ZEROFOX®

Weekly Intelligence Brief

Classification: TLP:GREEN

May 16, 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 May 14, 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 Flash Report - ShinyHunters' Campaign Against the Education Sector

ShinyHunters is very likely in the midst of an ongoing campaign of escalatory attacks. This campaign almost certainly includes intentional targeting of the education sector—most recently Canvas and Houghton Mifflin Harcourt. The group's targeting of the education sector is almost certainly due to the large amount of user, employee, and customer data housed by educational institutions and within learning management systems. Data retrieved from the attack on Canvas is very likely to be used for further attacks against companies and institutions that use the learning management system for corporate and online training. ShinyHunters is very likely employing escalatory tactics: using data stolen in one breach to attack the next organization in a ladder of escalation. Further attacks exploiting lax access token protocols—and empowered by sophisticated phishing attacks—will almost certainly occur in the coming weeks and months.

ZEROFOX

Weekly Intelligence Brief

Cyber and Dark Web Intelligence

Cyber and Dark Web Intelligence Key Findings

RubyGems Repository Abused for Covert Storage of Scraped Government Portal Data

What we know:

- A campaign dubbed "GemStuffer" has reportedly scraped data from UK local government portals and covertly exfiltrated it through more than 150 malicious gems uploaded to the RubyGems repository.

Background:

- The campaign reportedly abuses the RubyGems ecosystem as a storage and retrieval mechanism for scraped government-related data while testing large-scale package registry abuse techniques.
- The packages embedded the collected portal data into valid .gem archives and republished them using hardcoded Application Programming Interface (API) credentials, instead of distributing traditional malware strains.

Analyst note:

- The scraped data is likely to be archived, indexed, and reused for intelligence collection, profiling of government operations, or future targeting activity.
- Even if much of the information is publicly accessible, aggregating it at scale is likely to enable threat actors to map government structures, as well as identify personnel and operational patterns.

Threat Actors Used AI to Identify a Zero-Day Vulnerability

What we know:

- Threat actors reportedly used artificial intelligence (AI) to find a zero-day vulnerability in what is described as the world's first such instance.
- However, researchers were able to quietly identify and patch the flaw before it was released into the wild.

Background:

- It was a two-factor authentication bypass flaw targeting a popular open-source web-based administration platform.
- The flaw's Python script reportedly showed signs of Large Language Model (LLM) influence, including educational docstrings, a fake Common Vulnerability Scoring System (CVSS) score, and polished textbook-style code.

Analyst note:

- Threat actors are very likely using AI to enhance the technical complexity, speed, and scale of cyberattacks, making capabilities such as zero-day discovery and exploitation, once limited to advanced actors, more accessible to lower-level threat actors.

Strategic Terrain and Infrastructure Data Stolen from Russian Targets

What we know:

- A threat group known as "HeartlessSoul" is reportedly targeting aerospace firms and drone operators via phishing and malvertising campaigns to distribute malware disguised as legitimate aviation software.
- HeartlessSoul is suspected to be focused on stealing this data from Russian government and enterprise systems.

Background:

- HeartlessSoul reportedly stole geospatial mapping and GPS-related data, including Geographic Information System (GIS) shapefiles, digital terrain and elevation data, and proprietary mapping files from systems used for geographic and infrastructure analysis.

Analyst note:

- HeartlessSoul is very likely building detailed operational awareness of Russian infrastructure, terrain, and aviation activity to improve future targeting and mission planning based on the type of data stolen.

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 two vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalog on May 8, 2026 and May 14, and 23 Industrial Control System (ICS) advisories on May 12 and May 14. Researcher Chaotic Eclipse has released proof-of-concept (PoC) exploits for two unpatched Windows vulnerabilities, dubbed YellowKey and GreenPlasma, which are reportedly a BitLocker bypass and a privilege-escalation flaw. Microsoft released patches for 137 vulnerabilities, with no zero-days disclosed or exploited this month. The majority of the vulnerabilities involved elevation of privilege flaws, with notable critical remote code execution (RCE) flaws. Fortinet released security updates to address two critical vulnerabilities in FortiSandbox and FortiAuthenticator, tracked as CVE-2026-44277 and CVE-2026-26083, that could enable attackers to run commands or arbitrary code on unpatched systems. SAP also released its May 2026 security updates, patching 15 vulnerabilities across multiple products, including two critical flaws in Commerce Cloud and S/4HANA, tracked as CVE-2026-34263 and CVE-2026-34260. Adobe patched 52 vulnerabilities, including two critical flaws tracked as CVE-2026-34659 and CVE-2026-34660, which can lead to arbitrary code execution and privilege escalation, respectively. Chipmakers Intel and AMD addressed 70 vulnerabilities combined in their May 2026 Patch Tuesday. A critical use-after-free (UAF) vulnerability in Exim enables an unauthenticated remote attacker to execute arbitrary code without credentials.

HIGH

CVE-2026-6973

What happened: This is an improper input validation flaw in Ivanti Endpoint Manager Mobile (EPMM), reportedly exploited as a zero-day against Ivanti's customers at the time of disclosure.

- **What this means:** The flaw requires authenticated administrative access to execute arbitrary code remotely. The flaw can reportedly be chained with previously patched vulnerabilities (CVE-2026-1281 and CVE-2026-1340) for exploitation.
  - **Affected products:** EPMM versions 12.8.0.0, 12.7.0.0, 12.6.1.0, and earlier

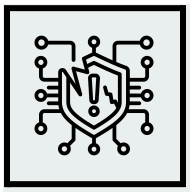

CRITICAL

CVE-2026-42208

What happened: This is a critical Structured Query Language (SQL) injection vulnerability in BerriAI's LiteLLM. The flaw resides in the proxy's API key verification process, where unsanitized user input is passed directly into database queries.

- **What this means:** Successful exploitation enables an unauthenticated attacker to manipulate those queries through a crafted request to any AI API endpoint.
  - **Affected products:** BerriAI's LiteLLM versions 1.81.16 to 1.83.6

Ransomware and Breach Intelligence

Ransomware and Breach Intelligence Key Findings

Ransomware Groups, Industry and Regional Trends

Last week in ransomware: In the past week, The Gentlemen, Qilin, Akira, INC Ransom, and SafePay were the most active ransomware groups. ZeroFox observed close to 150 ransomware victims, most of whom were located in North America. The Gentlemen ransomware group accounted for the largest number of attacks, followed by Qilin.

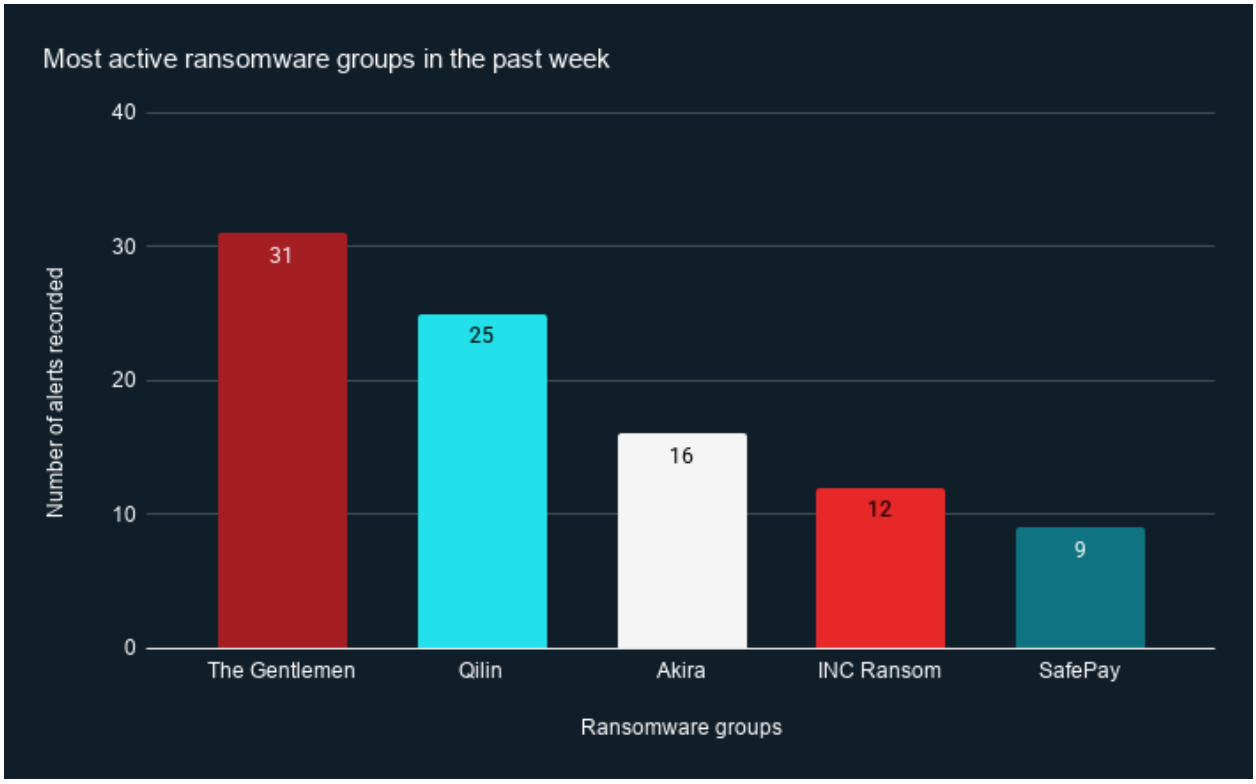

Source: ZeroFox Internal Collections

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

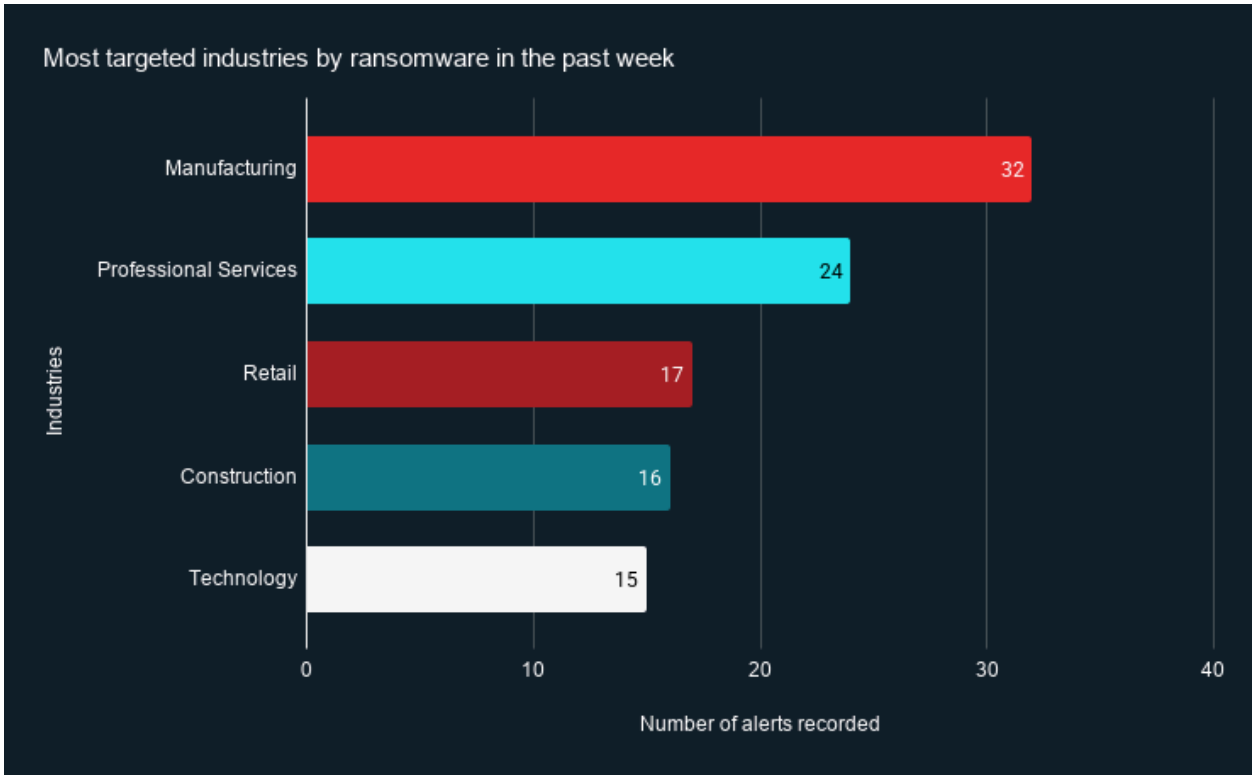

Source: ZeroFox Internal Collections

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 73 ransomware attacks observed in North America, while Europe and Russia accounted for 45, Asia-Pacific (APAC) for 17, Middle East and Africa for seven, and South America for five.

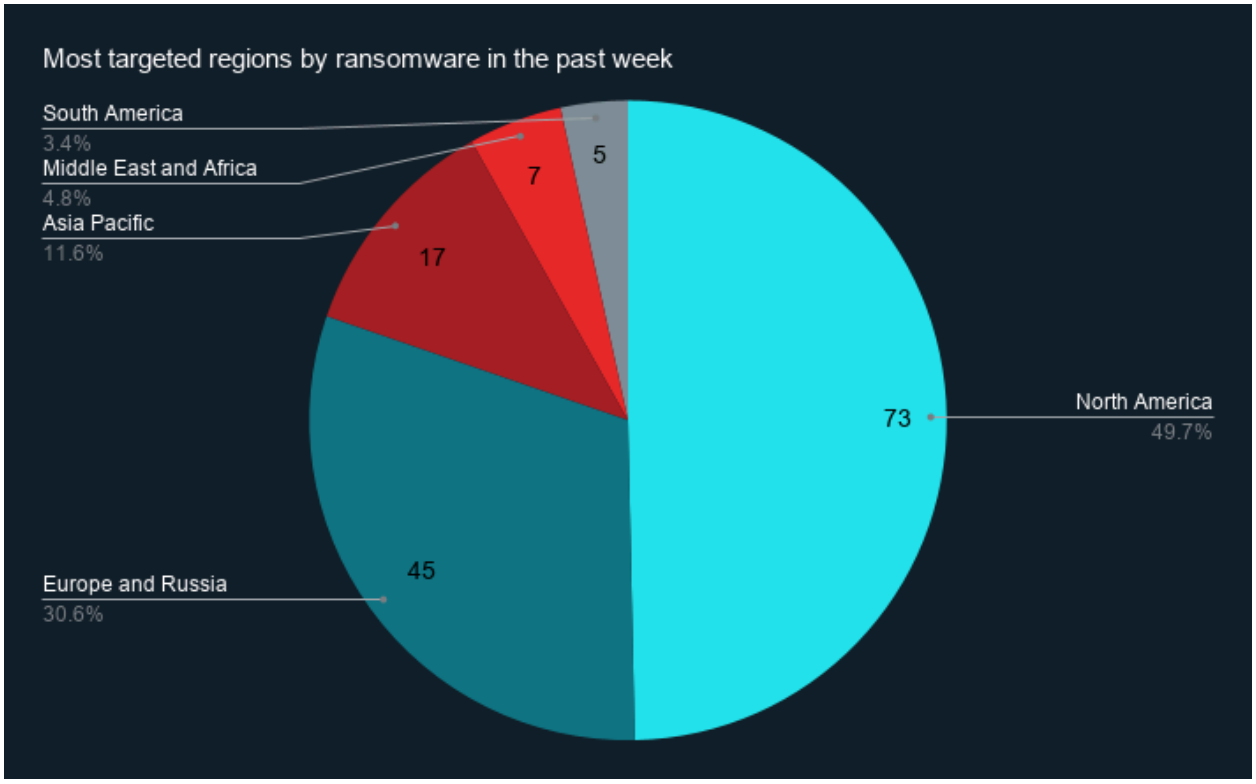

Source: ZeroFox Internal Collections

Significant Data Breaches Reported over the Past Week

Targeted Entity	OpenLoop Health	BWH Hotels	Škoda
Compromised Entities/Victims	716,000 individuals	BWH guests	Customers of Škoda Auto importer in Germany
Compromised Data Fields	Personally identifiable information (PII) and medical data	PII, including home addresses and reservation details	PII, order information, and login credentials
Suspected Threat Actor	Stuckin2019	N/A	N/A
Country/Region	United States	United States	Germany
Industry	Healthcare	Hospitality	Transportation
Possible Repercussions	Phishing, identity theft, medical insurance scams, and credential stuffing	Social engineering attacks, physical security risks, reputational damage, and payment scams	Phishing, account takeover attempts, identity theft, and invoice scams

Three major breaches observed in the past week

Physical and Geopolitical Intelligence Key Findings

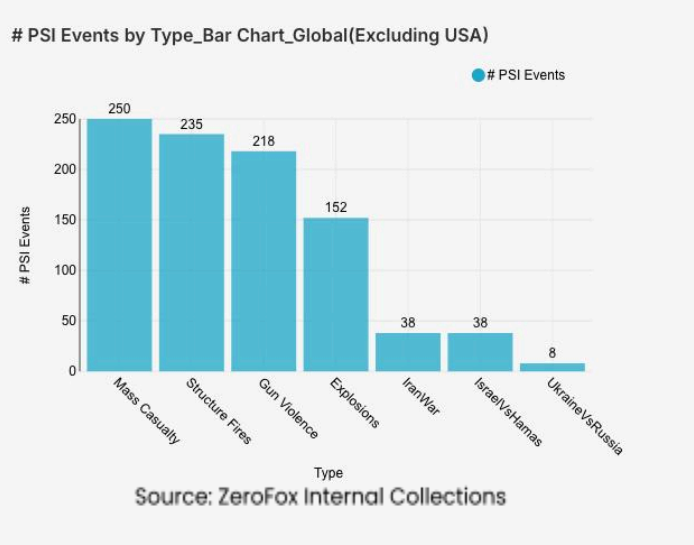

Physical Security Intelligence: Global

What happened: Excluding the United States, there was an 8 percent increase in mass casualty events this week from the previous week, with the top contributing countries or territories being Mexico, Lebanon, and India, in that order. Approximately 61 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 31 percent from the previous week, and alerts related to the war in Iran increased by 52 percent. Events related to Russia's war in Ukraine decreased by 38 percent. The top three most-alerted subtypes were structure fires, which saw a 2 percent increase from the previous week; gun violence, which decreased by 8 percent; and explosions, which increased by 6 percent.

- **What this means:** Global mass casualty activity rose this week, driven by persistent instability across several high-threat regions. In Mexico, industrial incidents contributed significantly to the count; a powerful explosion and fire struck the Pemex refinery in Salina Cruz, Oaxaca, on May 11, injuring at least six workers. In Lebanon, instability around the United Nations Interim Force in Lebanon (UNIFIL) continues to generate explosive incidents as well; on May 12, a drone explosion near the UNIFIL headquarters in Naqoura caused damage to a facility used by Malaysian military personnel. Alerts tied to the Israel-Hamas conflict increased this week, as a Palestinian man was killed and several others injured when an Israeli drone struck a motorcycle near the Jabalia refugee camp. Additionally, alerts for the Iran war had a significant increase, consistent with the ceasefire's fragility, as U.S. President Donald Trump warned on May 11 that the ceasefire with Tehran is on "life support." Russia-Ukraine alerts fell overall due to a three-day ceasefire; however, after the end of the ceasefire, Russia launched over 200 drones into Ukraine, damaging energy infrastructure, apartment buildings, and a kindergarten. Collectively, these trends reflect a persistently volatile global security environment, in which armed conflict and civil unrest continue to drive elevated mass casualty risk across multiple regions simultaneously.

Physical Security Intelligence: United States

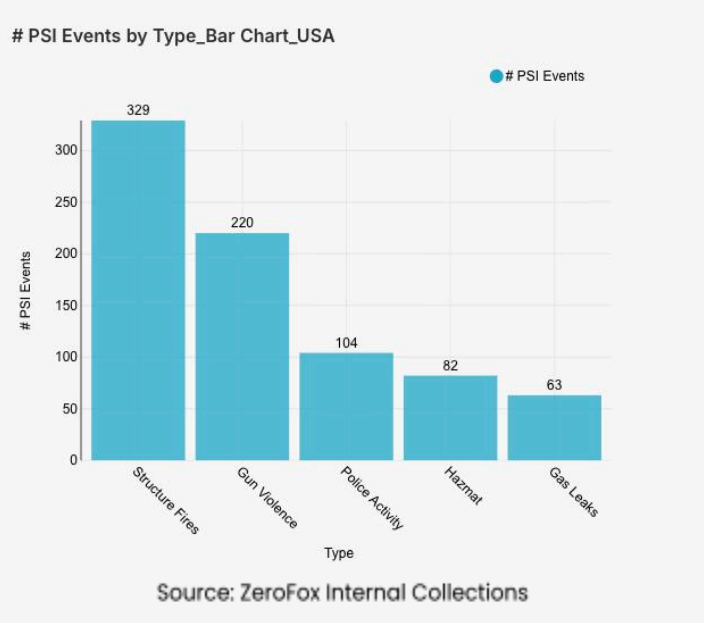

Source: ZeroFox Internal Collections

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 Ohio and Illinois, which together made up 17 percent of this week's nationwide total. Gun violence across the United States overall decreased by 17 percent from the week prior. Police activity alerts also decreased by 17 percent, and the top contributing states were California and Florida. Structure fires increased by 4 percent, and the top two states for this subtype were New York and California. Notably, hazmat incidents increased by 41 percent.

- **What this means:** Over the past week, structure fires, gun violence, and police activity were the most frequently alerted incident types across the United States. Gun violence remained prevalent despite an overall decrease, with eight mass shootings occurring within the last seven days; in a suspected drive-by shooting on May 10 in Paterson, New Jersey, a gunman opened fire on a group of people, killing two and injuring five others. Structure fires increased this week, with New York and California leading activity; on May 11, a fire broke out at a six-story apartment building in the Fordham section of the Bronx, New York City, New York, killing a one-year-old child and leaving two other young children in critical condition, with more than 80 fire and EMS personnel responding to the scene. The notable spike in hazmat alerts is exemplified by the man-made incident in North Carolina: on May 13, electric workers struck a large natural gas line in Matthews while drilling for a utility pole, triggering a major rupture and fire that required roughly 100 emergency personnel, damaging homes and displacing nearby residents. Overall, this week's incident data reflects an active and varied domestic threat environment, with simultaneous pressures across crime, fire, and hazardous materials underscoring the sustained demand on emergency response resources nationwide.

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%