When the Attackers Bring AI: 2026 Healthcare Threats in Plain View

Breaking down the **emerging AI-driven cyber threats reshaping healthcare in 2026.**

WEEK 3: VIRTUAL SUMMER SERIES
JULY 8 | 12:00 - 1:00 PM CT

Clearwater

Meeting Logistics

Microphones

All attendees are on mute.

Questions

Type your questions in the Q&A box.

Resources

Upcoming events, slides & resources linked.

Recording

Recording will be provided after event.

Survey

Survey Poll will prompt towards end of webinar.

Clearwater

When the Attackers Bring AI: 2026 Healthcare Threats in Plain View




WEEK 3: VIRTUAL SUMMER SERIES

JULY 8 | 12:00 - 1:00 PM CT

Dave Bailey
Vice President
Consulting Solutions & Strategy

Philip Burnham
Consultant
Penetration Tester, Technical Testing Services

Steve Akers
CTO/CSO
Corporate CISO

Justin Sun
Director
Security Operation Center

WEEK 3 · THE ESCALATION

One Year Later: Building on 2025

2025 · THE EMERGENCE		2026 · OPERATIONAL REALITY
WormGPT and uncensored LLMs made headlines	▸	**Named actors are deploying AI — operations confirmed, not theoretical**
Phishing volume surged — AI-assisted generation emerged	▸	**AI-driven BEC accounts for the majority of high-value fraud losses**
Voice cloning first appeared as a documented threat vector	▸	**Voice cloning has crossed the indistinguishable threshold**
Nation-state actors began experimenting with AI in operations	▸	**Polymorphic malware with AI mutation detected in client environments**
Healthcare sector identified as a high-value target	▸	**Prompt injection tested in clinical AI tools — 94.4% success rate**

AI Is Moving in Three Directions at Once

A New Attack Surface

The chatbots, copilots, and AI agents healthcare organizations already deploy for productivity are being manipulated, chained, or tricked into acting against the very systems they're meant to protect.

A Force Multiplier for Nation-States

Nation-state actors are using AI to scale target identification, credential theft, and social engineering — increasing the speed, reach, and precision of state-sponsored operations against the sector.

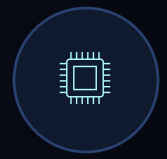

A Capability Outpacing Defenses

Frontier AI vulnerability-discovery capability is advancing faster than most healthcare organizations can adapt their defenses, compressing the time between disclosure and exploitation.

Individually, each of these developments is manageable. Together, they compound into a threat trajectory healthcare leadership cannot address with yesterday's playbook — reinforcing the case for building AI governance, vulnerability management, and incident response capacity now.

Frontier AI Model Day: The Clock Healthcare Is Racing

Frontier AI Model Day — the day frontier, Mythos-class AI vulnerability-discovery capability reaches the adversary groups that target the U.S. Healthcare and Public Health sector — and they begin using it.

1

Mythos NSA Red-Team Exercise

Anthropic's frontier model reportedly found vulnerabilities across classified systems in hours, not weeks.

2

Alibaba-Linked Distillation of Claude

28.8M exchanges used to extract vulnerability-hunting capability from a frontier model.

3

AI-Compressed Attacker Tradecraft

Average breach-to-lateral-movement breakout time is down to just 48 minutes.

!

FRONTIER AI MODEL DAY

The day this capability reaches the adversary groups targeting healthcare — at scale.

"The timeline is not years, it is months."
— Five Eyes intelligence alliance (US, UK, Canada, Australia, New Zealand), joint statement on frontier AI cyber capabilities, June 22, 2026. Signed by NSA's Director of the Cybersecurity Directorate and the acting CISA Director.

U.S. Policy Response: Review, Restriction, and Reversal

A Formal Review Process

- Jun 2, 2026 — White House executive order “Promoting Advanced AI Innovation and Security” creates a voluntary “covered frontier model” review process.
- Directs CISA to expand access to AI-enabled defense tools for critical infrastructure, including rural hospitals.
- Orders DOJ to prioritize AI-enabled cybercrime enforcement.

An Export Ban, Then a Reversal

- Jun 12, 2026 — Claude Fable 5 and Mythos 5 suspended worldwide under a Commerce Department export-control order — the first real-world government-mandated shutdown of a deployed frontier model.
- The trigger: a reported jailbreak surfaced exploit-adjacent output from the models.
- Jun 30 – Jul 1, 2026 — Commerce lifted the export controls; Anthropic restored access to both models.

The lesson for healthcare isn't the specific model or the specific order — it's that frontier AI is now squarely a matter of national security policy, with direct implications for AI vendor dependency, contract risk, and continuity planning.

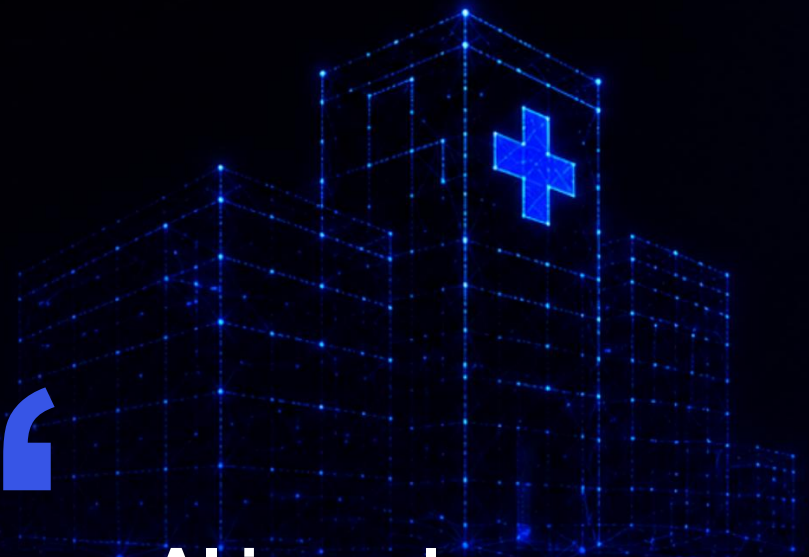

AI is no longer a weapon threat actors tinker with.

It is becoming the first, second and third lines of attack

The shift from research and script writing to full weaponization of AI by Threat Actors is happening faster than security programs can adapt.

Steve Akers | Corporate CISO

Clearwater

The Commoditization Shift: Nation-State Capability at Street Prices

CAPABILITY	PRICE / MONTH
Deepfake-as-a-Service video + audio	$50–$300
WormGPT / FraudGPT dark web LLM	$90–$200
Voice Cloning API clone from seconds of audio	$20
Phishing Kit Generator AI-personalized at scale	$30–$100
Polymorphic Payload Builder self-rewriting malware	$150–$300
Prompt Injection Templates targeting clinical AI	FREE – traded

WHAT THIS MEANS

- Best Player Capabilities for Minor League Pricing
- Deepfakes are only going to gain in believability and success
- More creative threat actors (better scenarios and prompts) will do more damage
- AI Tools ran by threat actors make each attack cost pennies – unfortunately, the inverse is not true

BARRIER TO ENTRY · HEALTHCARE TARGETING

< $200 / month

Sources: Kaspersky · Secureworks · SlashNext · AHA / Anthropic · Clearwater Cyber Briefing

JUSTIN SUN | DIRECTOR, SECURITY OPERATIONS CENTER

From the SOC:

What We're Actually Seeing in Client Environments

Real incidents. Real environments. No hypotheticals.

Clearwater

SEGMENT 3A · JUSTIN SUN · SOC

Case Study: AI-Generated Polymorphic BEC

WHAT CAUGHT IT

Behavioral anomaly detection flagged the **self-modifying payload — not the content**. Signature-based controls never triggered.

What the Payload Looked Like

- Arrived as BEC across multiple client environments
- Complexity far beyond human-authored phishing
- Self-modifying: rewrote itself mid-delivery
- Evaded signature + filter controls

Why We Believe It Was AI-Generated

- Cadence, syntax & exec voice inconsistent with humans
- Adaptive branching — responded to environment feedback
- Polymorphic self-modification = AI hallmark
- Sophistication exceeded the originating infrastructure

Source: Clearwater SOC internal incident data

SEGMENT 3A · JUSTIN SUN · SOC

AI-Enabled Social Engineering: Easy to Fake, Hard to Detect

The capability is trivial now. The defense has not changed.

3 SEC OF AUDIO
to clone a voice (FBI IC3)

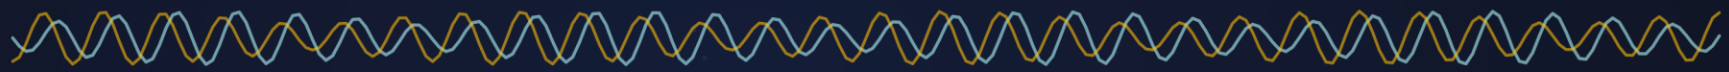

real ≈ clone
indistinguishable

Why It's Trivial Now

- FBI: AI voice cloning needs only ~3 seconds of audio
- FBI PSA (2025): AI-voice impersonation of senior U.S. officials, updated into 2026
- We proved it ourselves — an AI-generated video of our own CISO
- Help-desk vishing now breaches SaaS platforms (UNC6040 / ShinyHunters)

The Defense That Holds

- Verify identity over a DIFFERENT channel than the request arrived on
- Voice call in? Don't confirm by voice. Teams message in? Don't confirm by Teams.
- Out-of-band checks for wire transfers and credential / MFA resets
- Social engineering is back in healthcare's top-3 breach patterns (DBIR 2026)

THE BOTTOM LINE

The control is old and it still works: a second, independent channel. AI changes how convincing the request looks, not how you verify it.

Justin Sun | Director, Security Operations Center

Sources: FBI IC3 PSA I-051525 · Google Threat Intelligence (UNC6040) · Verizon 2026 DBIR

SEGMENT 3B · JUSTIN SUN · SOC

Dark Web & Threat Intelligence

What is actually being traded and targeted against healthcare

WHAT THE INTELLIGENCE SHOWS

- Hospital VPN admin access advertised for sale, tied to U.S. hospital environments (Jan 2026)
- Vishing crews impersonate IT help desk, then exfiltrate and extort (Google Threat Intelligence)
- Healthcare PHI commands premium prices; operational urgency drives fast ransom payment
- Communities run reputation systems, peer reviews, and technical support

CASE IN POINT

UNC6040 / ShinyHunters voice-phished help-desk staff into approving a malicious connected app, then stole Salesforce data for extortion. Google's own instance was hit — roughly 2.55 million records exposed.

The same help-desk playbook maps directly onto hospital IT support.

Source: Google Threat Intelligence Group (GTIG), 2025.

WHAT IS BEING TRADED · HEALTHCARE-RELEVANT

Hospital network access
VPN / RDP admin access to U.S. hospitals
ADVERTISED

Stolen PHI & patient records
Sold per record; healthcare premium
PER-RECORD

Ransomware-as-a-Service
Qilin affiliate model, double extortion
RaaS

Vishing / help-desk impersonation
IT-support social engineering for access
ACTIVE

AI phishing-kit generators
Personalized lures at scale
AI-ENABLED

Sources: Google Threat Intelligence (UNC6040 / ShinyHunters) · Clearwater Cyber Briefing · KELA

The Offensive Lens

What Pen Testers Find When AI Is in Scope

PHILIP BURNHAM | CONSULTANT, PENETRATION TESTER, TECHNICAL TESTING SERVICES

The attackers changed their tools. Have your pen tests changed with them?

Clearwater

AI as a Force Multiplier — Attackers and Pen Testers

What changed in offensive security — and what it means for defense

FOR ATTACKERS

- Recon in minutes: OSINT, target profiling, org-chart mapping
- Auto-adapted scripts iterate on environment feedback
- Skill floor drops — mid-tier actors reach expert-level output
- Payload generation and evasion accelerated at scale

FOR PEN TESTERS

- We set the scope — AI validates and deepens our findings
- AI drafts impact analysis; testers write the actual findings
- Social engineering content AI-assisted; campaign design stays human
- Prompt injection and API configuration now standard scope items

THE DIFFERENCE

Attackers iterate with no rules of engagement. We operate under written authorization, defined scope, and a full evidence trail.

Philip Burnham | Consultant, Penetration Tester, Technical Testing Services

Source: Clearwater Technical Testing Services

Shadow AI: The Inventory Gap

What's actually running in your environment — and what it means for testing

SHADOW AI IN THE WILD

Employees adopt AI tools outside IT visibility — without security review, data classification, or policy coverage.

Every new Windows machine ships with Copilot enabled by default. If your policy says use Claude — has Copilot been disabled?

You cannot govern, test, or defend what you haven't catalogued.

What we find in scope

AI tools deployed without security review · LLM API keys exposed in code repositories · Prompt injection vulnerabilities in internal AI tools

The governance gap

Traditional controls don't catch AI model drift or data poisoning. Any AI with read/write access to records, devices, or communications is a new attack surface.

Where it belongs in scope

Prompt injection testing for any AI with write access · API configuration review — a prompt that can't be injected means nothing if the API endpoint is open and exposed

IMMEDIATE ACTION

Build your AI asset inventory first. You cannot test, govern, or defend what you haven't catalogued.

Philip Burnham | Consultant, Penetration Tester, Technical Testing Services

Source: Clearwater Technical Testing Services

SEGMENT 4C · PHILIP BURNHAM · OFFENSIVE

How We Use AI on Offense

The AI-augmented penetration test, and what it means for your defense

01

RECON

AI validates findings & surfaces patterns. We define scope — the AI doesn't.

02

VULN DISCOVERY

AI scans code & configs for vulnerability families. Testers confirm.

03

EXPLOIT DEV

AI assists with payload drafting & iteration. Testers adapt to environment.

04

SOCIAL ENGINEERING

AI generates phishing & voice content. Campaign design stays human.

05

REPORT

AI polishes impact analysis. Findings & recommendations written by tester.

THE HONEST TAKE

AI accelerates specific phases. The tester's judgment drives every step — from scoping through to sign-off.

What it means for defenders

- Supplement point-in-time tests with continuous validation
- Every AI asset in your environment needs to be in scope
- The 43-day patch cycle is still a real window — AI just finds it faster

How we keep it safe

- Written authorization and defined scope before any AI use
- Every action is logged and evidence-based
- Adversaries use the same tooling with none of these guardrails

Philip Burnham | Consultant, Penetration Tester, Technical Testing Services

Sources: Clearwater Technical Testing Services · Anthropic (GTG-1002)

Defensive Priorities

What Do You Do Now?

For 2026–2027: Field-Tested Recommendations from the Panel

Detection · Response · Testing · Governance

Clearwater

SEGMENT 5 · ALL SPEAKERS · DAVE BAILEY MODERATES

Defensive Priorities: Field-Tested Recommendations

Each speaker's top two actions for healthcare security leaders

90
DAYS TO ACT

SA — Steve Akers
Corporate CISO

1. Treat AI Risk as a named risk category and update governance and compliance posture to match
2. Deploy AI-powered detection and response to match AI-speed attacks — but keep a human in the loop

JS — Justin Sun
Director, Security Operations Center

1. Add AI-specific threat-intel feeds — dark web monitoring for healthcare-targeted tooling
2. Build an IR playbook for AI-generated BEC and voice-clone scenarios — your current plan won't cover it

PB — Philip Burnham
Penetration Tester, Technical Testing

1. Audit pen-test contracts: add prompt-injection testing and AI-specific scope language
2. Inventory every AI-enabled device and map its write-access permissions

DB — Dave Bailey
VP, Consulting Solutions & Strategy

1. Integrate AI risk into your risk framework — model AI threat scenarios in IRM|Pro
2. Brief the board with AI threat data (1,210%, $893M, 94.4%) — not just compliance metrics

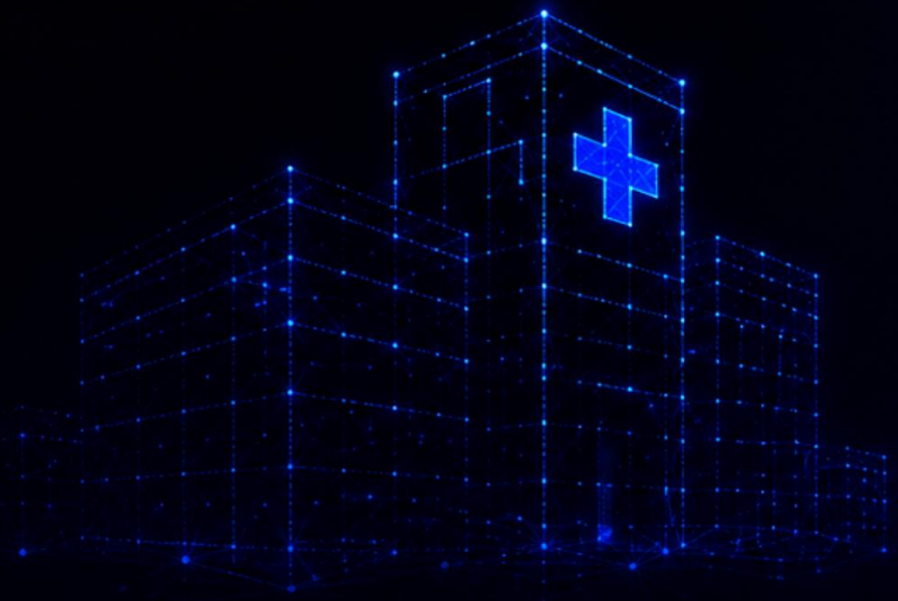

Q&A

Healthcare's Cyber Briefing

CRITICAL THREAT BRIEF FOLLOWED BY

RECOVERY UNDER PRESSURE: WHAT SEPARATES WEEKS FROM MONTHS IN HEALTHCARE INCIDENT RESPONSE

JULY 9 @ 12PM CT!

FEATURING:

David Bailey, VP of Consulting Solutions & Strategy

Heather Hanson, Manager, Resiliency Services

Coming Next

From Commitment to Execution:
Operationalizing AI Governance in Healthcare

WEEK 4: VIRTUAL SUMMER SERIES
JULY 15 | 12:00 - 1:00 PM CT

State of AI in Healthcare:
Where the Industry Actually Is

WEEK 5: VIRTUAL SUMMER SERIES
JULY 22 | 12:00 - 1:00 PM CT

We are here to help.

Moving healthcare organizations to a more secure, compliant, and resilient state so they can achieve their mission.

Clearwater

Healthcare – Secure, Compliant, Resilient

www.ClearwaterSecurity.com

800.704.3394

LinkedIn | linkedin.com/company/clearwater-security-llc/

Legal Disclaimer

Although the information provided by Clearwater Security & Compliance LLC may be helpful in informing customers and others who have an interest in data privacy and security issues, it does not constitute legal advice. This information may be based in part on current federal law and is subject to change based on changes in federal law or subsequent interpretative guidance. Where this information is based on federal law, it must be modified to reflect state law where that state law is more stringent than the federal law or other state law exceptions apply. This information is intended to be a general information resource and should not be relied upon as a substitute for competent legal advice specific to your circumstances. YOU SHOULD EVALUATE ALL INFORMATION, OPINIONS AND RECOMMENDATIONS PROVIDED BY CLEARWATER IN CONSULTATION WITH YOUR LEGAL OR OTHER ADVISOR, AS APPROPRIATE.

Copyright Notice

*The existence of a link or organizational reference in any of the following materials should not be assumed as an endorsement by Clearwater Security & Compliance LLC.