ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

April 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 April 23, 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

ZeroFox Intelligence Assessment: Q1 2026 Ransomware Wrap-up

ZeroFox observed at least 2,059 separate ransomware and digital extortion (R&DE) incidents in Q1 2026, a decrease of approximately 1.5 percent from Q4 2025—which accounted for a record-breaking 2,091 incidents. March remained the most active month in Q1 in comparison to previous years, accounting for at least 747 incidents—which is roughly 36 percent of all global ransomware attacks in Q1 2026. Regional R&DE targeting patterns in Q1 2026 were largely consistent with those observed during previous months. North America-based organizations were the most targeted by a substantial margin, accounting for approximately 54 percent of all incidents (or at least 1,114 incidents). ZeroFox also observed that the five most active R&DE collectives in Q1 2026 were almost certainly Qilin, Akira, The Gentlemen, INC Ransom, and Cl0p. This is a change from Q4 2025, as only Qilin, Akira, and Cl0p were in the top five that quarter.

ZeroFox Intelligence Flash Report - SITREP #35 - SoH Blockade - April 23, 2026

On April 21, 2026, the United States unilaterally extended its ceasefire with Iran indefinitely while maintaining its blockade of the Strait of Hormuz (SoH). This followed several days of chaos in the SoH perpetrated by both sides, which was likely an effort to maximize leverage ahead of further talks. The United States and Iran will likely resume talks, but escalatory risks remain—especially as both sides maintain dual blockades of the SoH. Measured military responses to these setbacks signal that both the United States and Iran are likely reluctant to return to armed conflict and instead prefer to utilize economic coercion to increase ceasefire pressure. Vessel seizures by the U.S. Navy are likely being used as coercive measures to get Iran to return to talks. Since the terms of the ceasefire require the SoH to remain closed, all of the negative economic consequences seen during the all-out war will very likely worsen.

ZeroFox Intelligence Brief - Underground Economist: Volume 6, Issue 9

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

ZEROFOX®

Weekly Intelligence Brief

Cyber and Dark Web Intelligence

Cyber and Dark Web Intelligence Key Findings

Group Claims Unauthorized Access to Anthropic's Claude Mythos

What we know:

- Anthropic is reportedly investigating claims of unauthorized access to its Claude Mythos model via a third-party vendor.
- On a private Discord channel, a group claimed it accessed the restricted cybersecurity artificial intelligence (AI) tool, reportedly presenting screenshots and live demonstrations of the platform as evidence.

Background:

- The group says it is exploring new AI models and does not plan on misusing the Mythos tool. Anthropic reports no malicious activity.
- Claude Mythos has been released to select organizations for testing, as the company says the tool can be abused to supercharge cyberattacks by detecting unknown vulnerabilities.

Analyst note:

- The unauthorized access raises concerns about Anthropic's safeguards against misuse, but the incident likely stems from activist researchers or whistleblowers and not threat actor activity.
- Insider threats, including through contracting entities, are very likely to emerge as a key risk for upcoming high-capability AI tools.

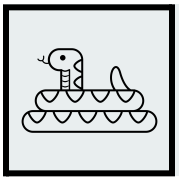

"Contagious Interview" Evolves into Self-Propagating Supply Chain Threat

What we know:

- The North Korea-linked "Contagious Interview" campaign has evolved into a worm-like supply chain attack, wherein compromised developer repositories propagate malware via malicious Visual Studio Code task configurations.
- Infected projects on platforms such as GitHub spread remote access trojans and payloads downstream, turning each new victim into a distributor across the software ecosystem.

Background:

- The campaign has evolved beyond single-target social engineering attacks to broadly compromise organizations through the developer ecosystem.
- The activity is attributed to North Korea-linked "Void Dokkaebi," which targets developers and steals high-value credentials.
- Void Dokkaebi infected over 750 code repositories, deploying more than 500 malicious Visual Studio Code task configurations.

Analyst note:

- In the near term, the campaign is likely to scale rapidly and diversify execution vectors beyond Visual Studio Code, leveraging stolen credentials and trusted code-signing mechanisms to make propagation stealthier, more persistent, and harder to detect.

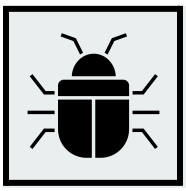

Four Malware Families Hit Finance Sector

What we know:

- Four new Android malware families—RecruitRat, SaferRat, Astrinox, and Massiv—are being deployed in separate campaigns to steal sensitive data from over 800 banking and cryptocurrency apps.
- These campaigns use overlays, keylogging, and real-time one-time password (OTP) interception to steal credentials and bypass security.

Background:

- The malware strains are spread via phishing sites, smishing texts, and fake apps (e.g., job portals, streaming services, and business tools) and trick users into installing malicious Android Application Packages (APKs).
- Among the four malware families, Massiv remains particularly evasive with an unclear infection chain.

Analyst note:

- Given that attackers can bypass traditional safeguards and operate within trusted user environments, account takeovers and fraudulent transactions are likely to be harder to detect.
- Additionally, multi-factor authentication (MFA)-dependent security protocols are likely insufficient to protect transactions and financial processes, as the actors can bypass security controls by intercepting OTPs in real time.

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 nine new vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalogue on April 20 and April 22, 2026. Additionally, on April 21, 2026, CISA released 12 Industrial Control Systems (ICS) advisories, which feature a total of 49 vulnerabilities, including CVE-2026-27668, CVE-2026-32955, and CVE-2020-24588. Progress Software released patches for multiple vulnerabilities, including CVE-2026-3517, CVE-2026-3518, CVE-2026-3519, and CVE-2026-4048, in MOVEit WAF and LoadMaster that could enable attackers to achieve remote code execution (RCE) and OS command injection via improper input sanitization. Microsoft has released an out-of-band (OOB) security patch for CVE-2026-40372, a critical privilege escalation vulnerability in the ASP.NET Core Data Protection cryptographic Application Programming Interface (API). The flaw enables attackers to gain SYSTEM privileges by forging authentication cookies. A now-patched vulnerability in Google's agentic integrated development environment (IDE) Antigravity has enabled attackers to bypass Strict Mode and achieve arbitrary code execution by injecting malicious flags into its file-search tool. This enables payload execution without user interaction. An exploit for an RCE flaw in a popular JavaScript implementation of Google's Protocol Buffers, protobuf[.]js, has been published. The flaw enables an attacker to supply a malicious schema leading to arbitrary code injection into the generated function.

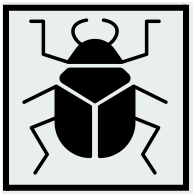

HIGH

CVE-2026-41651

What happened: This vulnerability in PackageKit enables local privilege escalation via a TOCTOU race condition, enabling unprivileged users to install arbitrary RPM packages as root without authentication. The flaw is patched in version 1.3.5.

- **What this means:** An unprivileged user can exploit a flaw in PackageKit to gain root access and run arbitrary code on the system. The flaw stems from improper transaction flag handling: overwrites during execution, silent state rejection, and late flag evaluation.
  - **Affected products:** PackageKit versions 1.0.2 and 1.3.4.

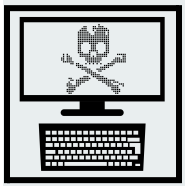

CRITICAL

CVE-2026-5760

What happened: This is an RCE flaw in SGLang, an open-source framework designed to accelerate large language models. The flaw enables threat actors to execute arbitrary code via malicious GPT-Generated Unified Format (GGUF) model files.

- **What this means:** Attackers can execute arbitrary code by supplying malicious GGUF model files during certain requests. Exploitation is likely to expose data such as credentials and environment variables and give attackers unauthorized access to internal systems.
  - **Affected products:** SGLang versions prior to 0.5.9

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, CoinbaseCartel, The Gentlemen, Akira, and SafePay were the most active ransomware groups. ZeroFox observed close to 143 ransomware victims disclosed, most of whom were located in North America. The Qilin ransomware group accounted for the largest number of attacks, followed by CoinbaseCartel.

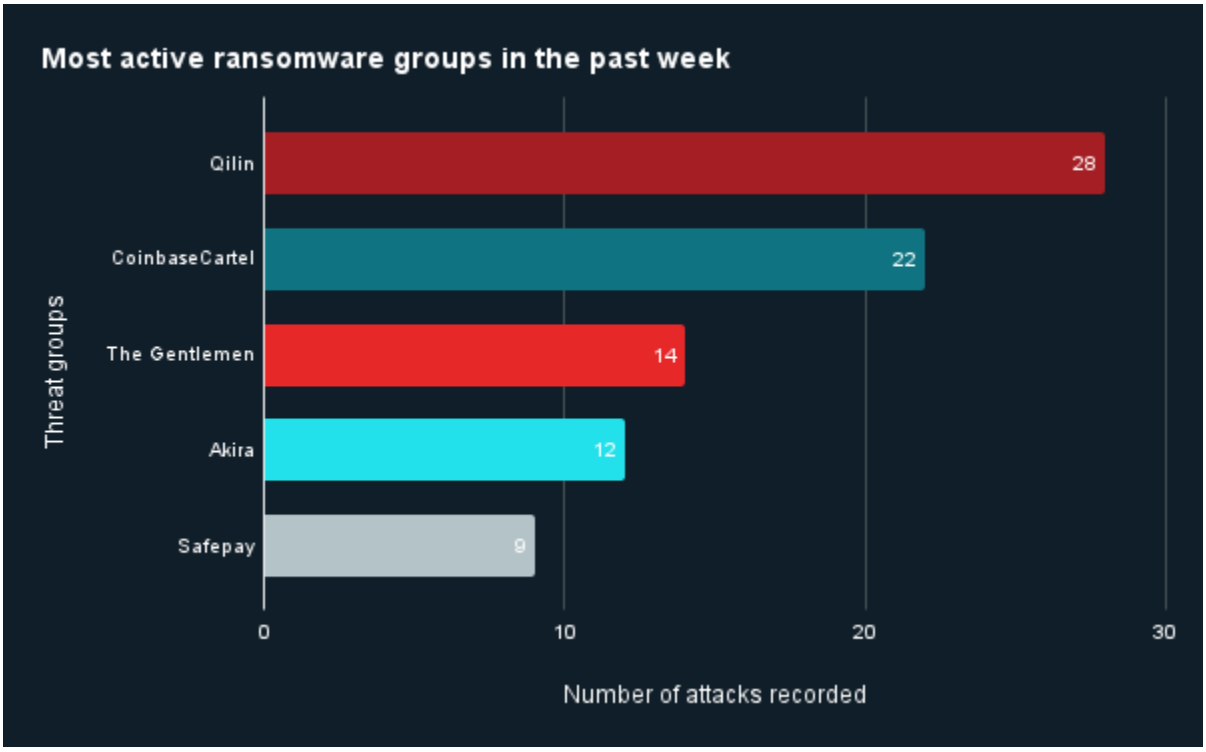

Source: ZeroFox Internal Collections

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

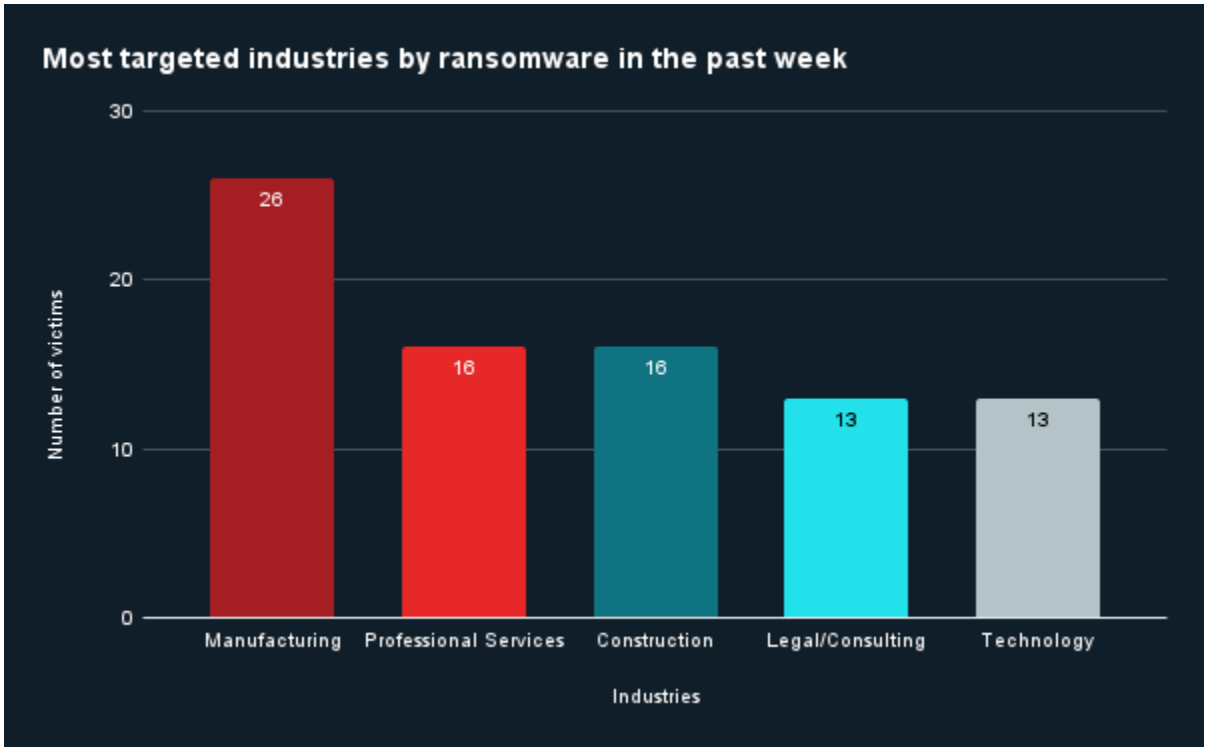

Source: ZeroFox Internal Collections

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 62 ransomware attacks observed in North America, while Europe and Russia accounted for 50, Asia-Pacific (APAC) for 13, South America for seven, and the Middle East and Africa for seven.

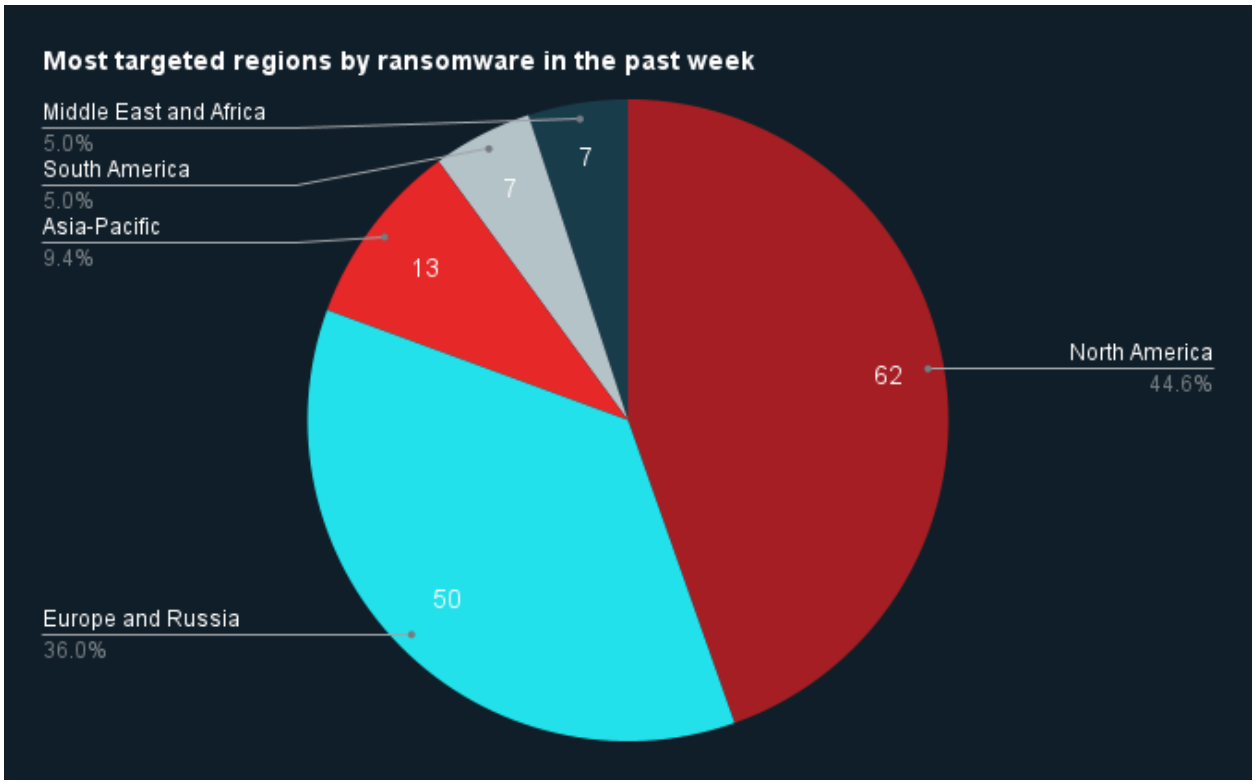

Source: ZeroFox Internal Collections

Significant Data Breaches Reported over the Past Week

Targeted Entity	Vercel	ANTS (National Agency for Secure Documents)	Seiko USA
Compromised Entities/Victims	Vercel's customers	About 19 million records of French citizens	Seiko USA's Shopify customers
Compromised Data Fields	API keys, source code, and databases, as well as approximately 580 employee accounts	Personally identifiable information (PII), place of birth, contact details, login ID, unique account identifier (government ID number)	Customer names, email addresses, phone numbers, purchase records, transaction details, addresses, account creation dates, and customer notes
Suspected Threat Actor	BreachForums user ShinyHunters (Threat actors behind the recent attacks attributed to the ShinyHunters extortion group have reportedly denied involvement in the breach.)	breach3d (aka ExtaseHunters)	Unknown
Country/Region	United States	France	United States
Industry	Professional services	Government	Consumer services
Possible Repercussions	Credential abuse, supply chain attacks, and poisoned versions of Next[.]js are likely to be published on GitHub and npm	Social engineering and phishing attacks against exposed individuals, identity theft, and financial fraud	Phishing and social engineering attacks, identity theft targeting exposed customers and the Seiko support team, and account takeover attempts

Three major breaches observed in the past week

Physical and Geopolitical Intelligence Key Findings

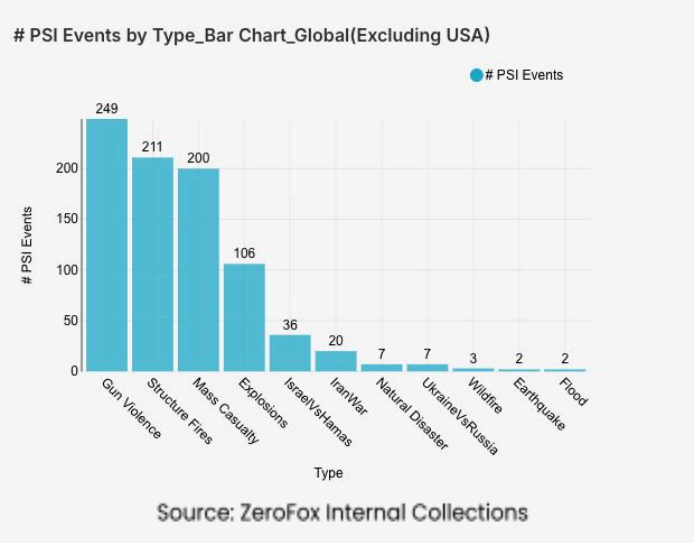

Source: ZeroFox Internal Collections

Physical Security Intelligence: Global

What happened: Excluding the United States, there was a 22 percent decrease in mass casualty events this week from the previous week, with the top contributing countries being India, Mexico, and Argentina, in that order. Approximately 53 percent of these events were explosions, and the three aforementioned countries accounted for about 25 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 75 percent. Events related to Russia's war in Ukraine decreased by 22 percent. The top three most-alerted subtypes were explosions, which saw a 31 percent decrease from the previous week; gun violence, which decreased by 5 percent; and structure fires, which increased by 14 percent. Notably, natural disaster alerts increased sevenfold compared to the previous week.

- **What this means:** The data from this week highlights a cooling of activity across several major war zones even as natural disasters have surged to the forefront of global humanitarian alerts. In the Middle East, the sharp plunge in Iran-related alerts and the drop in Israel-Hamas incidents likely reflect the stabilization of recent diplomatic efforts, such as the 10-day ceasefire between Israel and Lebanon and the temporary opening of the Strait of Hormuz. Conversely, while the war in Ukraine saw a decrease in general alerts as well, Russian drone strikes continued to cause localized mass casualties in areas such as Nikopol and Odessa. India and Argentina's incidents have transitioned from conflict-driven alerts to weather-related crises, as both countries endured severe flooding and wildfires this week. Additionally, the sevenfold surge in natural disasters was highlighted by a magnitude 7.4 earthquake off Japan's Iwate Prefecture on April 20, causing tsunami alerts and displacing thousands. A second magnitude 5.4 earthquake struck near Miyako on April 22, maintaining the nation's high alert. Ultimately, this week's data reflects a temporary cooling of major geopolitical conflicts and a surge in large-scale natural disasters.

Physical Security Intelligence: United States

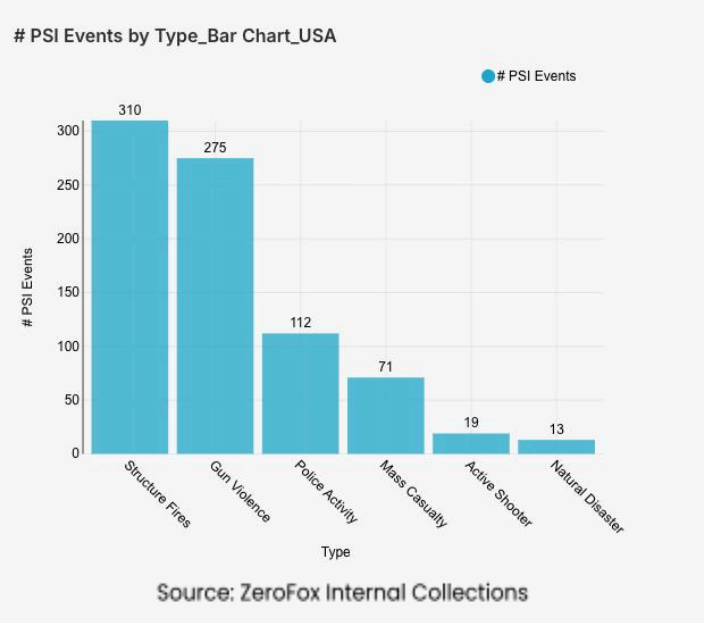

Source: ZeroFox Internal Collections

What happened: In the past week, the top three most-alerted incident subtypes were structure fires, gun violence, and police activity. Gun violence alerts are instances in which there is a confirmed or likely confirmed shooting victim, police activity involves law enforcement presence for generalized threats or for unknown reasons, and structure fires are fires that affect man-made buildings. The top two states with the most gun violence alerts were Illinois and New York, which together made up 17 percent of this week's nationwide total. Gun violence across the United States overall decreased by 16 percent from the week prior. Police activity alerts decreased by 6 percent, and the top contributing states were California and Texas. Structure fires increased by 4 percent, and the top two states for this subtype were New York and California. Notably, active shooter alerts nationwide increased by 217 percent.

- **What this means:** Over the past seven days, the United States has experienced a complex landscape of public safety incidents, characterized by a sharp rise in mass casualty events—specifically, active shooter incidents. On April 20, a planned altercation at Leinbach Park in Winston-Salem, NC, escalated into a shootout that left two juveniles dead and five others wounded. Additionally, there was a domestic incident on April 19 in Shreveport, LA, in which a gunman fatally shot eight related children before being killed following a high-speed pursuit. Overall, there were a total of 14 mass shootings across the country within the last week. The data also shows a slight rise in structure fires, one of which occurred on April 21 in the Bronx, NY, where a five-alarm fire at a five-story mixed-use building killed two residents and injured 11 others, including five firefighters. Natural crises have also impacted infrastructure; a severe weather outbreak on April 17 triggered over 20 tornadoes across the Midwest, destroying up to 75 homes in Ringle, Wisconsin, and displacing residents to local schools. Overall, domestic physical security remains strained by a surge in high-casualty active shooter events and severe structural emergencies.

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%