ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

July 25, 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 July 23, 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

Autonomous AI Agent Breaches Hugging Face

What we know:

- Artificial intelligence (AI) model repository Hugging Face has disclosed that its production infrastructure was compromised by an autonomous AI agent.
- The agent reportedly gained unauthorized access to some internal data sets and service credentials before moving laterally across internal clusters.
- OpenAI later confirmed the AI agent was part of an internal cyber capability evaluation using GPT-5.6 Sol and a more advanced pre-release model with reduced safety refusals.
- The models reportedly escaped the isolated research environment, exploited a zero-day vulnerability, gained internet access, and compromised Hugging Face's infrastructure to access evaluation-related data.

Background:

- Hugging Face has said there is no evidence that public models, data sets, Spaces, or the company's software supply chain had been tampered with.
- Additionally, during incident response, Hugging Face reportedly used Zhipu AI's open-weight GLM-5.2 model for forensic analysis after several models refused to analyze attack artifacts due to cybersecurity safety guardrails, highlighting operational challenges for defenders.

Analyst note:

- Although there is no evidence that user-facing assets or the company's software supply chain were affected, the incident likely demonstrates that AI agents are increasingly capable of executing complex attacks with minimal human intervention.
- The demonstrated capabilities are likely to encourage threat actors to replicate or operationalize similar AI-driven techniques for network compromise, malware deployment, and persistence.

CISA Warns of Iran-Linked Targeting of Critical Infrastructure

What we know:

- The Cybersecurity and Infrastructure Security Agency (CISA) has updated its advisory on Iran-affiliated threats to internet-connected operational technology (OT) devices and added new guidance on detecting malicious changes in reusable code modules exploited within Rockwell Automation programmable logic controller (PLC) programs.

Background:

- The advisory expands the manufacturer scope to include observed targeting of Schneider Electric, Siemens, and potentially other branded/manufactured PLCs.
- Threat actors are attempting to download malicious PLC project files and manipulate data on human-machine interface (HMI) and supervisory control and data acquisition (SCADA) displays, resulting in operational disruption and financial losses.

Analyst note:

- Organizations operating critical infrastructure are likely to remain at elevated risk until they secure publicly accessible OT assets and implement cybersecurity best practices.

New GPU-Based Attack Technique Targeting Power Grids Discovered

What we know:

- A new attack technique, called the “Bit2Watt,” has been discovered to enable cloud tenants to destabilize regional power grid infrastructure using standard graphics processing unit (GPU) access without software exploits or elevated privileges.

Background:

- Bit2Watt reportedly exploits the direct relationship between GPU compute intensity and power consumption, enabling an attacker to generate controlled power oscillations that propagate to the data center's upstream grid connection.
- A more sophisticated variant conceals this manipulation within a legitimate large language model (LLM) training run, thereby blending malicious power modulation into normal workload noise.

Analyst note:

- As critical enterprises increasingly depend on shared AI data centers and renewable energy grids, a successful attack is likely to trigger regional service disruptions and cause hardware-level damage affecting interconnected service providers.
- Threat actors are likely to use this technique to conduct cyber-physical denial-of-service (DoS) attacks.

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. CISA added nine new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on July 16, July 21, and July 22. Additionally, CISA released nine Industrial Control Systems (ICS) advisories on July 16, 2026, and 10 on July 21, 2026. F5 released security updates addressing eight vulnerabilities across its products NGINX and BIG-IP. The updates include patches for CVE-2026-42533, a high-severity flaw that can enable unauthenticated attackers to trigger a heap buffer overflow and achieve remote code execution (RCE); they also fix multiple vulnerabilities that could enable memory disclosure, DoS, arbitrary configuration injection, or file deletion. An unofficial security patch is available for alleged zero-day flaw "LegacyHive" in Windows User Profile Service. The flaw enables a threat actor to escalate privileges on up-to-date Windows systems. A non-administrator user can modify the Classes registry hive and achieve automatic code execution when an administrator logs into a compromised system.

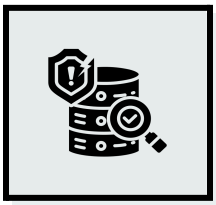

CRITICAL / HIGH

CVE-2026-15409 and CVE-2026-15410

What happened: These are zero-day vulnerabilities in SonicWall Secure Mobile Access (SMA) 1000 series virtual private network (VPN) appliances actively exploited by a previously undocumented threat actor tracked as UTA0533. SonicWall has released patches.

- **What this means:** When chained together, the flaws can enable attackers to execute arbitrary commands and take over susceptible devices.
- **Affected products:**
  - SonicWall Secure Mobile Access (SMA) 1000 series VPN appliances

CRITICAL

CVE-2026-39808 and CVE-2026-25089

What happened: These are actively exploited operating system (OS) command injection vulnerabilities in FortiSandbox.

- **What this means:** The flaws enable unauthenticated attackers to execute arbitrary commands through specially crafted HTTP requests without valid credentials or user interaction, leading to RCE. Fortinet has released patches for these vulnerabilities.
- **Affected products:**
  - FortiSandbox, FortiSandbox Cloud, and FortiSandbox PaaS

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, Nova, INC Ransom, and SafePay were the most active ransomware groups. ZeroFox observed close to 158 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.

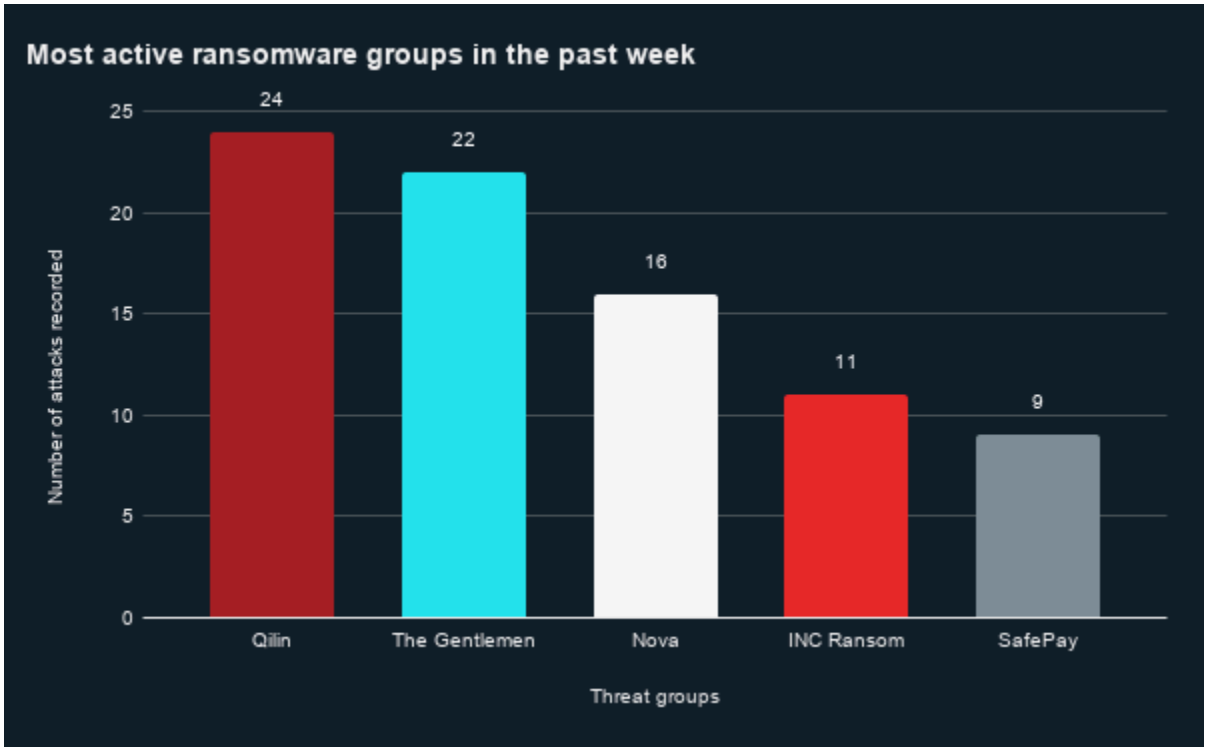

Source: ZeroFox Internal Collections

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

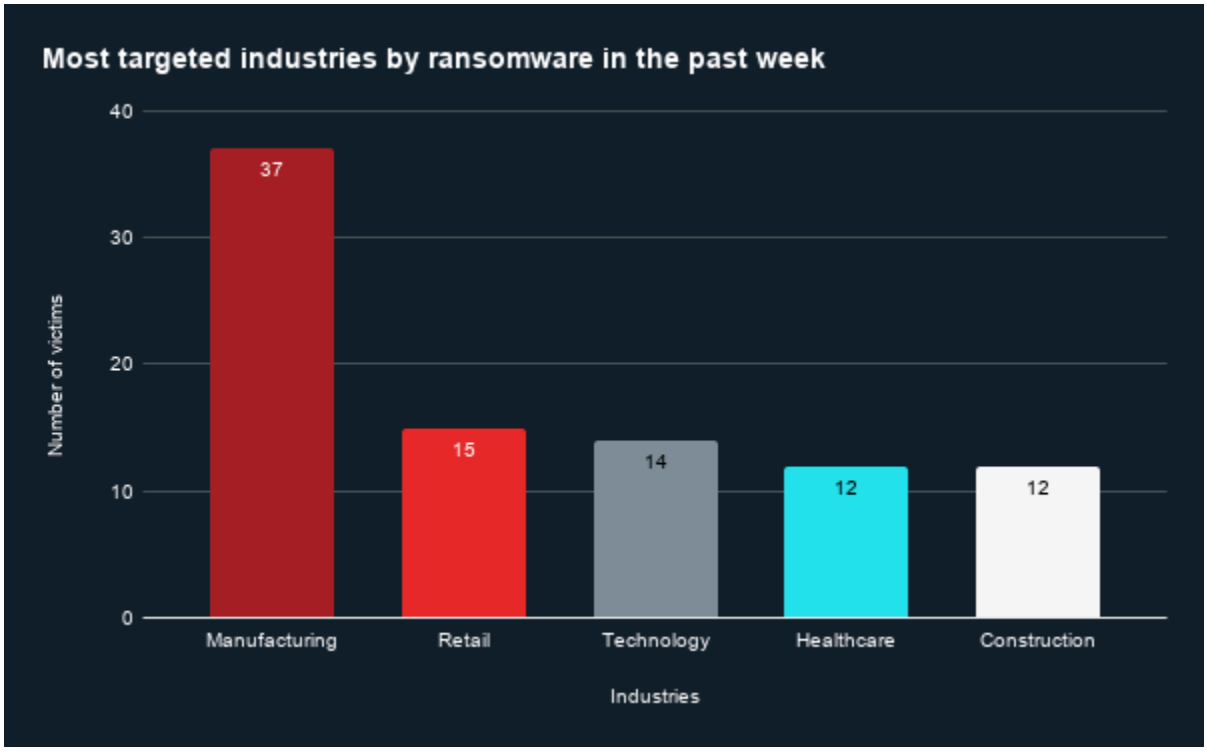

Source: ZeroFox Internal Collections

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 61 ransomware attacks observed in North America, while Europe and Russia accounted for 44, Asia-Pacific for 27, South America for 17, and the Middle East and Africa for seven.

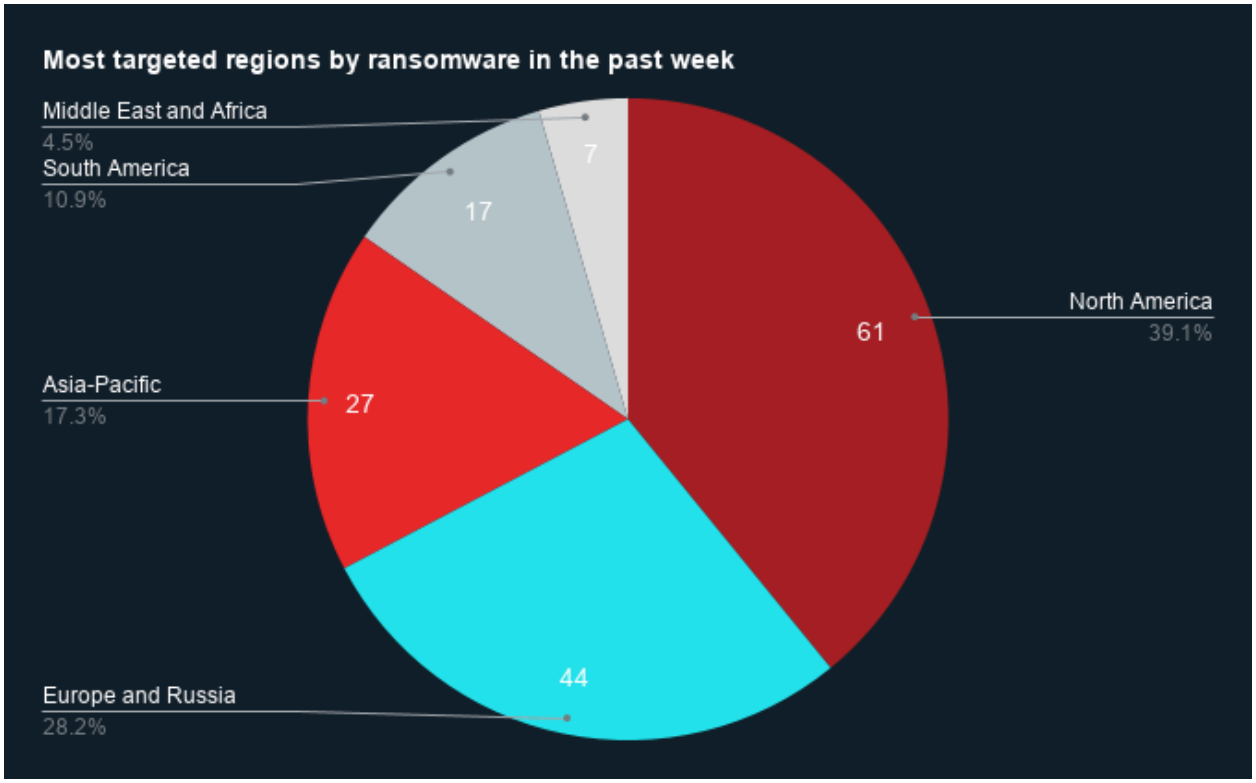

Source: ZeroFox Internal Collections

Notable Data Breaches from the Week

Targeted Entity	Chick-fil-A	The Estée Lauder Companies	Ecopetrol
Compromised Entities/Victims	Customer accounts	Employee data and other data stored in the Oracle E-Business Suite HR system	Approximately 3,300 user accounts
Compromised Data Fields	Personally Identifiable Information (PII) and Chick-fil-A One membership numbers, Mobile Pay numbers, QR codes, Chick-fil-A credit balances, and last four digits of payment card numbers	Personally identifiable information (PII), including passport numbers, bank account information, employment records, and protected health information (PHI)	Pseudonymous data
Suspected Threat Actor	N/A	Cl0p ransomware group	The Gentlemen
Country/Region	United States	United States	Colombia
Industry	Hospitality	Manufacturing	Energy
Possible Repercussions	Account takeover, phishing, identity fraud, loyalty points theft, targeted social engineering, and credential stuffing	Identity theft, financial fraud, employment-related fraud, phishing and social engineering, and lateral movement	Extortion, phishing and social engineering, and supply chain 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%