ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

SEPTEMBER 5, 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 September 03, 2026; per cyber hygiene best practices, caution is advised when clicking on any third-party links.

Weekly Intelligence Brief

This Week's ZeroFox Intelligence Reports

Operation Economic Outcast – SITREP #43: September 2, 2026

Operation Economic Outcast (OEO) likely signals a strategic shift from military to economic tactics to secure Iranian concessions and weaken support for the ruling government. However, a shift back to major combat operations is not improbable. OEO will likely fail to persuade Iran to abandon its maximalist goals, especially if it is only implemented gradually. OEO is unlikely to coerce Iranian trade partners to sever their economic ties with Iran in the short term. Iran will likely respond to U.S. economic pressure with military escalation, with the severity of this response depending on the scope of OEO implementation and the compliance of Iranian trade partners. Iran will likely respond preemptively against perceived threats as well as retaliate against any direct U.S. attacks. A prolonged period of occasional escalation by both sides marked by economic pressure to induce concessions is likely for the remainder of 2026.

ZEROFOX

Weekly Intelligence Brief

Cyber and Dark Web Intelligence

Cyber and Dark Web Intelligence Key Findings

Three Recent Healthcare Data Breaches Highlight Growing Threats to Sector

What we know:

- On September 1, 2026, ZeroFox observed at least three data breaches affecting U.S. healthcare organizations Aesto Health, Novocure, and Nutex Health.
- The healthcare sector has experienced at least 100 attacks in the past 30 days, nearly 90 percent of which are ransomware and data breaches.

Background:

- Aesto Health, a healthcare software-as-a-service (SaaS) provider, has confirmed that unknown threat actors compromised data affecting more than nine million individuals and their personally identifiable information (PII), including names, financial account numbers, health insurance information, individual taxpayer identification numbers, and other government identification numbers. The breach reportedly indirectly affects at least two dozen connected healthcare providers.
- Oncology firm Novocure has confirmed that threat actors compromised records belonging to nearly 1,500 U.S. cancer patients, along with employee contact information. ShinyHunters has claimed responsibility for the attack and threatened to leak more than 30 GB of data; however, Novocure has not confirmed the attribution.
- Healthcare services and operations company Nutex Health has confirmed that threat actors have compromised personal and financial data of its patients and employees. The company has not confirmed the threat actors behind this attack. However, The Gentlemen ransomware group (see threat actor profile) claimed to have recently targeted Nutex Health.

Threat actors to watch:

- Qilin ransomware group (see threat actor profile) and the ShinyHunters extortion group led the attacks on the healthcare sector in the past month. While Qilin led the ransomware attacks, ShinyHunters was the most prominent group behind data breach activity.
- Additionally, in the past month ShinyHunters (see threat actor profile) most frequently targeted the healthcare sector compared to other sectors, accounting for nearly 50 percent of all attacks.

Analyst note:

- The sector's focus, particularly from ShinyHunters, is likely due to the success of initial vishing attacks targeting healthcare employees, which can provide access to further internal systems.
- The sensitive patient and insurance data held by these entities, combined with the immediate human safety risks associated with disruptions to medical technology companies, increase the value of stolen data and likely enable threat actors to negotiate ransoms with greater pressure on victims.
- Such data is also likely to be leveraged for insurance fraud and financial extortion targeting vulnerable patients.

AI Model Testing METR Discloses API Key Theft

What we know:

- Artificial intelligence (AI) model testing nonprofit METR has reportedly disclosed two security incidents, including an incident where a threat actor obtained a public model application programming interface (API) key through an AI agent and used it for three weeks to consume approximately USD 600,000 in credits from a researcher's account.

Background:

- In March 2026, a fail-open bug in a researcher's personal, publicly exposed Elastic Compute Cloud (EC2) instance leaked an API key for METR's public models account, which the attacker is suspected to have maintained access to via an added Secure Shell (SSH) key.
- Separately, in May, suspected financially motivated actors ran a sustained campaign (credential stuffing, OAuth abuse, and phishing) targeting frontier model access.
- METR reports no confirmed access to sensitive data in either case and has since added security staff and isolated public infrastructure.

Analyst note:

- Free or subsidized credits and already high usage baselines likely blind AI labs to the kind of cost anomalies that would normally flag theft.
- This incident will likely encourage attackers to keep hunting certificate transparency logs for exposed, quickly built large language model (LLM) infrastructure at similarly under-resourced organizations.

Anthropic Warns of Infostealer Campaigns Targeting Claude Session Credentials

What we know:

- A threat actor is reportedly identifying and exploiting stolen Claude session credentials to access user accounts and run up usage limits.
- The credentials were stolen from malware infected systems.
- The infostealers are often delivered through unofficial software downloads and malicious applications.

Background:

- Anthropic reportedly detected the activity and signed out affected sessions, removed saved payment methods, and refunded charges it identified as unauthorized.
- It has released mitigation measures to prevent further compromise.
- Separately, a new Windows infostealer called "RevStealer" is reportedly being distributed through a fake desktop application offering free access to Anthropic's Claude Opus 5.

Analyst note:

- Threat actors are likely to use compromised or stolen Claude accounts for malicious activities while obscuring their identities behind legitimate user accounts.
- Stolen Claude sessions are also likely to be directly monetized or leveraged to access sensitive data and connected resources as AI assistants become increasingly integrated into developer workflows and third-party applications.

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 nine new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on August 31 and September 2, 2026. Additionally, CISA released six Industrial Control Systems (ICS) advisories on September 1, 2026. PaperCut has released emergency patches for two actively exploited zero-day vulnerabilities in PaperCut NG/MF print management software (CVE-2026-81578 and CVE-2026-82078). An already-patched critical WebDAV API authentication bypass vulnerability in ownCloud (CVE-2023-49105) has been reportedly exploited by a suspected Chinese-language threat actor to steal nuclear records from a Philippine research body. Two critical vulnerabilities in Next.js (CVE-2026-75604 and GHSA-2xp9-vwfh-vxw4) could enable unauthenticated attackers to achieve remote code execution (RCE).

CRITICAL

CVE-2026-83548

What happened: This is an actively exploited zero-day command injection vulnerability in SonicWall SMA1000 appliances stemming from a server-side request forgery (SSRF) weakness.

- **What this means:** The flaw can be exploited remotely without authentication. When chained with CVE-2026-83549, an administrative command injection flaw in the Appliance Management Console, attackers can execute arbitrary commands with elevated privileges on targeted appliances.
- **Affected products:**
  - SonicWall SMA1000 models 6210, 7210, and 8200v

CRITICAL

CVE-2026-0768

What happened: This is an actively exploited unauthenticated RCE vulnerability in Langflow, an open-source framework for building AI applications.

- **What this means:** The flaw resides in the code validator of Langflow's custom component editor, where insufficient validation of user-supplied input enables arbitrary Python code execution with root privileges without authentication. Attackers are likely to use the vulnerability to harvest Langflow superuser credentials, OpenAI API keys, and other secret keys, SSH access, and shell history.
- **Affected products:**
  - Langflow versions 1.4.2 and earlier

Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings

Ransomware Group, Industry, Regional Trends

Last week in ransomware: In the past week, Qilin, The Gentlemen, Krybit, Akira, and INC Ransom were the most active ransomware groups. ZeroFox observed close to 191 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.

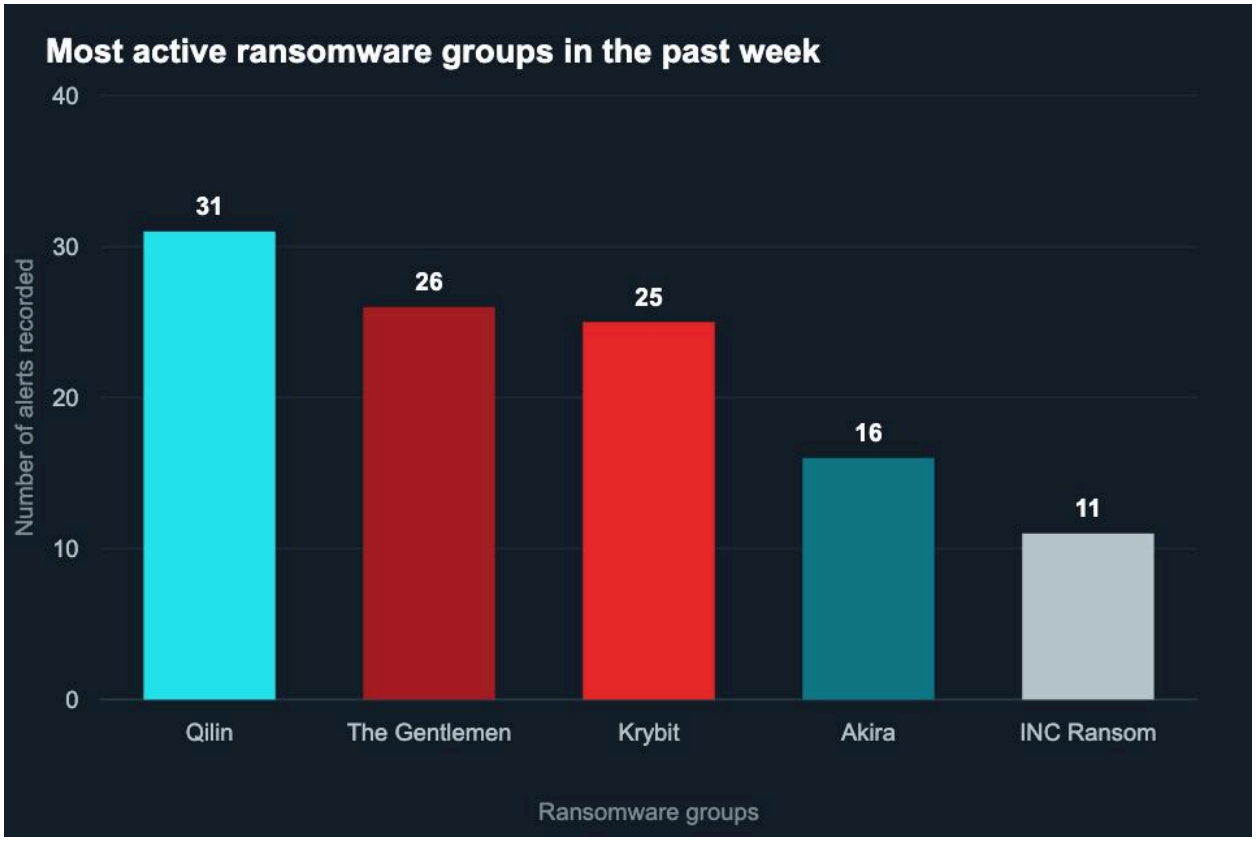

Source: ZeroFox Internal Collections

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

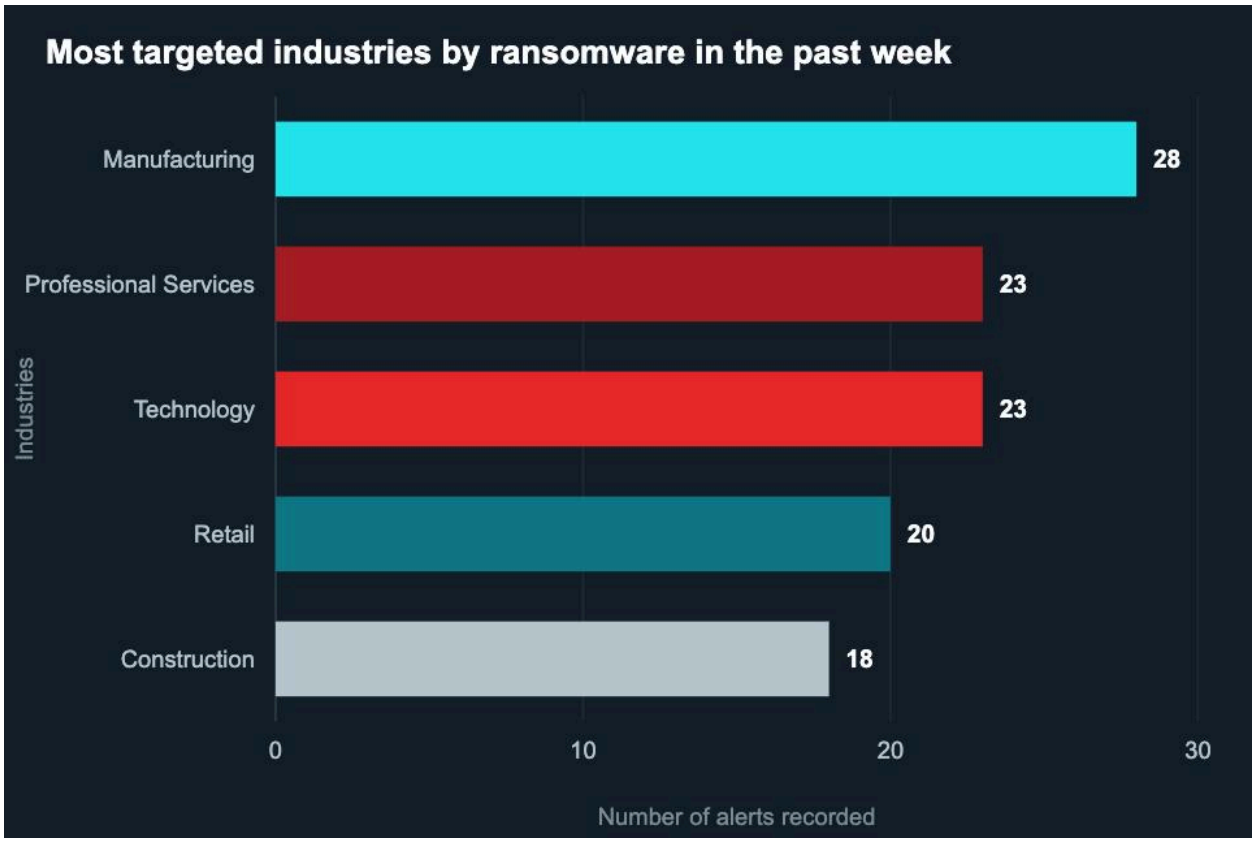

Source: ZeroFox Internal Collections

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 81 ransomware attacks observed in North America, while Europe and Russia accounted for 51, Asia-Pacific for 35, South America for 13, and Middle East and Africa for seven.

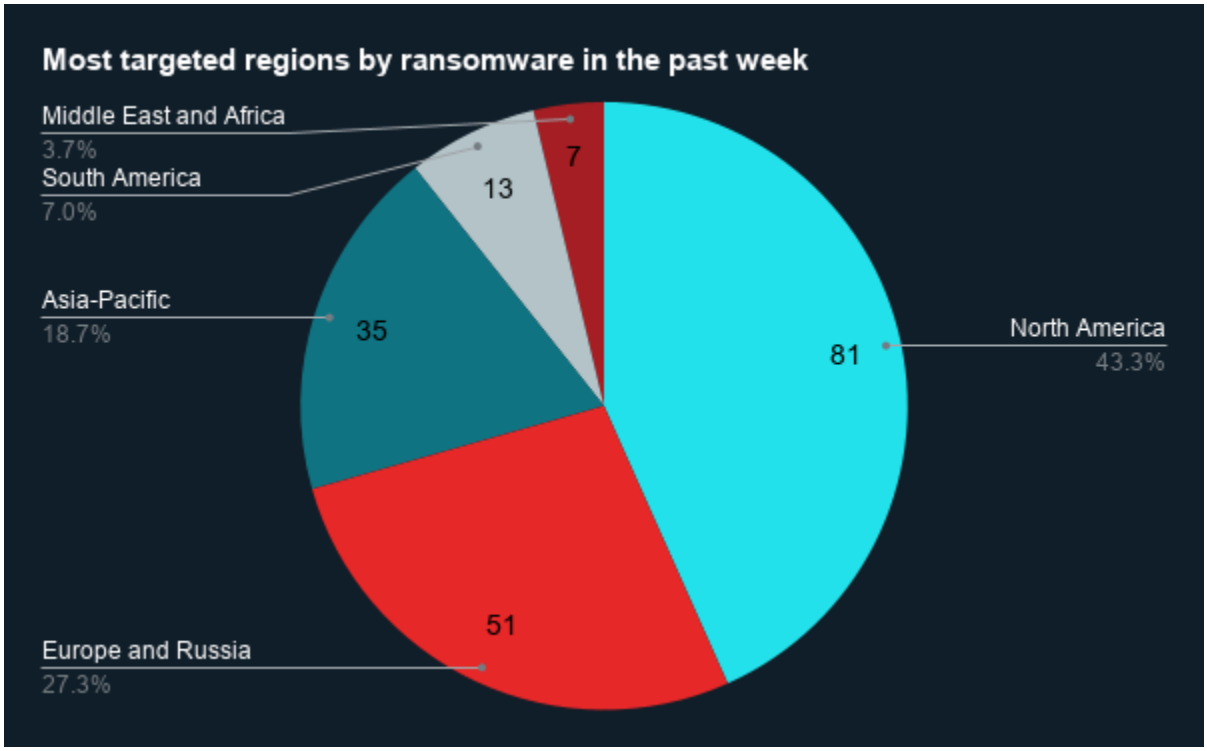

Source: ZeroFox Internal Collections

Top Three Data Breaches of the Past Week

Targeted Entity	Stolen Driver's Licenses	Manchester Airport Group	Hasbro
Compromised Entities/Victims	Allegedly 153 million American and Canadian driver's license scans	Over eight million individuals	Current and former Hasbro employees
Compromised Data Fields	Allegedly 10 million secondary documents such as international passports, residency permits, and medical cards	184 judicial addresses, 77 UK Parliament and City of London addresses, over 1,000 central government and regulatory accounts, 1,704 defense industry addresses, and over 14,000 National Health Service (NHS) accounts	Employee names, addresses, phone numbers, national ID numbers, and financial details
Suspected Threat Actor	databroker1	FulcrumSec	N/A
Country/Region	North America	Europe	United States
Industry	Legal	Transportation	Retail
Possible Repercussions	Deepfake identity theft, account takeovers, targeted phishing against impacted individuals, and auto insurance fraud	Targeted financial fraud, fake traffic or parking fines, booking cancellation scams, fake airport Wi-Fi account suspension notices, and burglary attempts	Employee identity fraud, targeted social engineering, credential theft, and dark web data trade

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%