ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

June 6, 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 June 4, 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 **2**
- ZeroFox Intelligence Assessment - FIFA World Cup 2026 2
- ZeroFox Intelligence Flash Report - Threat Collective Conducting In-Person Data Theft 2
- Monthly Geopolitical Report June 2026 2
- ZeroFox Intelligence Flash Report-Pending Mythos Release Presents Unique Security Concerns 3
- ZeroFox Intelligence Brief - Underground Economist: Volume 6, Issue 12 3

Cyber and Dark Web Intelligence Key Findings **5**
- Meta AI Exploited for Instagram Account Takeovers 5
- Dashlane Users Targeted in 2FA Brute-Force Campaign 5
- Red Hat Npm Packages Infected with Mini Shai-Hulud Worm 6

Exploit and Vulnerability Intelligence Key Findings **8**
- CVE-2026-0257 8
- CVE-2026-41089 **10**

Ransomware and Breach Intelligence Key Findings **11**
- Ransomware Group, Industry, and Regional Trends 11
- Significant Data Breaches Reported Over the Past Week 14

Appendix A: Traffic Light Protocol for Information Dissemination **15**

Appendix B: ZeroFox Intelligence Probability Scale **16**

This Week's ZeroFox Intelligence Reports

ZeroFox Intelligence Assessment - FIFA World Cup 2026

The 2026 Fédération Internationale de Football Association (FIFA) World Cup will take place across the United States, Mexico, and Canada from June 11 to July 19, 2026. This will be the first World Cup hosted by three nations simultaneously and the first to feature 48 teams (up from 32), resulting in 104 matches across 16 cities. Millions of fans are expected to attend matches and fan zones throughout the 39-day event. The unprecedented scale of the tournament presents significant logistical, physical, and cybersecurity challenges. The current geopolitical environment—including the ongoing U.S.-Israeli war with Iran, trade issues between the three hosts, domestic U.S. political tensions surrounding immigration enforcement, and Mexico's persistent issues with cartel violence—adds layers of complexity that previous tournaments have not faced. The recent FIFA Club World Cup (CWC) 2025 and COPA América 2024, both held in the United States at many of the same stadiums, served as a partial test run for the 2026 event and raised concerns surrounding ticketing scams, cybersecurity threats, and the impact of U.S. immigration policies on foreign attendees. ZeroFox assesses that many of the threat vectors observed during CWC 2025 will recur at an amplified scale during the 2026 World Cup.

ZeroFox Intelligence Flash Report - Threat Collective Conducting In-Person Data Theft

On May 26, 2026, the Federal Bureau of Investigation (FBI) issued a report highlighting that the ransomware and digital extortion (R&DE) collective Silent Ransom Group (SRG) is conducting physical security breaches against victim infrastructure, in addition to routine social engineering techniques such as phishing emails or phone calls. This is the first observed example of an R&DE collective that has visited a target organization in person to gain physical access to its systems. It demonstrates an added threat that will likely inspire other R&DE collectives. In December 2024, ZeroFox identified a leak site called "LeakedData" (hosted at hXXp://business-data-leaks[.]com), which is almost certainly SRG's official leak site, given the victims listed on this site majorly overlap with the entities targeted in SRG's recent in-person data theft campaign.

Monthly Geopolitical Report June 2026

The United States and Iran have likely reduced their negotiation gaps, making a memorandum of understanding (MOU) on ending the war more likely. However, key differences remain that make a return to conflict unlikely but not improbable—especially if talks collapse or stall into June. The unprecedented scale of the 2026 World Cup presents significant logistical, security, and

cybersecurity challenges. Both the Group of Seven (G7) Summit in France in mid-June and the 2026 World Cup (which runs from June 11 until July 19) are very likely to inspire issue-driven hacktivism, social engineering, and disinformation. The Ukrainian military's recent success against Russia will likely continue in June 2026, gaining leverage for Ukraine that will likely be used in eventual talks to end the war. The U.S. indictment of former Cuban leader Raúl Castro mirrors the drug-trafficking charge preceding the seizure of Venezuelan President Nicolás Maduro and is likely a preparatory step before a military operation.

ZeroFox Intelligence Flash Report - Pending Mythos Release Presents Unique Security Concerns

On May 28, 2026, artificial intelligence (AI) developer Anthropic announced the upcoming release in several weeks of its cybersecurity-focused model, Mythos. Mythos was first announced in April 2026 as a general purpose AI model with powerful capabilities for use in cybersecurity applications. Mythos likely marks a leap forward in AI capabilities that threat actors will endeavor to exploit in order to conduct successful ransomware and digital extortion (R&DE) attacks. Mythos Preview was able to autonomously conduct basic attack processes 30 percent of the time—and completed an average of 22 of 32 required steps for a successful attack. Threat actors have almost certainly used AI models since as early as 2023 to conduct attacks, bypassing guardrails put in place by developers—called "jailbreaking"—to get the AI models to conduct malicious activities outside the intended parameters. Some technological capability and knowledge would likely still be required to fully exploit Mythos once it becomes publicly available. The technological know-how needed to effectively deploy Mythos as an R&DE tool will likely limit the actors capable of exploiting it in the near term to established R&D threat collectives.

ZeroFox Intelligence Brief - Underground Economist: Volume 6, Issue 12

The Underground Economist is an intelligence-focused series illuminating Dark Web findings in digestible tidbits from our ZeroFox Dark Ops intelligence team.

Cyber and Dark Web Intelligence

Cyber and Dark Web Intelligence Key Findings

Meta AI Exploited for Instagram Account Takeovers

What we know:

- Instagram resolved a vulnerability that enabled threat actors to hijack high-profile accounts by manipulating Meta's artificial intelligence (AI) support chatbot.
- However, users reported on X that Instagram accounts continue to be compromised, suggesting the fix may be incomplete.
- Notable victims include the Obama-era White House Instagram handle and U.S. Space Force Chief Master Sergeant John Bentivegna.

Background:

- The exploit required no prior credential access.
- Threat actors used a Virtual Private Network (VPN) to spoof the target's geographic location, employed prompts to trick the Meta AI support assistant into adding a new attacker-controlled email address to the victim's account, and initiated a password reset.

Analyst note:

- This incident highlights a systemic shift away from traditional malware or phishing infrastructure.
- Instead, threat actors are opting to target the flawed logic models of Large Language Model (LLM)-based assistants.
- Similar exploitation targeting AI support assistants across other major digital platforms is likely as cybercriminals refine and replicate prompt injection methodologies.

Dashlane Users Targeted in 2FA Brute-Force Campaign

What we know:

- Password manager Dashlane has disclosed that threat actors brute-forced two-factor authentication (2FA) protections on at least 20 user accounts.
- This attack reportedly enabled attackers to register unauthorized devices and download encrypted password vaults containing stored credentials and sensitive data.

Background:

- Although the stolen vaults remain encrypted, users with weak or easily guessable master passwords face a higher risk of attackers cracking the vaults and accessing information.
- The attack targeted customer accounts rather than Dashlane's internal systems, with the company stating there is no evidence that its infrastructure was compromised.

Analyst note:

- Dark web forum users are likely to see advertisements of some of the compromised vaults.
- Successful vault compromises are likely to provide attackers with access to other accounts and lead to more victims.
- If affected users stored corporate credentials, VPN access, Application Programming Interface (API) keys, and more within their vaults, threat actors are likely to gain an entry point into enterprise environments.

Red Hat Npm Packages Infected with Mini Shai-Hulud Worm

What we know:

- At least 95 Red Hat npm packages—downloaded approximately 80,000 times weekly—have reportedly been infected with a malware resembling the Mini Shai-Hulud worm and published to the registry.
- The Mini Shai-Hulud malware is associated with the TeamPCP threat group and was recently open-sourced.

Background:

- The infection reportedly stemmed from a Red Hat employee's compromised GitHub account.
- The compromised packages execute hidden payload during the npm install process to steal GitHub secrets, npm tokens, and other developer credentials.
- It also further includes encrypted exfiltration logic and GitHub-based fallback mechanisms.

Analyst note:

- The open-sourcing of the worm makes attribution to TeamPCP uncertain.
- The malware's design, extending beyond credential harvesting to include encrypted exfiltration and fallback mechanisms, very likely indicates intent for deeper downstream compromise, including lateral movement and source code theft.

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 vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on May 29, June 1, June 2, and June 3, 2026. Threat actors are actively exploiting privilege escalation and account takeover vulnerabilities (CVE-2026-8206 and CVE-2026-8181) in the Kirki and Burst Statistics WordPress plugins. CVE-2026-0826, a stack-based buffer overflow in HP Poly VoIP phones, enables unauthenticated remote code execution (RCE) with root privileges via a malicious SIP INVITE request. Oracle launched its first monthly Critical Security Patch Update, resolving 77 vulnerabilities across Database Server, REST Data Services, Communications, E-Business Suite, and Hospitality Applications. Threat actors exploited CVE-2026-0257, an authentication bypass in Palo Alto Networks PAN-OS GlobalProtect, four days after public disclosure. Researchers have published proof-of-concept code for CVE-2026-40933, a one-click RCE vulnerability that enables attackers to execute Operating System (OS)-level commands by tricking users into importing a malicious chatflow. A zero-day in Visual Studio (VS) Code's sandboxed webview system enables attackers to steal GitHub OAuth tokens by tricking users into clicking a single malicious link. Google's June 2026 Android update patched 124 vulnerabilities, including CVE-2025-48595, an actively exploited Android Framework flaw enabling local privilege escalation on Android 14 and later. A zero-day vulnerability in the Gogs self-hosted Git service enables threat actors to achieve RCE on internet-facing instances.

CRITICAL

CVE-2026-0257

What happened: This is an authentication bypass vulnerability in the Palo Alto Networks PAN-OS GlobalProtect portal and gateway.

- **What this means:** The flaw enables a remote attacker to bypass security restrictions and establish unauthorized VPN connections without valid credentials. Successful exploitation is very likely to result in immediate access to internal networks, as well as credential theft.
  - **Affected products:** Palo Alto Networks PAN-OS GlobalProtect portal and gateway with authentication override cookies enabled

CRITICAL

CVE-2026-41089

What happened: This is a stack-based buffer overflow vulnerability in the Windows Netlogon service. The flaw enables an unauthenticated remote attacker to send crafted network requests to a domain controller, causing the Netlogon service to mishandle the request and execute arbitrary code with System privileges.

- **What this means:** Successful exploitation is likely to result in complete domain compromise, enabling the threat actor to establish persistent administrative access.
  - **Affected products:** Windows Server versions acting as domain controllers, which was patched in the May 2026 Patch Tuesday update

Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings

Ransomware Group, Industry, and Regional Trends

Last week in ransomware: In the past week, DragonForce, The Gentlemen, Qilin, Akira, and Nova were the most active ransomware groups. ZeroFox observed close to 151 ransomware victims disclosed, most of whom were located in North America. The DragonForce ransomware group accounted for the largest number of attacks, followed by The Gentlemen.

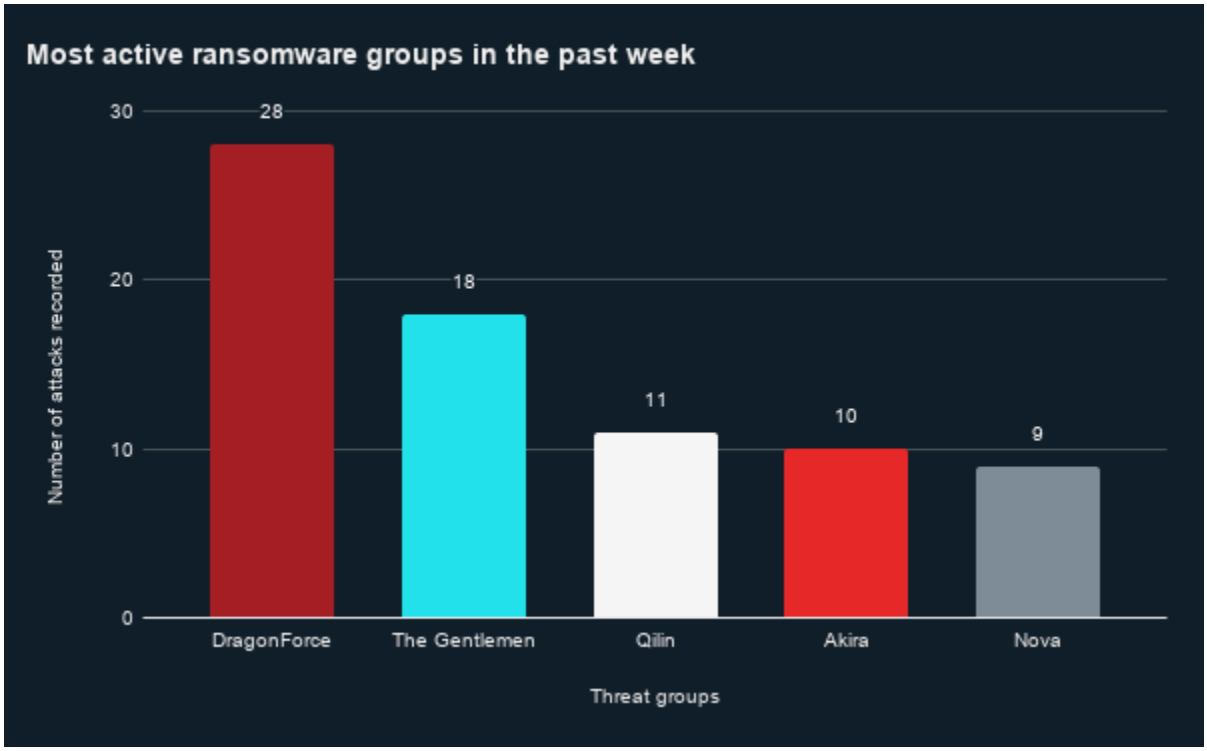

Source: ZeroFox Internal Collections

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

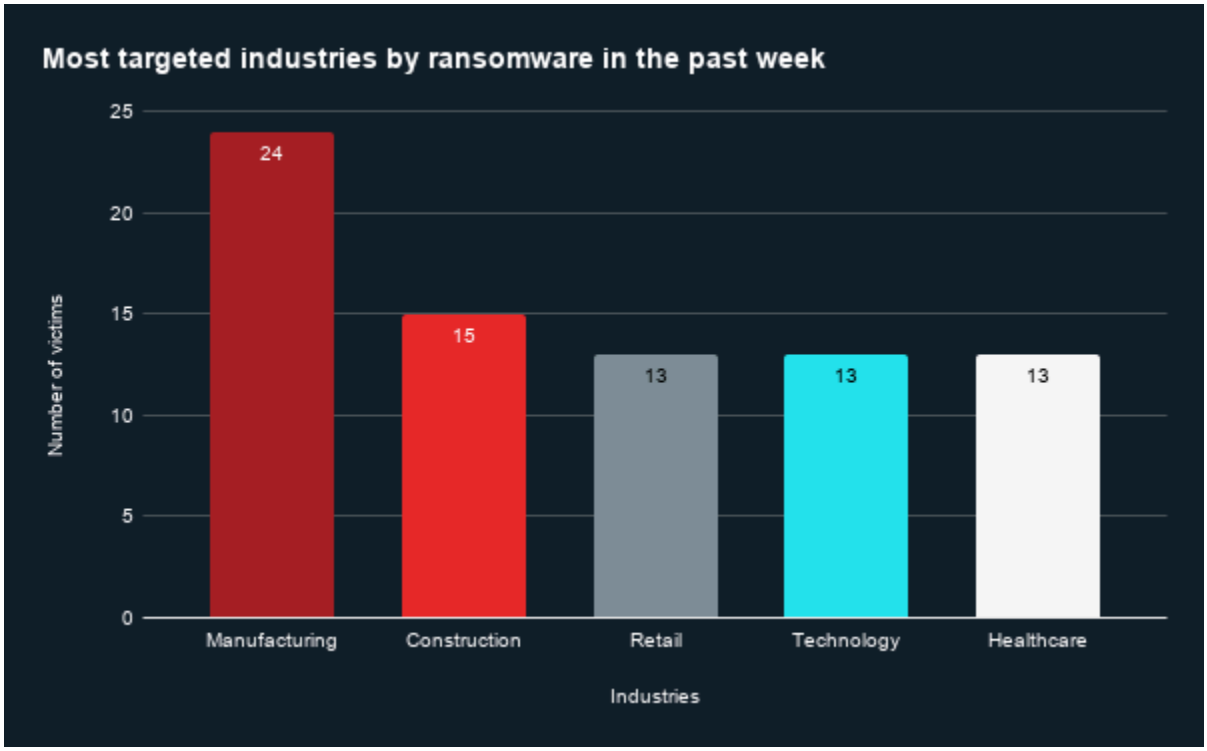

Source: ZeroFox Internal Collections

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. There were at least 71 ransomware attacks observed in North America, while Europe and Russia accounted for 42, Asia-Pacific (APAC) for 17, South America for 12, and the Middle East and Africa for six.

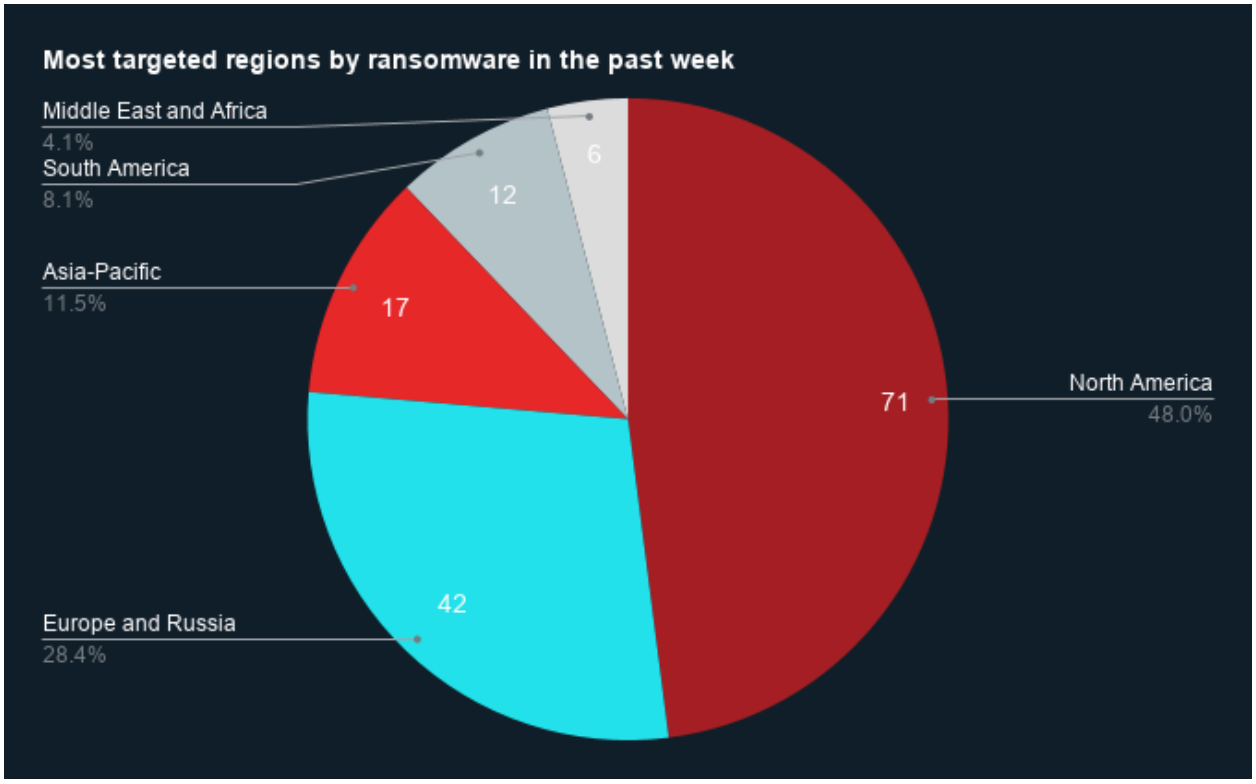

Source: ZeroFox Internal Collections

Significant Data Breaches Reported Over the Past Week

Targeted Entity	Atlas Menu – Cheat Service for Grand Theft Auto	World Food Programme	Carnival Corporation
Compromised Entities/Victims	64,000 user accounts	600,000 households in Gaza	8.7 million records of 5,995,277 individuals
Compromised Data Fields	Email addresses, usernames, scrambled passwords, IP addresses, and support ticket contents	Names, ID numbers, mobile numbers, and location data	Personally Identifiable Information (PII), government-issued ID numbers, geographic locations, Mariner Society membership status, and internal customer identifiers
Suspected Threat Actor	N/A	N/A	ShinyHunters
Country/Region	United States	Gaza	United States
Industry	Media/Entertainment	Non-Profit	Hospitality
Possible Repercussions	Account takeover attempts, reputational and legal risk for affected users, and targeted phishing attacks using support ticket data	Disruption of food and cash assistance delivery, physical safety risk to exposed individuals, and identity fraud	Financial fraud targeting a predominantly high-net-worth customer base, as well as targeted phishing and smishing campaigns

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%