DATEC INCORPORATED

Daily Security Bulletin

Friday, July 31, 2026

Today's most quotable data point comes from Microsoft's Digital Defense Report: the company now detects more than 600 million identity attacks every day, putting a hard number behind the identity-centric attack patterns this bulletin has tracked for weeks. We're also watching Laundry Bear, the Russian espionage group behind the Zimbra half-click attacks, expand the same technique to Microsoft Exchange OWA, Amazon's formal attribution of four poisoned npm packages to a North Korean crew, and a severe Azure Cosmos DB flaw dubbed CosmosEscape that exposed a platform-wide signing key capable of reaching any database on the service. And two survey findings are worth sitting with together: only 47% of companies know all the AI agents running on their own networks, while 67% of CISOs say their organization has experienced attackers actively working to undermine defenders mid-incident.

NEED-TO-KNOW

Why Identity Visibility Is Becoming Cybersecurity's Top Priority

Microsoft's Digital Defense Report finds the company now detects more than 600 million identity attacks every day, reflecting how identity has become one of the primary targets for cybercriminals rather than a secondary concern behind traditional malware or network-based attacks. That volume puts a concrete number behind what this bulletin has tracked as a dominant theme for weeks, credential stuffing, OAuth client ID spoofing, AitM phishing kits, service desk social engineering, all identity-centric attack patterns rather than classic malware delivery. This is a strong data point for any conversation about where security investment should actually be prioritized right now: identity visibility and control, not just endpoint or network defense, is where attack volume is concentrated.

Datec recommends Okta for identity visibility and control given Microsoft's own telemetry now confirms identity as the primary attack surface by volume.

READ FULL ARTICLE →

HEALTHCARE

Here's What to Know Before You Upload Your Medical Data to an AI Program

A piece on medical data and AI notes that hospitals and doctors' offices must follow strict codes of ethics and rules protecting sensitive patient health information, but AI tools themselves aren't bound by those same federal rules, and with no comprehensive federal privacy law in the US, protection for medical data uploaded to a consumer AI tool depends largely on that AI company's own policy promises rather than binding legal requirements. HIPAA governs covered entities and business associates, but a patient or provider uploading records directly into a general-purpose AI tool outside that formal relationship may not have the legal protections they'd assume apply.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

Agency's Push to Gather ER Data Sparks Privacy Clash

A federal agency charged with protecting the public against injuries associated with consumer products is asking large hospitals to give a government contractor the electronic health information of all emergency room patients, with some hospitals pushing back over patient privacy concerns and arguing the records requests are too sweeping. This creates a genuine tension between a legitimate public safety mission and the scope of data actually being requested. Healthcare clients receiving similarly broad government data requests should treat proportionality and scope as a legitimate basis for pushback.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

BUSINESS IMPACTS

ShinyHunters Claims Brinks Home Breach, Threatens to Leak Stolen Data

The ShinyHunters extortion gang claims to have breached Brinks Home, alleging theft of more than 4.9 million Salesforce records containing personally identifiable information, and told BleepingComputer directly that the breach occurred on July 13 through a Microsoft Entra voice phishing, or vishing, attack. This continues both the sustained ShinyHunters campaign covered repeatedly this month and the voice-phishing-against-identity-platforms pattern seen with the Qantas tech support scam, confirming vishing against Entra and Salesforce specifically remains a reliable technique for major extortion operations.

Datec recommends KnowBe4 for vishing-specific awareness training naming Entra and Salesforce directly, given ShinyHunters' own admission of the technique used here.

READ FULL ARTICLE →

Shadow AI, Leadership Resistance Make AI Governance Tough for Worried CISOs

An Okta survey finds shadow AI and leadership resistance are making AI governance genuinely difficult for CISOs, with only 47% of surveyed companies saying they knew all the AI agents actually running on their networks, and only 46% reporting they actually controlled those agents' access to corporate data. Fewer than half of organizations can even inventory their own AI agents, let alone control what they access, a stark visibility gap given how much this bulletin has covered about AI agent risk this month, all of which assumes organizations at least know what agents are running.

Datec recommends Wiz for AI agent discovery across cloud and development environments, since most organizations first need basic inventory before any governance policy can work.

READ FULL ARTICLE →

Key AI-Enabled Cyber Attack Trends to Watch Out For

A survey of CISOs identifies key AI-enabled attack trends: 67% report their firm experienced counter incident response, where attackers actively work to undermine defenders during live engagements, 55% report increased API-based attacks, and 37% confirm being targeted by island-hopping, where adversaries compromise one institution's infrastructure specifically to pivot against its customers. Two-thirds of CISOs reporting active attacker countermeasures mid-engagement suggests adversary sophistication now goes well beyond simply evading initial detection.

Datec recommends Reliquest for MDR built to withstand active attacker countermeasures mid-engagement, not just initial detection.

READ FULL ARTICLE →

US GOVERNMENT

The FCC Bans All Robot Vacuums Made Outside the U.S.

The FCC has banned all robot vacuums, and more broadly advanced robotic devices, made outside the United States, with the update preventing foreign-manufactured advanced robotic devices from being imported or sold domestically, explicitly naming limbed robots such as humanoids and quadrupeds. This is a significant supply chain and national security-driven trade restriction reflecting growing concern about embedded surveillance or remote-access capability in foreign-made connected robotic devices operating inside homes and businesses.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

K-12 Cyber Information Exchange, Incident Catalog Would Be Directed by CISA Under New Legislative Proposal

A new legislative proposal would direct CISA to establish a K-12 cyber information exchange and incident catalog, creating a Cybersecurity Information Exchange connecting federal, state, and local governments along with nongovernmental organizations to define and disseminate best practices specifically to schools. This extends the centralized, sector-specific information-sharing model seen with healthcare's Health-ISAC to K-12 education, a sector that's faced significant targeting but historically lacked comparable dedicated threat intelligence infrastructure.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

Navy Changes Title for Sailors Working in Cybersecurity to Reflect 'Warfighting' Ethos

The Navy has renamed sailors working in cybersecurity roles, with those previously known as information professionals now referred to as communications systems warfare officers, a title change intended to reflect a warfighting ethos even though the actual role continues focusing on securing network infrastructure and supervising critical data flow. This is primarily a cultural and organizational signal, reflecting a broader institutional shift toward treating cyber defense explicitly as a combat function.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

CISA Unveils a Six-Step Blueprint for Isolating Critical Infrastructure During Cyberattacks

CISA, along with several Five Eyes security agencies across the US, UK, Australia, Canada, and New Zealand, released a six-step blueprint for isolating critical infrastructure during cyberattacks, intended to help organizations isolate and separate vital systems during an active attack and allow for safe rebuild afterward. Given how much OT and critical infrastructure targeting this bulletin has covered over the past month, this blueprint is worth incorporating directly into any critical infrastructure client's incident response planning as an authoritative, multi-national reference framework.

Datec recommends Palo Alto for the segmentation architecture that makes this six-step isolation blueprint actually executable during a live incident.

READ FULL ARTICLE →

INTERNATIONAL

Amazon Links Four Poisoned npm Packages to One North Korean Crew

Amazon has linked four poisoned npm packages to a single North Korean threat crew, with AWS noting the strategy is as much about economics as espionage, since compromising a small number of widely-used packages offers a shot at thousands of downstream environments in one move rather than picking off individual victims. This formal attribution adds a name and motive to the steady drumbeat of npm supply-chain compromises covered throughout this bulletin this month, confirming at least some of this activity traces back to a state-sponsored actor.

Datec recommends Snyk for npm dependency security now that state-sponsored North Korean targeting of the package ecosystem is formally confirmed.

READ FULL ARTICLE →

Taiwan Needs an AI Cyber Shield Now to Defend Against Invasion

An analysis argues Taiwan needs an AI cyber shield now to defend against invasion, describing a system that would preemptively triage vulnerabilities, integrate real-time threat intelligence, detect anomalous activity, and automate compromise remediation with minimal human interaction. This closely mirrors the UK NCSC's own AI-powered Cyber Shield initiative covered earlier this month, suggesting this federated, AI-driven collective defense model is emerging as a template multiple nations are independently converging on.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

[Podcast] China, Russia, Iran and North Korea Form 'Axis of Aggressors' That Could Overwhelm U.S.

A podcast discussion covers a new book examining how China, Russia, Iran, and North Korea increasingly support one another across military, economic, cyber, and technology domains, arguing these four governments now function as a genuinely cooperative axis. This framing is directly reinforced by this bulletin's own coverage over recent weeks, North Korean IT worker revenue reportedly funding Russia's war effort, Iran potentially serving as a proxy for China and Russia, suggesting real, observable operational overlap rather than a purely theoretical grouping.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

Attack on U.S. Water Facilities Is Not Part of Iran's Military Doctrine

An Iranian source is rejecting attempts to attribute a Minnesota water facility attack to Iran, instead proposing that the attack may have been orchestrated by Israel or Israel-affiliated actors to escalate tensions and create grounds for blaming Iran, or alternatively that it was a false flag operation conducted by the United States itself. This should be read as Iran's own denial and counter-narrative rather than independently verified fact; clients should continue relying on the multi-agency US attribution rather than competing state-sourced denials when making security decisions.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

PRIVACY & SURVEILLANCE

Health Privacy Legislation Advances in the Senate

Health privacy legislation is advancing in the Senate that would limit health data collection, use, and sharing at the source, place limits on how any collected data can be used and shared, guarantee individuals access to and deletion rights over their own data, and specifically prohibit the government from purchasing people's health data. The government data-purchase prohibition closes a notable current loophole, since buying health or location data from commercial brokers rather than obtaining it through a warrant has been a recurring workaround to Fourth Amendment protections.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

DOGE Must Face Privacy Suit Over Access to Treasury, OPM Data

DOGE must face a privacy lawsuit over its access to Treasury and OPM data, with the Electronic Privacy Information Center and a federal employee alleging the administration unlawfully disclosed sensitive personal information belonging to millions of Americans by granting DOGE employees access to Treasury and OPM databases. This litigation proceeding establishes that federal agencies granting broad internal data access outside an agency's normal operational scope can face genuine legal exposure under existing privacy law.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

'Reverse Warrants' Should Worry All Americans Who Value Their Liberties

An opinion piece argues that reverse warrants, a technique that identifies suspects by first querying broad location or search data and working backward rather than starting with a specific suspect, should concern Americans across the political spectrum, noting a broad, non-partisan wave of opposition to government surveillance has been building for roughly two years. This connects to broader Fourth Amendment and third-party-doctrine themes covered elsewhere in this bulletin recently.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

THREATS & VULNERABILITIES

Russian Hackers Exploit Exchange OWA Zero-Day for Long-Term Mailbox Access

Russian hackers attributed to Laundry Bear, the same group behind the earlier Zimbra half-click attacks covered in this bulletin, are exploiting an Exchange OWA zero-day for long-term mailbox access, with the flaw, CVE-2026-42897, a cross-site scripting vulnerability that executes arbitrary JavaScript when a user simply opens a specially crafted email in OWA. This confirms Laundry Bear has expanded its half-click technique from Zimbra to Microsoft Exchange OWA as well, running parallel campaigns against multiple major email platforms with conceptually identical exploitation logic.

Datec recommends Qualys to confirm Exchange OWA patch status immediately given Laundry Bear's demonstrated pattern of long-term persistent mailbox access.

READ FULL ARTICLE →

Cisco FMC Zero-Day Actively Exploited, Static Credentials Could Expose Sensitive Data

A Cisco Firepower Management Center zero-day is being actively exploited, tracked as CVE-2026-20316, potentially permitting an unauthenticated remote attacker to log into an affected device using a low-privilege static account to access sensitive data. Static, hardcoded credentials embedded in a security product itself, one that manages Cisco's own firewall infrastructure, is a particularly consequential flaw, since the product responsible for enforcing security policy becomes the exposure point through a weakness baked into its own design.

Datec recommends confirming Cisco FMC remediation guidance directly given active exploitation of this static-credential flaw in a core security product.

READ FULL ARTICLE →

Word Worm Crawls Into Copilot, Spreads Chaos

Researchers found that an attacker can hide malicious instructions inside a Word document that, when included in Copilot for Word's context, may alter document output and copy those hidden instructions into newly created files that use the affected document as source material, without the victim noticing. This self-propagating, worm-like prompt injection means a single compromised document can silently corrupt every subsequent document Copilot generates using it as a reference, spreading through normal document workflows.

Datec recommends confirming Microsoft's Copilot mitigation guidance directly, given this worm-like propagation spreads silently through ordinary document workflows.

READ FULL ARTICLE →

Microsoft Teams Vishing Attacks Lead to Chaos Ransomware Attacks

Microsoft Teams vishing attacks are leading to Chaos ransomware deployments, with Sophos finding about 95% of observed attacks targeted organizations in Canada, 50%, and the United States, 45%, with services, manufacturing, energy, and construction and engineering experiencing the largest number of attacks. This quantifies the Teams-based IT support impersonation pattern covered in this bulletin two weeks ago with the EtherRAT campaign, now confirmed as a significant vector leading to full ransomware deployment.

Datec recommends Abnormal Security for Teams-based impersonation detection, given Sophos's hard data confirming this as a direct ransomware precursor.

READ FULL ARTICLE →

Russian Spies Take Their Half-Click Email Attack From Zimbra to Outlook

Additional reporting confirms this attack doesn't depend on persuading a victim to follow a link or download a file at all, instead executing attacker-controlled JavaScript directly inside the victim's authenticated mail session the moment a booby-trapped message is opened in OWA, while Exchange Online itself is notably not affected. This gives clients a clear, actionable distinction: cloud-only Exchange Online customers aren't at risk from this flaw, while any organization running on-premises Exchange with OWA access needs to treat this as an urgent, active threat.

Datec recommends Qualys to confirm on-premises Exchange OWA patch status specifically, since Exchange Online customers are confirmed unaffected by this flaw.

READ FULL ARTICLE →

Azure Cosmos DB Flaw Exposed Platform-Wide Key That Could Access Any Database

A flaw in Azure Cosmos DB, dubbed CosmosEscape by researchers at Wiz, exposed a platform-wide signing key that could potentially access any database on the platform, with the exploit chain beginning with a crafted query against an attacker-controlled Gremlin database, achieving code execution on a multi-tenant gateway that exposed both a platform-wide signing secret and a regional account directory. A platform-wide signing key exposure in a major cloud database service represents one of the more severe multi-tenant cloud vulnerabilities of the year.

Datec recommends Wiz, the same team that discovered CosmosEscape, to assess exposure across your own Azure Cosmos DB deployments.

READ FULL ARTICLE →

MEANWHILE

NY School Pauses $60K AI Robot Teacher Over Privacy Fears, Maker's Sex Doll Ties

A New York school has paused its $60,000 AI robot teacher, named "Sally," over privacy fears and concerns about the manufacturer's ties to sex doll production, with maker Realbotix responding by assuring families that student-linked information would be encrypted and remain under the district's control, while denying the robot was a repurposed sex doll. This illustrates that vendor due diligence for education technology needs to extend beyond data handling practices alone into basic reputational and manufacturing-history scrutiny.

Informational item. No specific Datec product recommendation applies.

READ FULL ARTICLE →

Want to talk through any of these threats?

Your Datec account manager can walk you through options that fit your environment and budget — no hard sell, just straight answers.