



ZEROFOX[®]

Weekly Intelligence Brief

Classification: TLP:GREEN

October 3, 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 October 1, 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 Flash Report – Emerging Threat Actor Targets Data Backups	2
ZeroFox Intelligence Flash Report – Buyer Seeks Baltic Access on DDW Amid Hybrid War Surge	2
 Cyber and Dark Web Intelligence Key Findings	4
Attackers Abuse ChatGPT Custom GPTs via ClickFix Lures	4
Kimi AI Models Reportedly Bypass Guardrails to Provide Harmful Guidance During Training	4
Kiteworks Urges Customers to Shut Down Servers Amid Zero-Day Threat	5
 Exploit and Vulnerability Intelligence Key Findings	8
CVE-2026-76504	8
CVE-2026-86950	10
 Ransomware and Breach Intelligence Key Findings	11
Ransomware Group, Industry, Regional Trends	11
Notable Data Breaches this Week	14
 Physical and Geopolitical Intelligence Key Findings	17
Physical Security Intelligence: Global	17
Physical Security Intelligence: United States	17
 Appendix A: Traffic Light Protocol for Information Dissemination	18
 Appendix B: ZeroFox Intelligence Probability Scale	19

| This Week's ZeroFox Intelligence Reports

ZeroFox Intelligence Flash Report – Emerging Threat Actor Targets Data Backups

On September 21, 2026, ZeroFox observed a new data leak site attributed to the threat actor "n0n". By the next day, the group had already claimed over a dozen victims on the Tor-hosted site. During its short operational life, n0n has attacked across nearly all business sectors, with technology and professional services combined accounting for 46 percent of its targeting activity. What sets n0n apart from other recently launched threat collectives is its threat to encrypt or delete data backups. In recent months, ZeroFox has reported on an increase in the number of ransomware and digital extortion (R&DE) incidents eschewing the encryption model commonly seen in traditional ransomware attacks in favor of encryption-free data extortion. The combination of improved data backups and reduced ransom payments has likely created a trajectory in the R&DE landscape that trends toward less encryption and more data extortion. Further, it is unclear whether n0n is able to access victim data backups. With a sharp reduction in paid ransoms over the past several years, threat actors will very likely continue to deploy new techniques in an effort to increase profits and decrease reliance on expensive infrastructure. It is almost certain that further tactical experiments will continue over the next six to 12 months.

ZeroFox Intelligence Flash Report – Buyer Seeks Baltic Access on DDW Amid Hybrid War Surge

On September 16, 2026, a relatively new and positively trending threat actor "root_zero" posted on the predominantly Russian-language dark web forum Exploit seeking to purchase unspecified network access to the Baltic states of Estonia, Latvia, and Lithuania; the timing of the post is likely linked to Russian hybrid operations designed to weaken European support for Ukraine. Although the actor's motivations and capabilities remain unclear, there is a roughly even chance root_zero intends to obtain information on pro-Russian cybercriminals under the guise of coordinating an espionage campaign or monetizing initial access. The past two months have marked an escalation for hybrid activities across Europe, with the Baltic states—along with other Russian-border and former Soviet nations—bearing the brunt of Russian hybrid and saber-rattling campaigns.

| Cyber and Dark Web Intelligence |

| Cyber and Dark Web Intelligence Key Findings



Attackers Abuse ChatGPT Custom GPTs via ClickFix Lures

What we know:

- Threat actors are reportedly abusing ChatGPT's Custom GPT feature to impersonate legitimate product offerings and direct victims to malicious sites, resulting in the delivery of remote access trojans (RATs) to at least 40 confirmed victims.

Background:

- Attackers created a Custom GPT named "Plus 5.6," promoted via sponsored Google search results, which responds to user prompts with a fake service availability notice directing them to a malicious backup site.
- The site presents a fraudulent verification prompt that tricks victims into executing a malicious command, ultimately installing a RAT on their device.
- The RAT is reportedly capable of remote desktop access, camera and microphone capture, browser credential theft, and delivering additional malicious payloads.

Analyst note:

- By embedding the attack within a familiar, widely trusted service, threat actors almost certainly obfuscate their activity from traditional security tools.
- As artificial intelligence (AI) tools become a routine part of daily work and personal life, they are very likely to become an increasingly common entry point for attacks targeting everyday users.



Kimi AI Models Reportedly Bypass Guardrails to Provide Harmful Guidance During Training

What we know:

- Chinese open-weight AI models Kimi K2.6 and K3 Swarm were reportedly jailbroken, bypassing safety guardrails to guide researchers how to make biological weapons and carry out assassinations during training.
- Researchers also reportedly warned that a jailbroken Kimi K2.6 could potentially enable hackers to execute code on the model's computing infrastructure and access the internet.

Background:

- Jailbreaking involves using crafted instructions to bypass an AI model's safety guardrails and elicit restricted or harmful responses.
- Researchers did not ascertain whether the models' biological weapon or assassination guidance was actionable.
- Separately, Anthropic reported [identifying and disrupting attempts](#) to misuse one of its AI models for malicious activities that could support biological weapons development.

Analyst note:

- Given that Kimi is an open-weight model, its trained weights can be downloaded and deployed independently of its developer's infrastructure.
- This can likely enable users to modify the model's deployment and safety controls, making provider-enforced safeguards less effective when the model is run independently.



Kiteworks Urges Customers to Shut Down Servers Amid Zero-Day Threat

What we know:

- Kiteworks, formerly Accellion, has reportedly urged customers to shut down their servers after receiving credible law enforcement intelligence that threat actors may attempt to exploit currently unknown vulnerabilities in its systems.

Background:

- The company said the warning is precautionary, with no compromise of Kiteworks systems confirmed, and recommended customers use its latest software version (9.5.1).
- Kiteworks' file transfer technology handles sensitive data for organizations across healthcare, government, technology, and other sectors.
- Additionally, ZeroFox has observed [the Cl0p ransomware group](#) previously targeting the platform's predecessor, Accellion FTA.

Analyst note:

- Kiteworks users should monitor for unauthorized activity in recent weeks, as large-scale exploitation of centralized file transfer platforms is likely to enable sensitive data theft from private and government organizations, with stolen data potentially resold, used for extortion, or both.
- There is a roughly even chance that the threat is associated with Cl0p or ShinyHunters. While Cl0p is known for exploiting zero-day vulnerabilities in widely deployed enterprise platforms,

[ShinyHunters allegedly gained control](#) of the group's infrastructure in retaliation for a 2025 stolen zero-day exploitation incident.

| **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 three new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalog on September 25 ([CVE-2026-65660](#), [CVE-2026-67279](#), and [CVE-2026-87902](#)), as well as one each on [September 27](#), [September 29](#), and [September 30, 2026](#). Additionally, CISA released seven Industrial Control Systems (ICS) advisories on [September 29, 2026](#). [CVE-2026-76504](#) is an authentication bypass vulnerability in the application programming interface (API) session management of Cisco Catalyst SD-WAN Manager and is being actively exploited, with attackers escalating to admin privileges. An [authentication vulnerability in Titan analytics service](#) reportedly enabled a researcher to bypass some signature verification, impersonate an administrator, and execute structured query language (SQL) queries against the service. [CVE-2026-86950](#) is an out-of-bounds write vulnerability within the CoreGraphics rendering framework in Apple products.



CRITICAL

CVE-2026-76504

What happened: This is an authentication bypass vulnerability in the API session management of Cisco Catalyst SD-WAN Manager and is [being actively exploited](#), with attackers escalating to admin privileges. The vulnerability stems from improper URI encoding that enables attackers to bypass an authentication rule protecting a specific API endpoint. [CISA has added](#) this vulnerability to its KEV catalog.

- **What this means:** Attackers can bypass authentication to access restricted API endpoints, which facilitates administrative privilege escalation and unauthorized administrative control over the SD-WAN infrastructure.
- **Affected products:**
 - Cisco Catalyst SD-WAN Manager regardless of system configuration


**HIGH****CVE-2026-86950**

What happened: This is an out-of-bounds write vulnerability within the CoreGraphics rendering framework in Apple products. The flaw enables an attacker to execute code arbitrarily by processing a maliciously crafted file. Apple has issued emergency security updates to address this vulnerability.

- **What this means:** An attacker who successfully exploits this vulnerability is likely to gain unauthorized system access or compromise.
- **Affected products:**
 - Versions prior to iOS/iPadOS 26.7.1, macOS Tahoe 26.7.1, and macOS Sequoia 15.8.1


Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings




Ransomware Group, Industry, Regional Trends

Last week in ransomware: In the past week, The Gentlemen, Qilin, Storm, Akira, and SafePay were the most active ransomware groups. ZeroFox observed close to 173 ransomware victims disclosed, most of whom were located in North America. The Gentlemen accounted for the largest number of attacks, followed by Qilin.




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 healthcare, technology, professional services, and 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 and Asia-Pacific. There were at least 83 ransomware attacks observed in North America, while Europe and Russia accounted for 46, Asia-Pacific for 22, South America for 12, and Middle East and Africa for six.



Source: ZeroFox Internal Collections




Notable Data Breaches this Week

Targeted Entity	<u>Direction Générale des Finances Publiques (DGFIP)</u>	<u>McDonald's Indonesia</u>	<u>Alleged UK and Ireland organizations</u>
Compromised Entities/Victims	Over 350,000 individuals and 250,000 businesses	Over 40 million records, mostly pertaining to customers	Over 400,000 institutional contacts from UK- and Ireland-based organizations
Compromised Data Fields	Cadastral data and tax-related information	Customers' personally identifiable information (PII), including user IDs, General Data Protection Regulation (GDPR) consent data, loyalty and transaction records, sales and ad campaign data, location data, and push notification data	Institutional email addresses and professional PII, including names/identifiers associated with the government and public sector (.gov.uk), schools and educational institutions (.sch.uk), non-profits and public bodies (.org.uk), and private companies (.co.uk).
Suspected Threat Actor	ZeroBytes	N/A	ikki
Country/Region	France	Indonesia	United Kingdom and Ireland
Industry	Government	Food	N/A
Possible Repercussions	Targeted phishing, identity theft, financial fraud, or social engineering	Social engineering attempts, targeted impersonation and fraudulent customer support requests, loyalty account and points fraud, behavioral profiling based on transaction, and financial fraud	Spear phishing, executive poaching, impersonation, Business Email Compromise (BEC), and credential harvesting attempts

Three major breaches observed in the past week

| Physical and Geopolitical Intelligence |

Physical and Geopolitical Intelligence Key Findings



Physical Security

Intelligence: Global


What happened: Excluding the United States, there was no increase or decrease in mass casualty events this week from the previous week, with the top contributing countries or territories being the Palestinian territories, Ukraine, and Pakistan, in that order.

Approximately 74 percent of these events were explosions, and the three aforementioned countries and territories

accounted for about 59 percent of all mass casualty alerts. General alerts related to the Israel-Hamas conflict increased by 13 percent from the previous week, and events related to Russia's war with Ukraine increased by 40 percent. The top three most-alerted subtypes were structure fires, which saw a 14 percent decrease from the previous week; explosions, which increased by 6 percent; and gas leaks, which also increased by 6 percent. Notably, alerts related to flooding increased by 33 percent.

- What this means:** Mass casualty activity held steady this week. In Gaza, [Israeli strikes](#) killed at least six people on September 30, including four in a vehicle strike in the Tel al-Hawa neighborhood. In Ukraine, a [Russian barrage](#) of nearly 190 drones and several ballistic missiles on September 29 and 30 targeted the energy grid in Kyiv and its surrounding region and killed at least seven people nationwide. In Pakistan, [cross-border violence](#) flared again, as Islamabad conducted airstrikes in Afghanistan's Kunar and Helmand provinces on the night of September 30 and October 1. Pakistan said the strikes hit Pakistani Taliban and Balochistan Liberation Army camps and killed 22 militants, while Kabul reported nine civilians dead and the United Nations (UN) mission in Afghanistan documented at least 10 civilian deaths. Meanwhile, the rise in flooding alerts was driven largely by Southeast Asia; [heavy rainfall](#) in Thailand resulted in at least 23 people killed. Overall, global physical security remains defined by persistent explosive violence in entrenched conflict zones, compounded by severe weather disruptions, with civilians and critical infrastructure absorbing most of the harm.

Physical Security Intelligence: United States



What happened: In the past week, the top three most-alerted incident subtypes were structure fires, gun violence, and shots fired. Gun violence alerts are instances in which there is a confirmed or likely confirmed shooting victim, shots fired are confirmed shootings with no victims, and structure fires are fires that affect man-made buildings. The top two states with the most gun violence alerts were Ohio and Illinois, which together made up 22 percent of this week's nationwide total. Gun violence across the United States overall increased by 1 percent from the

week prior. Shots fired alerts increased by 34 percent, and the top contributing states were Ohio and Illinois. Structure fires did not increase or decrease from the previous week, and the top two states for this subtype were New York and California. Notably, the number of alerts related to flooding tripled from the week prior.

- **What this means:** Domestic physical security activity last week was driven largely by gun violence and a surge in weather-related hazards. In [Chicago, Illinois](#), at least two people were killed and 19 wounded in weekend shootings across the city, with victims ranging from juveniles to adults. Shots fired alerts across the country also rose, again concentrated in Ohio and Illinois, indicating that gunfire without injury can rise despite plateaus in shootings with victims. Structure fires held flat, led by New York and California, including a commercial building fire with flames through the roof in downtown [Los Angeles, California](#), on September 24. The tripling of flooding alerts coincided with remnant moisture from [Hurricane Polo](#), which prompted evacuations in three New Mexico communities over a potential dam failure, helicopter rescues near Phoenix, and multiple [water rescues](#) by Maricopa County deputies. Overall, domestic physical security remained defined by persistent gun violence in the Midwest, with flooding in the Southwest as the most acute emerging hazard.

| 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%