TRINITY CYBER

Key Vulnerability Trends in 2026

Authors: Deepak Bellani, Jeremy Brown

Executive Summary

The vulnerability landscape in 2026 is increasingly defined by speed, automation, and scale. Attackers are exploiting vulnerabilities faster than ever, leveraging disposable infrastructure, spreading botnets, and AI-assisted tooling to expand operations across the internet. As 2026 hits its halfway point, three vulnerability trends emerge from our research:

01 Most exploitation originates from bulletproof hosting providers or compromised residential infrastructure.

02 Botnets remain the primary driver of high-volume exploitation campaigns.

03 AI-generated proof-of-concepts (POCs) are wasting both defender AND attacker resources.

These trends are not independent. Together, they reflect an ecosystem where exploitation is becoming increasingly automated, infrastructure is becoming more disposable, and defenders face growing challenges separating legitimate threats from misinformation.

The findings in this analysis highlight how attackers are operating today, and provide practical lessons for defenders responsible for protecting internet-facing applications, devices, and networks.

Background

Nearly everything about the vulnerability industry has changed within the past few years. Unfortunate cutbacks to NIST and its CVE enrichment program[1], the rise of alternative vulnerability tracking systems in both private and public sectors[2], and the overwhelming increase in AI-assisted[3] and AI hallucinated vulnerabilities is a staggering sign of the times. On average, vulnerabilities found in 2026 are exploited within 10 hours after disclosure, down from 56 days back in 2024[4]. These changes can, at times, lead to chaos.

This report examines vulnerability exploitation activity observed and prevented by Trinity Cyber's Full Content Inspection™ (FCI) platform between January and June 2026 - targeting organizations across multiple sectors, including healthcare, manufacturing, higher education, retail, defense, and critical infrastructure.

Because this dataset reflects activity prevented across Trinity Cyber's customer base, it should not be considered representative of all internet exploitation activity. Instead, it provides a focused view into the techniques, infrastructure, and operational behaviors most frequently used against these sectors during the first half of 2026.

All observations and statistics in this report are derived from exploitation attempts prevented by Trinity Cyber thus far in 2026.

Dataset Overview

OBSERVATION PERIOD:
January–June 2026

EXPLOITS PREVENTED:
2,396,919

DATA SOURCE:
Trinity Cyber Full Content Inspection™ (FCI) telemetry

INDUSTRIES TARGETED:
Healthcare, Manufacturing, Higher Education, Retail, Defense, Critical Infrastructure, and others

PURPOSE:
Identify recurring exploitation behaviors, infrastructure choices, and operational trends across large-scale internet exploitation campaigns.

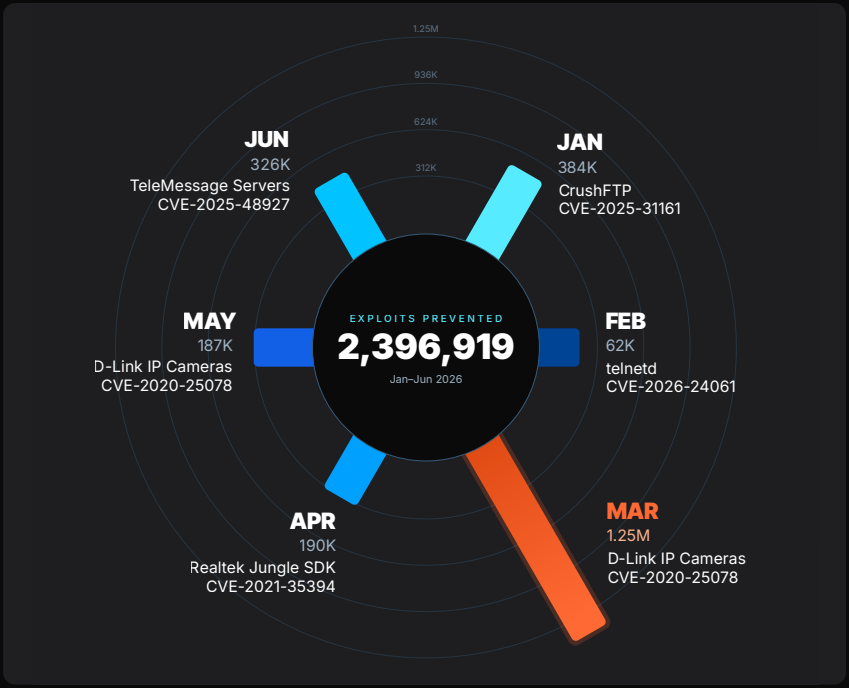

At a high level, the data illustrates how modern exploitation campaigns operate at internet scale. While attacker attention shifts rapidly toward newly disclosed vulnerabilities, older vulnerabilities with reliable exploit code continue to generate significant activity years after disclosure. This combination of opportunistic targeting and proven attack techniques creates a constantly evolving exploitation landscape.

The remainder of this report examines each trend and the representative exploit traffic behind some of the most prevalent campaigns prevented by Trinity Cyber during the first half of 2026.

TREND #1

Exploitation Infrastructure is Primarily Bulletproof Hosting and Residential Devices

Key Finding:

More than 80% of all vulnerability exploitation prevented by Trinity Cyber during the first half of 2026 originated from either bulletproof hosting providers (~50%) or compromised residential ISP devices (~30%).

Why It Matters:

Attackers increasingly rely on infrastructure that is inexpensive, difficult to attribute, and easy to replace. Rather than maintaining long-lived malicious servers, operators can rotate through residential proxy networks and permissive hosting providers at scale. This shift makes traditional IP-based blocking less effective. While attackers can quickly change infrastructure, they cannot easily change exploit behavior itself.

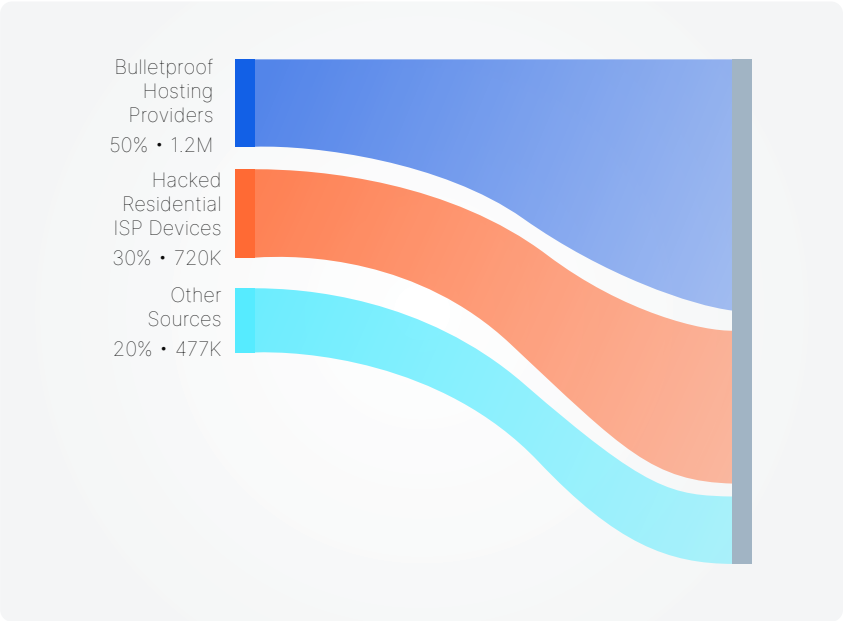

Behind the Trend:

Bulletproof hosting providers are notorious for enabling cyber-attacks. They provide infrastructure to cybercriminals, and very rarely engage in Know Your Customer (KYC) practices to thwart malicious activity. The threat is so severe that multiple governments have teamed up to provide defensive guidance against bulletproof hosting providers[5] for defenders who are overrun by the whack-a-mole infrastructure they provide.

Residential devices, usually those owned or leased by consumers to provide internet access at home (think ISP modems, routers, and WiFi gear) are also heavily abused by cyber actors. Bandwidth from these devices is unknowingly co-opted by hacking campaigns and residential proxy networks. Actors keep large amounts of them at the ready for non-attributable attacks across the internet. Once they're used up, actors pivot onto the next set of hacked gear.

There is also a considerable overlap with botnet activity in the residential ISP space, but not for the same reason: these proxies are usually fueled by IoT devices behind ISP gear, on home user's local networks[6]. The industry isn't being silent about it either — with security researchers and law enforcement agencies globally taking part in the identification and cleanup of hacked ISP gear[7]. Despite these efforts, attackers continue to thrive on both bulletproof hosting providers infrastructure and hacked residential devices.

Defenders must recognize bulletproof hosting providers and residential ISP infrastructure and stop exploitation coming from them at all costs. These days, anything that isn't legitimate business traffic from a residential ISP is probably exploitation. Inbound traffic from bulletproof hosting providers networks is rarely (if ever) legitimate.

TREND #2

Botnets Continue to Drive High-Volume Exploitation

Key Finding:

The largest exploitation spikes prevented by Trinity Cyber in the first half of 2026 were from botnet operations targeting vulnerable internet-facing devices. One campaign associated with the Mirai botnet generated more than 48% of exploit attempts leveraging CVE-2020-25078 against end-of-life D-Link devices.

Why It Matters:

Botnets continue to thrive because vulnerability exploitation remains a numbers game. Attackers automate discovery, exploitation, credential collection, and malware deployment at internet scale. As patching windows continue to shrink, defenders increasingly require preventative controls capable of stopping exploit traffic before vulnerable devices are reached.

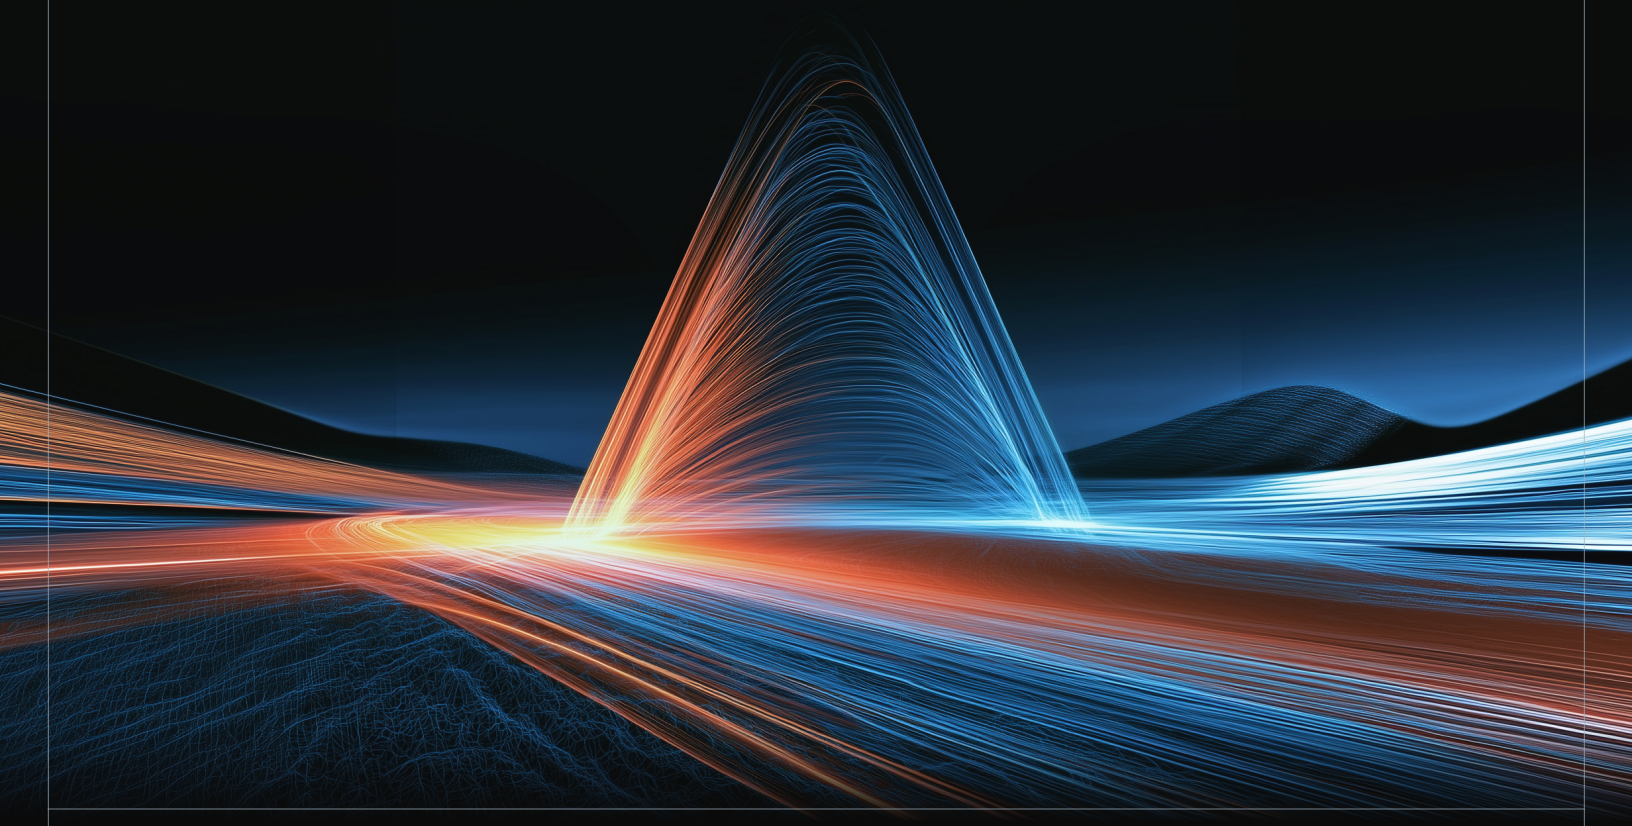

Behind the Trend:

Botnet ecosystems remain highly dynamic. Individual operations are routinely disrupted, rebranded, merged, or rebuilt, yet the overall volume of botnet-driven activity remains remarkably resilient. Many modern botnets now leverage residential proxy networks[8] and compromised consumer devices to distribute exploitation traffic across large numbers of seemingly legitimate IP addresses. As a result, the distinction between botnet infrastructure, bulletproof hosting providers, residential proxy networks, and compromised ISP devices is increasingly blurred.

This convergence is evident throughout Trinity Cyber's telemetry. The same infrastructure frequently served multiple purposes: hosting malicious payloads, proxying exploitation traffic, supporting botnet operations, and providing attacker-controlled relay points.

For defenders, attribution matters less than the operational reality that these systems collectively enable large-scale exploitation campaigns.

Botnets continue to succeed because exploitation remains a scale-driven activity. Automated scanning and exploitation frameworks continuously probe internet-facing services for known vulnerabilities, often targeting millions of hosts in search of a relatively small number of successful compromises. Information disclosure vulnerabilities are particularly attractive because they provide a low-cost path to collecting credentials, configuration data, and other sensitive information that can be monetized directly or leveraged in subsequent attacks.

Exploitation of information disclosure vulnerabilities offer three benefits to attackers and the botnets they operate:

01

Collected admin passwords can be sold to other cyber actors.

02

Botnets gain immediate admin level access to devices, using them for DDoS operations.

03

New attack infrastructure is gained and abused by attackers to launch further exploitation.

TREND #3

AI-Generated POCs Are Creating Detection Noise

Key Finding:

Trinity Cyber has observed a growing number of AI-assisted proof-of-concept (POC) repositories published shortly after vulnerability disclosure.Many of these POCs are incomplete, nonfunctional, or merely perform version detection while claiming successful exploitation.

During the first half of 2026, Trinity Cyber identified more than ten high-profile examples where publicly shared POCs failed to perform the exploitation they claimed to demonstrate.

Why It Matters:

Detection engineers frequently rely on newly released exploit code to develop signatures, analytics, and response playbooks. When those POCs are inaccurate, defenders risk wasting time building detections around activity that does not represent real exploitation.

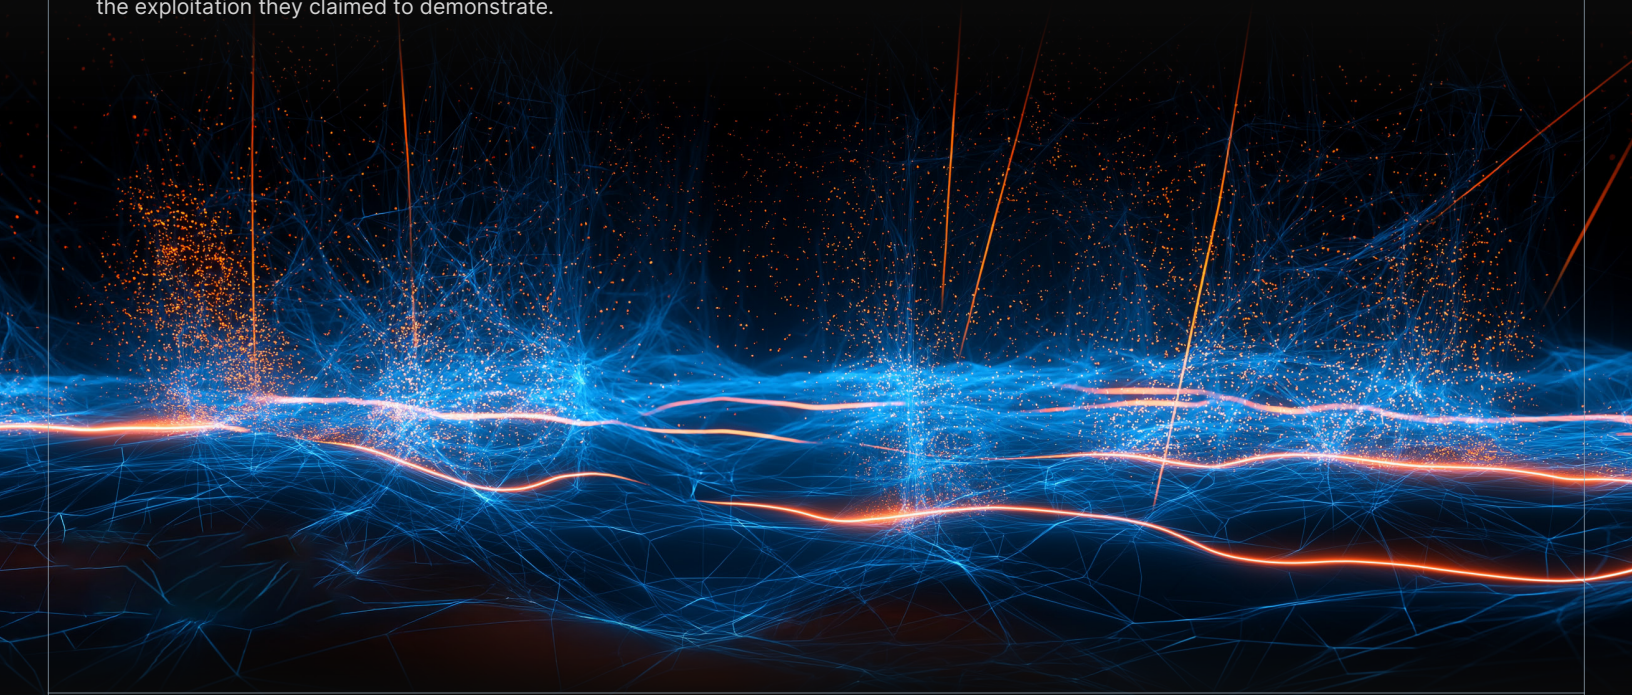

Behind the Trend:

Artificial Intelligence (AI) is a transformative technology with incredibly strong benefits to the cybersecurity industry. Vulnerability discovery is one of them. However, the technology is a double-edged sword. For every Project Glasswing[9] effort that benefits defenders, there are just as many attackers who use AI to develop exploits for vulnerabilities with limited or no public information.

Not all attackers are created equal, though. Trinity Cyber has observed a growing trend in rapid vulnerability exploitation where AI hallucinations are wasting both attacker and defender time. Essentially, low-skilled attackers are generating AI-assisted POCs that do not work. These POCs are often the only source of information that defenders (who are scrambling for detection guidance) can draw from. The resulting attacks do not work, and the resulting detections catch attacks that aren't real.

Our friends at GreyNoise touched on the problem back in 2025, too[10]. Fake POCs poison Detection Engineering efforts. This is happening at broader scale in 2026.

To frame this problem further, take the recent vulnerability in the WordPress UpdraftPlus Plugin (CVE-2026-10795) which was found and responsibly disclosed to Team Updraft in June 2026[11]. A few POCs came out on GitHub as early as June 11th that were clearly AI generated:

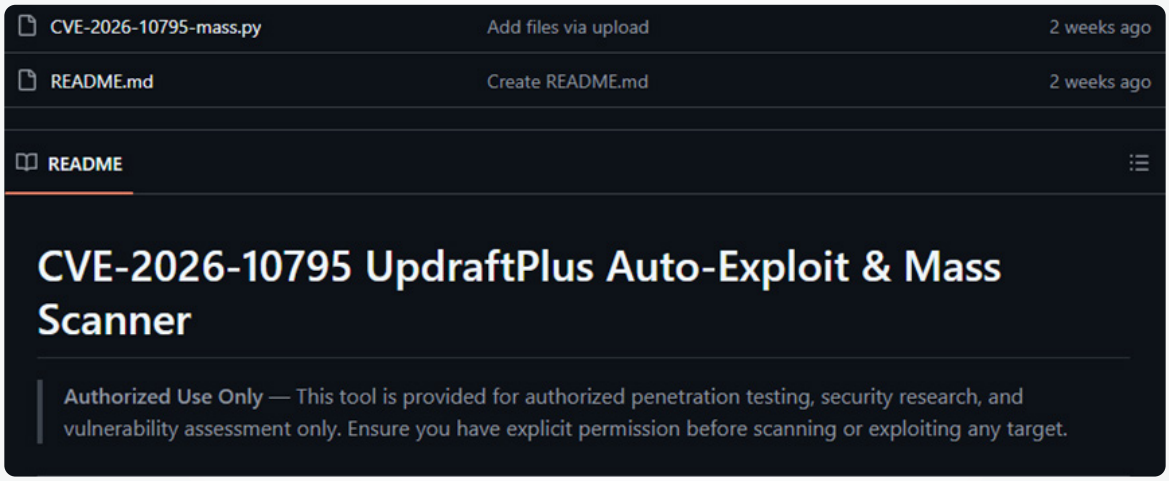

Figure 1. AI assisted POC for CVE-2026-10795 on GitHub found one day after disclosure.

The primary issue with many of these POCs is that they do not successfully exploit the vulnerability they claim to target. In several cases, repositories advertised as working exploits were merely version scanners, contained nonfunctional code, or redirected users to Telegram channels and other platforms promoting additional exploit sales. This pattern is becoming increasingly common following high-profile vulnerability disclosures.

The challenge for defenders is not simply that inaccurate POCs exist—it is the speed at which they are published and consumed. AI-assisted development has dramatically lowered the barrier to producing exploit code and accompanying research. As a result, repositories often appear within hours of disclosure, sometimes before researchers have fully validated the underlying vulnerability or exploitation methodology.

In practice, many of these POCs perform reconnaissance rather than exploitation. Determining whether a target is running a vulnerable version of software is a valuable step in an attack chain, but it is not evidence that exploitation has occurred or is even possible using the published code. Other repositories have been tied directly to actors attempting to monetize demand for exploit intelligence by advertising "exclusive" exploits in exchange for cryptocurrency[12].

For defenders, this creates a growing intelligence validation problem. Detection engineers, threat hunters, and incident responders frequently rely on public exploit code to develop detections and assess risk. When those POCs are inaccurate or AI-generated without sufficient validation, organizations can spend valuable time building detections around activity that does not represent real-world exploitation.

To address this challenge, our team maintains an internal repository of authors and sources associated with nonfunctional, misleading, or otherwise unreliable exploit releases. Defenders should similarly evaluate the credibility of exploit sources, validate claims through independent testing whenever possible, and treat newly released POCs with healthy skepticism until their functionality has been verified.

AI can be a valuable tool for identifying, classifying, and prioritizing exploit research. However, as with any AI-assisted workflow, human-in-the-loop validation is essential for separating credible exploit intelligence from distracting noise.

Trending Campaigns

The following vulnerabilities generated some of the highest exploitation volumes prevented in the first half of 2026:

CVE-2020-25078

D-Link IP Cameras

- 1,100,000+ exploit attempts prevented
- Most used exploit by Mirai botnet
- Dumps administrative credentials

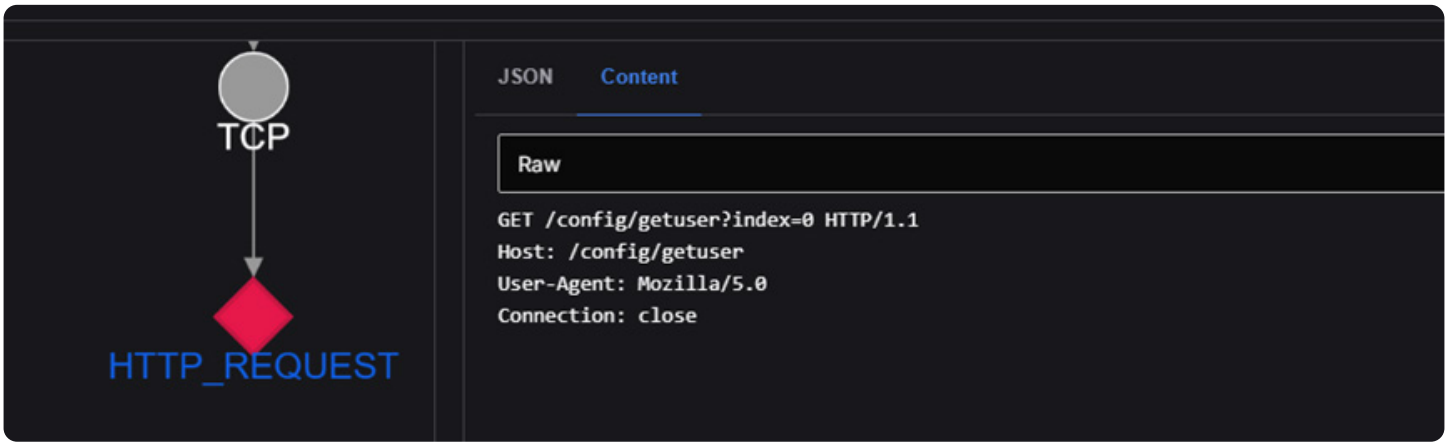

Figure 2. D-Link Information Disclosure Exploit (CVE-2020-25078)

CVE-2026-24061

GNU inetutils telnetd

- 73,000+ exploit attempts prevented
- Unauthenticated root access
- Evasive because Telnet protocol is frequently ignored

```
00000000  ff fa 27 00 00 55 53 45 52 01 2d 66 20 72 6f 6f  |ÿú'..USER.-f roo|
00000010  74 ff f0                                        |tÿð|
```

Figure 3. GNU inetutils telnetd Argument Injection Vulnerability (CVE-2026-24061)

CVE-2025-31161

CrushFTP

- 20,000 exploit attempts prevented
- Administrative takeover
- Rapidly targeted after disclosure

```
GET /WebInterface/function/?command=getUserList&serverGroup=MainUsers&c2f=1346 HTTP/1.1
Host: [redacted]:49176
User-Agent: Mozilla/5.0 (Kubuntu; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/133.0.0.0 Safari/537.36
Connection: close
Authorization: AWS4-HMAC-SHA256 Credential=crushadmin/
Cookie: CrushAuth=6702046443664_sEpAkJyvVrAzbWKulmATaaCugCFg1346; currentAuth=1346
Accept-Encoding: gzip
```

Figure 4. CrushFTP Authentication Bypass (CVE-2025-31161)

CVE-2021-35394

Realtek Jungle SDK

- 129,000+ exploit attempts prevented
- Supply-chain exposure across many IoT devices
- Heavily used by Mirai and Mozi botnets

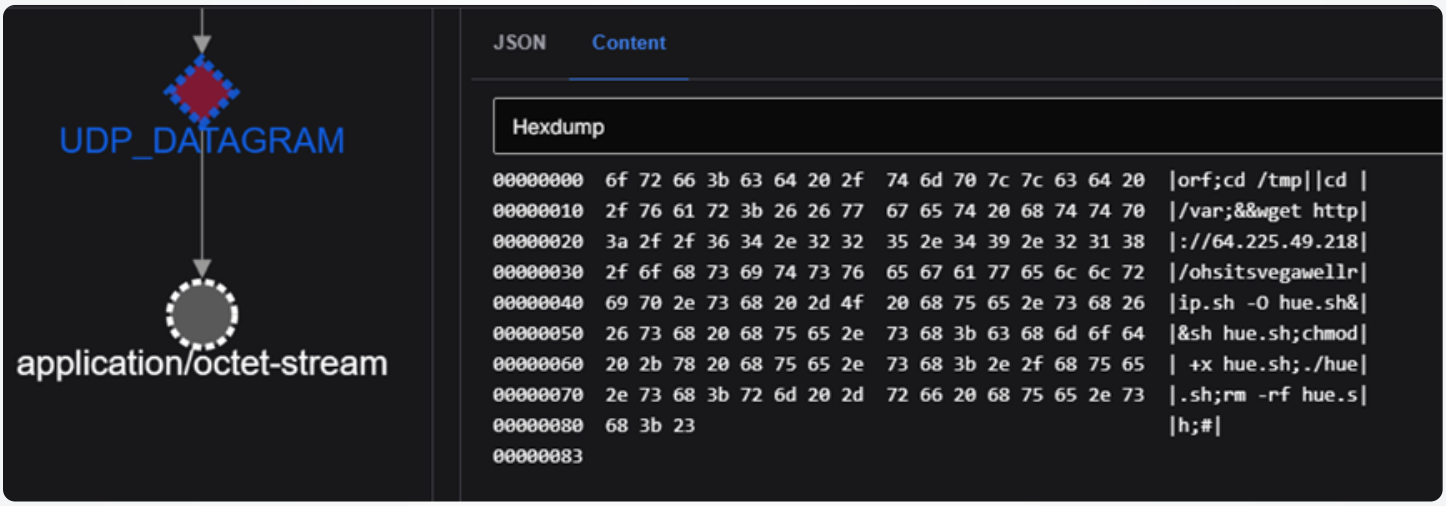

Figure 5. Realtek Jungle SDK Command Injection (CVE-2021-35394)

CVE-2025-48927

TeleMessage Infrastructure

- 47,000+ exploit attempts prevented
- Information disclosure involving credentials and sensitive communications

```
GET /actuator/heapdump HTTP/1.1
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/120.0.0.0 Safari/537.36
Cache-Control: no-cache
accept-encoding: gzip
Accept-Language: en-US,en;q=0.9
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
X-Forwarded-Proto: http
Host: [redacted]
connection: Keep-Alive
```

Figure 6. TeleMessage Spring Boot Actuator Information Disclosure (CVE-2025-48927)

Looking Ahead

Based on exploitation activity observed during the first half of 2026, Trinity Cyber expects several trends to continue accelerating throughout the year:

- Increased use of residential proxy infrastructure and compromised ISP devices.
- Continued dominance of botnet-driven exploitation campaigns.
- Faster weaponization of newly disclosed vulnerabilities.
- Growing volumes of AI-generated exploit content, including both legitimate researchand hallucinated POCs.
- Increased targeting of embedded software and IoT supply chains.

While attacker tooling continues to evolve, the underlying patterns remain consistent: automation, scale, and infrastructure designed to frustrate traditional detection approaches.

Summary

From Trinity Cyber's perspective, the first half of 2026 revealed three consistent trends across internet exploitation: attackers increasingly rely on disposable infrastructure, botnets continue to automate exploitation at scale, and AI is accelerating both legitimate vulnerability research and exploit-related noise.

These trends all point to the same challenge for defenders. Infrastructure changes. Payloads evolve. New vulnerabilities emerge daily. Attempting to keep pace through indicators, blocklists, and reactive detection alone becomes increasingly difficult as exploitation volume continues to grow.

The most effective defensive strategy is to focus on what attackers cannot easily change: the exploit traffic itself.

By inspecting and preventing exploitation before it reaches vulnerable systems, organizations can reduce risk regardless of where an attack originates, whether it comes from a botnet, a bulletproof hosting provider, or newly compromised infrastructure.

During the first six months of 2026, Trinity Cyber's Full Content Inspection™ (FCI) platform prevented more than 2.3 million exploitation attempts against customer infrastructure. As exploitation continues to scale, prevention-first approaches that stop attacks before compromise will become increasingly important for defending modern networks.

References

01 https://www.nist.gov/news-events/news/2026/04/nist-updates-nvd-operations-address-record-cve-growth

02 https://db.gcve.eu

03 https://daniel.haxx.se/blog/2026/01/26/the-end-of-the-curl-bug-bounty

04 https://thehackernews.com/2026/05/your-purple-team-isnt-purple-its-just.html

05 https://www.cisa.gov/resources-tools/resources/bulletproof-defense-mitigating-risks-bulletproof-hosting-providers

06 https://www.first.org/blog/20260424-Infrastructure-Nobody-Owns

07 https://www.justice.gov/opa/pr/justice-department-conducts-court-authorized-disruption-dns-hijacking-network-controlled

08 https://www.nokia.com/blog/one-year-later-the-residential-proxy-botnet-problem-got-bigger-not-smaller

09 https://www.anthropic.com/glasswing

10 https://www.labs.greynoise.io/grimoire/2025-07-30-ai-poc

11 https://teamupdraft.com/blog/important-security-update-for-updraftplus-and-updraftcentral-users

12 https://www.uptycs.com/blog/threat-research-report-team/fake-poc-repositories-malicious-code-github

TRINITY CYB3R™

Want to see how Full Content Inspection defeats attacks like this before they reach your users?

Schedule a live demo with Trinity Cyber.

trinitycyber.com