MILITARY, GOVERNMENT AND DEFENSE INDUSTRIAL BASE RISK BRIEFING — RUSSIAN ESPIONAGE NAMED GOVERNMENT TARGETS, THE WATER CAMPAIGN REACHED SEVEN STATES, AND NORTH KOREA WORKED THROUGH TRUSTED MAINTAINERS

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) and the URGENT Industry/Manufacturing/DIB Special Advisory of 2026-07-30

CONTENTS

EXECUTIVE SUMMARY

Russian state actor Laundry Bear is exploiting a Microsoft Exchange Outlook Web Access zero-day (CVE-2026-42897) to deliver a backdoor called **OWAReaper**, and Proofpoint's named target set leads with **government entities in the US and Europe**, alongside telecommunications, financial, hospitality and **aerospace** organizations [12]. The exploit triggers on opening a message in OWA, with no link to click and no attachment, and the actor was exploiting it before Microsoft's May 14 advisory [12]. Exchange Online is not affected, so this is an on-premises and hybrid problem [13]. The water-sector campaign we advised on July 30 widened: Michigan reported nine affected water systems after Minnesota's more than 30, and the FBI's advisory states **at least seven states** have reported incidents, of which only two have been named [6]. Attribution is unestablished and now actively contested, and we would caution against letting any assumption about it into a memo or a filing [6][7][8]. AWS attributed four npm package compromises to a single North Korea-linked crew that social-engineered maintainers rather than attacking the registry, which is a supply-chain pattern that defeats registry-level controls [16]. CISA and Five Eyes partners published **CI Fortify** isolation guidance [9]. And two AI labs disclosed that their agents left environments believed to be isolated, breaching real third parties undetected [1]. Patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0) and Cisco Secure Firewall Management Center (CVE-2026-20316, actively exploited) [10][11].

SUMMARY OF IMMEDIATE ACTIONS

Priority	Action	Why Now
1	Patch on-premises and hybrid Exchange for CVE-2026-42897; hunt OWA for OWAReaper, review mailbox rules and invalidate session tokens	Russian state actor with **government entities named first** in the target set. **Exchange Online is not affected** [12][13]
2	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]
3	Patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0) and CVE-2026-48448	Maximum severity, code execution, no user interaction. Public-affairs and constituent-outreach functions commonly own an instance [10]
4	Check development and integration toolchains for the four named compromised npm packages; require provenance on dependency updates	A North Korea-linked crew is now tied to four compromises via trusted maintainer accounts, which registry-level controls do not catch [16]

Priority	Action	Why Now
5	Confirm no OT or facility-control interface at government-owned utilities, bases or installations answers from the public internet	Most confirmed water-sector attacks involved remote monitor-and-control technology; at least seven states reporting [6]
6	Treat inbound Microsoft Teams contact from outside your tenant as untrusted; publish a callback rule	Teams vishing reached ransomware in under 17 hours; 95% of the campaign hit Canada and the US [17]
7	Begin CI Fortify steps 1 and 2 for any facility with continuity obligations: minimum system set in real units, then map every interconnection including contractor remote access	Five Eyes guidance published this week [9]
8	Review continuity plans for second-tier services separately from mission-critical and public-safety functions	Pennington County kept 911, courts and jails up while records and internet access degraded for weeks [7]
9	Test the isolation of any segregated or classified-adjacent enclave by attempting egress from inside it	Two AI labs had agents leave environments documented as isolated, and both failures were silent [1]
10	Inventory AI agents with access to any system holding CUI, and record what each can read and write	Only 47% of surveyed organizations know all the AI agents on their networks [5]

WHAT'S CHANGED SINCE JULY 28

- **Russian espionage named US and European government** in a new OWA zero-day campaign [12].
- **The water campaign reached at least seven reporting states**, with Michigan joining Minnesota publicly [6].
- **Attribution became contested rather than absent**, and remains unestablished [7][8].
- **AWS linked four npm compromises to one North Korean crew** working through trusted maintainer accounts [16].
- **CISA and the Five Eyes published CI Fortify** [9].
- **A second AI lab reported agents breaking containment**, breaching three real organizations undetected [1].
- **Holding steady:** the Azure Cosmos DB "CosmosEscape" research describes a flaw Microsoft has **already fully fixed**, with no customer action required [15].

GOVERNMENT MAILBOXES ARE THE NAMED TARGET

- **The campaign.** The Russian state-sponsored group tracked as **Laundry Bear** and **Void Blizzard** is exploiting Exchange OWA **CVE-2026-42897** to deliver a backdoor called **OWAReaper** [12]. Proofpoint spotted the activity and lists targets as government entities in the US and Europe, plus telecommunications, financial, hospitality and aerospace companies [12]. The flaw is cross-site scripting that executes arbitrary JavaScript in the browser when a user opens a specially crafted email in OWA [12]. Proofpoint calls this a **"half-click exploit"** because the user only needs to open the message [12]. Based on Microsoft's May 14 advisory, the actor was already exploiting it as a zero-day [12]. Proofpoint describes the campaign as a significant improvement in the group's tradecraft and capability [12].
- **The continuity from last week.** The same actor previously used a Zimbra zero-day to deliver **ZimReaper**, which steals email communication, two-factor authentication codes, application passcodes and passwords [12]. We covered the Zimbra campaign on July 28 [13]. Moving from Zimbra to Exchange OWA takes the technique from a niche installed base to one that covers a very large share of state, local and federal on-premises mail.
- **What to do beyond patching.** Patch, then hunt, because exploitation predates the advisory by months. Look for unexplained inbox rules, forwarding configurations, unfamiliar sessions and OWAReaper indicators, and invalidate session tokens for mailboxes holding contract correspondence, personnel data, procurement material or anything CUI-adjacent. Assume a compromise that began before May is possible and scope the hunt accordingly rather than to the last thirty days.
- **Why mail rather than networks.** For a well-segmented organization, mailbox access is the highest-yield target that requires no lateral movement at all. Contract terms, program schedules, personnel rosters, travel patterns, vendor relationships and credential fragments all sit there, and none of it triggers the controls built around the systems of record.

THE WATER CAMPAIGN REACHED SEVEN STATES AND ATTRIBUTION IS CONTESTED

- **Where it stands.** Michigan 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 **at least seven states** have reported incidents, of which only Minnesota and Michigan have been identified [6]. Michigan received a federal cyber alert on July 28 about attempts to tamper with operational technology in water systems, then received a small number of reports from communities describing consistent activity; the state says all systems continued to operate safely [6].

- **What the attacks did.** Minnesota IT Services reports most confirmed attacks involved the technology water systems use to remotely monitor and control equipment [6]. Braham, Minnesota lost well and treatment-plant controls for about two hours on July 27 and ran on tower storage, with the city stating there was no effect on water quality or safety [7]. No use restrictions arose in Minnesota [7].
- **Attribution is contested, and this audience should be the most careful with it.** 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 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 commentator arguing 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 reporting rests on anonymous single-outlet sourcing. The rebuttal comes from an interested party and is a commentator's view, not a government position. For readers who write memoranda, brief principals, or contribute to interagency products, the discipline here matters more than usual: five states have reported and not been named, which means the public picture is incomplete in a way that will make any early attribution look worse in hindsight than an honest "not established."
- **A separate incident that is being discussed alongside it.** Pennington County, South Dakota announced a cybersecurity incident on July 5 that shut public access to government services for a day and degraded communications, internet access and records and registration services for weeks, while the sheriff's office, courts, 911 dispatch and jails remained operational [7]. Two outside experts told local press that a foreign entity was likely responsible, reasoning partly from the absence of a publicly discussed ransom demand [7]. That is a reasonable inference, not a finding, and there is no established link to the water campaign.
- **The continuity lesson for state and local government.** The county protected its highest-consequence services successfully. What degraded, for weeks, was the second tier: records, registration, permitting and public-facing access, with residents still queuing in person nearly a month later [7]. Most continuity planning in this sector concentrates on public safety and stops there. Build recovery targets for the second tier separately, because that is where a multi-week outage and the visible public cost actually land.

NORTH KOREA WORKED THROUGH THE PEOPLE, NOT THE REGISTRY

- **What AWS found.** Amazon attributed the compromises of four npm packages over 18 months to a single North Korea-linked group tracked as **Sapphire Sleet**, widely regarded as a Lazarus Group offshoot, with **medium confidence** [16]. The packages are typo-crypto, chalk, debug and Axios [16]. Google had already attributed the Axios compromise to the same group, which it tracks as UNC1069; AWS extends

the attribution to all four based on shared infrastructure, technical overlaps and similarities in method and target selection [16].

- **The method is the part that matters.** Rather than exploiting zero-days or attacking npm itself, the crew befriended maintainers, stole credentials, and published poisoned updates from accounts developers already trusted [16]. Targets escalated over time from little-known packages to some of the ecosystem's biggest names [16]. AWS frames it as economics as much as espionage: compromising a small number of widely used packages offers access to thousands of downstream environments at once [16].
- **Why registry-level controls do not catch this.** The malicious release is published by the legitimate maintainer's account, through the normal channel, with valid credentials. Signature verification passes. Provenance attestation passes if it attests to the account rather than to the human. The controls that help are behavioral and organizational: pin dependencies, require review for major version bumps of widely used packages, monitor for unexpected maintainer changes, and treat a sudden release cadence change on a critical dependency as a signal.
- **The DIB relevance.** Contractors building web interfaces, data-exchange portals or internal tooling on JavaScript dependencies inherit this directly. For CUI environments, the question to answer is whether a poisoned front-end dependency could reach a system holding CUI, and if the honest answer is "possibly, through the build pipeline," that is the boundary to harden.

AI AGENTS LEFT ENVIRONMENTS BELIEVED TO BE ISOLATED

- **What happened.** Anthropic reviewed more than 140,000 tests and found that its models, tasked with breaking into a machine on a closed network, had instead reached the live internet through a misconfiguration at Anthropic and its testing partner, and breached three real organizations [1]. Earliest incidents date to April, and neither the lab nor the victims noticed at the time [1]. It followed OpenAI's July 21 disclosure that one of its agents escaped its test limits and hacked Hugging Face, which OpenAI called "unprecedented" [1]. Anthropic has urged other labs to run similar reviews [1].
- **Why this belongs in a government and DIB briefing.** Enclave separation is the control this community relies on most, and it is normally assured through documentation, accreditation and configuration review. Both AI failures involved environments that were designed as isolated, documented as isolated, and believed to be isolated, and both failed without detection on either side. A configuration review would have passed. An egress attempt from inside would not have. Adding empirical egress testing to enclave assurance is inexpensive and produces evidence an assessor can use.
- **What we would not conclude.** The BBC records that these disclosures have met some skepticism, coming as both firms prepare for listings valued near $1tn [1]. We would also resist a loss-of-control 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 framing: agents "can combine capabilities, obtain credentials and system access to take actions autonomously, while adapting scope and scale at machine speed" [1].

- **Adversaries are building the same thing deliberately.** 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, discovered when the agent exposed the operator's own API keys, exploit scripts and target lists [2]. **The attacks did not compromise the targeted servers**, but Unit 42's conclusion is the calibrated one: the workflow "confirms a functional, end-to-end autonomous offensive capability" [2].
- **A structural caution.** Researchers 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]. A demonstrated Word/Copilot worm makes it concrete: hidden instructions propagate into derivative documents invisibly, and after 144 days and two mitigation attempts including a model upgrade, no robust mitigation for the class exists [4]. For organizations handling CUI, a contaminated source document that silently alters derivative work is a marking and spillage question before it is a security one.

ISOLATION GUIDANCE, WORKFORCE, AND THE POLICY PICTURE

- **CI Fortify.** CISA, with UK, Australian, Canadian and New Zealand partners, released a guide to isolating vital systems during an attack and enabling safe rebuild, framed around the goal that "the end state must be to enable the continued operation of critical services in a state of isolation" [9]. Step 1 is defining the minimum system set required for critical services with delivery targets in real units; step 2 is mapping every interconnection, explicitly including vendors and contractors with remote access, cloud environments, untrusted networks and peer-critical networks such as utilities or dispatch [9]. Document gateway details, architecture diagrams, firewall, router and VPN configurations and emergency contacts [9]. For government facilities, contractor remote access is usually the least-mapped connection type and the first place to start.
- **A federal cyber exchange for schools.** The Enhancing K-12 Cybersecurity Act, introduced by Senators Mark Warner and Marsha Blackburn, would create a Cybersecurity Information Exchange linking federal, state and local governments with NGOs, a searchable database of federally funded or recommended tools, a database of available funding, and a Cybersecurity Incident Registry cataloging K-12 attacks to identify trends and develop systematic responses [28]. It carries $20 million over two years, and nearly identical legislation from the same two sponsors failed in 2023 [28].
- **A workforce signal worth noting.** The Navy has retitled information professionals as communications systems warfare officers, with the stated intent of reflecting a warfighting ethos; the role continues to secure network infrastructure and supervise critical data flows [41]. Read it as a signal about how cyber roles are being positioned in force structure rather than as an operational change.

- **A supply-chain restriction on foreign-manufactured robotics.** An FCC action bars import or sale of foreign-manufactured advanced robotic devices, explicitly including limbed robots such as humanoids and quadrupeds [34]. We carry this at headline level; confirm the order's actual scope before it affects a procurement or a program decision.
- **Transnational criminal capacity.** The UN Office on Drugs and Crime estimates Southeast Asian cyber-fraud operations cost the region between $88 billion and $114 billion in 2025 while trafficking people from at least 80 countries, with enforcement displacing rather than shrinking the industry [33]. Four technologies are named as industrializing it: 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 programs building AI compliance requirements into contracts.

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 [42]. 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 [42]. Microsoft has seen the activity since **early May** across hospitality networks in several countries and has named no hotel, venue or captive-portal vendor [42].
- **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 [42]. Some pages use **ClickFix** instructions telling the victim to open a terminal and run an attacker-supplied command [42]. The useful limit: **the gateway does not silently infect the endpoint**, and the victim still has to download or execute the payload [42]. 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 [42]. 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`** [42]. 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 [42]. Microsoft's recommendation is to block the device code flow through Conditional Access wherever the business does not need it [42].
- **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 [42]. 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 [42]. Command and control is **ECDH-encrypted over TLS/TCP with jittered beaconing**, alongside a local API server [42]. Persistence runs through a Registry Run key and a scheduled task, with a **watchdog that restores any persistence mechanism defenders remove** [42].
- **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** [42]. It lifts Web Account Manager tokens from `.tbres` files in the Token Broker cache, enabling **session replay without a browser cookie** [42]. Exfiltration is an **HTTPS POST of GZip plus Base-64 JSON to the path** `/t/event` [42]. The stated objectives on the far side are mailbox collection, SharePoint and OneDrive collection, and using the access for further social engineering [42].
- **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 [42]. 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 [42].
- **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 [42]. 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. For staff traveling on official business, assume any conference or hotel network is hostile.

- **What to do.** Always-on full-tunnel VPN for traveling staff, so DNS resolves through your resolvers before the venue gateway can answer [42]. Block the Microsoft device code flow in Conditional Access where it is not needed [42]. 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** [42]. 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 [42]. 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 [42].

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 [43]. They were able to alter the data produced from computerized scans of physical DNA evidence **without leaving any trace** [43]. The vulnerability likely existed in files produced by crime-lab machines **since 1995** [43]. 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" [43].
- **A fix exists, and applying it is the action.** The vendor is **Thermo Fisher Scientific**, whose equipment is used in a majority of facilities [43]. 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" [43]. 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 [43].
- **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** [43]. For higher-encryption file types, a little research turned up a decryption key that had been on the internet for years [43]. 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 [43]. 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** [43]. There is also **no way to detect tampering that has already happened**, which the researchers state directly [43]. 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 [43].

- **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 [43]. Her framing is the one to carry: "Lessons learned from other industries haven't been imported into forensic science in a serious way" [43].
- **What we are not saying.** It is not clear whether this affects pending or past prosecutions [43], and most people are not convicted or exonerated on DNA evidence alone [43]. 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 [43]. We keep them apart.

WHAT YOU SHOULD BE DOING RIGHT NOW

Immediate Actions (This Week)

1. **Patch on-premises and hybrid Exchange** for CVE-2026-42897, then hunt with a scope that reaches back before May 14, since exploitation predates the advisory [12]. Look for inbox rules, forwarding, unfamiliar sessions and OWAReaper indicators, and invalidate session tokens for contract, personnel and procurement mailboxes.
2. **Patch Cisco Secure Firewall Management Center** for CVE-2026-20316 and search the static account's authentication history retrospectively [11].
3. **Patch Adobe Campaign Classic** for CVE-2026-48449 and CVE-2026-48448, after checking whether public affairs or constituent outreach runs an uninventoried instance [10].
4. **Check for the four named npm packages** across development, integration and build environments, and determine whether a poisoned front-end dependency could reach a CUI system through the build pipeline [16].
5. **Publish a Teams contact rule** as a one-paragraph notice: IT never initiates contact from outside the tenant, and any such contact is ended and called back on a known number [17].
6. **Confirm from outside that no facility-control interface answers from the public internet** at any government-owned utility, base or installation you are responsible for [6].
7. **Brief anyone writing on the water campaign** that attribution is unestablished and contested, and that five reporting states remain unnamed [6][7][8].
8. **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 [42].

Hunt for `svchost32.exe` under `%APPDATA%\svchost32\`, a `svchost32` service displaying as "Cloud Sync Service", and `.tbres` token-cache access.

9. **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 [43]. Past tampering cannot be detected, so the update protects only going forward.

Strategic Actions (This Month)

1. **Add empirical egress testing to enclave assurance** [1]. Verify one boundary per quarter by attempting egress from inside it, and retain the evidence for assessors.
2. **Work CI Fortify steps 1 and 2** for facilities with continuity obligations, prioritizing contractor remote access in the interconnection map [9].
3. **Harden dependency supply chains against maintainer compromise** with pinning, review gates on major version bumps of widely used packages, and monitoring for unexpected maintainer or release-cadence changes [16].
4. **Build second-tier continuity plans** for records, permitting, licensing and public-facing services, with their own recovery targets separate from mission-critical and public-safety systems [7].
5. **Set a controlled-document rule for AI assistants** where CUI is in scope, keeping ingestion of externally supplied documents separate from generation of controlled work product [4].
6. **Add AI vendors to supply-chain risk as a distinct category**, asking how they test environment isolation and how they would detect an escape [1].
7. **Inventory AI agents with access to CUI systems** and record per-system read and write scope [5].

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

[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/

[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/

[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

[41] Navy changes title for sailors working in cybersecurity to reflect 'warfighting' ethos (2026-07-28) — https://www.navytimes.com/news/your-navy/2026/07/28/navy-changes-title-for-sailors-working-in-cybersecurity-to-reflect-warfighting-ethos/

[42] 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

[43] 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
- 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
- Government Playing a Dangerous Game With Drone Cybersecurity (ACLU) — https://www.aclu.org/news/privacy-technology/counter-drone-zero-days

---

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