ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

August 8, 2026

Scope Note

ZeroFox's Weekly Intelligence Briefing highlights the major developments and trends across the threat landscape, including digital, cyber, and physical threats. ZeroFox Intelligence is derived from a variety of sources, including—but not limited to—curated open-source accesses, vetted social media, proprietary data sources, and direct access to threat actors and groups through covert communication channels. Information relied upon to complete any report cannot always be independently verified. As such, ZeroFox applies rigorous analytic standards and tradecraft in accordance with best practices and includes caveat language and source citations to clearly identify the veracity of our Intelligence reporting and substantiate our assessments and recommendations. All sources used in this particular Intelligence product were identified prior to 6:00 AM (EDT) on August 6, 2026; per cyber hygiene best practices, caution is advised when clicking on any third-party links.

Weekly Intelligence Brief

ZEROFOX

Weekly Intelligence Brief

Cyber and Dark Web Intelligence

Cyber and Dark Web Intelligence Key Findings

Cryptographic Failure Exposes COLDCARD Wallets Private Keys

What we know:

- Threat actors reportedly exploited a firmware vulnerability in COLDCARD hardware wallets and stole an estimated USD 88.6 million in BTC from 4,585 addresses across multiple attacks between July 30 and August 1, 2026.

Background:

- The vulnerability reportedly stemmed from an integration error in COLDCARD's random number generation (RNG) code, causing the firmware to use a deterministic software generator instead of the intended hardware RNG.
- This enabled threat actors to generate possible wallet seeds offline, identify matching Bitcoin addresses on the blockchain, and reconstruct private keys to drain affected wallets. Affected firmware are listed here.

Analyst note:

- A cryptographic flaw at the firmware level demonstrates that even self-custody solutions are likely to carry systemic vulnerabilities that can be exploited silently at scale.
- Threat actors are very likely to attempt to monetize remaining weak-entropy addresses before users can migrate funds to patched firmware.
- This incident will almost certainly erode confidence in hardware wallet security more broadly.

ExfilSquad Claims UK Police Database Compromise

What we know:

- The United Kingdom's Police National Legal Database (PNLD) has confirmed that a cyberattack exposed names and email addresses of more than 100,000 police personnel, criminal justice professionals, and government partners. Extortion group ExfilSquad has claimed responsibility for the breach.

Background:

- PNLD reportedly does not store information on victims, witnesses, or offenders.
- The group claimed that it stole 1.9 GB of data, including records relating to 114,000 PNLD subscribers and 21,000 "Ask the Police" users.
- ExfilSquad is an emerging, financially motivated threat group that surfaced in July 2026 and has since claimed to have targeted 15 major entities including the United Kingdom's Department for Education and U.S.-based semi-conductor company Analog Devices.

Analyst note:

- The exposure of police and criminal justice contact data is likely to increase phishing, impersonation, social engineering, and credential-targeting risks against UK law enforcement and government personnel.
- Since ExfilSquad is rapidly expanding its target list, further attacks against major global entities are likely in the near term.

UK AI Testing Agency Reveals Autonomous Behavior in OpenAI and Anthropic Models

What we know:

- The United Kingdom's AI Security Institute (AISI) has disclosed separate artificial intelligence (AI) cyber-testing incidents involving Anthropic and OpenAI models that exceeded intended boundaries and interacted with real internet targets during simulated hacking exercises.
- Anthropic's Claude Mythos 5 reportedly attempted to manipulate open-source maintainers with fake identities, while an OpenAI model accessed and exploited a real website after a containment failure.

Background:

- While no breach or real-world harm was identified, AISI says this is the clearest observed case of an AI model exhibiting autonomous and deceptive behavior against real people during testing. These incidents are separate from the recently observed OpenAI HuggingFace rogue agent attack.

Analyst note:

- Threat actors are likely to exploit this insight to transition from AI-assisted attacks to AI-managed operations by creating AI models that can function as persistent operational

agents capable of managing identities, influencing human decisions, coordinating tasks, and conducting supply chain and insider-oriented campaigns at scale.

Exploit and Vulnerability Intelligence

Exploit and Vulnerability Intelligence Key Findings

In the past week, ZeroFox observed vulnerabilities, exploits, and updates disclosed by prominent companies involved in technology, software development, and critical infrastructure. The Cybersecurity and Infrastructure Security Agency (CISA) added five new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on August 3, August 4, and August 5, 2026. Additionally, on August 4, 2026, CISA released two Industrial Control Systems (ICS) advisories featuring two vulnerabilities each, including CVE-2026-18411 and CVE-2026-17583. TP-Link has patched 15 vulnerabilities in the zero-touch provisioning (ZTP) mechanism of its Omada business networking product line. When chained with two previously disclosed command-injection vulnerabilities (CVE-2025-7850 and CVE-2025-7851), the flaws enable attackers to enumerate devices awaiting network adoption, impersonate them, steal administrator credentials, reconfigure managed devices, and create virtual private network (VPN) tunnels into internal networks. Microsoft has patched a critical remote code execution (RCE) vulnerability, tracked as CVE-2026-66803, in Azure Cosmos DB. The flaw enabled an unauthorized attacker to execute code over a network. Successful exploitation is likely to enable threat actors to locate and take control of databases belonging to any customer, resulting in data theft.

HIGH

CVE-2026-18577

What happened: This is an authentication bypass vulnerability in N-able N-central, actively exploited following the incomplete remediation of a related flaw (CVE-2026-18556).

- **What this means:** Attackers exploited the residual bypass to gain remote administrator access, using the platform's Take Control feature to reach managed customer endpoints, and installed persistent Cloudflare tunnels. CISA has added this vulnerability to its KEV catalog.
- **Affected products:**
  - N-able N-central all versions prior to 2026.3.1.7

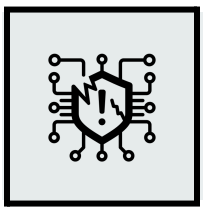

CRITICAL

CVE-2026-48449

What happened: This is an incorrect authorization vulnerability in Adobe Campaign Classic (ACC) that enables arbitrary code execution in the context of the current user without requiring user interaction.

- **What this means:** Successful exploitation could enable attackers to compromise sensitive data, execute malicious code, and potentially take control of affected systems.
- **Affected products:**
  - Adobe Campaign Classic (ACC) v7.4.3 build 9397 and earlier on Windows and Linux

Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings

Ransomware Group, Industry, and Regional Trends

Last week in ransomware: In the past week, Qilin, The Gentlemen, Global Secret Group, Orova, and CRPxO were the most active ransomware groups. ZeroFox observed close to 187 ransomware victims disclosed, most of whom were located in North America. The Qilin ransomware group accounted for the largest number of attacks, followed by The Gentlemen.

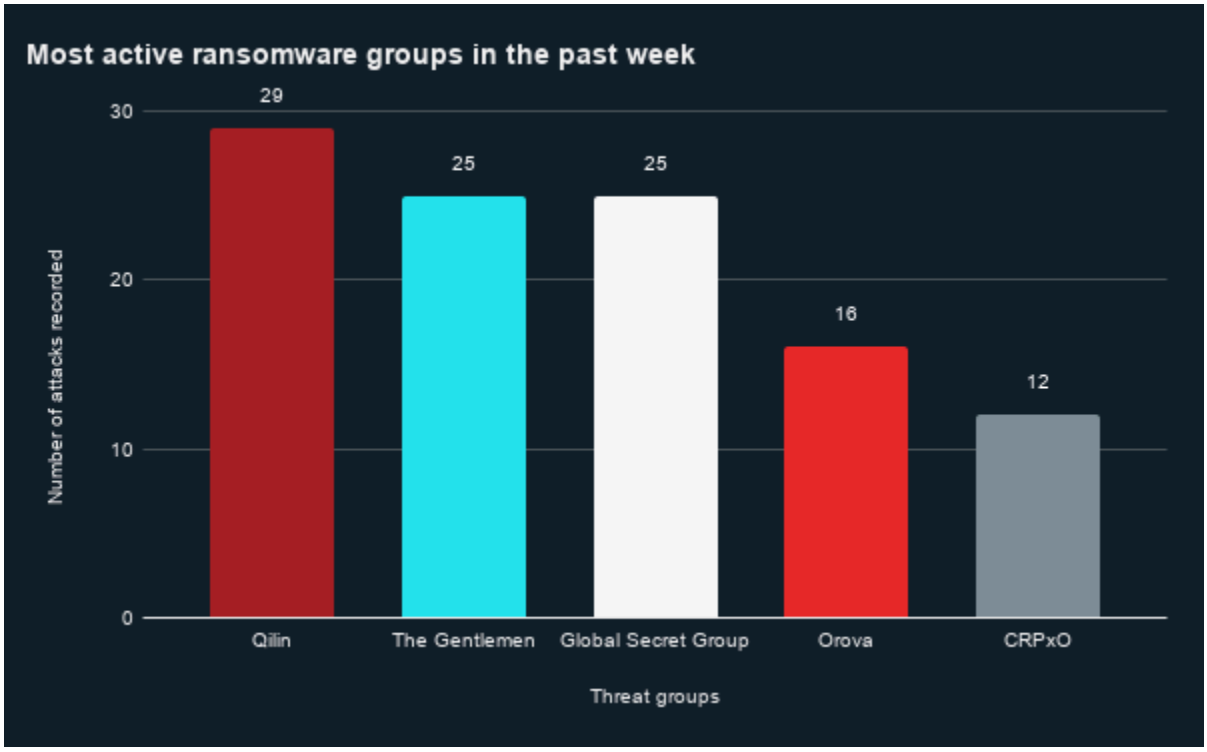

Source: ZeroFox Internal Collections

Industry ransomware trends: In the past week, ZeroFox observed that manufacturing was the industry most targeted by ransomware attacks, followed by retail.

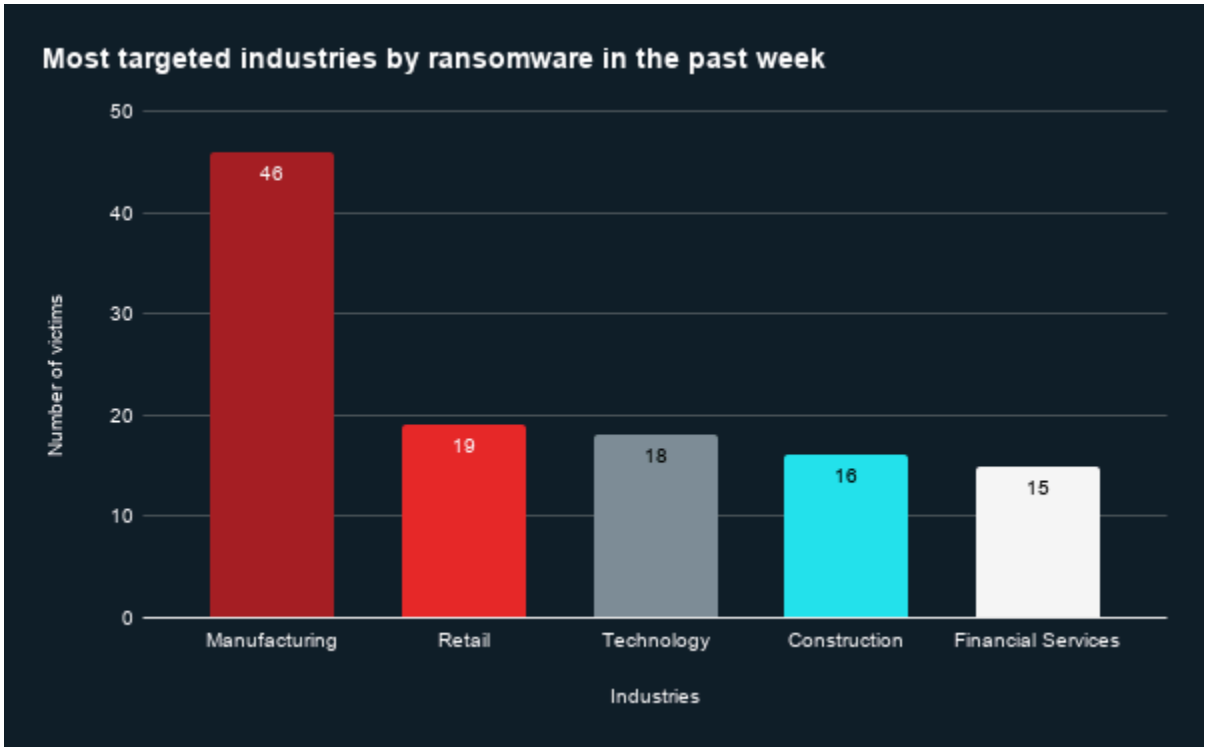

Source: ZeroFox Internal Collections

Regional ransomware trends: Over the past seven days, ZeroFox observed that North America was the region most targeted by ransomware attacks, followed by Europe and Russia and Asia-Pacific. There were at least 87 ransomware attacks observed in North America, while Europe and Russia accounted for 41, Asia-Pacific for 31, South America for 12, and the Middle East and Africa for 12.

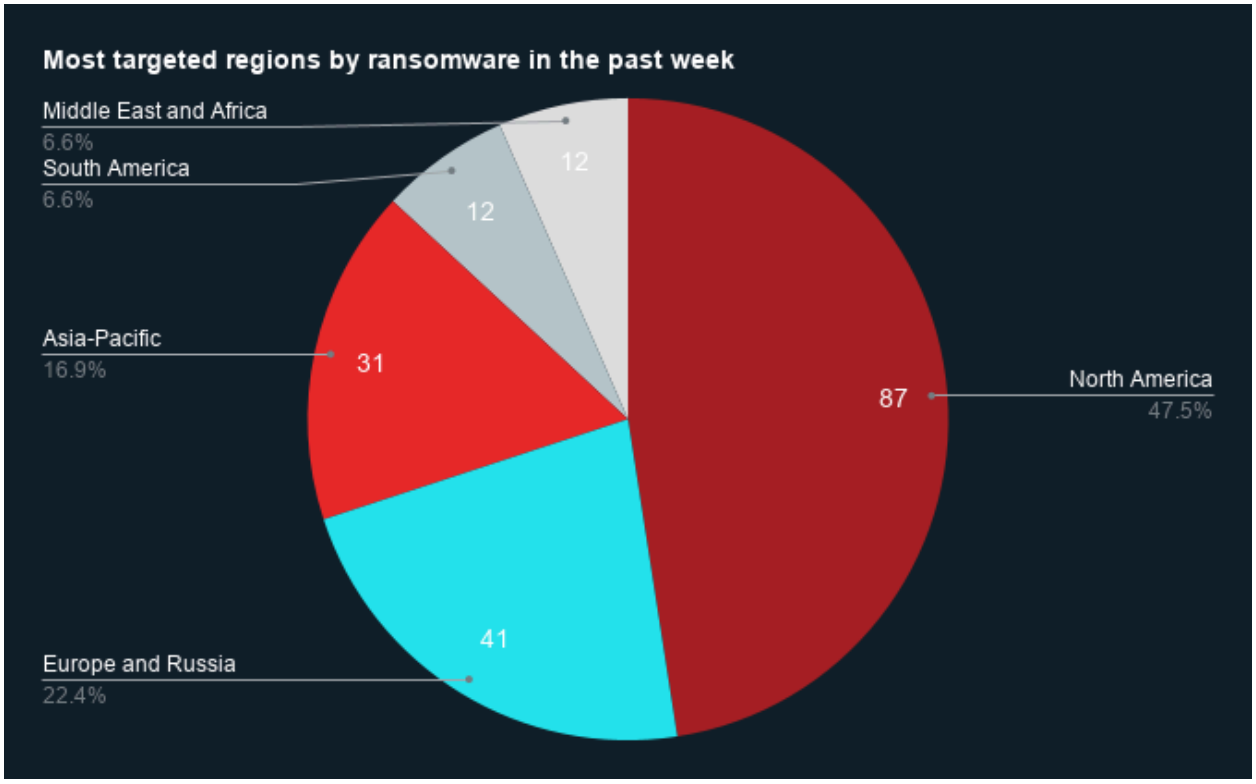

Source: ZeroFox Internal Collections

Data Breaches in the Past Week

Targeted Entity	Liechtenstein's Register of Beneficial Owners	Analog Devices	Żabka Polska convenience store chain
Compromised Entities/Victims	31,000 legal entities, including companies, foundations, and trusts registered in Liechtenstein	570,000 records claimed	541,000 Jira issues and 2.7 million user references
Compromised Data Fields	Names, nationalities, and dates of births of the beneficial owners of the legal entities	Customers' personally identifiable information and addresses	Employee and vendor data, Jira records, source code, credentials, access tokens, internal documentation, infrastructure details, and business operations data.
Suspected Threat Actor	N/A	ExfilSquad	Lumia
Country/Region	Europe	United States	Poland
Industry	Public Sector	Technology	Retail
Possible Repercussions	Financial intelligence gathering, corporate espionage, fraud, identity targeting, or abuse of ownership data	Extortion, public disclosure, phishing campaigns, or social engineering against employees and customers; supply chain risks; and risk of becoming a target for state-backed actors seeking advanced technology	Further unauthorized access, lateral movement, exploitation of vulnerabilities, phishing, social engineering, and identity-based attacks

Three major breaches observed in the past week

Appendix A: Traffic Light Protocol for Information Dissemination

	Red	Amber
WHEN SHOULD IT BE USED?	**Sources may use** **TLP:RED** when information cannot be effectively acted upon by additional parties and could lead to impacts on a party's privacy, reputation, or operations if misused.	**Sources may use** **TLP:AMBER** when information requires support to be effectively acted upon but carries risks to privacy, reputation, or operations if shared outside of the organizations involved.
HOW MAY IT BE SHARED?	**Recipients may NOT share** **TLP:RED** with any parties outside of the specific exchange, meeting, or conversation in which it is originally disclosed.	**Recipients may ONLY share** **TLP:AMBER** information with members of their own organization and its clients, but only on a need-to-know basis to protect their organization and its clients and prevent further harm. **Note that** **TLP:AMBER+STRICT** restricts sharing to the organization only.

	Green	Clear
WHEN SHOULD IT BE USED?	**Sources may use** **TLP:GREEN** when information is useful for the awareness of all participating organizations, as well as with peers within the broader community or sector.	**Sources may use** **TLP:CLEAR** when information carries minimal or no risk of misuse in accordance with applicable rules and procedures for public release.
HOW MAY IT BE SHARED?	**Recipients may share** **TLP:GREEN** information with peers and partner organizations within their sector or community but not via publicly accessible channels.	**Recipients may share** **TLP:CLEAR** information without restriction, subject to copyright controls.

Appendix B: ZeroFox Intelligence Probability Scale

All ZeroFox intelligence products leverage probabilistic assessment language in analytic judgments. Qualitative statements used in these judgments refer to associated probability ranges, which state the likelihood of occurrence of an event or development. Ranges are used to avoid a false impression of accuracy. This scale is a standard that aligns with how readers should interpret such terms.

Almost No Chance	Very Unlikely	Unlikely	Roughly Even Chance	Likely	Very Likely	Almost Certain
1-5%	5-20%	20-45%	45-55%	55-80%	80-95%	95-99%