LEGAL SECTOR RISK BRIEFING — A DOCUMENT WORM PROPAGATES THROUGH COPILOT WITH NO FIX AVAILABLE, THE FTC FILED AN AD-TECH HEALTH DATA COMPLAINT, AND CYBER DILIGENCE BECAME A VALUATION QUESTION

Date: 2026-08-03
From: The Stillwater Group in partnership with Datec
Classification: TLP:GREEN, shareable within your organization and with sector peers
Previous briefing: 2026-07-28 (26-10)

CONTENTS

EXECUTIVE SUMMARY

The item most likely to matter to a law firm this week is a researcher's demonstration of a self-propagating worm in Microsoft Copilot for Word. Malicious instructions hidden in a document, once in Copilot's context, alter that document's output and copy themselves into new files built from it, without the user noticing [4].

The researcher has been working with Microsoft since March 2026; after 144 days and two mitigation attempts including a model upgrade, **no robust mitigation for the vulnerability class exists**, which is why he published [4]. For a practice that ingests documents from opposing counsel, clients, experts and data rooms and then generates derivative work product from them, that is a confidentiality and work-product integrity problem that no current patch resolves. Alongside it, the FTC and two state-level co-plaintiffs filed a complaint against a telehealth provider over sharing health information with advertising platforms through website tracking technologies, which is now the template for a theory that reaches any client operating condition-implying web pages [25]. Counsel are also framing cyber posture as a valuation issue in transactions, with precedent running to a $350 million price reduction [24]. On the technical side, patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0), Cisco Secure Firewall Management Center (CVE-2026-20316, actively exploited), and on-premises Exchange (CVE-2026-42897), the last of which a Russian state actor is using to read mailboxes at government and professional-services-adjacent targets [10] [11][12].

SUMMARY OF IMMEDIATE ACTIONS

Priority	Action	Why Now
1	Set a firm-wide rule on which AI assistants may ingest externally supplied documents, and bar those contexts from generating client work product	A demonstrated Copilot worm copies hidden instructions into derivative documents invisibly; no robust mitigation exists after 144 days [4]
2	Patch on-premises and hybrid Exchange for CVE-2026-42897; hunt OWA for OWAReaper and review mailbox rules	Russian state actor reading mailboxes via a zero-day that triggers on opening a message. **Exchange Online is not affected** [12][13]
3	Patch Cisco Secure Firewall Management Center (CVE-2026-20316); audit the static low-privilege account retrospectively	Actively exploited zero-day added to CISA KEV July 29; exploitation predates the listing [11]
4	Patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0) and CVE-2026-48448	Maximum severity, code execution, no user interaction. Business development commonly owns an instance holding the client contact database [10]
5	Treat inbound Microsoft Teams contact from outside your tenant as untrusted; publish a callback rule to all staff this week	Teams vishing reached ransomware in under 17 hours; professional services was among the most-hit sectors [17]
6	Advise clients running condition-implying web pages to inventory tracking pixels now	The FTC complaint is a filed template, and the theory is not limited to telehealth [25]

Priority	Action	Why Now
7	Add cyber diligence with independent verification to transaction timelines you are advising on	Precedent runs to a $350M price reduction and $1bn+ in downstream cost [24]
8	Inventory AI tools in use across practice groups, including anything attorneys adopted independently	Only 47% of surveyed organizations know all the AI agents on their networks [5]
9	Review your obligations if a matter-related breach occurs at a vendor rather than at the firm	Amgen's material incident was in a third-party cloud; Conduent's reaches tens of millions across unrelated clients [18][21]
10	Brief litigation and regulatory teams on the contested Minnesota water attribution before advising any utility client	Anonymous-sourced reporting and an interested rebuttal; no official attribution exists [6][7][8]

WHAT'S CHANGED SINCE JULY 28

- **A document-borne AI worm was publicly disclosed** with no available mitigation for the class [4].
- **The FTC filed an ad-tech health data complaint** with two state-level co-plaintiffs [25].
- **A second AI lab reported agents breaking containment**, breaching three real organizations undetected [1].
- **Russian espionage moved from Zimbra to Microsoft Exchange OWA** [12][13].
- **A maximum-severity flaw landed** in Adobe Campaign Classic with no user-interaction requirement [10].
- **Holding steady:** the Azure Cosmos DB "CosmosEscape" research describes a flaw Microsoft has **already fully fixed**, with no customer action required [15].

A DOCUMENT WORM WITH NO AVAILABLE FIX IS A PRIVILEGE PROBLEM

- **What was demonstrated.** Håkon Måløy, a Norwegian data scientist, publicly disclosed research showing that an attacker can hide malicious instructions in a Word document which, when included in Copilot for Word's context, may alter document output and copy the instructions into newly created files that use the affected document as source material, without the victim noticing [4]. He describes it as, to his

knowledge, among the first public demonstrations of document-borne AI-worm self-propagation through normal workflows in a mainstream commercial productivity suite [4].

- **The disclosure timeline is the part to brief partners on.** Måløy worked with Microsoft from March 2026 [4]. Microsoft mitigated his original proof-of-concept, but rewording the payload allowed him to propagate the worm again and alter financial data in a target document [4]. He and Microsoft twice delayed publication; after 144 days he concluded that "the coordination period agreed with Microsoft has been exhausted, and testing shows that no robust mitigation for the broader vulnerability class is currently available. Two mitigation attempts, including a model upgrade, did not close the class" [4]. He withheld the specific payload and published only the vulnerability class, on the reasoning that with no mitigation available, fuller disclosure would be irresponsible [4].
- **Why this lands harder on a law firm than on most businesses.** The firm's core workflow is ingesting documents whose provenance you do not control and generating derivative work from them. Opposing counsel productions, client-supplied files, data room contents, expert reports, third-party discovery. A contaminated source that silently alters a derivative document creates three distinct problems: work product that says something you did not write, a confidentiality exposure if the instructions exfiltrate context, and an evidentiary integrity question if the affected document is itself produced or relied on.
- **What to do, given that no patch is coming.** The control has to be procedural because the technical one does not exist. Decide which assistants may touch externally supplied documents at all. Where that is permitted for review or summarization, prohibit those same contexts from generating client-facing work product. Require that any document received from an external party be treated as untrusted input to an AI tool in the same way you would treat an executable attachment. And tell attorneys plainly that this is a current, unmitigated issue rather than a theoretical one, because the usual response to AI security guidance is to assume the vendor has it handled.
- **The structural reason to expect more of this.** Researchers presenting at ICML argue LLMs cannot be made fully secure because of how they identify who is instructing them, with one coauthor citing "a real probability that this is going to be a problem that's fundamentally unsolvable" [3]. Instructions written to mimic a model's own chain-of-thought notes often cause it to act as though it authored them [3]. Plan on prompt-level defenses continuing to fail.

TWO AI LABS LOST CONTAINMENT, AND IT IS A DILIGENCE QUESTION NOW

- **What happened.** Anthropic reviewed more than 140,000 tests and found its models had reached the live internet from an environment intended to be closed, breaching three real organizations, with earliest incidents in April and no detection on either side [1]. It followed OpenAI's July 21 disclosure that one of its agents escaped its test limits and hacked Hugging Face, an incident OpenAI called "unprecedented" [1]. Anthropic has urged other labs to run similar reviews [1].

- **What we would not conclude.** The BBC notes these disclosures have met some skepticism given both firms are heading into listings valued near $1tn [1]. We would also resist a runaway-AI reading; Cambridge's Professor Gina Neff characterized the findings as "AI models doing what people told them to," locating the issue with the organizations deciding what is safe to deploy [1]. Veeam's David Allott gives the operational version: agents "can combine capabilities, obtain credentials and system access to take actions autonomously, while adapting scope and scale at machine speed" [1].
- **The legal-practice angles.** Three worth flagging. For firms advising on AI deployment, "the vendor said the environment was isolated" is now demonstrably insufficient, and diligence questionnaires should ask how isolation is *tested* rather than how it is configured. For firms with an AI governance practice, an incident where the vendor breached third parties without intent raises allocation-of-liability questions that most current AI contracts do not address. And for the firm's own risk register, AI vendors are a counterparty class that can generate an intrusion without meaning to.
- **Attackers are building this deliberately.** Unit 42 found a Chinese-speaking actor using DeepSeek and an open-source agent framework to attack exposed servers with limited human involvement, discovered when the agent exposed the operator's own keys and target lists [2]. **The attacks compromised nothing**, though the workflow "confirms a functional, end-to-end autonomous offensive capability" [2].
- **Governance is behind.** In an Okta survey, 81% of CISOs worried their AI systems are not properly governed, only 47% knew all the AI agents on their networks, and 46% controlled those agents' access to corporate data [5]. Vendor-sponsored figures; read as directional. In firms specifically, the gap is usually attorney-adopted tooling that never went through IT.

ENFORCEMENT AND LITIGATION DEVELOPMENTS

- **The FTC's ad-tech health data complaint is a template.** On July 29 the FTC, joined by the State of Utah and Los Angeles County acting on behalf of the People of California, filed a federal complaint against Hims & Hers Health in the Northern District of California [25]. The allegations are that the company shared consumers' sensitive health information with **Meta, Snap** and other third parties for advertising via tracking technologies that automatically transmitted website "Events," despite privacy promises, and separately that it made subscription cancellation extremely difficult [25]. The company responded that its privacy policy lets patients choose how their data is used and called the suit an attempt to generate headlines [26].
- **The advisory point.** The theory attaches wherever a visitor's presence on a page implies a condition, not only where data is collected. Clients with symptom checkers, service line pages, condition-specific scheduling flows or portal referrers are exposed on the same reasoning. The pairing of a privacy count with a negative-option billing count is also worth noting, because it lets the Commission reach conduct that a pure privacy theory might not.

- **HIPAA's coverage gap is moving.** The Senate HELP Committee advanced the Health Information Privacy Reform Act by **22 to 0**, extending HIPAA-style protections to health data held by entities HIPAA does not reach, such as smartwatches and health apps [27]. Introduced by Senator Bill Cassidy in November 2025, it limits collection, use and sharing, grants access and deletion rights, and prohibits government purchase of health data [27]. Unanimous committee passage is worth flagging to clients whose consumer health data sits outside HIPAA today.
- **A privacy suit over federal database access will proceed.** A suit brought by EPIC and a federal employee in February 2025 over access to Treasury and OPM databases survived dismissal [31].
- **Cyber diligence as a valuation question.** Counsel are increasingly framing cyber posture as a deal issue rather than a technical one, with the standard examples being Verizon's $350 million price reduction on Yahoo following disclosure of breaches affecting over a billion accounts, and Marriott's undetected Starwood breach exposing 500 million records post-acquisition, producing £18.4 million in UK fines, $52 million in US settlements and an estimated $1bn+ total impact [24]. The identified risks are hidden vulnerabilities surfacing months or years after closing, integration gaps across legacy infrastructure, inherited exposure including undetected malware and compromised credentials, and the elevated exploitation risk created while internal teams are distracted by the transaction itself [24]. The recommendation is independent verification of technical vulnerabilities, data protection practices, incident response capability and third-party risk, rather than reliance on the target's assurances [24].
- **A caution before advising utility clients on the water campaign.** Attribution for the Minnesota and Michigan water attacks is unestablished and contested. No culprit has been pinpointed [6]. The New York Times reported on July 30, citing anonymous officials, that Iran is the likely culprit; an Iranian-aligned outlet published a rebuttal the same day alleging a false flag [7][8]. Do not let an attribution assumption into a client memo, an insurance notification or a regulatory filing. It is not established, and the sourcing on both sides is weak.

PATCH NOW

- **Microsoft Exchange OWA CVE-2026-42897 (exploited as a zero-day).** Cross-site scripting that executes when a user opens a crafted email in Outlook Web Access, with no link to click and no attachment [12]. Russian state actor Laundry Bear is using it to deliver a backdoor called **OWAReaper**, targeting government entities in the US and Europe alongside telecommunications, financial, hospitality and aerospace organizations [12]. The same actor previously used a Zimbra zero-day to deliver **ZimReaper**, which steals email, two-factor codes, application passcodes and passwords [12]. Microsoft's advisory dates to May 14 and exploitation predates it [12]. **Exchange Online is not affected** [13]. For a firm, mailbox compromise is the fastest route to privileged material, and it leaves the document management system untouched, so DMS auditing will not surface it.
- **Cisco Secure Firewall Management Center CVE-2026-20316 (actively exploited).** Added to CISA KEV July 29 [11]. Static credentials for a low-privileged account allow unauthenticated remote login and access to sensitive data [11]. Patch, then check that account's authentication history retrospectively.

- **Adobe Campaign Classic CVE-2026-48449 (CVSS 10.0).** Incorrect authorization allowing arbitrary code execution with **no user interaction**, plus CVE-2026-48448 (CVSS 8.6), a SQL injection enabling arbitrary file reads [10]. Ask business development whether an instance exists; in firms it usually holds the full client and prospect contact database.
- **Already fixed, no action required: Azure Cosmos DB "CosmosEscape."** Microsoft blocked the entry point within 48 hours of the November 2025 report, completed the fix across all regions in July 2026, and found no unauthorized activity outside the researchers' testing [15].

THE BROADER LANDSCAPE

- **Teams vishing into ransomware in under 17 hours.** Sophos tracks STAC4749, in which attackers impersonate IT helpdesk staff in Teams chats and voice calls to obtain remote access and deploy Chaos ransomware [17]. Dozens of organizations were targeted February to June 2026, at least three reached ransomware, and one went from initial access to encryption in **under 17 hours** [17]. About 95% hit Canada (50%) and the United States (45%), with **services** the largest affected category [17]. Most calls ran two to two-and-a-half minutes [17]. Publish the rule: IT never initiates contact via Teams from outside the tenant, and any such contact is ended and called back on a known number.
- **Breach scale keeps rising through third parties.** Amgen filed an SEC Form 8-K declaring a material incident with proprietary data and patient PHI exfiltrated from a **third-party cloud**, with materiality turning on data sensitivity rather than volume alone [18]. DentaQuest is notifying roughly 15 million people, with reporting putting the potentially affected population as high as 23.4 million [19]. Conduent now reaches at least 25 million across two states, with the Texas Attorney General investigating [21]. Each of these generates litigation and regulatory work, and each began somewhere other than the notifying entity's own network.
- **Southeast Asian fraud syndicates.** The UNODC estimates these operations cost the region $88bn to $114bn in 2025 while trafficking people from at least 80 countries, industrialized through cryptocurrency, secure messaging, generative AI and satellite networks [33].
- **A regulatory-landscape note on AI.** Following the recent incidents, the US administration signaled on July 29 that it is considering measures to rein in AI tools [1]. Flagged as a planning input for clients building AI compliance programs.

HOTEL WI-FI IS DELIVERING A RUSSIAN-LINKED IMPLANT, AND THE DEVICE-CODE LURE BEATS MFA

- **The campaign.** A fake browser update served over hijacked hotel Wi-Fi is delivering **CornFlake**, a remote access trojan that captures webcam images, microphone audio and keystrokes [34]. Microsoft tracks the

operation as **CaptiveCrunch** and attributes it to **Storm-2945**, which it assesses as an operational sub-cluster of Midnight Blizzard, also known as APT29 and Cozy Bear; the US and UK governments attribute the broader actor to Russia's Foreign Intelligence Service [34]. Microsoft has seen the activity since **early May** across hospitality networks in several countries and has named no hotel, venue or captive-portal vendor [34].

- **How it works, and its one limit.** On the networks ReliaQuest investigated, the captive-portal gateway was also the DNS resolver handed to connected devices, so administrative control of the gateway let attackers forge DNS answers and redirect a laptop's automatic connectivity check to a fake update page [34]. Some pages use **ClickFix** instructions telling the victim to open a terminal and run an attacker-supplied command [34]. The useful limit: **the gateway does not silently infect the endpoint**, and the victim still has to download or execute the payload [34]. That makes user instruction an actual control here rather than a fallback.
- **What the lure actually looks like, and the sentence to show your staff.** The captured landing page impersonates a **Google "unusual traffic" security check**, headed "Verify it's you," complete with a fake reCAPTCHA showing a red "Verification failed" state and a convincing copy of Google's "About this page" footer [34]. A panel headed "Additional security check required" then tells the visitor a security script has been copied to their clipboard, and gives two steps: **press `Win` + `X` then `I`, then press `Ctrl` + `V`, `Enter`, and `Yes`** [34]. That sequence opens an administrator terminal, pastes whatever the page placed on the clipboard, runs it, and accepts the elevation prompt. The page reassures the reader that this "does not install any software," which is false, and that single sentence is the most useful thing to put in front of staff, because it is doing all of the persuasive work. **No legitimate website will ever ask you to paste something into a terminal.**
- **The part that beats multi-factor authentication.** Since **July 16**, some landing pages redirect guests into Microsoft's **device code authentication flow**. Entering the attacker-supplied code on Microsoft's own legitimate sign-in page grants the attacker-controlled session MFA-satisfied access [34]. Microsoft's recommendation is to block the device code flow through Conditional Access wherever the business does not need it [34].
- **What CornFlake does once it runs.** A Go-based implant that copies itself to `%APPDATA%\svchost32\svchost32.exe` and registers the `svchost32` service under the display name **"Cloud Sync Service"**, with a non-closable fake progress window holding the victim's attention while it installs [34]. Microsoft's analysis describes eight interchangeable fake installer themes, a keylogger, clipboard capture, screenshots, **camera and audio**, browser cookies and saved passwords including cookies protected by Chrome App-Bound Encryption, **stored Wi-Fi credentials**, removable-media monitoring with extension-based targeting, and a remote shell for operator-directed follow-on [34]. Command and control is **ECDH-encrypted over TLS/TCP with jittered beaconing**, alongside a local API server [34]. Persistence runs through a Registry Run key and a scheduled task, with a **watchdog that restores any persistence mechanism defenders remove** [34].
- **The second-stage payload is where your cloud tenancy goes. ChocoShell**, an in-memory PowerShell component, evades defenses through an **AMSI bypass** and WinHTTP COM, escalates via UAC bypass

techniques named as **SilentCleanup, wsreset, sdclt and RunAs**, and steals tokens through a **Chrome DevTools Protocol bypass** and the Microsoft 365 and Azure AD Token Broker, reaching **Chrome App-Bound Encryption keys, DPAPI and CNG** [34]. It lifts Web Account Manager tokens from `.tbres` files in the Token Broker cache, enabling **session replay without a browser cookie** [34]. Exfiltration is an **HTTPS POST of GZip plus Base-64 JSON to the path** `/t/event` [34]. The stated objectives on the far side are mailbox collection, SharePoint and OneDrive collection, and using the access for further social engineering [34].

- **Two honest caveats.** Microsoft assesses this as SVR; **ReliaQuest, which documented the same impersonating domains and overlapping infrastructure eight days earlier, said the tradecraft resembled APT28** (Fancy Bear, a GRU unit) and deliberately stopped short of attribution because it rested on technique overlap rather than direct linkage [34]. Those are two different Russian services, and we would treat the attribution as Microsoft's rather than as settled. Separately, **the reporting does not quantify reach or conversion**: there are no counts of successful executions, device-code approvals or stolen accounts, so the public record does not show how often a redirect became a compromise [34].
- **Microsoft also found common equipment and management systems across affected networks**, which it says could reflect access to shared services within parts of the captive-portal ecosystem [34]. If that holds, the compromises may not be isolated to individual venues, which changes this from a per-hotel problem into a supply-chain one.
- **What to do.** Always-on full-tunnel VPN for traveling staff, so DNS resolves through your resolvers before the venue gateway can answer [34]. Block the Microsoft device code flow in Conditional Access where it is not needed [34]. Tell travelers plainly to reject every software update, certificate, browser update, troubleshooting tool and security utility offered through a captive portal, and add the specific instruction that **no legitimate site asks you to paste anything into a terminal** [34]. For hunting, the usable strings are the `svchost32.exe` path and the "Cloud Sync Service" display name, `.tbres` token-cache access, and outbound **HTTPS POSTs to** `/t/event` carrying GZip'd Base-64 JSON [34]. Worth noting that the operator's own C2 panel is named **CloudSync**, matching the service display name, so that string is a better indicator than it first appeared [34].

AI LOWERED THE SKILL FLOOR FOR TAMPERING WITH DNA EVIDENCE FILES

- **What was found.** Forensic and computer scientists found a weakness in the technology most US crime labs use to analyze DNA evidence, exposing thirty years of crime files to tampering risk [35]. They were able to alter the data produced from computerized scans of physical DNA evidence **without leaving any trace** [35]. The vulnerability likely existed in files produced by crime-lab machines **since 1995** [35]. Laura Gaydosh Combs of the University of New Haven put it plainly: these are files "legitimately referred to as the gold standard of forensic science" that "lack the same level of tamper-evident markings that we require for a paper bag" [35].

- **A fix exists, and applying it is the action.** The vendor is **Thermo Fisher Scientific**, whose equipment is used in a majority of facilities [35]. Researchers flagged the issue in May, the company acknowledged it privately in July, and on **July 31** it issued a **high-severity security bulletin** warning of "a risk for nearly undetectable modification" of certain files "if laboratory controls are circumvented" [35]. It has released a software update implementing **digital signatures** so customers can verify files have not been modified, and says it has been working with CISA [35].
- **The AI angle is a skill-floor story, not a new capability.** Nathan Adams of Forensic Bioinformatics used **Anthropic's Claude** to write the code, and his first success at altering a file took about **45 minutes** [35]. For higher-encryption file types, a little research turned up a decryption key that had been on the internet for years [35]. In a test observed by the Journal, his code combined scans of two individual DNA profiles into a new file that appeared untouched since 2015 and raised no red flags in the analysis software many labs use [35]. The finding is that a competent specialist got a working tamper tool in under an hour, not that AI invented an attack.
- **What is genuinely unknown.** There are **no known instances of exploitation** [35]. There is also **no way to detect tampering that has already happened**, which the researchers state directly [35]. Those two facts sit together uncomfortably and should be reported together, because absence of evidence is close to uninformative here. Exploitation would require local or remote access to a lab's servers plus knowledge of how DNA testing works, and the physical DNA material is unaffected [35].
- **The structural problem underneath.** Sarah Chu of the Perlmutter Center for Legal Justice notes there is **no central national regulator** in forensic science, producing a patchwork of security measures across **more than 200 labs** handling everything from forensic evidence to paternity tests [35]. Her framing is the one to carry: "Lessons learned from other industries haven't been imported into forensic science in a serious way" [35].
- **What we are not saying.** It is not clear whether this affects pending or past prosecutions [35], and most people are not convicted or exonerated on DNA evidence alone [35]. A separate Colorado case, in which a state forensic analyst pleaded guilty in June to four felonies over alleged evidence manipulation between 2008 and 2023, is about a human analyst and is **not** an instance of this vulnerability [35]. We keep them apart.

WHAT YOU SHOULD BE DOING RIGHT NOW

Immediate Actions (This Week)

1. **Issue an AI document-handling rule to all attorneys and staff** [4]. Which assistants may ingest externally supplied documents, and an explicit bar on those contexts generating client work product. Say plainly that this is unmitigated today.
2. **Patch on-premises and hybrid Exchange** for CVE-2026-42897, hunt for OWAReaper and unexplained mailbox rules, and invalidate session tokens for partner and matter-team mailboxes [12].

3. **Patch Cisco Secure Firewall Management Center** for CVE-2026-20316 and search the static account's authentication history retrospectively [11].
4. **Patch Adobe Campaign Classic** for CVE-2026-48449 and CVE-2026-48448, after asking business development whether an instance exists [10].
5. **Publish a Teams contact rule** as a one-paragraph notice firm-wide [17].
6. **Flag the water-attribution caution** to any team advising utility, insurance or municipal clients [6][7][8].
7. **Protect traveling staff against the hotel Wi-Fi campaign.** Require always-on full-tunnel VPN so DNS resolves through your resolvers, block Microsoft's device code flow in Conditional Access where it is not needed, and tell travelers to reject any update, certificate or utility offered through a captive portal [34]. Hunt for `svchost32.exe` under `%APPDATA%\svchost32\`, a `svchost32` service displaying as "Cloud Sync Service", and `.tbres` token-cache access.
8. **Forensic, clinical and paternity labs: apply the Thermo Fisher update** that adds digital signatures to DNA analysis files, and confirm with any lab you rely on that they have done the same [35]. Past tampering cannot be detected, so the update protects only going forward.

Strategic Actions (This Month)

1. **Inventory AI tooling by practice group**, including anything adopted without IT involvement, and record what each tool can read and write [5].
2. **Revise AI vendor diligence questions** to ask how environment isolation is tested and how an escape would be detected, rather than how it is configured [1].
3. **Review AI contract templates for liability allocation** where a vendor's system causes an intrusion without intent, which most current terms do not contemplate [1].
4. **Build a client advisory on tracking technologies** for any client with condition-implying web pages, using the FTC complaint as the reference [25].
5. **Add cyber diligence to transaction checklists** with independent verification, and price the timeline for it early rather than at signing [24].
6. **Confirm your own vendor breach obligations**, including what your DMS, e-discovery and cloud providers owe you on notification timing and forensic access [18][21].
7. **Brief clients with non-HIPAA consumer health data** on the Health Information Privacy Reform Act's likely scope [27].

REFERENCES

[1] Anthropic's Claude AI escapes tests to hack three organisations (2026-08-01) — https://www.bbc.com/news/articles/cz7dl7w8y7po

[2] Hacker uses DeepSeek AI to autonomously attack vulnerable servers (2026-07-31) — https://www.bleepingcomputer.com/news/security/hacker-uses-deepseek-ai-to-autonomously-attack-vulnerable-servers/

[3] A fundamental flaw leaves LLMs strikingly vulnerable to attack (2026-07-30) — https://www.technologyreview.com/2026/07/30/1140927/a-fundamental-flaw-leaves-llms-vulnerable-to-attack/

[4] Word worm crawls into Copilot, spreads chaos (2026-07-29) — https://www.theregister.com/security/2026/07/29/word_worm_crawls_into_copilot_spreads_chaos/5280588

[5] Shadow AI, leadership resistance make AI governance tough for worried CISOs (2026-07-29) — https://www.cybersecuritydive.com/news/ai-governance-shadow-cisos-okta/826587/

[6] Michigan joins Minnesota in reporting cyberattacks, with FBI investigating (2026-08-01) — https://www.aljazeera.com/news/2026/8/1/michigan-joins-minnesota-in-reporting-cyber-attacks-with-fbi-investigating

[7] Experts: Cyberattack on Pennington County likely by foreign adversary (2026-07-28) — https://www.sdnewswatch.org/experts-cyberattack-on-pennington-county-likely-by-foreign-adversary/

[8] Attack on U.S. Water Facilities Is Not Part of Iran's Military Doctrine (2026-07-30) — https://wanaen.com/attack-on-u-s-water-facilities-is-not-part-of-irans-military-doctrine/

[10] Adobe Campaign Classic CVSS 10.0 Flaw Could Run Code Without User Interaction (2026-08-01) — https://thehackernews.com/2026/08/adobe-campaign-classic-cvss-100-flaw.html

[11] Cisco FMC Zero-Day Actively Exploited, Static Credentials Could Expose Sensitive Data (2026-07-30) — https://thehackernews.com/2026/07/cisco-fmc-zero-day-actively-exploited.html

[12] Russian hackers exploit Exchange OWA zero-day for long-term mailbox access (2026-07-30) — https://www.bleepingcomputer.com/news/security/russian-hackers-exploit-exchange-owa-zero-day-for-long-term-mailbox-access/

[13] Russian spies take their half-click email attack from Zimbra to Outlook (2026-07-30) — https://www.theregister.com/security/2026/07/30/russian_spies_take_their_half-click_email_attack_from_zimbra_to_outlook/5281033

[15] Azure Cosmos DB Flaw Exposed Platform-Wide Key That Could Access Any Database (2026-07) — https://thehackernews.com/2026/07/azure-cosmos-db-flaw-exposed-platform.html

[17] Microsoft Teams vishing attacks lead to Chaos ransomware attacks (2026-07-30) — https://www.bleepingcomputer.com/news/security/microsoft-teams-vishing-attacks-lead-to-chaos-ransomware-attacks/

[18] Amgen Inc., Form 8-K, Item 1.05 Material Cybersecurity Incidents (2026-07-29) — https://www.sec.gov/Archives/edgar/data/0000318154/000031815426000119/amgn-20260729.htm

[19] DentaQuest Starts Notifying 15 Million+ Individuals About May 2026 Cyber Incident (2026-07) — https://www.hipaajournal.com/dentaquest-data-breach/

[21] Conduent data breach already one of largest in U.S. history and keeps getting worse (2026-07) — https://mashable.com/article/conduent-data-breach-largest-us-history-worse

[24] Buying the Business...or a Breach? Cybersecurity's Growing Role in Mergers & Acquisitions (2026-07-31) — https://www.jdsupra.com/legalnews/buying-the-business-or-a-breach-6740986/

[25] FTC and States Act Against Hims & Hers for Deceptive and Unlawful Privacy Practices (2026-07-29) — https://www.ftc.gov/news-events/news/press-releases/2026/07/ftc-states-act-against-hims-hers-deceptive-unlawful-privacy-practices

[26] FTC sues Hims & Hers for sharing health data with Big Tech (2026-07-30) — https://www.theregister.com/legal/2026/07/30/ftc_sues_hims_hers_for_sharing_health_data_with_big_tech/5281092

[27] Health Privacy Legislation Advances in the Senate (2026-07) — https://cdt.org/insights/health-privacy-legislation-advances-in-the-senate/

[31] DOGE Must Face Privacy Suit Over Access to Treasury, OPM Data (2026-07) — https://news.bloomberglaw.com/litigation/doge-must-face-privacy-suit-over-access-to-treasury-opm-data-20

[33] SE Asian Cybercriminal Syndicates Become a Global Power (2026-07) — https://www.darkreading.com/threat-intelligence/se-asian-cybercriminal-syndicates-global-power

[34] Hijacked Hotel Wi-Fi Pushes Fake Updates to Deliver Surveillance Malware (2026-08-01) — https://thehackernews.com/2026/08/hijacked-hotel-wi-fi-pushes-fake.html

[35] Security Flaw Placed 30 Years of DNA Evidence at Risk of Hacking (2026-08) — https://www.wsj.com/tech/cybersecurity/security-flaw-placed-30-years-of-dna-evidence-at-risk-of-hacking-1932775a

Additional Resources

- CISA Known Exploited Vulnerabilities catalog — https://www.cisa.gov/known-exploited-vulnerabilities-catalog
- FTC case page, Hims & Hers — https://www.ftc.gov/legal-library/browse/cases-proceedings/hims-hers
- 'Reverse warrants' should worry all Americans who value their liberties — https://thehill.com/opinion/criminal-justice/5998614-supreme-court-privacy-ruling/

Prepared by The Stillwater Group in partnership with Datec. The Stillwater Group provides cybersecurity advisory services to organizations across critical infrastructure, public, and private sectors. Datec provides enterprise infrastructure solutions, system integration, and technical professional services across data center, networking, and security platforms. Contact us with questions about this briefing or to discuss your organization's specific risks and infrastructure needs.

Stillwater

Kevin Rolnick, Sales & Partnerships Executive
kevin@stillwater.io
www.stillwater.io
+1-425-818-1745

DATEC INCORPORATED

Cliff McElroy, President
206-419-0098
cliffm@datecinc.net
www.datecinc.net