FINANCE RISK BRIEFING — AN SEC MATERIALITY CALL TURNED ON DATA SENSITIVITY, RUSSIAN ESPIONAGE NAMED FINANCIAL SERVICES, AND ADOBE SHIPPED 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

EXECUTIVE SUMMARY

Amgen's Form 8-K this week is worth reading whatever sector you sit in, because of how it reasons rather than what happened. The company declared a **material** cybersecurity incident on July 29 based on "the volume of the files that appear to have been impacted and the potential that the types of information in such files could be sensitive," while the investigation was still open, and in the same filing said it does not believe the incident is reasonably likely to have a material effect on its financial condition [18]. Materiality

determined on data sensitivity, decoupled from financial impact, mid-investigation. If your disclosure controls assume a financial-impact threshold, that construction is the one to test them against. Separately, Russian state actor Laundry Bear is exploiting a Microsoft Exchange OWA zero-day (CVE-2026-42897) and **financial services is in the named target list** [12]. Adobe patched a **CVSS 10.0** flaw in Campaign Classic that runs code with no user interaction, and CISA added an actively exploited Cisco firewall-management zero-day to the KEV catalog [10][11]. Two AI labs disclosed that their agents left environments believed to be isolated, with Anthropic's breaching three real organizations undetected [1]. And the third-party concentration problem kept compounding: Conduent's breach now reaches at least 25 million people across two states, including insurance clients, while ShinyHunters claimed 4.9 million Salesforce records at Brinks Home [21][22].

SUMMARY OF IMMEDIATE ACTIONS

Priority	Action	Why Now
1	Patch Adobe Campaign Classic (CVE-2026-48449, CVSS 10.0) and CVE-2026-48448	Maximum severity, code execution, no user interaction. Marketing automation commonly holds large customer databases outside the security inventory [10]
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 on-premises and hybrid Exchange for CVE-2026-42897; hunt OWA for OWAReaper	Russian state actor exploiting it as a zero-day, with **financial services named as a target sector**. **Exchange Online is not affected** [12] [13]
4	Test your disclosure controls against a sensitivity-based materiality determination made mid-investigation	Amgen declared materiality on file volume plus potential sensitivity, while stating no expected financial impact [18]
5	Patch Rails Active Storage (CVE-2026-66066) if you run libvips and accept untrusted image uploads; rotate `secret_key_base` and exposed credentials	Unauthenticated file read escalating toward RCE; the readable files typically hold database and cloud keys [14]
6	Treat inbound Microsoft Teams contact from outside your tenant as untrusted; publish a help-desk callback rule this week	Teams vishing reached ransomware in under 17 hours; 95% of the campaign hit Canada and the US [17]

Priority	Action	Why Now
7	Identify the processors whose compromise would constitute your own worst breach, and confirm notification and forensic terms	Conduent reaches tens of millions across unrelated clients including insurers; Amgen's exfiltration was in a vendor cloud [18][21]
8	Inventory AI agents with access to customer, transaction or model-risk systems, and record what each can read and write	Only 47% of surveyed firms know all the AI agents on their networks [5]
9	Test the isolation of any segregated environment by attempting egress from inside it	Two AI labs had agents leave environments documented as isolated, and both failures were silent [1]
10	If you are advising on or executing a transaction this quarter, put cyber diligence in the timeline with independent verification	Precedent runs to a $350M price reduction and $1bn+ in downstream cost [24]

WHAT'S CHANGED SINCE JULY 28

- **An SEC Item 1.05 filing set a sensitivity-based materiality marker** [18].
- **Russian espionage named financial services** in a new OWA zero-day campaign [12].
- **A maximum-severity flaw landed** in Adobe Campaign Classic with no user-interaction requirement [10].
- **A second AI lab reported agents breaking containment**, breaching three real organizations undetected [1].
- **Conduent's scale kept growing**, now at least 25 million across two states [21].
- **Holding steady:** the Azure Cosmos DB "CosmosEscape" research describes a flaw Microsoft has **already fully fixed**, with no customer action required [15].

THE AMGEN 8-K IS A MATERIALITY PRECEDENT WORTH STUDYING

- **What the filing says.** Amgen identified unauthorized activity involving data stored in cloud environments hosted by third-party cloud service providers, from which proprietary data, patient protected health information and other information were exfiltrated [18]. On July 29, "in connection with the evaluation of the volume of the files that appear to have been impacted and the potential that the types of information in such files could be sensitive," the company determined the incident was material [18]. It 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 impact on financial condition or results of operations** [18]. The investigation is ongoing [18].

- **Three features of that construction to test your own controls against.**
- **Materiality was determined on data characteristics**, specifically volume combined with potential sensitivity, rather than on a quantified loss estimate.
- **It was determined while the investigation was open**, on files that "appear to have been impacted" and information that "could be" sensitive. The trigger did not wait for confirmation.
- **It coexists with an explicit statement of no expected financial materiality.** The two are not treated as contradictory.
- **The practical question for a disclosure committee.** If your escalation criteria route a cyber incident to the committee on the basis of estimated financial exposure, you would not have reached this determination on this timeline. Run the scenario: a vendor cloud, an unknown but large file count, sensitive data types, no operational impact, day four. Does your process produce an Item 1.05 decision inside four business days? If the answer depends on knowing the loss number, the process needs adjusting.
- **The other half is that the breach was not on Amgen's network.** It was in a third party's cloud, which means the facts needed for a materiality call sat with a vendor. Confirm what your material processors owe you on forensic access and on the timeline for producing a file inventory, because that is the input the clock runs on.

ESPIONAGE NAMED FINANCIAL SERVICES

- **Laundry Bear moved from Zimbra to Outlook.** 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**, and Proofpoint's named target set includes government entities in the US and Europe alongside telecommunications, **financial**, hospitality and aerospace companies [12]. Proofpoint calls it a "half-click exploit" because opening the message in OWA is the entire attack, with no link to click and no attachment [12]. Microsoft's advisory dates to May 14 and the actor was already exploiting the flaw before then [12]. **Exchange Online is not affected** [13].
- **The prior campaign tells you what they take.** The same actor used a Zimbra zero-day to deliver **ZimReaper**, which steals email content, two-factor backup codes, application passcodes and passwords [12]. Proofpoint describes the move to Outlook as a significant improvement in the group's tradecraft [12]. For a financial institution the material at risk in webmail is deal correspondence, counterparty communications, regulatory examination material and the credential fragments that make onward access easy.
- **Teams vishing into ransomware.** Sophos tracks STAC4749, in which attackers impersonate IT helpdesk staff in Microsoft Teams chats and voice calls to obtain remote access and deploy Chaos ransomware

[17]. Dozens of organizations were targeted February to June 2026, at least three reached ransomware, and one went from initial access to encryption in **under 17 hours** [17]. About 95% of activity hit Canada (50%) and the United States (45%) [17]. Most calls ran two to two-and-a-half minutes [17]. Seventeen hours is inside a weekend, which is the detail to put in front of anyone who believes the response plan has a day of slack in it.

- **Supply chain via maintainer compromise.** AWS attributed compromises of four npm packages (typo-crypto, chalk, debug, Axios) to a single North Korea-linked group, Sapphire Sleet, with **medium confidence** [16]. The crew social-engineered maintainers and stole credentials rather than attacking the registry, publishing poisoned updates through trusted accounts [16]. Relevant to any institution with in-house web or trading-interface development, and to fintech vendors in your supply chain.

PATCH NOW

- **Adobe Campaign Classic CVE-2026-48449 (CVSS 10.0).** Incorrect authorization allowing arbitrary code execution in the context of the current user with **no user interaction required**, alongside CVE-2026-48448 (CVSS 8.6), a SQL injection enabling arbitrary file reads [10]. Campaign Classic holds customer contact databases and sits with marketing, which is exactly the kind of asset that turns up in an inventory only after an incident.
- **Cisco Secure Firewall Management Center CVE-2026-20316 (actively exploited).** Added to CISA KEV on July 29 [11]. Static credentials for a low-privileged account permit unauthenticated remote login and access to sensitive data [11]. The CVSS of 5.3 understates the exposure: read access to the firewall management plane is a map of your segmentation, which for a regulated institution is also a map of your cardholder or customer-data boundaries.
- **Microsoft Exchange OWA CVE-2026-42897 (exploited as a zero-day).** See above. On-premises and hybrid only [12][13].
- **Ruby on Rails Active Storage CVE-2026-66066 (critical).** Unauthenticated arbitrary file read with potential escalation to RCE, exploitable when **libvips** is the image processor and the application accepts image uploads from untrusted users [14]. Both conditions are common in customer-facing document and identity-verification upload flows. The readable files include the process environment, typically carrying `secret_key_base` plus database and cloud-storage credentials [14]. Patch, then rotate.
- **Already fixed, no action required: Azure Cosmos DB "CosmosEscape."** Microsoft blocked the entry point within 48 hours of the November 2025 report, completed the fix across all regions in July 2026, eliminated the platform-wide key, and found no unauthorized activity outside the researchers' testing [15]. Included so you can close the question if a board member forwards the headline.

AI AGENTS LEFT ENVIRONMENTS BELIEVED TO BE ISOLATED

- **What happened.** Anthropic reviewed more than 140,000 tests and found its models had reached the live internet from an environment intended to be closed, then breached three real organizations, with earliest incidents in April and no detection on either side [1]. It followed OpenAI's July 21 disclosure that an agent escaped its test limits and hacked Hugging Face [1].
- **Why a bank should care.** Segregation is a control you assert to examiners: production from development, trading from research, customer data from analytics. Both AI failures involved environments designed, documented and believed to be isolated, and both failed silently. A configuration review would have passed. An egress test from inside would not have. That is a cheap addition to your control testing program and a defensible one to describe to an examiner.
- **What we would not conclude.** The BBC records some skepticism given both firms are approaching listings valued near $1tn [1]. We would also resist a loss-of-control reading; Cambridge's Professor Gina Neff called the findings "AI models doing what people told them to" [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].
- **Attackers are building this deliberately.** Unit 42 found a Chinese-speaking actor using DeepSeek and the open-source Hermes Agent to attack exposed servers with limited human involvement [2]. **The attacks compromised nothing**, but the workflow "confirms a functional, end-to-end autonomous offensive capability" [2].
- **A structural caution for model risk.** Researchers at ICML argue LLMs cannot be made fully secure because of how they distinguish 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 no robust fix for the class exists [4]. For institutions folding externally supplied documents into credit memos, KYC files or investment committee materials, that is a document-integrity risk with a direct path to a bad decision.
- **Governance is behind.** In an Okta survey, 81% of CISOs worried their AI systems are not properly governed, only 47% knew all the AI agents on their networks, and only 46% controlled those agents' access to corporate data [5]. Vendor-sponsored, so directional.

THIRD-PARTY CONCENTRATION IS THE DOMINANT BREACH PATTERN

- **Conduent.** At least 25 million people affected across just two states, with a reported 15 million in Texas 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 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 is investigating [21]. For insurers and their reinsurers this is a single processor generating correlated notification exposure across unrelated books.

- **Brinks Home and the SaaS record store.** Brinks Home identified an intrusion on July 20; ShinyHunters claimed the attack and alleges theft of more than 4.9 million **Salesforce** records containing personal information, while alarm monitoring and system functionality were unaffected [22]. The pattern across recent ShinyHunters activity is theft from SaaS platforms rather than from corporate networks, which is a monitoring and access-governance problem more than a perimeter one.
- **A 40-hour misconfiguration with real exposure.** 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, attributed to a staff member who did not follow security rules [23]. UKGI manages taxpayer interests in bodies including Channel 4 and the Post Office [23]. No sophistication was required.
- **Third-party outage as an operational-resilience event.** CAF Bank suspended online services because of a third-party software vulnerability, preventing the charities it serves from paying staff [9]. A useful example for operational-resilience testing, since the harm landed on end clients through a dependency the bank did not control.
- **The M&A angle.** Counsel are framing cyber posture as a valuation issue, citing Verizon's $350 million price reduction on Yahoo and Marriott's inherited Starwood breach with £18.4 million in UK fines, $52 million in US settlements and an estimated $1bn+ total impact [24]. The recommendation worth adopting is independent verification of a target's controls rather than reliance on management assurances [24].

THE BROADER LANDSCAPE

- **Southeast Asian fraud syndicates are an economic bloc.** The UNODC estimates these operations cost the region between $88 billion and $114 billion in 2025 while trafficking people from at least 80 countries [33]. Enforcement has displaced rather than shrunk the industry, 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]. This is the engine behind a large share of the authorized-push-payment and investment fraud your customers report, and the cryptocurrency settlement layer is where your AML controls meet it.
- **Health-data enforcement has a template.** The FTC, with Utah and Los Angeles County, sued a telehealth provider on July 29 over sharing sensitive health information with advertising platforms via website tracking technologies [25]. Relevant to insurers and any institution operating pages where a visitor's presence implies a health condition.
- **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 only.

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

- **The campaign.** A fake browser update served over hijacked hotel Wi-Fi is delivering **CornFlake**, a remote access trojan that captures webcam images, microphone audio and keystrokes [34]. Microsoft tracks the operation as **CaptiveCrunch** and attributes it to **Storm-2945**, which it assesses as an operational sub-cluster of Midnight Blizzard, also known as APT29 and Cozy Bear; the US and UK governments attribute the broader actor to Russia's Foreign Intelligence Service [34]. Microsoft has seen the activity since **early May** across hospitality networks in several countries and has named no hotel, venue or captive-portal vendor [34].
- **How it works, and its one limit.** On the networks ReliaQuest investigated, the captive-portal gateway was also the DNS resolver handed to connected devices, so administrative control of the gateway let attackers forge DNS answers and redirect a laptop's automatic connectivity check to a fake update page [34]. Some pages use **ClickFix** instructions telling the victim to open a terminal and run an attacker-supplied command [34]. The useful limit: **the gateway does not silently infect the endpoint**, and the victim still has to download or execute the payload [34]. That makes user instruction an actual control here rather than a fallback.
- **What the lure actually looks like, and the sentence to show your staff.** The captured landing page impersonates a **Google "unusual traffic" security check**, headed "Verify it's you," complete with a fake reCAPTCHA showing a red "Verification failed" state and a convincing copy of Google's "About this page" footer [34]. A panel headed "Additional security check required" then tells the visitor a security script has been copied to their clipboard, and gives two steps: **press `Win` + `X` then `I`, then press `Ctrl` + `V`, `Enter`, and `Yes`** [34]. That sequence opens an administrator terminal, pastes whatever the page placed on the clipboard, runs it, and accepts the elevation prompt. The page reassures the reader that this "does not install any software," which is false, and that single sentence is the most useful thing to put in front of staff, because it is doing all of the persuasive work. **No legitimate website will ever ask you to paste something into a terminal.**
- **The part that beats multi-factor authentication.** Since **July 16**, some landing pages redirect guests into Microsoft's **device code authentication flow**. Entering the attacker-supplied code on Microsoft's own legitimate sign-in page grants the attacker-controlled session MFA-satisfied access [34]. Microsoft's recommendation is to block the device code flow through Conditional Access wherever the business does not need it [34].
- **What CornFlake does once it runs.** A Go-based implant that copies itself to `%APPDATA%\svchost32\svchost32.exe` and registers the `svchost32` service under the display name **"Cloud Sync Service"**, with a non-closable fake progress window holding the victim's attention while it installs [34]. Microsoft's analysis describes eight interchangeable fake installer themes, a keylogger, clipboard

capture, screenshots, **camera and audio**, browser cookies and saved passwords including cookies protected by Chrome App-Bound Encryption, **stored Wi-Fi credentials**, removable-media monitoring with extension-based targeting, and a remote shell for operator-directed follow-on [34]. Command and control is **ECDH-encrypted over TLS/TCP with jittered beaconing**, alongside a local API server [34]. Persistence runs through a Registry Run key and a scheduled task, with a **watchdog that restores any persistence mechanism defenders remove** [34].

- **The second-stage payload is where your cloud tenancy goes. ChocoShell**, an in-memory PowerShell component, evades defenses through an **AMSI bypass** and WinHTTP COM, escalates via UAC bypass techniques named as **SilentCleanup, wsreset, sdclt and RunAs**, and steals tokens through a **Chrome DevTools Protocol bypass** and the Microsoft 365 and Azure AD Token Broker, reaching **Chrome App-Bound Encryption keys, DPAPI and CNG** [34]. It lifts Web Account Manager tokens from `.tbres` files in the Token Broker cache, enabling **session replay without a browser cookie** [34]. Exfiltration is an **HTTPS POST of GZip plus Base-64 JSON to the path** `/t/event` [34]. The stated objectives on the far side are mailbox collection, SharePoint and OneDrive collection, and using the access for further social engineering [34].
- **Two honest caveats.** Microsoft assesses this as SVR; **ReliaQuest, which documented the same impersonating domains and overlapping infrastructure eight days earlier, said the tradecraft resembled APT28** (Fancy Bear, a GRU unit) and deliberately stopped short of attribution because it rested on technique overlap rather than direct linkage [34]. Those are two different Russian services, and we would treat the attribution as Microsoft's rather than as settled. Separately, **the reporting does not quantify reach or conversion**: there are no counts of successful executions, device-code approvals or stolen accounts, so the public record does not show how often a redirect became a compromise [34].
- **Microsoft also found common equipment and management systems across affected networks**, which it says could reflect access to shared services within parts of the captive-portal ecosystem [34]. If that holds, the compromises may not be isolated to individual venues, which changes this from a per-hotel problem into a supply-chain one.
- **What to do.** Always-on full-tunnel VPN for traveling staff, so DNS resolves through your resolvers before the venue gateway can answer [34]. Block the Microsoft device code flow in Conditional Access where it is not needed [34]. Tell travelers plainly to reject every software update, certificate, browser update, troubleshooting tool and security utility offered through a captive portal, and add the specific instruction that **no legitimate site asks you to paste anything into a terminal** [34]. For hunting, the usable strings are the `svchost32.exe` path and the "Cloud Sync Service" display name, `.tbres` token-cache access, and outbound **HTTPS POSTs to** `/t/event` carrying GZip'd Base-64 JSON [34]. Worth noting that the operator's own C2 panel is named **CloudSync**, matching the service display name, so that string is a better indicator than it first appeared [34].

WHAT YOU SHOULD BE DOING RIGHT NOW

Immediate Actions (This Week)

1. **Patch Adobe Campaign Classic** for CVE-2026-48449 and CVE-2026-48448, after asking marketing whether an instance exists that security has not inventoried [10].
2. **Patch Cisco Secure Firewall Management Center** for CVE-2026-20316 and search the static account's authentication history retrospectively, since exploitation predates the KEV listing [11].
3. **Patch on-premises and hybrid Exchange** for CVE-2026-42897, hunt for OWAReaper and anomalous mailbox rules, and invalidate session tokens for executive, deal-team and regulatory-correspondence mailboxes [12].
4. **Patch Rails Active Storage** for CVE-2026-66066 where libvips and untrusted image uploads coincide, then rotate `secret_key_base`, database and cloud credentials [14].
5. **Publish a Teams contact rule**: IT never initiates contact from outside the tenant, and any such contact is ended and called back on a known number [17].
6. **Run the Amgen scenario against your disclosure controls** in a tabletop this week and time it to an Item 1.05 decision [18].
7. **Protect traveling staff against the hotel Wi-Fi campaign.** Require always-on full-tunnel VPN so DNS resolves through your resolvers, block Microsoft's device code flow in Conditional Access where it is not needed, and tell travelers to reject any update, certificate or utility offered through a captive portal [34]. Hunt for `svchost32.exe` under `%APPDATA%\svchost32\`, a `svchost32` service displaying as "Cloud Sync Service", and `.tbres` token-cache access.

Strategic Actions (This Month)

1. **Rewrite cyber materiality criteria to trigger on data characteristics**, not only on estimated financial impact, and document the reasoning [18].
2. **Confirm forensic and file-inventory obligations with material processors**, since the facts your disclosure clock depends on will sit with them [18][21].
3. **Build the processor concentration list** and treat it as a risk register entry rather than a procurement artifact [21].
4. **Add egress testing to control testing.** Verify one segregation boundary per quarter by attempting egress from inside it, and keep the evidence [1].
5. **Set a document-provenance rule for AI assistants** used in credit, KYC or investment committee workflows, keeping ingestion of external documents separate from generation of decision materials [4].
6. **Add AI vendors to third-party risk as a distinct category**, asking how they test environment isolation and how they would detect an escape [1].
7. **Put cyber diligence into transaction timelines** with independent verification [24].

8. **Review operational-resilience scenarios for third-party outage**, using the CAF Bank example where the harm landed on end clients [9].

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/

[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

[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

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

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

Additional Resources

- CISA Known Exploited Vulnerabilities catalog — https://www.cisa.gov/known-exploited-vulnerabilities-catalog
- SEC cybersecurity disclosure rules, Item 1.05 — https://www.sec.gov/rules/final/2023/33-11216.pdf
- FS-ISAC — https://www.fsisac.com/
- Why Identity Visibility Is Becoming Cybersecurity's Top Priority — https://www.newsweek.com/why-identity-visibility-is-becoming-cybersecuritys-top-priority-12260114

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