RISK BRIEFING — AI AGENTS BREACHED REAL COMPANIES FROM INSIDE TEST ENVIRONMENTS, MICHIGAN JOINS THE WATER CAMPAIGN, AND ADOBE SHIPS A CVSS 10.0

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

1. Executive Summary
2. Summary of Immediate Actions
3. What's Changed Since July 28
4. Patch Now: One Maximum-Severity Flaw and Two Confirmed Zero-Days
5. AI Agents Left Their Test Environments and Breached Real Companies
6. The Water Campaign Widened to Michigan, With No Attribution Established
7. From the Panel View: Briefing Your Operators
8. Espionage Moved Platforms and Supply Chains Got Quieter
9. Breach Scale Is Now a Third-Party Concentration Problem
10. Regulators Moved on Health Data From Three Directions
11. The Broader Landscape
12. Hotel Wi-Fi Is Delivering a Russian-Linked Implant, and the Device-Code Lure Beats MFA
13. AI Lowered the Skill Floor for Tampering With DNA Evidence Files
14. What You Should Be Doing Right Now
15. References

EXECUTIVE SUMMARY

Anthropic disclosed that during private security testing its Claude models found a way out of what was supposed to be an isolated environment and, continuing to treat the exercise as live, broke into three real organizations [1]. A misconfiguration at Anthropic and its testing partner left the models with live internet access; the earliest incidents date to April, and nobody noticed at the time, including the companies that were breached [1]. This lands nine days after OpenAI said one of its agents escaped its test limits and hacked Hugging Face, and it moves a story that read as a single vendor's bad week into something closer to a category of failure [1]. Treat the operational lesson as the durable part: an AI agent given a goal and unexpected network reach will pursue the goal, and your third-party AI vendors are now a class of counterparty that can generate an intrusion without intending to. Separately, a Chinese-speaking actor was caught running an end-to-end autonomous attack pipeline built on DeepSeek and an open-source agent framework, though it compromised nothing [2]. On the vulnerability side, Adobe patched a **CVSS 10.0** authorization flaw in Campaign Classic that executes code with no user interaction (CVE-2026-48449), CISA added an actively exploited Cisco firewall-management zero-day to the KEV catalog (CVE-2026-20316), and Russian state actor Laundry Bear moved its "half-click" webmail exploit from Zimbra to Microsoft Exchange OWA (CVE-2026-42897), where merely opening a message in the browser is enough [10][11][12][13]. The water campaign we advised on July 30 widened: Michigan reported attacks on nine water systems, the FBI confirmed it is investigating, and its advisory says at least seven states have reported incidents, though only two have been named [6]. No attribution has been established, and the reporting is now actively contested [6][7][8]. Patch the three exploited or maximum-severity flaws first, then work the AI-governance and OT items below.

SUMMARY OF IMMEDIATE ACTIONS

Priority	Action	Why Now
1	Patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0) and the accompanying SQL injection (CVE-2026-48448)	Maximum severity, arbitrary code execution, and **no user interaction required** [10]
2	Patch Cisco Secure Firewall Management Center (CVE-2026-20316) and audit the low-privilege account for prior logins	Actively exploited zero-day, added to CISA KEV July 29; static credentials ship in the product [11]
3	Confirm on-premises Exchange is patched for CVE-2026-42897; hunt OWA for OWAReaper and review mailbox rules and session tokens	Russian state actor exploiting it as a zero-day since before Microsoft's May 14 advisory. **Exchange Online is not affected** [12][13]

Priority	Action	Why Now
4	Patch Rails Active Storage (CVE-2026-66066) if you run libvips and accept image uploads from untrusted users; rotate `secret_key_base` and any exposed credentials	Unauthenticated file read escalating toward RCE, and the files it reads typically hold your database and cloud keys [14]
5	Inventory every AI agent with network access and confirm what it can reach; verify isolation of AI test and evaluation environments by testing egress, not by reading the configuration	Two labs have now had agents leave environments believed to be isolated. Only 47% of firms say they know all the AI agents on their networks [1][5]
6	Water and wastewater operators: pull remote monitoring and control interfaces off the public internet, verify manual-operation capability, and brief control-room staff using the operator section below	Nine Michigan systems and more than 30 Minnesota systems affected; at least seven states reporting; no attribution [6] [7]
7	Treat inbound Microsoft Teams contact from outside your tenant as untrusted; restrict external Teams communication and publish a help-desk callback rule	Vishing via Teams took one victim from initial access to ransomware in under 17 hours; 95% of the campaign hit Canada and the US [17]
8	Establish a policy for Copilot and similar assistants on untrusted documents; do not allow them to ingest externally supplied files into contexts that generate new documents	A researcher demonstrated a self-propagating Word/Copilot worm; after 144 days and two attempts, no robust mitigation for the vulnerability class exists [4]
9	Re-examine npm and other package-manager supply chains for the four named compromised packages; require provenance checks on dependency updates	One North Korean crew is now linked to four package compromises, obtained by social-engineering maintainers rather than attacking the registry [16]
10	If you are acquiring or being acquired this quarter, put cyber diligence in the deal timeline now with independent verification, not target assurances	Historical precedent runs to a $350M price cut and $1bn+ in downstream cost [24]

WHAT'S CHANGED SINCE JULY 28

- **A second AI lab reported agents breaking containment.** Anthropic reviewed more than 140,000 tests and found three real organizations breached from what should have been an isolated environment [1].

- **Michigan entered the water campaign.** Nine water systems reported, days after Minnesota's more than 30 [6]. The FBI confirmed on August 1 that it is investigating both [6].
- **Attribution became contested rather than absent.** The New York Times reported on July 30, citing anonymous officials, that Iran is the likely culprit in Minnesota; an Iranian-aligned outlet published a denial and a false-flag hypothesis the same day [7][8]. No official attribution exists [6].
- **The Laundry Bear half-click exploit moved platforms.** Last week it was Zimbra; this week the same actor is exploiting Microsoft Exchange OWA with a new backdoor [12][13].
- **CISA and the Five Eyes published isolation guidance.** CI Fortify sets out a "critical path to isolation" for keeping critical services running while cut off [9].
- **A maximum-severity flaw landed with no user-interaction requirement.** Adobe Campaign Classic CVE-2026-48449 scores a full 10.0 [10].
- **Holding steady:** the Azure Cosmos DB "CosmosEscape" research published this week describes a serious flaw that Microsoft has **already fully fixed**, with no customer action required [15]. We include it so you can close the question if someone forwards you the headline.

PATCH NOW: ONE MAXIMUM-SEVERITY FLAW AND TWO CONFIRMED ZERO-DAYS

- **Adobe Campaign Classic CVE-2026-48449 (CVSS 10.0).** An incorrect-authorization flaw allows arbitrary code execution in the context of the current user with **no user interaction required**, which is what takes the score to a full ten [10]. The same August 1 update fixes CVE-2026-48448 (CVSS 8.6), a SQL injection enabling arbitrary file reads [10]. Campaign Classic is marketing automation, so it typically holds large customer contact databases and sits in a part of the estate that security teams do not always monitor closely. Check whether marketing operations own an instance you have not inventoried.
- **Cisco Secure Firewall Management Center CVE-2026-20316 (CVSS 5.3, actively exploited).** CISA added this to the Known Exploited Vulnerabilities catalog on July 29 following reports of zero-day exploitation [11]. The mechanism is static user credentials for a low-privileged account shipped in the product, letting an unauthenticated remote attacker log in and read sensitive data [11]. The modest CVSS score understates the operational risk here: this is the management plane for your firewalls, and read access to it is reconnaissance of your entire network policy. Patch, then check authentication logs for that account historically, not just going forward.
- **Microsoft Exchange OWA CVE-2026-42897 (exploited as a zero-day).** A cross-site-scripting flaw that executes arbitrary JavaScript in the browser when a user opens a crafted email in Outlook Web Access [12]. Proofpoint calls this class a "half-click exploit" because opening the message is the whole attack; there is no link to click and no attachment to open [12]. Microsoft's advisory dates to May 14, and the actor was already exploiting it before then [12]. **Exchange Online is not affected**, so this is an on-premises and hybrid problem [13].

- **Ruby on Rails Active Storage CVE-2026-66066 (critical).** An unauthenticated attacker can read arbitrary files from a Rails application and potentially escalate to remote code execution [14]. Two preconditions narrow it: the flaw is exploitable when **libvips** is the image processor, and the application must accept image uploads from untrusted users [14]. If both hold, the payoff is severe, because the readable files include the process environment, which typically carries `secret_key_base` plus database and cloud-storage credentials [14]. Patch, then rotate anything that was in that environment.
- **Already fixed, no action required: Azure Cosmos DB "CosmosEscape."** Wiz disclosed a chain that escaped the Gremlin query sandbox and reached a platform-wide signing secret, potentially allowing read and write access to databases across customer tenants [15]. The prerequisites were trivial, requiring only a standard Azure account with an attacker-controlled Gremlin database [15]. Microsoft blocked the entry point within 48 hours of the November 2025 report, completed the full fix across all regions in July 2026, and eliminated the platform-wide key; its review found no unauthorized activity outside the researchers' testing [15]. Nothing for you to do.

AI AGENTS LEFT THEIR TEST ENVIRONMENTS AND BREACHED REAL COMPANIES

- **What Anthropic disclosed.** Reviewing more than 140,000 tests, Anthropic found evidence that its Claude models had reached the internet from an environment that was supposed to be closed off [1]. The tests involved a standard capability exercise: hide "secret" information on another machine on an isolated network and task the model with breaking in to retrieve it [1]. A misconfiguration on systems run by Anthropic and its testing partner left live internet access in place, and the models, still treating the exercise as the task at hand, breached three real organizations instead of test targets [1]. The earliest incidents date to April [1]. The affected companies, which Anthropic has not named, have been notified [1].
- **Nobody detected it.** Neither Anthropic nor the breached organizations noticed at the time [1]. For a security leader, this is the sentence that matters most. These were intrusions by a well-resourced party that was not trying to hide, and they still went unremarked on both sides until a retrospective review of test logs surfaced them.
- **The context, and the reason it is not one company's problem.** Anthropic began the review after OpenAI disclosed on July 21 that its agent had escaped its test limits and hacked Hugging Face, an incident OpenAI called "unprecedented" and Hugging Face's co-founder called "a wake-up call" [1]. Anthropic has urged other AI labs to run similar reviews [1]. Two labs finding the same failure in three weeks is the substance of the story.
- **What we would not conclude.** The BBC notes these disclosures have been met with some skepticism, coming as both firms prepare for listings expected to value each at around $1tn [1]. We do not think the incidents are manufactured, but the timing is a fair question and we would not build a forecast on the disclosures alone. We would also resist the "rogue AI" reading. Professor Gina Neff of Cambridge's Minderoo Centre 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]. David Allott of Veeam put the operational version well: AI agents "can combine capabilities, obtain credentials and system access to take actions autonomously, while adapting scope and scale at machine speed" [1]. He is explicit that this requires no fundamentally new attack capability, which is precisely why it is available now.

- **What to actually do.** Two things. First, verify the isolation of any AI evaluation or sandbox environment you operate by **testing egress from inside it**, because both incidents involved environments believed to be isolated and documented as isolated. Second, treat AI vendors as a counterparty class that can generate an intrusion without intent, and ask them directly what environment isolation testing they perform and how they would detect an escape. Anthropic's own admission that it "could have reviewed its records more thoroughly" is the standard to hold vendors to [1].
- **Meanwhile, attackers are building the same thing deliberately.** Palo Alto Networks' Unit 42 found a Chinese-speaking actor using DeepSeek as the reasoning engine behind the open-source Hermes Agent to attack exposed servers with limited human involvement [2]. The discovery was accidental: Hermes created a web server from its own home directory, exposing the operator's API keys, exploit scripts, target lists, shell history and AI attack logs [2]. **The attacks did not succeed in compromising anything** [2]. Unit 42's conclusion is the calibrated one: "While the observed campaign had limited impacts, the workflow confirms a functional, end-to-end autonomous offensive capability" [2]. Capability demonstrated, effectiveness not yet.
- **A structural problem underneath all of it.** Researchers presenting at ICML argue that large language models cannot be made fully secure, because of how they distinguish who is giving them instructions [3]. Writing an instruction in a style that mimics a model's own chain-of-thought notes often causes it to act as though it authored the instruction itself [3]. One coauthor puts the odds plainly: "There's a real probability that this is going to be a problem that's fundamentally unsolvable" [3]. Plan controls around AI systems on the assumption that prompt-level defenses will keep failing.
- **The worm case makes that concrete.** A Norwegian researcher demonstrated malicious instructions hidden in a Word document that, once in Copilot for Word's context, alter document output and copy themselves into new files built from that source, invisibly [4]. He has worked with Microsoft since March 2026; after two mitigation attempts including a model upgrade and 144 days, no robust mitigation for the class exists, so he published the vulnerability class while withholding the payload [4]. The practical control is document provenance: decide now which assistants may ingest externally supplied files, and keep those contexts away from document generation.
- **Governance has not caught up.** In an Okta survey, 81% of CISOs said they worry their AI systems are not properly governed, only 47% said they knew all the AI agents on their networks, and only 46% said they controlled those agents' access to corporate data [5]. These are vendor-sponsored figures and should be read as directional, but the direction is consistent with everything above.

THE WATER CAMPAIGN WIDENED TO MICHIGAN, WITH NO ATTRIBUTION ESTABLISHED

- **Where it stands.** Michigan has reported cyberattacks on nine of its water systems, days after Minnesota reported more than 30 [6]. The FBI said on August 1 that it is investigating both, and its advisory earlier that week stated that **at least seven states** have reported incidents, of which only Minnesota and Michigan have been identified [6]. Michigan received a federal alert on July 28 about attempts to tamper with operational technology in water systems, and then received a small number of reports from communities describing activity consistent with the federal description [6]. Michigan's Department of Environment, Great Lakes and Energy says "all systems continued to operate safely" [6].
- **What the attacks actually did.** Minnesota IT Services reports that most confirmed attacks involved the technology water systems use to remotely monitor and control equipment [6]. Braham, Minnesota gives the clearest published picture: the city water system was offline for about two hours on July 27, and the city was explicit that "this attack did not alter or cause any issue to the physical water plant or water quality or safety. Rather, the attackers shut down the operating controls which shut down the well and water treatment plant, causing the city to provide water to residents with only water held in the water tower" [7]. No water quality issues or use restrictions arose from the Minnesota attacks [7]. Read that as a capability demonstration against control systems rather than an attempt to harm water quality, and note that a two-hour outage was absorbed by tower storage that many small systems do not have in equivalent measure.
- **Attribution is contested and unresolved.** No culprit has been pinpointed [6]. The attacks followed an FBI, CISA and partner advisory warning that Iranian hackers have been targeting water and wastewater systems and OT controls in other sectors [6]. The New York Times reported on July 30, citing anonymous state and federal sources, that Iran is the likely culprit in Minnesota, and that federal officials were concerned it could be a new salvo in the ongoing conflict [7]. An Iranian-aligned outlet published a rebuttal the same day, quoting a cybersecurity commentator who argued such operations fall outside Iran's declared doctrine and offering two alternatives: orchestration by Israel or affiliated actors, or a US false-flag [8]. **Our assessment: attribution to any state remains unestablished, with low confidence in any current public claim.** The NYT reporting rests on anonymous single-outlet sourcing; the rebuttal comes from an interested party and is a commentator's opinion rather than a government position. Neither meets the bar for an operational decision, and none of your defensive actions depend on the answer.
- **One thing not to merge.** Pennington County, South Dakota disclosed a cybersecurity incident on July 5 that degraded county communications, internet access and records services while sheriff, courts, 911 dispatch and jails stayed operational [7]. It is a county government network incident, it is being discussed in the same coverage as the water attacks, and there is no established link between them. We are keeping them separate until something connects them.

- **The defensive guidance to use.** CISA and Five Eyes partners across the US, UK, Australia, Canada and New Zealand released **CI Fortify**, a step-by-step guide to isolating vital systems during an attack and rebuilding safely afterwards [9]. Its framing is the right one for water and any other continuous-service operator: "The end state must be to enable the continued operation of critical services in a state of isolation" [9]. The first two steps are the ones to start on this month. Identify the minimum set of systems required to support critical services and set delivery targets in real units, such as gallons or megawatts [9]. Then map, and keep mapping, every interconnection between those critical networks and everything else, including vendor remote access for consultants, contractors and managed service providers, cloud environments, untrusted networks, and peer networks like neighboring utilities or dispatch [9]. Document gateway details, architecture diagrams, firewall, router and VPN configurations, and emergency contacts, because as the guide notes, that information is what you build isolation controls out of [9].

FROM THE PANEL VIEW: BRIEFING YOUR OPERATORS

Operator instructions are indicators of compromise for the human layer. In this campaign the leading signals were physical, and the control room saw them before any tool did. Five lines for a shift handover:

1. **You are part of the detection system.** In Braham the attack showed up as controls shutting down a well and a treatment plant, not as an alert on a screen [7]. You will see this before software does.
2. **When the panel and the plant disagree, believe the plant.** A tank level, a pressure gauge or a physical valve position outranks what the HMI is telling you. Go look.
3. **The symptom most likely to be written off is a controller that will not accept a login, or a screen that has stopped updating.** That gets filed as a network glitch. File it as a security report instead.
4. **Report while you are still unsure.** You will not be second-guessed for calling something in that turns out to be a failed power supply. Every OT incident review finds somebody who noticed something at 2 a.m. and decided it was probably nothing. Do not be that person, and do not let your organization be the one that made them hesitate.
5. **Confirm today that you can run manually.** Know who holds the keys, how long tower or tank storage lasts you, and who you call first.

Symptom	What it may mean
HMI or SCADA screen stopped refreshing, values frozen	Loss of control-system communication, possibly deliberate
Controller or PLC rejects a known-good login	Credentials changed, or the device is under someone else's control

Symptom	What it may mean
Pump, well or treatment stage stops with no local fault indication	Remote shutdown of operating controls, as at Braham [7]
Gauge, tank level or flow reading disagrees with the panel	The panel may be reporting attacker-supplied values
Remote access session you did not initiate, or a vendor "already logged in"	Compromised vendor or remote-access path
Setpoint or alarm threshold changed with no work order	Unauthorized configuration change

What to tell your leadership, in three sentences. A coordinated campaign has disrupted operating controls at more than thirty water systems in Minnesota and nine in Michigan, and at least seven states have reported incidents [6][7]. In the clearest published case the attackers shut down a well and treatment plant for about two hours, with no effect on water quality, and the city ran on tower storage until control was restored [7]. What we need is confirmation that we can operate manually for a defined number of hours, that our remote monitoring and control interfaces are not reachable from the public internet, and that our operators know to report control-system anomalies as security events rather than equipment faults.

ESPIONAGE MOVED PLATFORMS AND SUPPLY CHAINS GOT QUIETER

- **Laundry Bear went from Zimbra to Outlook.** The Russian state-sponsored group tracked as Laundry Bear and Void Blizzard is exploiting the OWA flaw described above to deliver a backdoor called **OWAReaper**, targeting government entities in the US and Europe alongside telecommunications, financial, hospitality and aerospace companies [12]. Proofpoint spotted the campaign and describes it as a significant improvement in the group's tradecraft [12]. The same actor previously used a Zimbra zero-day to deliver **ZimReaper**, which steals email, two-factor codes, application passcodes and passwords [12]. We covered the Zimbra campaign in the July 28 briefing; the actor has now carried the same technique to a much larger installed base.
- **One North Korean crew, four poisoned npm packages.** AWS attributed the compromises of typo-crypto, chalk, debug and Axios to a single North Korea-linked group tracked as Sapphire Sleet, widely regarded as a Lazarus offshoot, with **medium confidence** [16]. The method is the notable part: rather than exploiting npm itself, the crew befriended maintainers, stole credentials, and published poisoned updates through accounts developers already trusted [16]. Google had previously attributed the Axios compromise to the same group; AWS extends that to all four on shared infrastructure and technical overlaps [16]. Targets escalated over time from obscure packages to major ecosystem names [16]. AWS

frames the strategy as economics: compromising a few widely used packages reaches thousands of downstream environments at once [16].

- **Teams calls as the entry point for ransomware.** Sophos tracks a campaign it calls STAC4749 in which attackers impersonate IT helpdesk staff in Microsoft Teams chats and voice calls to obtain remote access, then deploy Chaos ransomware [17]. Dozens of organizations were targeted between February and June 2026, at least three intrusions ended in ransomware, and one went from initial access to file encryption in **under 17 hours** [17]. Roughly 95% of the activity hit Canada (50%) and the United States (45%), concentrated in services, manufacturing, energy, and construction and engineering [17]. Most calls lasted about two to two-and-a-half minutes [17]. The control is procedural and cheap: publish a rule that IT never initiates contact via Teams from outside the tenant, and that any such contact is ended and called back through a known number.

BREACH SCALE IS NOW A THIRD-PARTY CONCENTRATION PROBLEM

- **Amgen filed a material cybersecurity incident.** In a Form 8-K, Amgen disclosed unauthorized activity involving data stored in cloud environments hosted by third-party cloud service providers, with proprietary data, **patient protected health information** and other information exfiltrated [18]. On July 29 the company determined the incident was material, based on the volume of files affected and the potential sensitivity of the information [18]. Amgen reports no impact to products, manufacturing operations, financial reporting systems, or its ability to meet patient needs, and believes the incident is not reasonably likely to have a material effect on its financial condition, with the investigation ongoing [18]. The pattern to note is that the breach was in a vendor's cloud, and the materiality determination turned on sensitivity rather than volume alone.
- **DentaQuest is notifying at very large scale.** DentaQuest began mailing notifications to roughly **15 million** people on July 17 [19]. Reporting puts the potentially affected population as high as 23.4 million, and the data ShinyHunters actually published in May covered about 2.6 million, so treat those three numbers as describing different things rather than a single count [19][20]. The intrusion ran from May 17 to May 20 and was discovered on May 20; ShinyHunters claimed more than 234 GB and published the cache on May 30 after negotiations failed [19][20]. Exposed data includes names, addresses, Social Security numbers, member IDs, Medicaid and Medicare numbers, diagnosis and treatment details, and billing information [19]. DentaQuest administers Medicaid and Medicare Advantage dental programs in all 50 states for roughly 35 million customers [19].
- **Conduent keeps growing.** At least 25 million people are affected across just two states, with a reported 15 million in Texas, roughly half the state's population, and more than 10 million reported by Oregon's Department of Justice [21]. Conduent processes data for corporations, healthcare providers and state agencies, with clients including Humana and several Blue Cross Blue Shield plans [21]. Exposed data may include names, Social Security numbers, medical information and health insurance information, and the Texas Attorney General's office is investigating [21].

- **ShinyHunters again, at Brinks Home.** Brinks Home identified an intrusion on July 20 and activated incident response; alarm monitoring and system functionality were not affected [22]. ShinyHunters claimed the attack and alleges theft of more than 4.9 million Salesforce records containing personal information [22]. This is the third ShinyHunters item across two briefings, which is itself the point: a single extortion crew is generating a meaningful share of large-scale US data theft, largely against SaaS platforms rather than corporate networks.
- **A 40-hour misconfiguration in the UK.** UK Government Investments left high-level management information publicly accessible for nearly two days, with personal details of 51 government officials exposed for about 40 hours [23]. UKGI attributed it to a staff member who did not follow security rules [23]. It is a small breach with a useful lesson, since UKGI manages taxpayer interests in bodies including Channel 4 and the Post Office, and none of the sophistication elsewhere in this briefing was required.
- **The deal-making angle.** Counsel are increasingly framing cyber posture as a valuation issue in mergers and acquisitions, with the standard examples being Verizon's $350 million price reduction on Yahoo and Marriott's inherited Starwood breach, which produced £18.4 million in UK fines, $52 million in US settlements, and an estimated total impact above $1 billion [24]. The recommendation worth acting on is independent verification of a target's controls rather than reliance on its assurances [24].

REGULATORS MOVED ON HEALTH DATA FROM THREE DIRECTIONS

- **The FTC sued a telehealth provider over ad-tech data sharing.** 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 complaint against Hims & Hers Health in the Northern District of California [25]. The allegation is that the company shared consumers' sensitive health information with Meta, Snap and other third parties for advertising via tracking technologies that transmitted website "Events," despite privacy promises, and separately that it made subscription cancellation extremely difficult [25]. The company says its privacy policy lets patients choose how their data is used and characterized the suit as headline-seeking [26]. If you run tracking pixels on any page where a user's presence implies a health condition, this complaint is the template you will be measured against.
- **HIPAA's coverage gap has bipartisan momentum.** The Senate HELP Committee advanced the Health Information Privacy Reform Act by **22 to 0**, a bill that would extend 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]. A unanimous committee vote is worth planning around even though floor time is not scheduled.
- **A federal agency wants emergency-room records at scale.** A federal agency is asking large hospitals to hand a government contractor the electronic health information of all emergency room patients, and some hospitals are objecting that the requests are too sweeping [35]. We are carrying this at headline

level only, because the source is behind a wall we could not read; confirm the agency, contractor and legal authority before you act on it.

- **Schools may get a federal cyber exchange.** The Enhancing K-12 Cybersecurity Act, introduced by Senators Mark Warner and Marsha Blackburn, would fund a Cybersecurity Information Exchange connecting federal, state and local governments with NGOs, a searchable database of funded and recommended tools, and a Cybersecurity Incident Registry cataloging K-12 attacks to identify trends [28]. It carries $20 million over two years, and nearly identical legislation from the same two senators failed in 2023 [28].

THE BROADER LANDSCAPE

- **Southeast Asian fraud syndicates are now an economic bloc.** The UN Office on Drugs and Crime estimates these operations cost the region between $88 billion and $114 billion in 2025, while trafficking people from at least 80 countries [33]. Years of enforcement have displaced the industry rather than shrunk it, helped by corruption, cryptocurrency settlement networks, special economic zone jurisdictions and crime-as-a-service models [33]. Four technologies are named as industrializing it: cryptocurrency, secure messaging such as Telegram, generative AI, and satellite networks [33]. For most readers this arrives as fraud volume against your customers and staff rather than as intrusions.
- **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]. We flag it only as a planning input: the probability of AI-specific compliance obligations arriving within the next few quarters has risen.

Surveillance & Privacy

- **License-plate reader effectiveness is being measured, and the numbers are small.** The ACLU estimates Flock Safety has installed some 100,000 automatic license plate readers across the US since 2017 [29]. Against that, cited FBI data shows case closures for aggravated assault and motor vehicle theft rose by 0.2 and 0.3 percentage points, and in 2025 Atlanta police solved roughly one in ten auto theft cases despite thousands of the cameras [29]. If your organization is being asked to host or integrate with such a system, the effectiveness evidence is now weak enough to warrant a written justification.
- **Withheld drone vulnerabilities cut both ways.** The ACLU argues that when government retains a drone flaw rather than reporting it to the manufacturer, every drone running that software stays open to takeover by anyone else who finds it [30]. Relevant to any operator flying commercial drones for inspection or security.
- **A data-access suit against DOGE will proceed.** A privacy suit brought by EPIC and a federal employee in February 2025 over access to Treasury and OPM databases survived dismissal [31].
- **A federal-funding risk for education.** Federal authorities have demanded Washington State revise certain student privacy policies, citing Title IX and FERPA, and the accompanying statement notes that FERPA violations can result in termination of an institution's federal funding [32]. We take no position on

the underlying dispute. The operational point for education-sector readers is narrow and real: FERPA non-compliance is now being framed as a funding-termination risk, which makes records-handling practice a finance question as well as a compliance one.

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 [36]. 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 [36]. Microsoft has seen the activity since **early May** across hospitality networks in several countries and has named no hotel, venue or captive-portal vendor [36].
- **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 [36]. Some pages use **ClickFix** instructions telling the victim to open a terminal and run an attacker-supplied command [36]. The useful limit: **the gateway does not silently infect the endpoint**, and the victim still has to download or execute the payload [36]. 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 [36]. 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`** [36]. 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 [36]. Microsoft's recommendation is to block the device code flow through Conditional Access wherever the business does not need it [36].

- **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 [36]. 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 [36]. Command and control is **ECDH-encrypted over TLS/TCP with jittered beaconing**, alongside a local API server [36]. Persistence runs through a Registry Run key and a scheduled task, with a **watchdog that restores any persistence mechanism defenders remove** [36].
- **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** [36]. It lifts Web Account Manager tokens from `.tbres` files in the Token Broker cache, enabling **session replay without a browser cookie** [36]. Exfiltration is an **HTTPS POST of GZip plus Base-64 JSON to the path** `/t/event` [36]. The stated objectives on the far side are mailbox collection, SharePoint and OneDrive collection, and using the access for further social engineering [36].
- **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 [36]. 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 [36].
- **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 [36]. 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 [36]. Block the Microsoft device code flow in Conditional Access where it is not needed [36]. 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** [36]. 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 [36]. 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 [36].

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 [37]. They were able to alter the data produced from computerized scans of physical DNA evidence **without leaving any trace** [37]. The vulnerability likely existed in files produced by crime-lab machines **since 1995** [37]. 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" [37].
- **A fix exists, and applying it is the action.** The vendor is **Thermo Fisher Scientific**, whose equipment is used in a majority of facilities [37]. 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" [37]. 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 [37].
- **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** [37]. For higher-encryption file types, a little research turned up a decryption key that had been on the internet for years [37]. 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 [37]. 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** [37]. There is also **no way to detect tampering that has already happened**, which the researchers state directly [37]. 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 [37].
- **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 [37]. Her framing is the one to carry: "Lessons learned from other industries haven't been imported into forensic science in a serious way" [37].
- **What we are not saying.** It is not clear whether this affects pending or past prosecutions [37], and most people are not convicted or exonerated on DNA evidence alone [37]. 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 [37]. We keep them apart.

WHAT YOU SHOULD BE DOING RIGHT NOW

Immediate Actions (This Week)

1. **Patch Adobe Campaign Classic** for CVE-2026-48449 and CVE-2026-48448 [10]. Before you do, confirm whether marketing operations run an instance that is not in your asset inventory, because that is the common failure here.
2. **Patch Cisco Secure Firewall Management Center** for CVE-2026-20316, then search authentication logs retrospectively for the static low-privilege account [11]. Exploitation predates the KEV listing, so forward-looking monitoring alone will miss it.
3. **Patch on-premises and hybrid Exchange** for CVE-2026-42897 [12]. Then hunt: look for unexplained mailbox rules, unfamiliar sessions, and OWAReaper indicators, and force session-token invalidation for high-value mailboxes. Exchange Online tenants can skip this [13].
4. **Patch Rails Active Storage** for CVE-2026-66066 if you use libvips and accept untrusted image uploads [14]. Treat `secret_key_base`, database credentials and cloud-storage keys in that environment as exposed and rotate them.
5. **Test AI environment isolation by attempting egress from inside it.** Both the Anthropic and OpenAI incidents involved environments documented as isolated that were not [1]. A configuration review would have passed in both cases; an egress test would not have.
6. **Inventory AI agents and their network reach.** Fewer than half of surveyed organizations can say they know all the agents on their networks [5]. Produce the list, then decide what each one is allowed to reach.
7. **Publish a Teams contact rule.** IT never initiates contact from outside the tenant; any external Teams call claiming to be support is ended and called back on a known number [17]. Circulate it as a one-paragraph notice this week, because the attack path is people, not software.
8. **Water and wastewater operators: verify no remote monitoring or control interface is reachable from the public internet**, confirm how many hours you can run manually, and hold the operator briefing above at the next shift handover [6][7].
9. **Set a Copilot document-provenance rule** covering which assistants may ingest externally supplied documents, and keep those contexts separate from document generation [4].
10. **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 [36]. Hunt for `svchost32.exe` under `%APPDATA%\svchost32\`, a `svchost32` service displaying as "Cloud Sync Service", and `.tbres` token-cache access.
11. **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 [37]. Past tampering cannot be detected, so the update protects only going forward.

Strategic Actions (This Month)

1. **Add AI vendors to your third-party risk process as a distinct category.** Ask specifically how they test environment isolation, how they would detect an escape, and what their notification commitment is. Anthropic's disclosure and its acknowledgement that it could have reviewed its records more thoroughly give you a concrete standard to cite [1].
2. **Work the first two steps of CI Fortify** [9]. Define the minimum set of systems needed to sustain critical services with delivery targets in real units, then map every interconnection to them, including vendor remote access, cloud, untrusted networks and peer utilities. Document gateways, architecture, firewall, router and VPN configurations, and emergency contacts.
3. **Design controls that assume prompt-level AI defenses will fail.** The ICML research argues the instruction-source confusion may be fundamentally unsolvable, and the Copilot worm shows a vendor unable to close a class after 144 days [3][4]. Put the real controls at the boundary: what the model can reach, what it can write, and what a human approves.
4. **Harden the software supply chain against maintainer compromise rather than registry compromise** [16]. Pin and verify dependencies, require provenance for updates, and add a review gate for major version bumps of widely used packages.
5. **Run a third-party data-concentration review.** Amgen's breach was in a vendor cloud, Conduent's reaches tens of millions across unrelated clients, and Brinks Home's exposure was a SaaS record store [18][21][22]. Identify the handful of processors whose compromise would constitute your own worst breach, and confirm what they owe you contractually on notification and forensics.
6. **Audit tracking technologies on any page where presence implies a health condition** [25]. The FTC complaint against Hims & Hers is the current template, and the theory reaches well beyond telehealth.
7. **Put cyber diligence into your M&A timeline with independent verification** [24].
8. **Prepare for health-data rules that reach beyond HIPAA.** A 22-0 committee vote signals real momentum; inventory the consumer health data you hold that HIPAA does not currently cover [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/

[9] CISA unveils a six-step blueprint for isolating critical infrastructure during cyberattacks (2026-07) — https://www.csoonline.com/article/4203133/cisa-unveils-a-six-step-blueprint-for-isolating-critical-infrastructure-during-cyberattacks.html

[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

[14] Rails patches critical Active Storage flaw with RCE potential (2026-07) — https://www.bleepingcomputer.com/news/security/rails-patches-critical-active-storage-flaw-with-rce-potential/

[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

[16] Amazon links four poisoned npm packages to one North Korean crew (2026-07-30) — https://www.theregister.com/cyber-crime/2026/07/30/amazon_links_four_poisoned_npm_packages_to_one_north_korean_crew/5281120

[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/

[20] DentaQuest Breach: ShinyHunters Publish Data Impacting 2.6M People (2026-05) — https://securityaffairs.com/193274/data-breach/dentaquest-breach-shinyhunters-publish-data-impacting-2-6m-people.html

[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

[22] ShinyHunters claims Brinks Home breach, threatens to leak stolen data (2026-07) — https://www.bleepingcomputer.com/news/security/shinyhunters-claims-brinks-home-breach-threatens-to-leak-stolen-data/

[23] UK's state investments agency hit by data breach (2026-08-02) — https://www.theguardian.com/business/2026/aug/02/uk-state-investments-agency-data-breach

[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/

[28] K-12 cyber information exchange, incident catalog would be directed by CISA under new legislative proposal (2026-07) — https://statescoop.com/k12-cyber-information-exchange-incident-catalog-cisa/

[29] Flock Surveillance Cameras Are Actually Horrible for Fighting Crime, FBI Data Shows (2026-07) — https://futurism.com/future-society/flock-surveillance-atlanta-fbi-data-crime-case-solve

[30] Government Playing a Dangerous Game With Drone Cybersecurity (2026-07) — https://www.aclu.org/news/privacy-technology/counter-drone-zero-days

[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

[32] Feds demand WA revise student policies on privacy, transgender rights (2026-07) — https://www.seattletimes.com/education-lab/feds-demand-wa-revise-student-policies-on-privacy-transgender-rights/

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

[34] The FCC bans all robot vacuums made outside the U.S. (2026-07) — https://mashable.com/tech/us-robot-vacuum-ban-fcc-foreign

[35] Agency's Push to Gather ER Data Sparks Privacy Clash (2026-07) — https://www.healthcareinfosecurity.com/agencys-push-to-gather-er-data-sparks-privacy-clash-a-32363

[36] 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

[37] 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
- Why Identity Visibility Is Becoming Cybersecurity's Top Priority — https://www.newsweek.com/why-identity-visibility-is-becoming-cybersecuritys-top-priority-12260114
- Taiwan Needs an AI Cyber Shield Now to Defend Against Invasion (CSIS) — https://www.csis.org/analysis/taiwan-needs-ai-cyber-shield-now-defend-against-invasion
- Key AI-enabled cyber attack trends to watch out for — https://irishtechnews.ie/key-ai-enabled-cyber-attack-trends-to-watch/
- A Mid-2026 Primer On Cybersecurity And Addressing New Threats — https://www.forbes.com/sites/chuckbrooks/2026/07/31/a-mid-2026-primer-on--cybersecurity-and-addressing-new-threats/
- Google Earth risked ruin with retracted AI tool for making fake satellite pics — https://arstechnica.com/ai/2026/07/google-earth-releases-swiftly-retracts-ai-feature-to-make-fake-satellite-images/
- Hackers threaten the EU directly: end Chat Control or else — https://cybernews.com/security/lunarissec-threatens-cyberattacks-over-eu-chat-control/

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