterns. Each exploits the browser session layer. Each was invisible to the existing security stack at time of delivery.

90%	35%	20%
Of Evasive Phishing Attempts Neutralized by Menlo HEAT Shield Were 'Zero-Days'	**Of Evasive Threats Originated from Sites Already Deemed 'Safe'**	**Of Phishing Links Actively Clicked by Users Went Undetected by Legacy URL Filtering**
(Menlo Threat Research, Q1 2026)	In 2024 alone — 758,000 exposed every single day (Healthcare IT Today)	(Menlo Threat Research, Q1 2026)

Attack Pattern 1: Credential Phishing and Session Token Theft

Credential phishing accounts for 58% of browser-based attacks in the healthcare sector, which is one of the highest of any vertical. *Source: Menlo Threat Research, Q1 2026*

In 2026, the dominant technique is no longer simple password theft. Adversary-in-the-Middle (AiTM) phishing intercepts the authenticated session token after MFA has already completed, which means that the attacker arrives as a fully verified user. MFA does not prevent this. Resetting the password afterward does not help because the token, not the password, is what the attacker holds.

Healthcare providers are particularly exposed because clinical workflows demand rapid authentication across multiple systems, including EHR platforms, patient portals, billing systems, and imaging interfaces. For payers, the exposure is equally significant: claims payment workflows, billing reconciliation, and customer service operations all require continuous authentication across web-based platforms throughout the day. Across both, the volume of authentication events creates proportionally more opportunities for credential harvesting, and the urgency of the work means staff are psychologically primed to complete verification steps quickly, without scrutiny.

Attack Pattern 2: Ransomware via Browser-First Initial Access

Ransomware groups understand that healthcare organizations are uniquely sensitive to operational downtime. Qilin, the most active ransomware group globally in Q1 2026 with 113 victims in February alone, uses browser-delivered credential theft as its primary initial access vector.

Their method: phishing delivers the initial foothold, and then Qilin harvests passwords by extracting credentials directly from the operating system's memory — the same process Windows uses to manage active logins — without needing to crack encrypted files or wait for a user to type a password. This gives them immediate access to every credential stored on a compromised clinical workstation. Stolen sessions then enable lateral movement through legitimate remote management tools that generate no behavioral alerts.

54% of ransomware victims had their credentials in stealer marketplaces before the attack began. Time from credential sale to encryption is sometimes under 48 hours. *Source: Verizon DBIR 2025*

Why Healthcare Is the Primary Target

Ransomware groups prioritize healthcare because operational downtime creates direct patient safety pressure that accelerates payment decisions. Healthcare was the costliest sector for data breaches for the **14th consecutive year, with an average breach cost of $7.42M globally and $10.22M in the United States**. Healthcare breaches also take the longest to detect and contain: an average of 279 days, five weeks longer than the global average.

Source: IBM Cost of a Data Breach Report 2025 (Ponemon Institute)

Attack Pattern 3: Exploitable Files and Password-Protected Archives

Menlo disarmed 433,314 exploitable files in Q1 2026, 110,357 of which were concealed behind password protection. This technique specifically defeats automated scanning tools, which cannot inspect encrypted content without the key. In healthcare environments where encrypted file exchange is standard practice for PHI protection, this creates a direct blind spot: the same encryption that protects patient data in transit also prevents security tools from inspecting files that may be weaponized.

Scale of the Patient Data Exposure Problem

More than 276 million patient records were compromised in 2024 alone — more than 758,000 records exposed every single day. More than half of health systems are now increasing cybersecurity investment in response.

Source: Healthcare IT Today, 2026 Health IT Predictions.

Healthcare Customer Case Studies

Case Study 1: How Behavioral, Intent-Based Prevention Stopped an Attack Every Other Tool Missed

In February 2026, an attack targeted a user at one of the nation's largest healthcare organizations — a major integrated health system processing millions of patient records across a large network of clinical facilities. Like many organizations of its size, the health system had invested in a layered security stack: a secure web gateway, endpoint detection, and email filtering. Each tool was performing as designed. None of them saw what was coming. The case demonstrates exactly why the browser session layer cannot be left unprotected.

CUSTOMER PROFILE

Large Integrated Health System— National Scale

Attack type: Social engineering with trojanized Remote Management tool delivery

Technique: User clicked an email link to a fake Adobe secure document portal hosted on a compromised WordPress site. The apparent PDF download was actually a trojanized remote management executable.

VirusTotal at time of click: Zero vendors flagged the domain as malicious.

Why every tool missed it: Domain was clean. Page rendered as a legitimate Adobe portal. No static payload. No known-bad signature. No tool in the stack, including SWG, EDR, or email gateways, saw anything wrong.

How it was stopped: Behavioral intent analysis identified that the page was delivering an executable disguised as a PDF. Menlo blocked the file before download executed, based on what the page was doing, not where it came from

Patient safety implication: Persistent remote access to a clinical workstation would have given an attacker the ability to observe, modify, or disrupt workflows connected to patient care, without triggering a single alert in the existing security stack.

CASE STUDY 2

How One of Georgia's Largest Nonprofit Health Systems Secured More Than 40,000 Users Without Replacing Their Browser

One of Georgia's largest nonprofit health systems, operating more than 400+ medical offices and hospitals with more than 40,000+ staff, faced a challenge common to large integrated health systems. The organization needed to secure high-risk web traffic across a distributed clinical workforce against sophisticated phishing and evasive ransomware, without replacing the browser clinical staff already use.

CUSTOMER PROFILE:

One of Georgia's Largest Nonprofit Health Systems

Challenge: Securing high-risk web traffic against sophisticated phishing and evasive ransomware across a distributed clinical workforce. Needed policy-based data protection controls without replacing clinical browser workflows.

What triggered investment: Internal dedicated personnel to quantify the risk associated with human error across the business and used that quantifiable risk to justify browser security budget to the board. This process is a model for risk-driven security investment in healthcare.

Results: 40,000+ staff now secured. Significantly reduced exposure to phishing and ransomware across 400+ medical offices. Gained visibility and traceability with forensics capability.

Key differentiator: No new or replacement browser required. Menlo delivers seamless deployment into existing clinical workflows with an abrasion-free user experience

Competitors evaluated: Palo Alto Networks, Island, Cloudflare, Conceal.io, LayerX

Clinical staff reaction: Provisioned within hours. Flexible configurations. Transparent policy controls with no disruption to clinical operations.

Source: Menlo Security win wire and customer success documentation, Wellstar Health System

CASE STUDY 3:

How an Orange County Nonprofit Health System Closed Its Browser Security Gap

A leading nonprofit health system in Orange County operating multiple hospitals, medical groups, imaging centers, and surgical centers across a major metropolitan region had a security posture that looked strong on paper — firewalls, endpoint protection, email filtering — but had no ability to govern web traffic at the browser session layer. Clinical workstations had open access to the web with only basic content filtering standing between staff and browser-based threats. For a health system of this size, that exposure represented a significant ransomware risk and a growing compliance gap as regulators increasingly scrutinized web-layer controls.

CUSTOMER PROFILE:

One of Georgia's Largest Nonprofit Health Systems

Challenge: Prime ransomware target with no ability to secure web traffic outside basic content filtering through firewalls.

Results: Full isolation of all web traffic with granular controls, visibility, and no impact to clinical performance or workflow.

Key benefits: Full isolation for all sites including documents. Granular control over file downloads and file types. Coverage for both authenticated and non-authenticated users. Seamless end-user experience for clinical staff.

Competitors displaced: Proofpoint, McAfee

Why Menlo won: Ability to reduce attack surface while maintaining business-as-usual clinical operations. Integration with existing tech stack. No endpoint agent required.

CASE STUDY 4:

Long-Term Care Pharmacy Governed PHI Across a Distributed Workforce

A long-term care pharmacy serving senior care facilities across multiple states processes sensitive PHI across a distributed workforce that accesses clinical systems, vendor portals, and patient records from dozens of locations. That distributed profile, combined with the volume of web-based clinical workflows, created elevated browser-based risk at every point of data access. Phishing attacks targeting vendor and customer portals were a recurring pattern, and the organization had limited visibility into what was happening inside browser sessions across its workforce.

CUSTOMER PROFILE:

Long-Term Care Pharmacy / Healthcare Services

PrimaryThreat Driver: Phishing attacks targeting vendors and customer portals — fake Microsoft login pages on clean domains were a recurring pattern found.

Key Finding: In a 30-day reporting period: 56M traffic requests analyzed, 20 phishing detections, multiple zero-hour attacks blocked by Menlo that no other tool in the stack caught. 75% of phishing attacks observed were hosted on uncategorized sites or sites with benign reputations.

GenAI Exposure: 276,000 file uploads to GenAI tools in a single 30-day period — with no prior visibility, policy, or governance over that data flow

HIPAA, HHS Enforcement, and the Browser Layer

The Department of Health and Human Services Office for Civil Rights has shifted HIPAA Security Rule enforcement from whether a risk analysis exists to how organizations act on its findings — moving from verifying the presence of an assessment to evaluating how identified risks are actually managed over time. HHS also proposed the most significant updates to the HIPAA Security Rule since 2013, introducing mandatory technical controls including encryption of ePHI, required multi-factor authentication, and 72-hour incident reporting requirements. In 2026, enforcement focus areas include:

Source: HHS OCR Security Risk Analysis Initiative; Medcurity, June 2026

- **Technical safeguards** — whether PHI transmitted or accessed through browser sessions is adequately protected
- **Access controls** — whether organizations can demonstrate that PHI access via web-based clinical systems is monitored and governed
- **Audit controls** — whether organizations maintain logs of activity in information systems containing PHI, including browser sessions
- **Transmission security** — whether PHI transmitted over networks, including browser sessions, is encrypted and monitored

The browser session layer is directly implicated in each of these focus areas. A clinician accessing Epic or Oracle Health, submitting data via a web-based claims system, or pasting PHI into a browser-based AI tool is creating PHI transmission events that most organizations cannot monitor, govern, or log at the session layer.

Healthcare Compliance Pattern

A mid-size health system's cyber-insurance audit recently surfaced a gap that most security teams don't find until something forces them to look: internal DLP tools had no visibility into browser session activity involving PHI. Data typed into web forms, pasted into AI tools, and submitted through browser-based clinical platforms was generating no logs, no alerts, and no compliance record. The audit requirement, rather than a breach, was what brought it into view. The coverage gap existed before the audit. It will exist after the audit in any organization that hasn't addressed it. The question isn't whether this gap is present in your environment. It's whether something has surfaced yet.

Healthcare Security Evaluation Framework

Five questions determine whether existing security controls address the browser-based threat landscape documented in this brief. Each maps directly to a documented attack pattern and a specific HIPAA compliance gap.

#	Question	What a 'No' Means	HIPAA Implication
1	Can your tools block a zero-day exploit before the patch is deployed to clinical endpoints?	6-day minimum exposure window on most zero-days — with more than 90 zero-days actively exploited in 2025.	Technical safeguards gap: clinical workstations processing PHI are unprotected during the patch window
2	Can your tools detect a page impersonating an EHR or clinical portal, based on behavior, not domain reputation?	AiTM and credential phishing on clean domains pass undetected, 58% of healthcare attacks arrive this way	Access control gap: unauthorized session token harvesting from clinical staff credentials
3	Can your tools inspect content assembled in the browser's JavaScript engine, not delivered as a file?	HTML smuggling and in-memory payload assembly are invisible to every network inspection tool	Transmission security gap: malicious content processed within systems containing PHI
4	Can your tools govern what PHI is pasted into AI tools, clinical web forms, or browser-based collaboration platforms?	PHI exposure via ChatGPT, browser forms, and clinical collaboration tools is undetected and unlogged	Audit control gap: no log of PHI transmitted via browser session — directly implicates HIPAA audit requirements
5	Can your tools govern AI agents browsing on behalf of clinical users?	Gemini in Chrome and Copilot in Edge are already operating in clinical environments without governance or audit trail	Access control gap: autonomous PHI access by AI agents with no audit trail or access controls

The Path Forward for Healthcare Security Teams

Browser security is not a replacement for the tools you already have. Your network controls, endpoint protection, and email gateway are doing their jobs. The browser session layer is the one layer none of them were designed to protect — and it is the initial access point for every attack pattern in this brief. Five specific capabilities close that gap:

Isolate the Clinical Endpoint from the Web	→	Execute web content away from clinical workstations, so that zero-day exploits and malicious payloads have no target. When an attack detonates, it detonates somewhere the patient data isn't.
Behavioral Intent Detection	→	Analyze what a page is trying to do, not where it came from. Behavioral, intent-based prevention is the only model that catches zero-day phishing and AiTM attacks on clean domains that reputation tools cannot flag.
Credential Input Protection	→	Block credential entry on suspicious or spoofed pages before session tokens can be harvested. MFA does not provide this — the attacker steals the authenticated session, not the password.
Session-Layer PHI Governance	→	Govern copy-paste, file uploads, and text entry inside browser sessions — covering GenAI tools, clinical web forms, and collaboration platforms where traditional DLP has no visibility.
AI agent Governance	→	Apply the same access controls and audit logging to AI agents already operating in your existing browsers such as Chrome and Edge that you apply to clinical staff — because agents can access and move PHI at machine speed without human oversight.

Three Actions for Healthcare Security Teams This Quarter

1. Ask your EHR and clinical portal vendors one question: "If a clinician's session token was stolen from our Epic or Cerner login page today — not the password, the authenticated session — would any tool in our stack detect it?" If the answer is no or I don't know, you have a credential defense gap directly implicated in 58% of healthcare attacks.

2. Run a PHI exposure audit for GenAI tools. Pull your proxy logs and identify how much data is flowing from clinical workstations to ChatGPT, Copilot, and other AI tools. If you don't have logs for that activity, that absence is your audit finding. HIPAA audit controls require that you can account for PHI transmission. Browser sessions are transmission events.

3. Test your detection against your own clinical infrastructure. Ask your security vendors: would your tools flag an attack hosted on a domain serving Epic, Cerner, or Change Healthcare traffic? If detection depends on a domain being flagged, it will miss attacks hosted on infrastructure your stack has already decided to trust, which is where 35% of the threats we blocked in Q1 2026 originated.

About Menlo Security

Menlo Security eliminates evasive threats and protects productivity with the Menlo Cloud. Menlo delivers on the promise of cloud-based security—enabling zero trust access that is simple to deploy. The Menlo Cloud prevents attacks and makes cyber defenses invisible to end users while they work online, reducing the operational burden on security teams.

Menlo protects your users and secures access to applications, providing a complete enterprise browser solution. With Menlo, you can deploy browser security policies in a single click, secure SaaS and private application access, and protect enterprise data down to the last mile. Secure your digital transformation with trusted and proven cyber defenses, on any browser.

MENLO SECURITY

Learn more: https://www.menlosecurity.com
Contact us: ask@menlosecurity.com