



ZEROFOX[®]

Weekly Intelligence Brief

Classification: TLP:GREEN

SEPTEMBER 19, 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 17, 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 – ShinyHunters Continues Shift to Encryption-Less Data Theft	2
Houthi's, Iran-Oman Deal, and Conflicting Military Doctrines – SITREP #44: September 15, 2026	2
 Cyber and Dark Web Intelligence Key Findings	5
FBI Disrupts NightmareStresser DDoS-for-Hire Service	5
AI Coding Assistant Session Hijacked to Spread Shai-Hulud Worm	5
Spain Reports First Confirmed AI Agent-Linked Data Breach	6
 Exploit and Vulnerability Intelligence Key Findings	9
CVE-2026-58704	9
CVE-2026-85706	11
 Ransomware and Breach Intelligence Key Findings	12
Ransomware Group, Industry, Regional Trends	12
Notable Data Breaches this Week	15
 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 – ShinyHunters Continues Shift to Encryption-Less Data Theft

On September 10, 2026, threat actor group ShinyHunters leaked data from McKesson Corporation (McKesson) on its dark web leak site. The data allegedly came from an August 2026 breach of 6.4 million personally identifiable information (PII) records associated with McKesson. Neither the public disclosures nor the ShinyHunters claim mentioned ransomware or encryption, making it very likely the group did not encrypt McKesson data; instead, ShinyHunters very likely extracted the data, issued a ransom demand, and then almost certainly leaked the data when the victim refused to pay. ZeroFox assesses that the lack of encryption in the McKesson breach was almost certainly intentional and very likely represents a tactical shift by ShinyHunters from double extortion to encryption-less data theft. The ShinyHunters data breach of McKesson and the subsequent release of sensitive data almost certainly reflect the continuation of the shift in the ransomware and digital extortion (R&DE) ecosystem from encryption-based attacks to “encryption optional” data extortion.

Houthi's, Iran-Oman Deal, and Conflicting Military Doctrines – SITREP #44: September 15, 2026

Houthi actions are not directly coordinated by Iran but almost certainly benefit Iran by increasing energy prices and stretching military resources. Houthi advances will likely persist unless Saudi Arabia offers concessions or increases military support to retake Houthi-controlled parts of Yemen. The U.S. military has reportedly implemented a tanker-for-tanker strategy in response to Iranian attacks on U.S. warships, to run alongside an economic pressure campaign. As part of this, the United States has very likely sought to expand commercial shipping through the Strait of Hormuz (SoH). A resumption in U.S. attacks likely reflects concern that Iran will target vessels transiting the SoH outside Iranian terms. Occasional U.S. strikes are likely necessary to degrade Iran's ability to control the SoH while pressuring it to make concessions. Iran's political-military leadership very likely views SoH control as key to deterring future conflict as well as an economic opportunity, which constrains the effectiveness of the U.S. approach. Iranian leadership has consistently stated its new military doctrine is offensive and based on retaliation and preemptive attacks. U.S. military and economic escalation is unlikely to restrain Iran, and there is a roughly even chance that Iran escalates further if U.S. tactics prove successful. This risks maintaining a new cycle of escalation with the U.S. military. Continued progress on the Iran-Oman SoH transit agreement, alongside reports of Gulf and U.S. support, indicate that Iran is unlikely to surrender control of the strait and that continued conflict is likely

costlier than accepting Iranian terms. Even if the deal is agreed upon, obstacles to normal SoH passage remain, as Iran has made clear that reopening still depends on further U.S. concessions linked to the June Memorandum of Understanding (MOU). If the accord never materializes or is rejected by the Trump administration, escalation is likely. Iran's hardline leadership is very likely continuing to drive policy decisions, reducing the likelihood that Iran adapts its positions.

| Cyber and Dark Web Intelligence |

| Cyber and Dark Web Intelligence Key Findings



FBI Disrupts NightmareStresser DDoS-for-Hire Service

What we know:

- The Federal Bureau of Investigation (FBI) has seized domains associated with NightmareStresser, a distributed denial-of-service (DDoS)-for-hire platform reportedly used by threat actors to facilitate attacks against targets worldwide.

Background:

- The domains were placed under law enforcement control as part of the disruption, with no arrests or charges announced against the service's administrators.
- NightmareStresser was reportedly active since 2022, targeting educational institutions, government agencies, and gaming platforms.

Analyst note:

- Although the seizure disrupts NightmareStresser's current infrastructure, the operators are likely to attempt to reconstitute the service using alternative infrastructure.
- The disruption is also likely to damage NightmareStresser's long-standing reputation and customer base if users perceive the service as compromised.



AI Coding Assistant Session Hijacked to Spread Shai-Hulud Worm

What we know:

- A threat actor reportedly hijacked an active artificial intelligence (AI) coding assistant session at an unnamed software-as-a-service (SaaS) provider and deployed the Shai-Hulud worm across approximately 100 internal code repositories to exfiltrate repository secrets and proprietary source code.

Background:

- The actor poisoned a PyPI package that the AI assistant recommended to a developer; upon acceptance, an infostealer was installed, and GitHub OAuth tokens were stolen.

- The worm self-propagated across internal repositories, and a second employee pulling the compromised package from the company's official namespace triggered a secondary infection.
- The Shai-Hulud worm family has featured in several recent developer-targeted campaigns.

Analyst note:

- The threat actor is almost certainly assessing the exfiltrated source code and credentials for further exploitation.
- Organizations relying on AI-assisted development are very likely to face increased targeting via similar vectors as threat actors mature techniques for abusing AI tool trust relationships.



Spain Reports First Confirmed AI Agent-Linked Data Breach

What we know:

- Spain's Data Protection Agency (AEPD) has reported the first confirmed personal data breach allegedly carried out by an AI agent.
- A third party reportedly used an autonomous system powered by a large language model (LLM) to identify vulnerabilities, gain unauthorized access, modify personal data, and access invoices—all with limited human intervention.

Background:

- The AI agent autonomously logged into the target system, searched for application weaknesses, exploited a vulnerability, altered personal information, and accessed billing records.
- The agency has not disclosed the identity of the targeted organization or the LLM involved.

Analyst note:

- This incident very likely signals that AI agent-driven intrusions are transitioning from theoretical risk to an emerging operational reality.
- As AI agents become more accessible and capable, threat actors are very likely to leverage them to establish persistent footholds, exfiltrate data, and conduct fraud and impersonation at a speed and scale that outpaces traditional human-speed incident response.

| **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. Additionally, CISA released [eight](#) Industrial Control Systems (ICS) advisories. ZeroFox observed active exploitation of a Linux local privilege escalation vulnerability, [CVE-2026-87886](#), in Acronis Backup plugins for cPanel, WHM, and Plesk. TP-Link released firmware updates for Tapo C200 home security cameras to address authentication bypass and service crash vulnerabilities. The Dutch National Cyber Security Centrum warned of impending exploitation of Check Point VPN vulnerabilities [CVE-2026-85102](#) and [CVE-2026-85103](#), which enable remote code execution on security gateways. Additionally, threat actors were observed chaining JFrog Artifactory vulnerabilities [CVE-2026-42018](#) and [CVE-2026-42016](#) to gain administrator access and deploy backdoors, alongside actively exploiting [CVE-2026-82329](#).



HIGH

CVE-2026-58704

What happened: A known exploited improper authorization zero-day flaw in the Android Modem component of Pixel devices caused by a logic error.

- **What this means:** Enables an attacker with adjacent network access and basic privileges to escalate privileges without user interaction.
- **Affected products:**
 - Google Pixel devices running affected Android Modem firmware



CRITICAL

CVE-2026-85706

What happened: A path traversal vulnerability in GitLab Community Edition (CE) and Enterprise Edition (EE) stemming from improper path confinement and missing authentication checks in the repository commits application programming interface (API).

- **What this means:** Enables unauthenticated attackers to read arbitrary files, secrets, and Secure Shell (SSH) configurations from a vulnerable GitLab server.

➤ **Affected products:**

- GitLab Community Edition and Enterprise Edition prior to versions 19.3.2, 19.2.6, and 19.1.8


Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings




Ransomware Group, Industry, Regional Trends

Last week in ransomware: In the past week, Qilin, Krybit, SafePay, Akira, and LockBit 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 Krybit.




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, construction, technology, and transportation.



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 75 ransomware attacks observed in North America, while Europe and Russia accounted for 34, Asia-Pacific for 25, South America for 13, and Middle East and Africa for eight.



Source: ZeroFox Internal Collections




Notable Data Breaches this Week

Targeted Entity	CenterPoint Energy	Japan's Digital Agency	Revolut
Compromised Entities/Victims	CenterPoint Energy customers	Government Solution Service (GSS) personnel and users	Revolut customers
Compromised Data Fields	Over 7.49 million records containing names, phone numbers, billing and service addresses, email addresses, driver's license numbers, and partial Social Security numbers (SSNs)	246,000 records including names, email addresses, phone numbers, and physical addresses	Customer identity and contact information, identity documents, verification selfies, account statements, and transaction histories
Suspected Threat Actor	Hex_4d722e4d656f77	N/A	Mister777
Country/Region	United States	Japan	United Kingdom
Industry	Energy & Utilities	Government	Financial Services
Possible Repercussions	Targeted phishing, identity fraud, and unauthorized account access	Targeted phishing and impersonation of government personnel and organizations	Social engineering, identity theft, and financial fraud

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 a 4 percent decrease in mass casualty events this week from the previous week, with the top contributing countries or territories being the Palestinian territories, Mexico, and Yemen, in that order. Approximately 73 percent of these events were explosions, and the three aforementioned countries and territories accounted for about 29

percent of all mass casualty alerts. General alerts related to the Israel-Hamas conflict increased by 20 percent from the previous week, and alerts related to the war in Iran decreased by 25 percent. Events related to Russia's war with Ukraine decreased by 32 percent. The top three most-alerted subtypes were structure fires, which saw a 2 percent increase from the previous week; gas leaks, which increased by 9 percent; and gun violence, which decreased by 1 percent.

- **What this means:** Global physical security conditions this week reflected a landscape still shaped by active armed conflict and organized violence, even as overall event volume ticked slightly downward. In Gaza, a six-story residential building in the Tal al-Hawa area collapsed after [Israeli strikes](#) on September 16, killing at least 21 people, including 11 children, and leaving dozens missing beneath the rubble; this is consistent with the rise in general Israel-Hamas conflict alerts. Yemen's contribution was driven by a [major Houthi advance](#) this month; the Iran-backed group seized the Red Sea port of Mokha on September 10 and the strategic island of Perim (Mayun) the following day, extending their reach toward the Bab el-Mandeb Strait and heightening threats to Saudi oil shipments. Mexico's contribution was anchored by a [fireworks explosion](#) at a festival in Temascalcingo, which killed at least 10 people and injured more than 60. Meanwhile, alerts tied to Russia's war in Ukraine dropped overall, even as attacks continued; a [Russian missile and drone barrage](#) overnight into September 17 struck Kyiv, Zaporizhzhia, and Odessa regions, injuring at least 19 people. Taken together, the data reflects a global security landscape that remains highly volatile, with conflict-driven violence in the Middle East and Eastern Europe continuing to generate significant risk.

Physical Security Intelligence: United States



What happened: In the past week, the top three most-alerted incident subtypes were structure fires, gun violence, and knife attacks. Gun violence alerts are instances in which there is a confirmed or likely confirmed shooting victim, knife attacks include confirmed stabbings or slashings, and structure fires are fires that affect man-made buildings. The top two states with the most gun violence alerts were Illinois and California, which together made up 21 percent of this week's nationwide total. Gun violence across the United States overall decreased by 4 percent

from the week prior. Knife attack alerts decreased by 11 percent, and the top contributing states were California and New York. Structure fires increased by 16 percent, and the top two states for this subtype were also California and New York. Notably, wildfires nationwide increased by 125 percent.

- What this means:** Gun violence remained concentrated in Illinois, where at least three people were killed and 17 others injured in [shootings across Chicago](#) over the weekend of September 11–14. In New York, two people were stabbed following a dispute at [SUNY Westchester Community College in Valhalla](#) on September 16, with two suspects taken into custody. In California, a fire burned through the top floor of a four-story apartment building in [Oakland](#) on September 15, requiring about two dozen firefighters to bring it under control. The sharp rise in wildfire activity was mirrored on the ground, with the fast-moving [Lucas Fire](#) along the Lake Mendocino county line triggering mandatory evacuations, emergency shelter activations, and major road closures, as of September 13. Taken together, this week's data points to a domestic physical security landscape marked by persistent, geographically concentrated violent crime and a rapidly intensifying wildfire threat, even as some incident categories showed modest declines.

| Appendix A: Traffic Light Protocol for Information Dissemination

WHEN SHOULD IT BE USED?

Red

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.

Amber

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

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.

Clear

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%