# ZEROFOX®

## *Weekly Intelligence Brief*

**Classification: TLP:GREEN**

**June 20, 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 18, 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

## The United States and Iran Agree to End War – SITREP #40: June 18, 2026

U.S. President Donald Trump and Iranian President Masoud Pezeshkian signed a Memorandum of Understanding (MOU) on ending the war. The final terms strongly align with those previously outlined in ZeroFox's SITREP report dated May 28, 2026. The MOU establishes a timeline for future talks on Iran's nuclear program and sanctions relief, despite both sides almost certainly holding fundamentally different positions on most issues. It mandates a 60-day ceasefire extension to facilitate talks and the reopening of the Strait of Hormuz (SoH)—a timeline almost certain to be extended given the anticipated difficulty of negotiations. A durable solution regarding Israeli operations in Lebanon remains very unlikely as it necessitates that Iran restrain Hezbollah and the United States curb Israel. Along with divergent interpretations of the SoH reopening and sanctions relief, this impasse poses the primary short-term risk of conflict resurgence during the 60-day negotiation period.

## ZeroFox Intelligence Flash Report - Spamming Package Targeting the U.S. SSA Advertised on DDW

On June 6, 2026, untested threat actor "mailerborn" advertised a spam distribution package targeting the U.S. Social Security Administration (SSA) on the predominantly Russian-language deep and dark web (DDW) forum Exploit. Mailerborn joined Exploit on May 21, 2026, and has made nearly 30 posts as of reporting but has garnered only one reputation point. The claimed features of the package—including a command loader that can evade common security controls, access to about 500 corporate Simple Mail Transfer Protocol (SMTP) servers, and per-recipient email generator—are likely to enhance phishing capabilities. Although ZeroFox has previously observed similar spam-related services—including email bombing and SMS spamming tools—on underground forums, this is likely a dedicated offering built around SSA-themed lures.

## ZeroFox Intelligence Flash Report - AI-Ransomware Toolkit Automates Operations

On June 2, 2026, security researchers discovered an unknown threat actor was almost certainly using commercially available artificial intelligence (AI) technologies to develop and iteratively test Endpoint Detection and Response (EDR) evasion techniques within a post-exploitation framework that was presented as a "red team" exercise. The threat actor reportedly used AI to accelerate tool development and testing, but the operation remained human-driven. AI was very likely used

primarily to coordinate workflows and support experimentation, while the EDR-bypass work followed a structured engineering test cycle that included human review and iteration. ZeroFox assesses the framework was very likely built for criminal use rather than legitimate security testing. The activity is linked to known ransomware deployment and data theft operations, and the red team framing was likely a pretext to circumvent the AI model's safety guardrails. ZeroFox assesses that the use of AI to accelerate tooling and test evasion techniques likely lowers the barrier to entry for sophisticated, red team-style intrusions but does not change defensive priorities. Fundamentals such as timely patching, multi-factor authentication (MFA), modern authentication (such as passkeys), and broad EDR deployment likely remain the primary mitigations.

## ZeroFox Intelligence Flash Report - LLM Jailbreak Chatter on the Deep and Dark Web

ZeroFox researchers have continuously observed discussions on the DDW regarding alleged jailbreaks of various AI and large language model (LLM) tools, including some pertaining to the June 9, 2026, release of Anthropic's Claude Mythos 5 and Fable 5. Researchers have also identified discussions and offers for Claude 4.6 and Opus 4.8 jailbreaks on the Telegram channel "GANOSECTEAM COMUNITY." On June 10, 2026, Russian-language threat actor "d4rm3an" announced on the dark web forum ReHub that Anthropic had released its Fable 5 model to the public one day earlier; the actor claimed to have successfully extracted and bypassed the model's system prompt. Offerings of jailbreaks for LLMs pose significant security concerns and demonstrate an ever-evolving cybercrime space almost certainly comprising highly motivated threat actors seeking to identify new and novel ways to conduct malicious activity targeting potential victims.

# Cyber and Dark Web Intelligence

# Cyber and Dark Web Intelligence Key Findings

## Cal Water Allegedly Hacked by Iran-Linked Handala Hack Team

**What we know:**

- Iran-linked hacktivist group Handala Hack Team has claimed responsibility for hacking California Water Service (Cal Water) in retaliation for U.S. actions against Iran.
- The group alleges it has the capability to disrupt water access but has opted not to.

**Background:**

- Handala has also published 5 GB of data allegedly linked to the U.S. water utility. Cal Water's Chico District has reportedly been affected in the attack.
- The leaked data reportedly contains personally identifiable information (PII) of customers, administrative credentials for the RTKBase platform, and a mountpoint-level NTRIP source password.

**Analyst note:**

- The RTKBase platform compromise is likely to enable threat actors to manipulate GPS data used for mapping underground water infrastructure, guiding repair or construction work, and monitoring ground movement near the infrastructure.
- However, the accesses listed so far are unlikely to directly disrupt customer water supply. Exposed individuals are likely to be targeted in identity theft, phishing, and social engineering attacks.

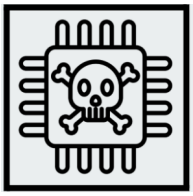

## "FortiBleed" Campaign Targets Over 320K Fortinet Devices Globally

**What we know:**

- An ongoing credential-harvesting campaign, dubbed "FortiBleed," has reportedly compromised approximately 75,000 Fortinet firewall and Virtual Private Network (VPN) devices across 194 countries.
- Fortune 500 companies and government agencies in more than 15 countries have reportedly been affected.

**Background:**

- The primary exposed server contains usernames, email addresses, and plaintext passwords for 73,932 unique firewall URLs.
- Attackers launched roughly 1.16 billion credential-stuffing attempts against more than 320,000 FortiGate targets.
- The campaign does not exploit a known vulnerability; instead, it relies on previously exposed Fortinet infrastructure to deploy credential-stuffing and brute-force techniques.

**Analyst note:**

- Organizations that have not rotated passwords or enforced MFA almost certainly face immediate risk.
- A credential leak of this scale is also likely to fuel a significant surge in Initial Access Broker (IAB) activity, enabling cybercriminals to sell network access to ransomware syndicates.

## Exposed Database Sourced from Infostealer Logs Taken Offline

**What we know:**

- Researchers have reportedly discovered a leaked database of 24 billion records originating from an Elasticsearch cluster, containing over 8 TB of data.
- The data set is an aggregation of data and not the result of a recent breach. The database is now reportedly offline.

**Background:**

- The database reportedly contained data sourced from infostealer logs with usernames, emails, passwords, and login URLs, originating from 36 sources, including Telegram channels and breach compilations.
- It is yet to be determined how many of these records are duplicates.

**Analyst note:**

- The exposure is a data leak as a result of a server misconfiguration and almost certainly does not involve threat actor malicious activity.
- Although the database was taken offline, the briefly exposed information is likely to have been archived or copied.
- Organizations are advised to rotate credentials and enable MFA).

# 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 four new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on June 12, June 15, and June 16, 2026. Additionally, CISA released five Industrial Control Systems (ICS) advisories on June 16. A vulnerability in Google Cloud Vertex AI SDK for Python, dubbed "Pickle in the Middle," reportedly enables threat actors to hijack a victim's machine learning model upload and run code inside Google's serving infrastructure. Threat actors reportedly only require the victim's project ID, which is often public, to exploit the flaw. A pre-authentication remote code execution (RCE) vulnerability in Splunk Enterprise, tracked as CVE-2026-20253, enables any network-reachable user to invoke arbitrary file operations without credentials. Oracle has released emergency mitigations for CVE-2026-35273, a critical PeopleSoft Suite vulnerability exploited as a zero-day in the ShinyHunters data theft campaign targeting the education sector.

**CRITICAL**

## CVE-2026-48907

**What happened:** CISA has ordered federal agencies to patch this actively exploited vulnerability in the Widget Factory Joomla Content Editor (JCE) plugin by Friday. The vulnerability enables unauthenticated attackers to upload and execute malicious PHP code via automated, low-complexity attacks targeting Joomla deployments that use the JCE WYSIWYG editor plugin.

- **What this means:** Successful exploitation is likely to give threat actors full control over the victims' web servers, leading to operational disruptions and/or data theft. Compromised servers also likely risk being used in cybercriminal infrastructure.
  - **Affected products:** Versions before JCE Pro 2.9.99.6

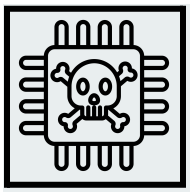

**MEDIUM**

## CVE-2026-20262

**What happened:** This is an actively exploited vulnerability affecting Cisco Catalyst SD-WAN Manager (formerly SD-WAN vManage). An attacker with at least a lower-privileged, single-task user account credential can exploit this vulnerability by sending a crafted HTTP request to an affected Application Programming Interface (API) endpoint of the affected system.

- **What this means:** A successful exploit is likely to enable the attacker to create or overwrite any file on the underlying operating system.
  - **Affected products:** The affected products are listed in this advisory.

# 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, DragonForce, LockBit and SafePay were the most active ransomware groups. ZeroFox observed close to 145 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.

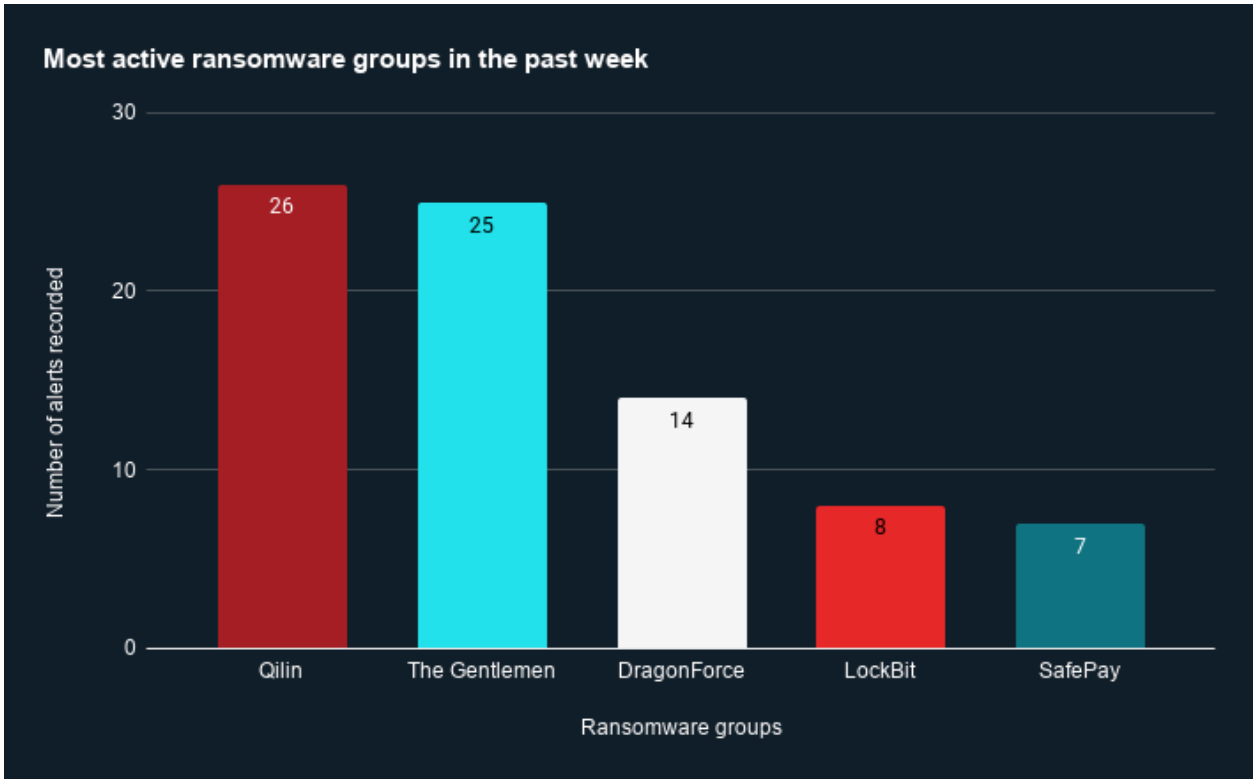


Source: ZeroFox Internal Collections

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

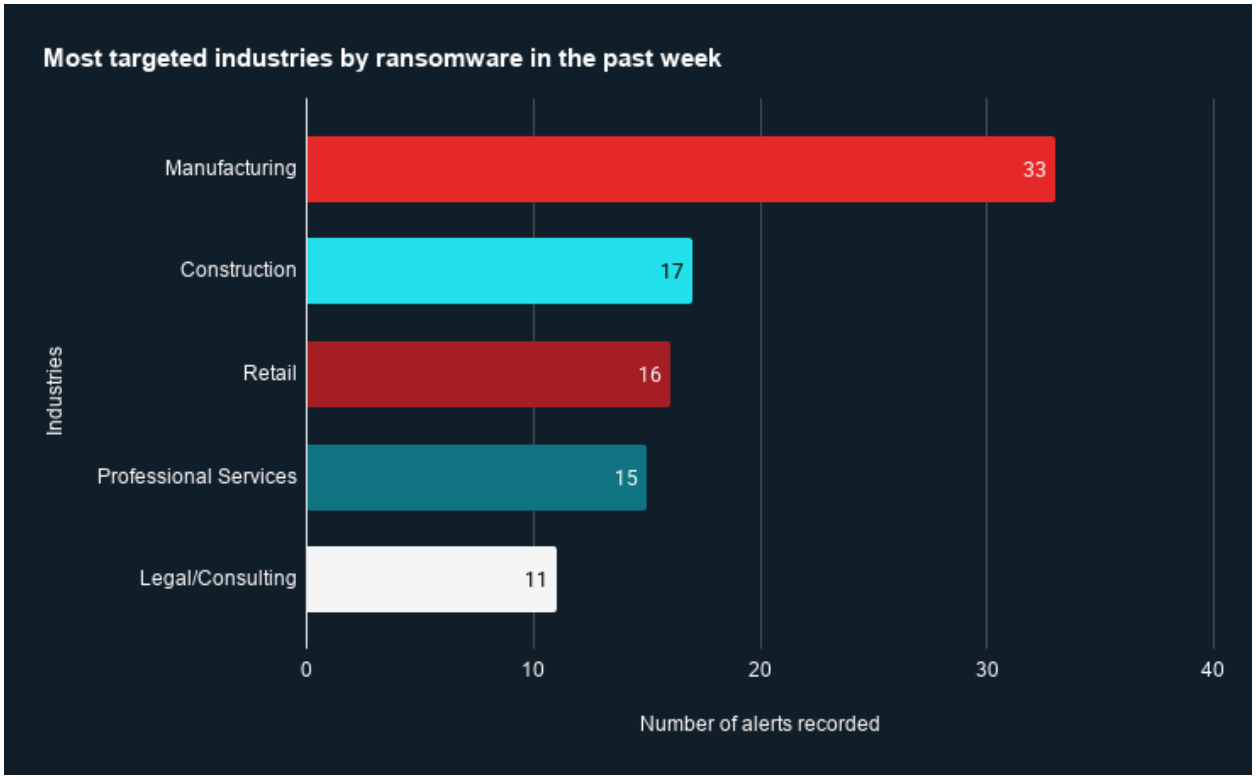


Source: ZeroFox Internal Collections

**Regional ransomware trends:** Over the past seven days, ZeroFox observed that North America was the region most targeted by ransomware attacks, followed by Asia-Pacific and Europe and Russia. There were at least 62 ransomware attacks observed in North America, while Asia-Pacific accounted for 33, Europe and Russia for 30, South America for 11, and the Middle East and Africa for six.

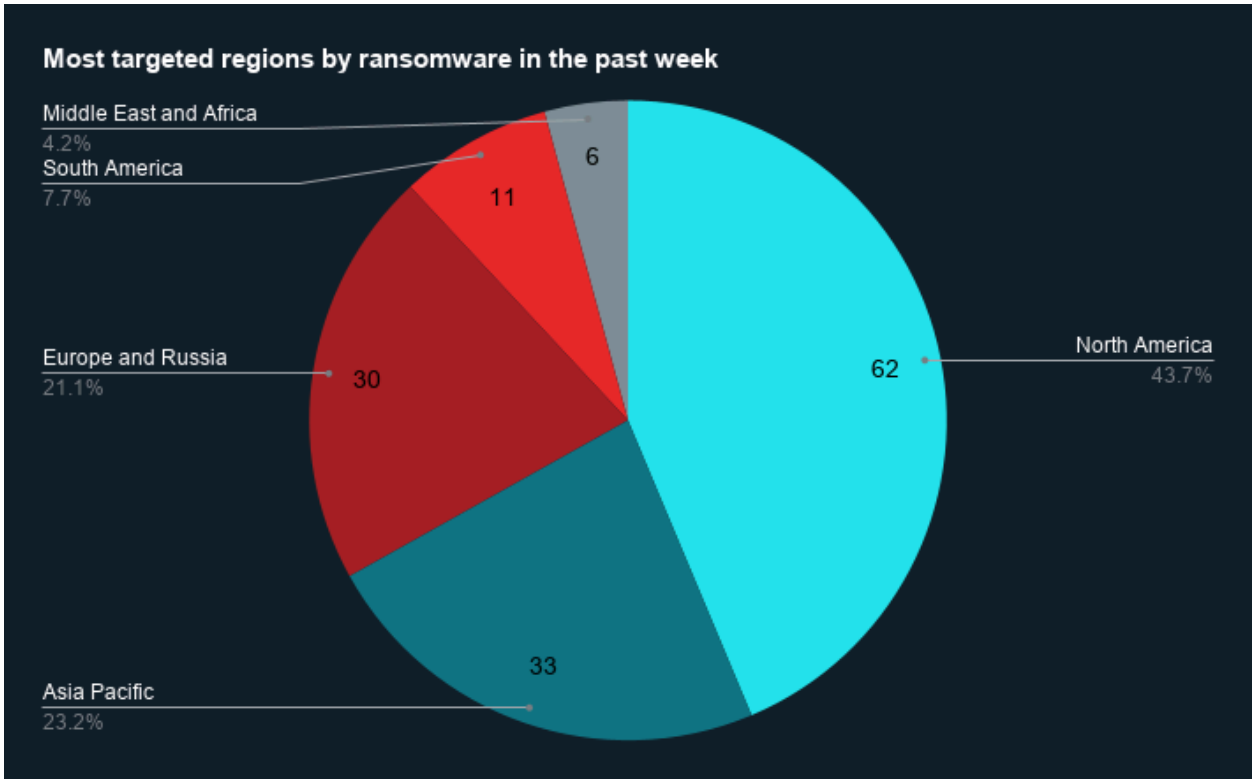


Source: ZeroFox Internal Collections

## Significant Data Breaches Reported over the Past Week

| Targeted Entity | iRhythm | Kodak | JeVeuxAider.gouv[.]fr |
|---|---|---|---|
| Compromised Entities/Victims | Patients | Kodak customers, employees, and business partners | Official French government portal JeVeuxAider.gouv[.]fr users |
| Compromised Data Fields | Proprietary data, patient protected health information, and other personal information | Over 2.2 million records containing PII, financial and identification data for some individuals, and internal corporate documents | Names, emails, phone numbers, dates of birth, and physical addresses |
| Suspected Threat Actor | Unknown | ShinyHunters | misere |
| Country/Region | United States | United States | France |
| Industry | Healthcare | Technology | Government |
| Possible Repercussions | Medical identity fraud, phishing, and healthcare-focused social engineering | Credential abuse, phishing, identity theft, and business email compromise (BEC) | Targeted phishing, smishing, identity fraud, and impersonation of French government services |

**Three major breaches observed in the past week**

# Physical and Geopolitical Intelligence Key Findings

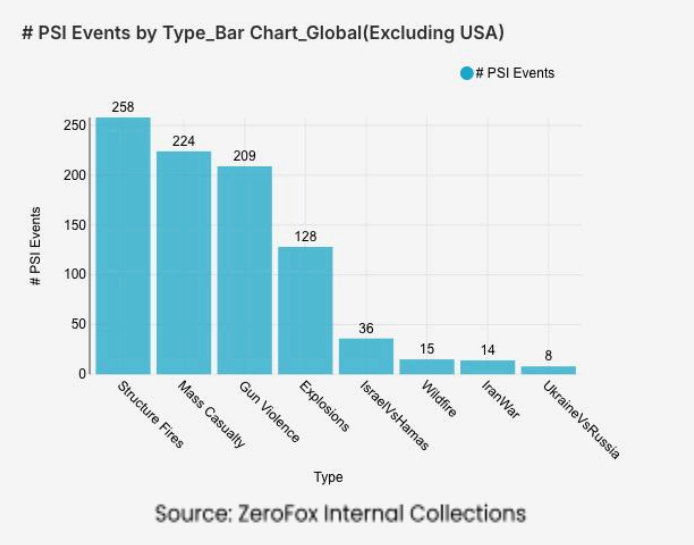


## Physical Security Intelligence: Global

**What happened:** Excluding the United States, there was a 28 percent decrease in mass casualty events this week from the previous week, with the top contributing countries or territories being Mexico, India, and France, in that order. Approximately 57 percent of these events were explosions, and the three aforementioned countries and territories accounted for about 32 percent of all mass casualty alerts. General alerts related to the Israel-Hamas conflict decreased by 35 percent from the previous week, and alerts related to the war in Iran decreased by 71 percent. Events related to Russia's war in Ukraine decreased by 33 percent. The top three most-alerted subtypes were structure fires, which saw a 5 percent increase from the previous week; gun violence, which did not increase or decrease from the previous week; and explosions, which decreased by 38 percent. Notably, wildfires increased by 150 percent.

- **What this means:** Global mass casualty alerts declined meaningfully this week, shaped by diplomatic breakthroughs and shifting conflict dynamics. The sharp decline in Iran-related alerts is attributable to the June 15 U.S.-Iran memorandum of understanding, which extended the ceasefire for 60 days and announced the reopening of the Strait of Hormuz, a major de-escalation after months of direct conflict. The goal in upcoming talks will be a permanent end to the war, although the fate of Iran's nuclear program remains unresolved for now. Mexico remains the top contributing country to mass casualty alerts, driven by waves of coordinated cartel violence. France is currently at a Level 2 terrorism advisory amid a pattern of knife and explosive attacks. The 150 percent spike in wildfire alerts reflects the early arrival of an intense global fire season across multiple continents, with Greece and Turkey having the highest numbers this week. While active conflicts and fragile ceasefires reduce violence in some areas, environmental crises and structural instability drive a continued rise in global alerts.

## Physical Security Intelligence: United States

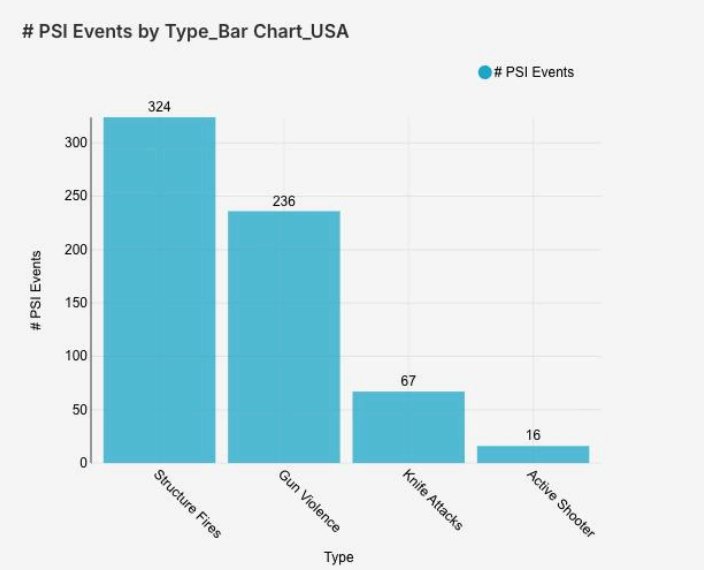

Source: ZeroFox Internal Collections

**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, structure fires are fires that affect man-made buildings, and knife attacks are confirmed stabbing or slashing instances. The top two states with the most gun violence alerts were Illinois and Ohio, which together made up 27 percent of this week's nationwide total. Gun violence across the United States overall decreased by 3 percent from the week prior. Knife attack alerts decreased by 7 percent, and the top contributing states were California and New York. Structure fire numbers did not increase or decrease from the previous week, and the top two states for this subtype were also California and New York. Notably, the number of active shooter incidents increased by 100 percent nationwide from the week prior.

- **What this means:** This past week, structure fires, gun violence, and knife attacks continued to dominate the threat landscape across the United States, reflecting persistent and widespread public safety challenges. Gun violence remained especially acute in Illinois, underscored by the violence in Chicago, where at least 20 people were shot and seven killed across 15 incidents in a single weekend. There were eight total mass shootings across the country within the last seven days, with a significant increase of active shooter events. This was exemplified by the June 12 shooting in Midland, Texas, where a fugitive opened fire on passing cars, killing one person and injuring 10 others before barricading themself and dying during a standoff with police. Structure fire alerts held flat week-over-week, with California and New York again leading in volume; on June 18 in Los Angeles, a blaze erupted at a warehouse, prompting a shelter-in-place order for nearby residents. Knife attacks also remained a concern this week, including a mass stabbing on a charter bus in Bessemer, Alabama, on June 15 that left one person dead and four others injured. Taken together, this week's incidents underscore that violent crime and public safety threats remain a persistent and nationwide challenge.

# 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% |