CYBER THREAT ANALYSIS

Recorded Future®

Insikt Group®

September 3, 2026

H1 2026 Malware and Vulnerability Trends

Threat actors increasingly blended malicious activity into legitimate tools, trusted services, and routine workflows, making behavioral detection and stronger identity, exposure, and third-party controls essential.

Insikt Group identified 215 actively exploited CVEs in H1 2026, up 34% year over year, with network-accessible, unauthenticated RCE flaws posing the greatest operational risk to exposed systems.

AI-enabled threats grew more visible in H1 2026, but most activity augmented existing tradecraft rather than replacing it; RATs, supply chain abuse, and NFC-enabled mobile fraud remained prominent.

Executive Summary

H1 2026 activity showed a continued adversary preference for abusing legitimate tools, trusted platforms, and routine workflows already present in enterprise and consumer environments. Threat actors used exposed software, developer tools, remote access utilities, payment workflows, and third-party services to gain access, steal credentials, move laterally, and monetize intrusions while blending into expected activity. This emphasis on evasion through normalcy rather than technical novelty increases the risk that malicious activity will progress through approved tools and trusted services before defenders recognize it, reinforcing the need for stronger exposure management, identity and credential governance, behavioral detection, developer-environment security, backup resilience, mobile fraud monitoring, and third-party oversight.

AI-enabled cyberattacks became more visible in H1 2026, but remained mostly additive to established intrusion tradecraft. In the vulnerability ecosystem, AI-assisted research increased the volume of vulnerability reports, which, moving forward, could further compress remediation timelines by accelerating exploit-path analysis and lowering exploit-development costs for skilled operators. In malware activity, observed AI-enabled capabilities largely aligned with lower-to-mid-level stages of Recorded Future's AI Malware Maturity Model (AIM3), where AI supported discrete functions such as persistence, user interface (UI) interaction, malware development, and delivery rather than fully autonomous operations.

Vulnerability exploitation broadened across enterprise operating systems, application frameworks, and network and security management products. Insikt Group identified 215 actively exploited common vulnerabilities and exposures (CVEs), and the most consequential cases combined network reachability, few access prerequisites, and code execution. Campaign reporting also showed that threat actors reused established post-exploitation playbooks across both newly disclosed and long-standing vulnerabilities, making exposure and impact more informative indicators of operational risk than vendor ranking or severity score alone. Across phishing and malware-delivery operations, threat actors repeatedly relied on familiar execution, obfuscation, discovery, and payload-transfer techniques rather than novel capabilities. Supply-chain compromises targeted package managers and developer environments, including AI-enabled tooling, where compromised credentials, trusted integrations, and software distribution channels enabled propagation into downstream cloud and software ecosystems. Other prominent activities included mobile malware enabling payment fraud through Near Field Communication (NFC) abuse and early AI-assisted workflows, as well as Magecart campaigns leveraging trusted third-party services and checkout manipulation.

The common risk across these threats is that malicious activity can progress through legitimate tools, trusted services, and routine workflows before defenders recognize it as part of a broader intrusion. Defenders should therefore prioritize vulnerabilities that can be exploited remotely or enable code execution, focus detection on suspicious sequences of behavior rather than isolated events, and strengthen controls protecting developer credentials, backup infrastructure, company-owned mobile devices, and payment environments.

Key Findings

- Insikt Group identified 215 actively exploited CVEs in H1 2026, up 34% from 161 in H1 2025. 142 of the 146 vulnerabilities that could be exploited without prior authentication were also network-accessible, and 60 of the 82 remote code execution (RCE) vulnerabilities combined network access with no authentication requirement.
- Similar to H1 2025, Recorded Future Malware Intelligence data showed the continued prominence of remote access trojans (RATs) throughout H1 2026. RATs led Insikt Group malware reporting, and Recorded Future Malware Intelligence submissions identified AsyncRAT as the top submitted malware family by total unique hashes and command-and-control (C2) configurations. AsyncRAT, Cobalt Strike, XWorm, Stealc, and REMCOS RAT remained in the top ten across H1 2025 and H1 2026.
- H1 2026 AI-enabled malware activity was concentrated in AIM3 Levels 1 to 3, with threat actors using AI to augment existing malware and intrusion workflows rather than conduct fully autonomous attacks.
- Ransomware operators continued to refine payload capabilities while relying on established intrusion methods, including ClickFix-style social engineering; exploitation of public-facing applications; abuse of legitimate tools such as s5cmd, PsExec, and AnyDesk; and techniques intended to reduce victims' recovery options.
- Android NFC malware was the most notable mobile malware trend in H1 2026, with families such as NFCShare and NGate abusing device NFC functionality to steal payment card data, relay contactless transactions, and facilitate ATM cash-outs.

Table of Contents

AI-Enabled Cyberattacks

AI-enabled cyber threats became more visible in H1 2026, but the available evidence indicates that most observed activity still augmented existing tradecraft rather than replacing it with fully autonomous operations.

In the vulnerability ecosystem, the release of Anthropic's Claude Mythos Preview reinforced the growing relevance of frontier models to offensive and defensive security research. Broader vulnerability reporting also increased following the model's release under Project Glasswing: June National Vulnerability Database (NVD) disclosures were 43% above the previous six-month average, while Mozilla reported that Mythos Preview identified 271 vulnerabilities that it fixed in Firefox 150, compared with 22 security flaws fixed following its earlier testing with Claude Opus 4.6. Vendors, AI developers, and bug-bounty platforms like Microsoft, Anthropic, and HackerOne have also reported rising vulnerability discovery or submission volumes alongside broader adoption of AI-assisted research, adding pressure to validation, disclosure, and remediation workflows.

Despite the increase in vulnerability reports, AI has not changed the fundamentals of vulnerability management: Attackers still need to identify, validate, weaponize, and operationalize vulnerabilities that offer reach, reliability, and return on investment. This means that only a small proportion of newly discovered vulnerabilities are likely to be a threat.

Nonetheless, AI-enabled vulnerability research can still increase defender workload in three ways: It can produce more credible vulnerability reports that require triage; reduce the time available to mitigate exploitable vulnerabilities by accelerating exploit-path analysis and weaponization; and lower the cost of exploit development by helping skilled operators produce proof-of-concept (PoC) code, test attack paths, and iterate toward weaponizable exploits more quickly. For defenders, the near-term issue is therefore not a sudden flood of fully autonomous exploitation, but a narrower window to determine which vulnerabilities matter most and remediate them before threat actors operationalize them. Additionally, early H2 2026 reporting on the July 2026 Hugging Face incident demonstrated that autonomous agents can perform discovery, validation, weaponization, and operationalization with limited human intervention. Defenders should prepare by automating vulnerability enrichment, prioritization, and mitigation to reduce the gap between machine-speed attack development and defensive response.

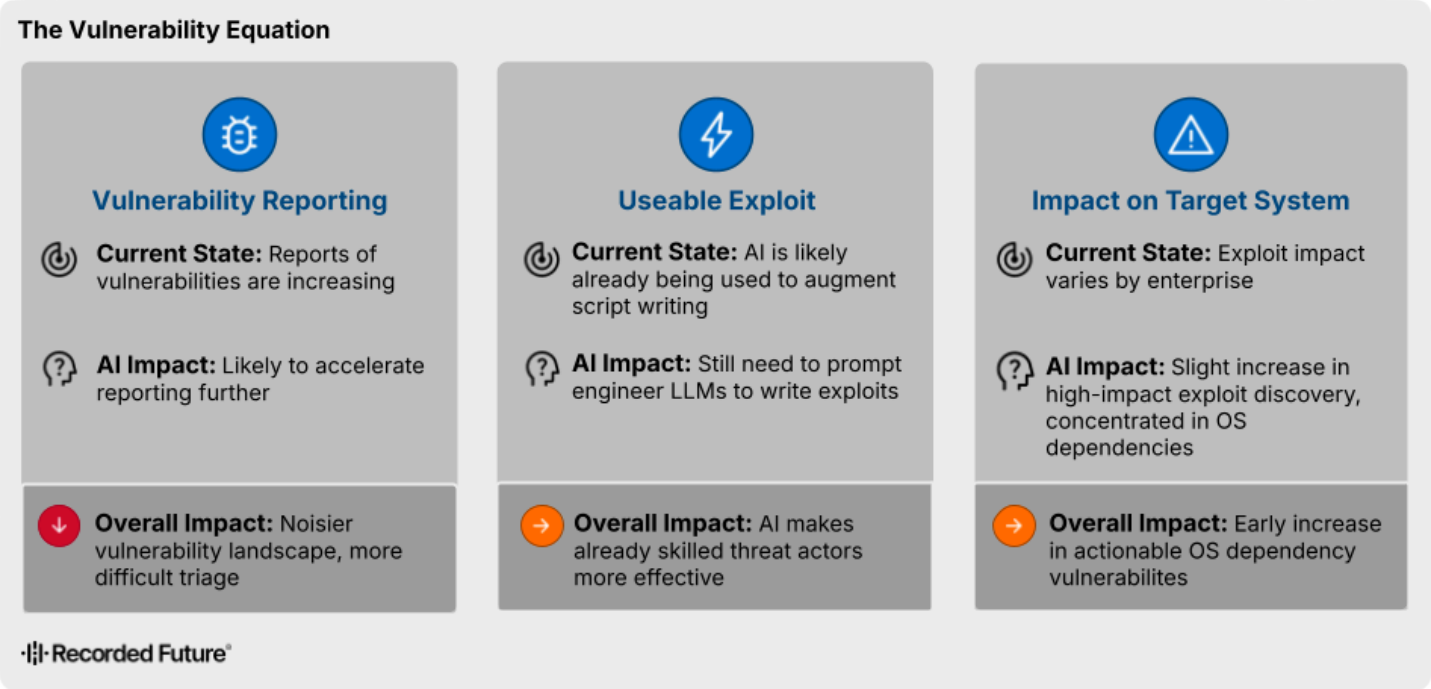

Figure 1: *How automated capabilities will likely impact reporting, exploit development, and impact (Source: Recorded Future)*

Malware reporting in H1 2026 showed more direct experimentation with AI-enabled capabilities compared to H1 2025. ESET identified PromptSpy as the first known Android malware to use generative AI (GenAI), using Google's Gemini to interpret on-screen UI elements and generate step-by-step instructions to improve persistence across device layouts. Reporting on CANFAIL also showed threat actors using LLM-generated decoy logic or AI-assisted development artifacts to complicate analysis and facilitate malware delivery against Ukrainian organizations. These examples suggest that threat actors are testing AI in ways that can solve narrow operational problems: adapting to user interfaces, generating code or decoy logic, improving obfuscation, and supporting analyst confusion.

Threat actors also used AI-related products and ecosystems as delivery mechanisms. In February 2026, VirusTotal reported malicious OpenClaw skills disguised as useful automation for a local AI agent ecosystem, while Malwarebytes reported fake OpenClaw installers hosted on GitHub and surfaced through search results to deliver infostealers and proxy malware. This activity shows that attackers are exploiting user interest in AI tools in the same way they have historically abused popular software brands, distributing trojanized installers, malicious extensions, fake repositories, and dependency-based payloads through otherwise familiar channels.

As of H1 2026, Insikt Group assesses that most observed AI-enabled malware activity aligns with the low-to-mid-level stages of Recorded Future's AI Malware Maturity Model (AIM3), specifically experimentation, adoption, and optimization, rather than fully autonomous AI-driven malware operations. Per the Recorded Future's AIM3 framework, most publicly observed "AI malware" remains concentrated in Levels 1 to 3, where AI supports discrete operational tasks such as UI interpretation, persistence guidance, transaction timing, or operator guidance. Therefore, the near-term risk is not

primarily self-directed malware operating independently, but rather the use of AI to make existing intrusion workflows faster, more convincing, more adaptive, and harder to analyze. Defenders should not expect a single control, model, or endpoint detection layer to reliably identify all AI-enabled malware activity. Instead, organizations should use defense-in-depth controls that can detect or disrupt different parts of the attack chain, including endpoint detection for script or installer activity tied to suspicious AI-themed downloads, extensions, repositories, or packages, mobile device management for accessibility service abuse and automated UI interaction, and restrictions on unapproved AI tools and installers.

Vulnerability Exploitation Trends

Key Takeaways

- Network reachability and low access requirements compounded risk, as 142 of the 215 exploited CVEs were network-accessible and could be exploited without prior authentication; 60 of those also enabled RCE, while public exploits were available for 66 (31%).
- Defenders faced both newly exploited flaws and persistent patch backlogs, as 162 CVEs were disclosed in 2025 or 2026, 53 predated 2025, and seventeen were dated from 2020 or earlier.
- Threat actors reused post-exploitation playbooks across different initial vulnerabilities; StrikeShark applied the same six-tool stack across thirteen CVEs, while Storm-1175 linked credential theft, remote execution, data transfer, and ransomware tooling across ten. More broadly, stealware was the most common malware category, followed by offensive security tools, backdoors, remote access trojans, and ransomware.

Microsoft Remains the Leading Vendor for Exploited Vulnerabilities

In H1 2026, Insikt Group identified 215 actively exploited vulnerabilities, up 34% from 161 in H1 2025. From January to June, the number of exploited vulnerabilities increased by an average of seven vulnerabilities per month. Microsoft remained the vendor most frequently associated with exploited vulnerabilities, accounting for 40 unique CVEs, up from 28 in H1 2025, a 43% year-over-year increase. Red Hat ranked second with fifteen CVEs, Cisco third with thirteen, Vercel fourth with eleven, and Fortinet fifth with nine. By comparison, H1 2025's top affected vendors after Microsoft were Apple with eight CVEs, Ivanti with seven, Linux with six, and SonicWall, Google, Fortinet, and Craft CMS with four each.

The exploited vulnerabilities affected products from 98 vendors, 67 of which were associated with only one CVE. This indicates that exploitation was not confined to the most frequently affected vendors. For defenders, this supports maintaining risk-based remediation across the full software inventory, including less common products that may receive less monitoring or slower patching than widely deployed enterprise platforms.

At the product-family level, Windows and Windows Server accounted for the largest H1 2026 concentration, with twenty unique CVEs, followed by Red Hat Enterprise Linux with thirteen, Vercel Next.js with eleven, Cisco Catalyst SD-WAN Manager with eight, and Apple iOS and iPadOS with seven. While Microsoft exploitation remained prominent, H1 2026 activity also increasingly involved enterprise Linux, application frameworks, and network management products. Several vendors' vulnerabilities were also driven by concentrated exploitation of one product family rather than uniform interest across the vendor's portfolio: all eleven Vercel CVEs centered on Next.js, eight of Cisco's thirteen affected Catalyst SD-WAN, and thirteen of Red Hat's fifteen affected Red Hat Enterprise Linux.

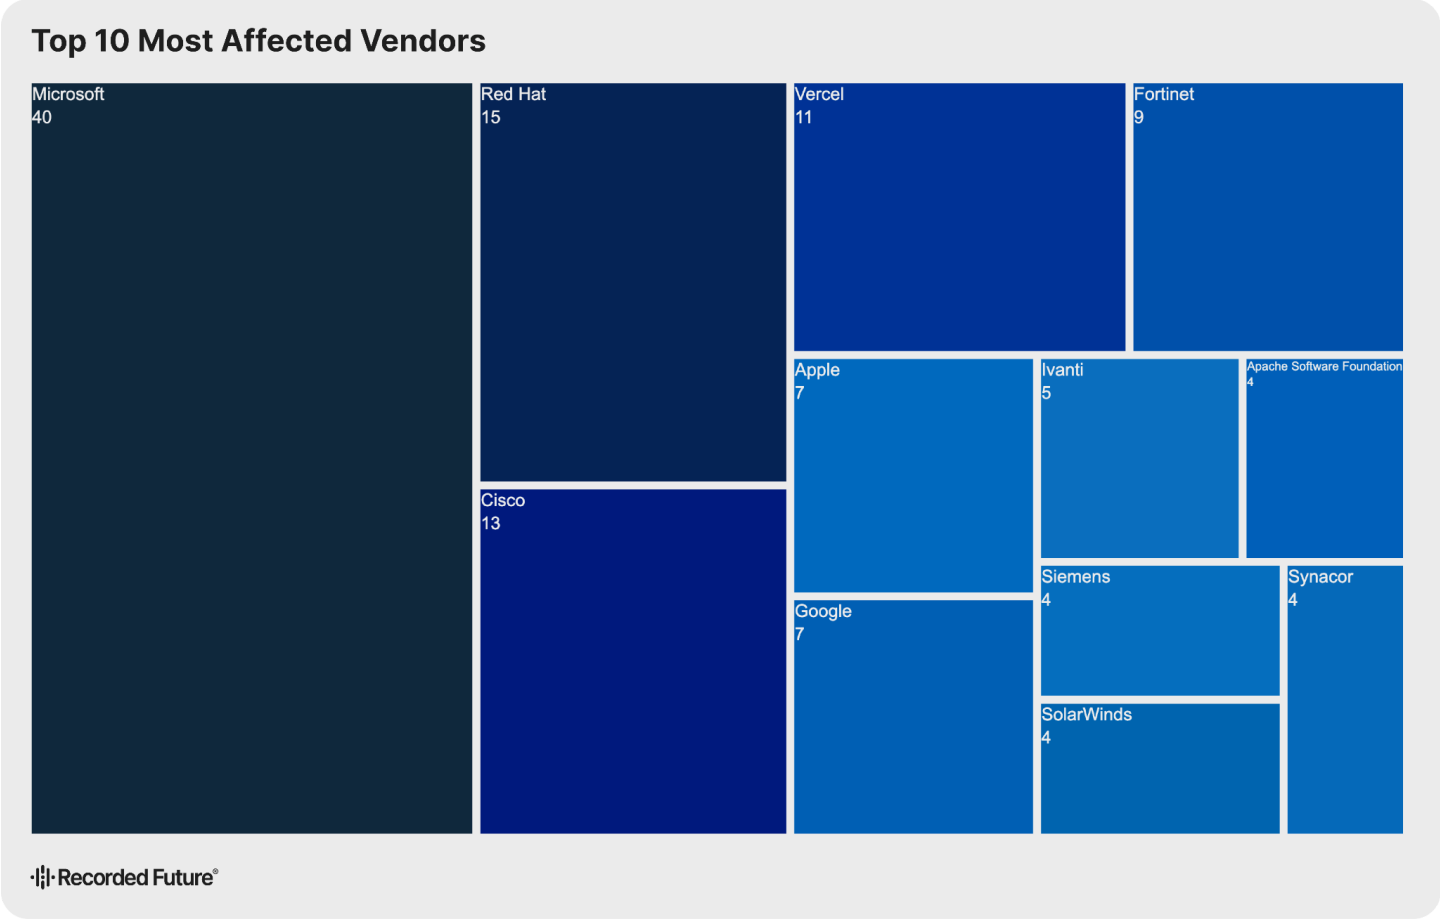

Figure 2: *Most affected vendors (top ten, including ties) by number of actively exploited vulnerabilities in H1 2026 (Source: Recorded Future)*

Of the 215 unique exploited CVEs, 66 had public proof-of-concept (PoC) exploits. Public exploits can reduce the time and expertise required to test or operationalize a vulnerability, but availability alone does not mean that the exploit was independently tested or reliably weaponized. Teams managing vulnerabilities should validate PoC exploits, prioritize affected internet-facing assets, and use interim mitigations when patches cannot be applied before exploit activity accelerates.

Across the 215 exploited vulnerabilities, exploitation was reported or detected an average of 564 days after disclosure, although the median was 31 days; this difference reflects the continued exploitation of

older vulnerabilities alongside the rapid targeting of newly disclosed flaws. Of the total, 81 (38%) were exploited before, on, or within seven days of NVD disclosure, including 20 whose honeypot activity preceded public CVE disclosure. Specifically, exploitation evidence preceded NVD publication for 23 CVEs (11%), occurred on the publication date for 20 (9%), and within the following seven days for 38 (18%).

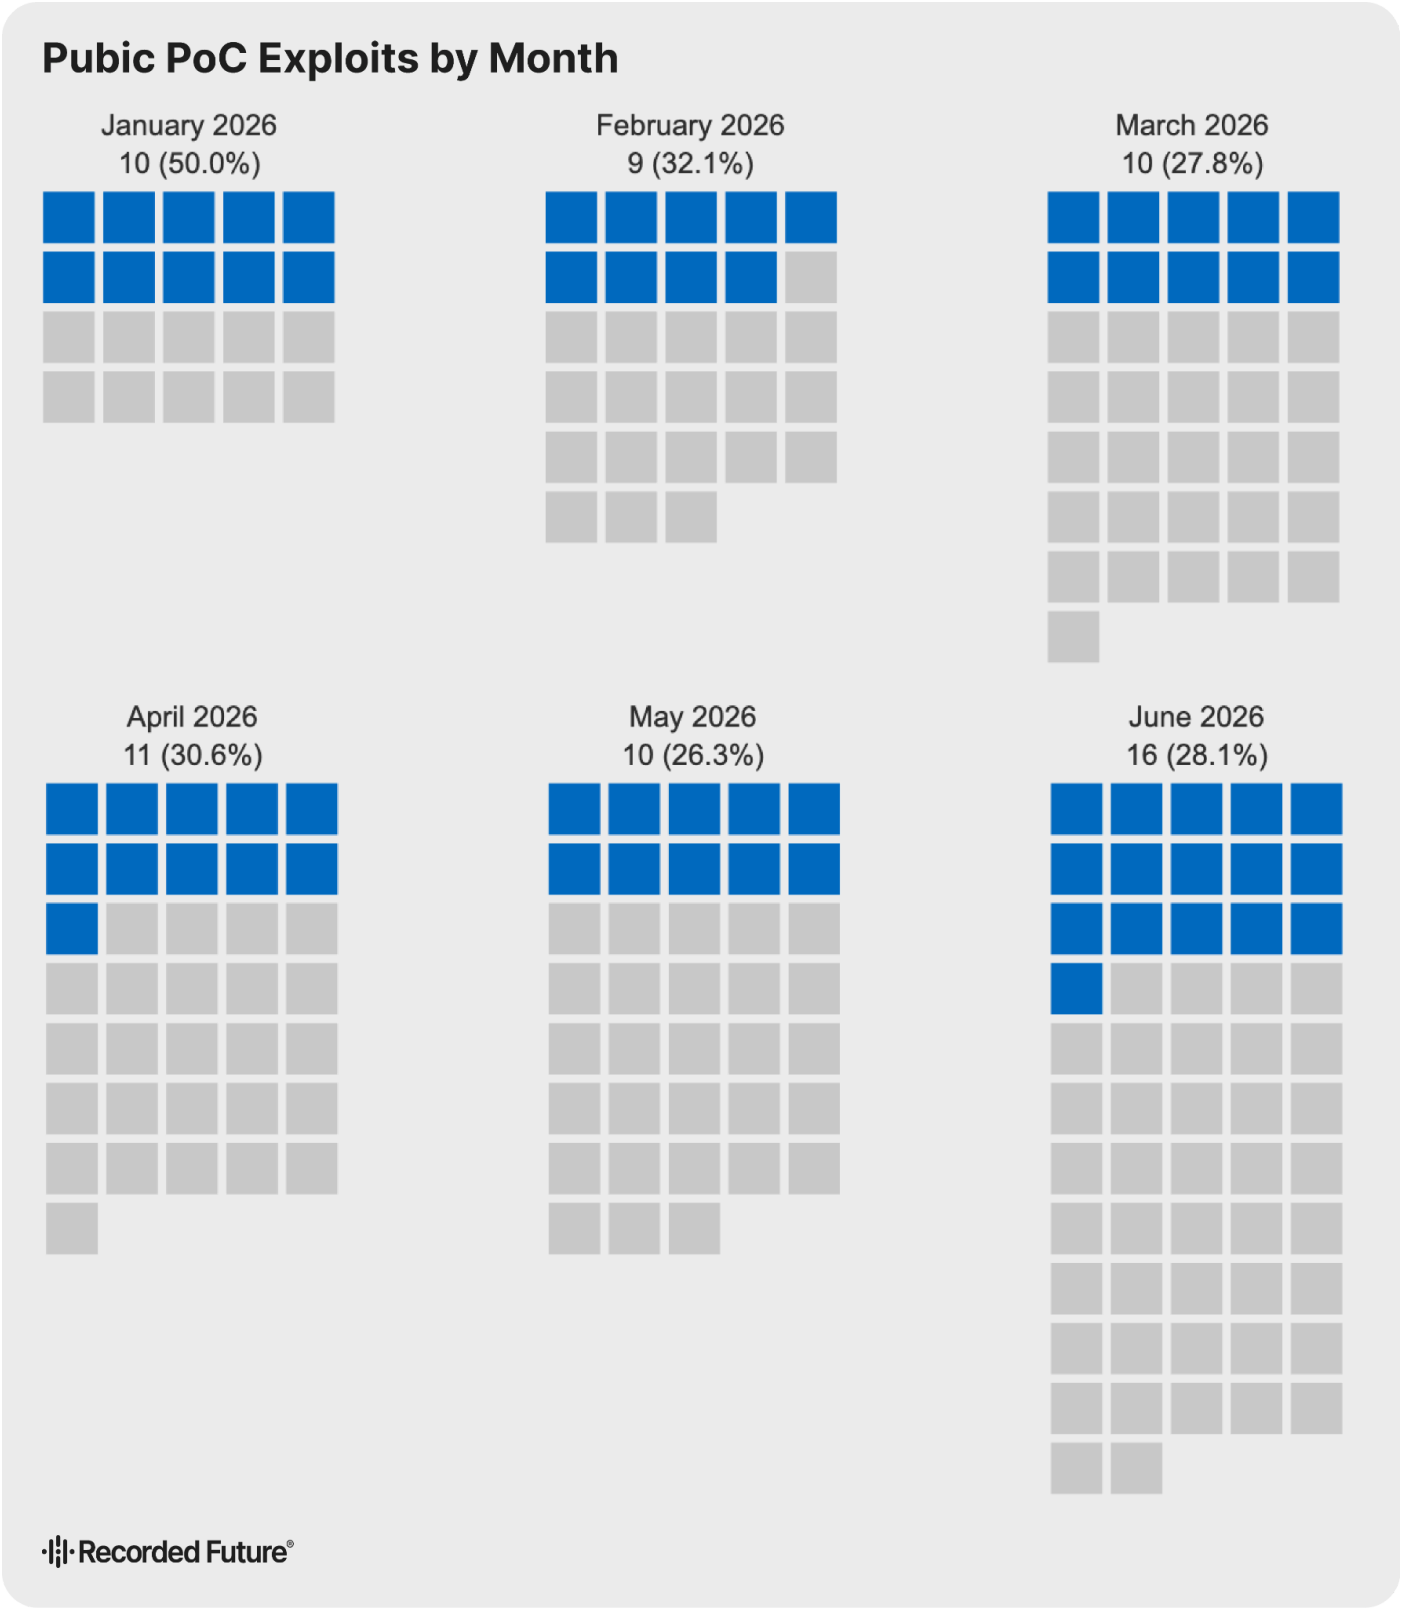

Figure 3: *Actively exploited vulnerabilities with a public exploit in H1 2026, by month (Source: Recorded Future)*

In terms of access requirements and impact, 176 (82%) of the 215 CVEs were network-accessible, and 146 (68%) could be exploited without prior authentication. More significantly, 142 of those 146 were also network-accessible, meaning that almost every vulnerability requiring no existing access could be reached over a network. Across H1 2026, 82 CVEs involved RCE, 27 involved privilege escalation, and six involved denial-of-service (DoS). Sixty of the 82 RCE vulnerabilities were also network-accessible and could be exploited without prior authentication. Defenders should prioritize these 60 RCE vulnerabilities on exposed systems because they combine remote reachability, unauthenticated exploitation, and code execution.

These combinations matter more than any one attribute in isolation. A local privilege-escalation flaw generally requires an attacker to establish access first, while a network-reachable, unauthenticated RCE vulnerability can provide the foothold itself. Defenders should begin with exposed management interfaces, application frameworks, security appliances, and other systems where exploitation can directly create server-side execution.

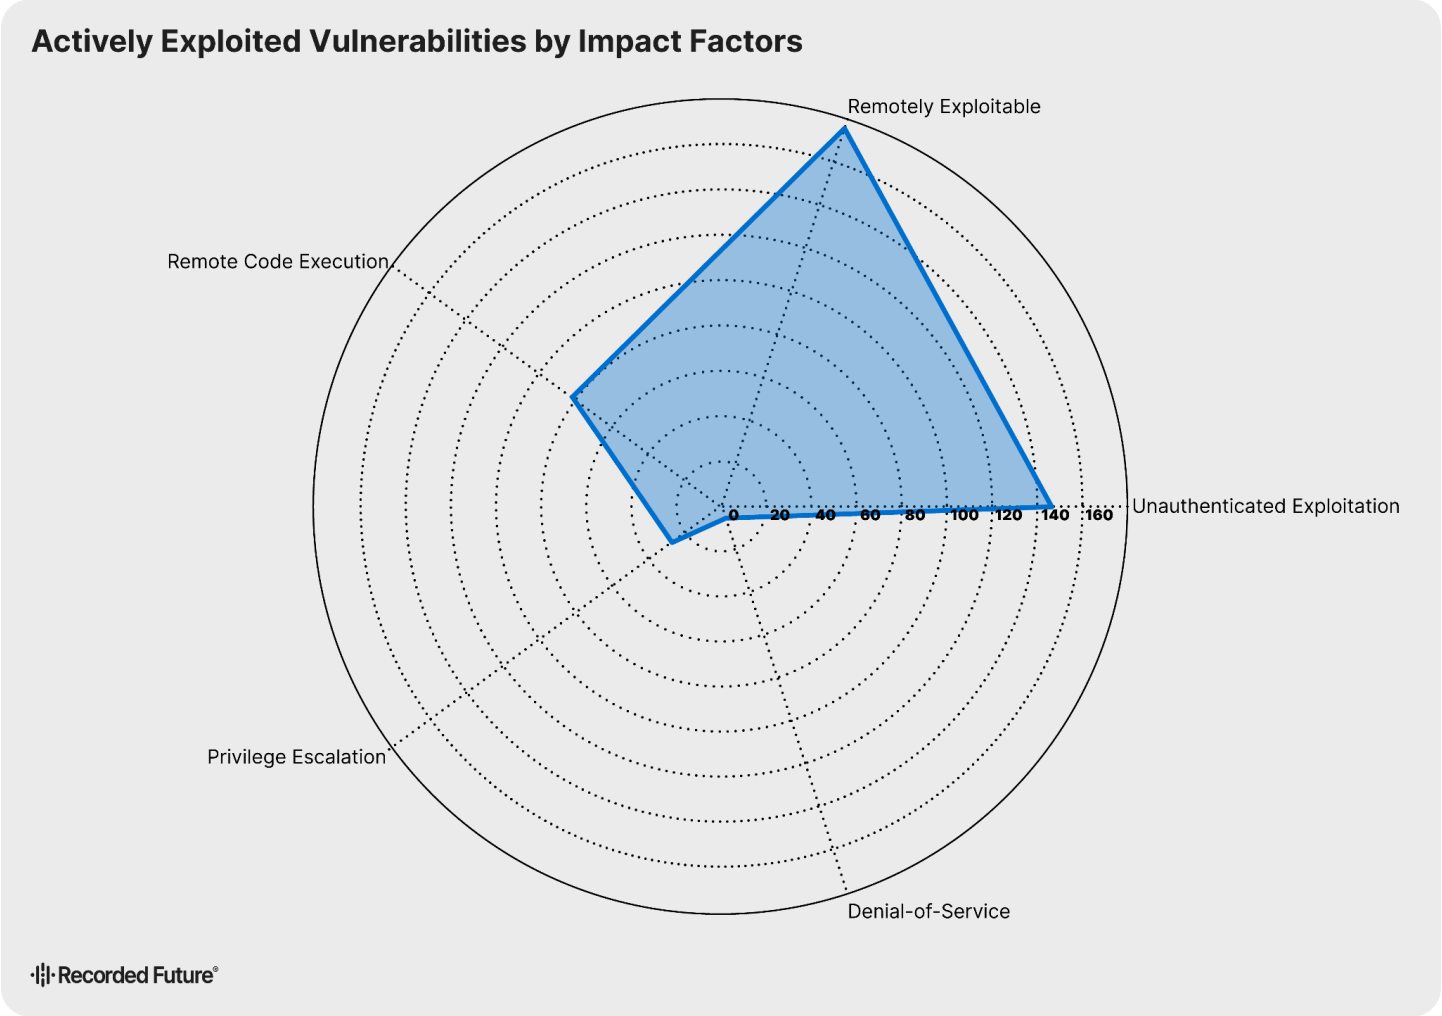

Figure 4: *Actively exploited vulnerabilities by impact factors in H1 2026 (Source: Recorded Future)*

The weakness distribution was highly fragmented: 212 of the 215 CVEs mapped to 83 distinct common weakness enumeration (CWE) classes, and no individual class accounted for more than 6% of the total. The most frequently represented weakness classes were code injection (CWE-94) across thirteen CVEs, deserialization of untrusted data (CWE-502) across twelve, and path traversal (CWE-22), improper authentication (CWE-287), and OS command injection (CWE-78) across eleven each. These flaw types can turn reachable application and management interfaces into paths for code execution, unauthorized file access, or authentication bypass. Additionally, the age distribution showed that exploitation pressure came from newer and older vulnerabilities: 162 CVEs were disclosed in 2025 or 2026, 53 predated 2025, and seventeen were from 2020 or earlier.

Cobalt Strike Beacon and Stealware Most Frequently Associated with Post-Exploitation Activity

The leading tool counts were driven by two campaigns, StrikeShark and a Storm-1175 operation, which were associated with a consistent post-exploitation tool stack across multiple initial vulnerabilities. In

the StrikeShark campaign, SharkLoader, Cobalt Strike Beacon, FScan, Searchall, Pillager Stealer, and SharpGPOAbuse were linked to the same thirteen CVEs. Those vulnerabilities spanned 2016 through 2025 and affected Microsoft Exchange and SharePoint, Fortinet FortiOS, Cisco IOS XE, F5 BIG-IP, GeoServer, Apache Shiro, and other public-facing technologies. Defenders should hunt for Cobalt Strike Beacon indicators of compromise (IoCs) and maintain behavioral detections for C2 beaconing, process injection, credential access, and remote execution.

Storm-1175 produced a second multi-vulnerability cluster. Ten CVEs were associated with the group's use of Mimikatz, Impacket, PsExec, Rclone, and Medusa ransomware. Separately, Mimikatz was also associated with the China-linked SHADOW-EARTH-053's exploitation of CVE-2021-26855, which affected Microsoft Exchange Server. This combination shows a repeatable procedure across credential dumping, remote execution via Windows administration mechanisms, data transfer, and ransomware deployment. Defenders should therefore not only monitor attempts to exploit, but also subsequent credential access, remote service execution, unusual bulk transfer activity, and ransomware behavior, rather than treating the initial vulnerability alert as an isolated event.

Of the 215 actively exploited vulnerabilities, 43 (20%) were linked with known post-exploitation malware and tooling. Stealware was the most common malware category in post-exploitation activity, appearing 29 times, followed by offensive security tools at 26, backdoors at 21, remote access trojans (RATs) at fifteen, ransomware at thirteen, broader hacking tools with twelve, loaders with five, mobile malware with four, exploit kits and web shells with three each, and botnets and malware packers with two each.

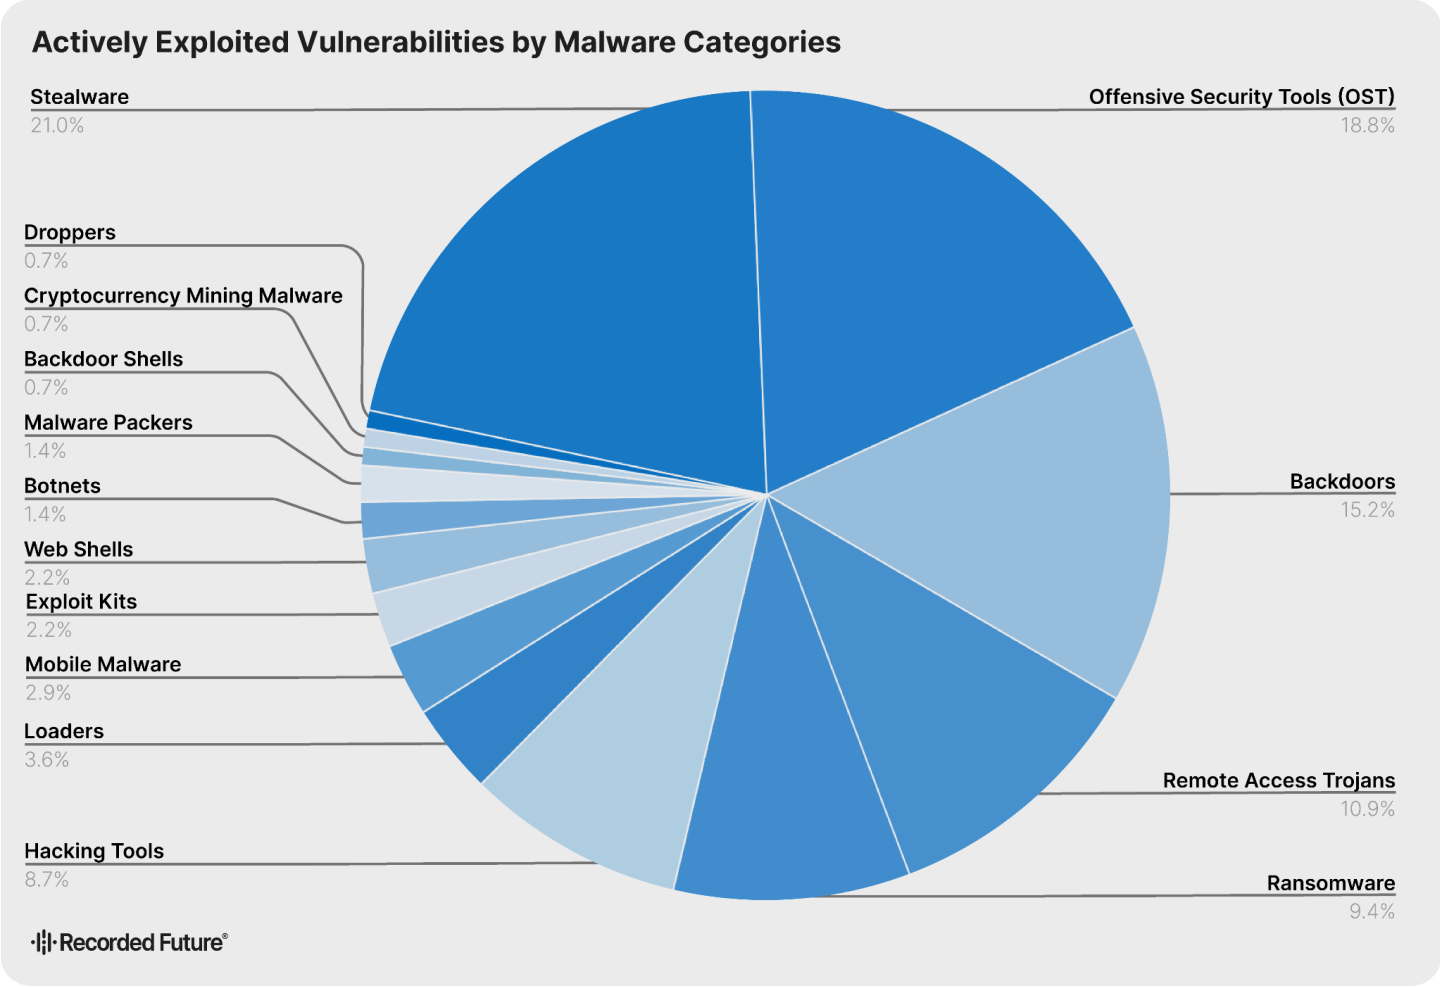

Figure 5: *Actively exploited vulnerabilities by malware category in H1 2026 (Source: Recorded Future)*

Insikt Group has reported on 114 of the 215 CVEs, mapping exploitation details to the MITRE ATT&CK Framework. Based on these mappings, Exploit Public-Facing Application (T1190) was the most common MITRE ATT&CK technique, associated with 77 CVEs (68%). It was followed by System Information Discovery (T1082), with 39 (34%); Data from Local System (T1005), with 36 (32%); and Ingress Tool Transfer (T1105), with 35 (31%). PowerShell (T1059.001) was associated with 32 CVEs; Windows Command Shell (T1059.003) and Unix Shell (T1059.004) with 31 each; Exfiltration Over C2 Channel (T1041) and Web Protocols (T1071.001) with 30 each; and Web Shell (T1505.003) with 28. The prominence of T1190 is consistent with the high number of network-accessible vulnerabilities overall, whereas the remaining techniques are frequently associated with discovery, collection, command execution, tool transfer, C2, and exfiltration.

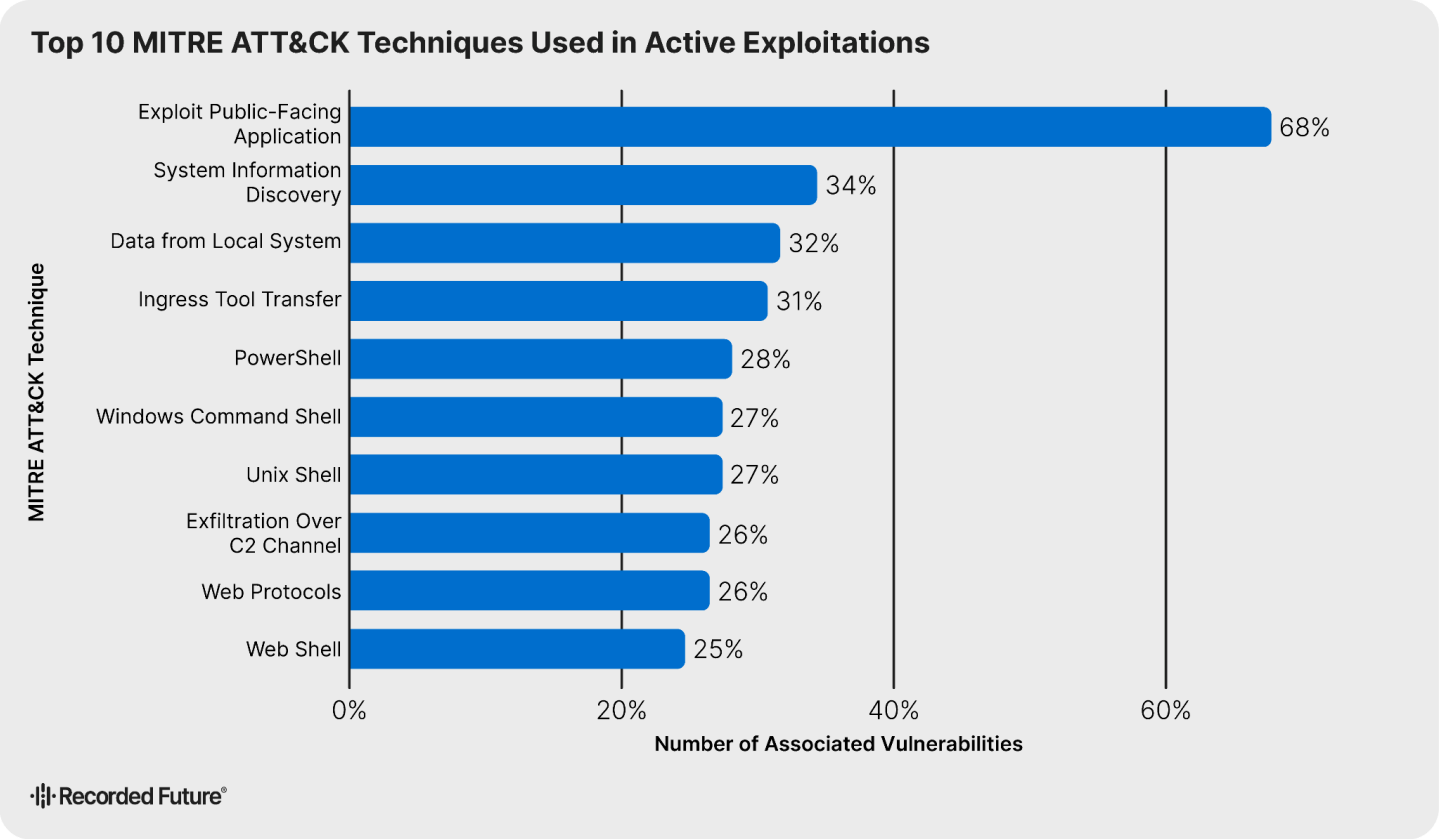

Figure 6: *Most frequent MITRE ATT&CK techniques (top ten) associated with actively exploited vulnerabilities in H1 2026 (Source: Recorded Future)*

The co-occurrence of these techniques provides a more useful defensive signal than the ranking alone. 50 of the 77 CVEs associated with the exploitation of public-facing applications were also linked to PowerShell, Windows Command Shell, or Unix Shell, and all 28 web-shell-associated CVEs also included T1190. This indicates that detailed exploitation reports frequently documented a transition from the exposed application to shell-based execution or to persistent server access. Monitoring should connect exploitation telemetry from public-facing systems to process creation, command execution, payload transfer, web shell creation, discovery, and outbound C2, rather than relying on isolated exploit signatures.

Malware Trends

Key Takeaways

- In Recorded Future Malware Intelligence submissions, AsyncRAT combined high sample volume with high C2 configuration diversity, while Cobalt Strike and Gh0st RAT generated large submission volumes but reused comparatively small sets of C2 configurations.
- RATs remained the most prominent malware type in H1 2026 Insikt Group reporting, with AsyncRAT, Cobalt Strike, REMCOS RAT, and XWorm ranking highly across both Recorded Future Malware Intelligence submissions and Insikt Group reporting.
- In the mobile malware threat landscape, threat actors moved further into real-time fraud workflows, as NFCShare and NGate abused NFC functionality for payment-card theft, unauthorized contactless payments, and ATM cash-outs, while PromptSpy and PixRevolution demonstrated early AI-assisted approaches to persistence, UI navigation, and transaction timing.

Malware Family Insights

Recorded Future® Malware Intelligence

Insikt Group reviewed malware samples captured by Recorded Future® Malware Intelligence in H1 2026 to determine the top malware families based on the total number of unique hashes and configuration diversity. This excludes malware families that are polymorphic, defunct, or seized, particularly where no recent public reporting exists to justify inclusion. This methodology reduced the risk that automated mutation, legacy detections, or inactive malware ecosystems would artificially inflate a malware family's ranking. The resulting list prioritizes families more relevant to the H1 2026 threat landscape and better reflects the malware families that defenders are more likely to encounter in current analysis workflows. **Figure 7** shows the top ten malware families by total unique hashes along with their total unique C2 configurations based on Recorded Future Malware Intelligence data captured in H1 2026.

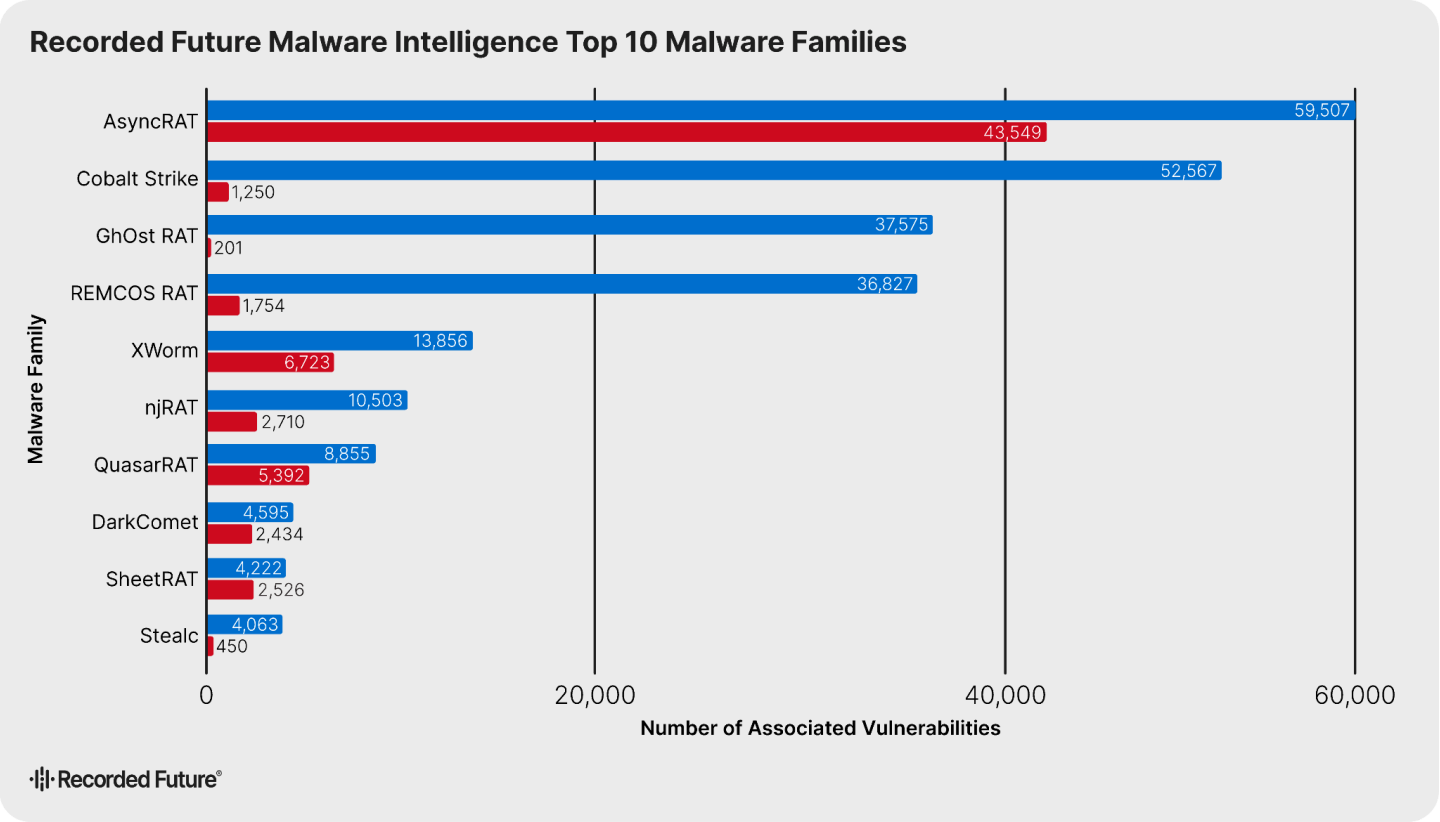

Figure 7: *Top malware families based on total unique hashes of samples submitted in H1 2026, along with corresponding unique C2 configurations (Source: Recorded Future Malware Intelligence)*

Recorded Future Malware Intelligence sample submissions for H1 2026 showed heavy concentration among RATs and offensive security tools, with AsyncRAT, Cobalt Strike, Gh0st RAT, REMCOS RAT, and XWorm comprising the five most prevalent families. Five families of the top ten families — AsyncRAT, Cobalt Strike, XWorm, Stealc, and REMCOS RAT — also ranked in the H1 2025 top ten, underscoring the continued utility of commodity RATs, stealware, and offensive security tools across different intrusion types.

AsyncRAT ranked first with 59,507 hashes and 43,549 unique C2 configurations, representing the strongest combination of sample count and infrastructure diversity among the top families. Cobalt Strike followed with 52,567 hashes but only 1,250 unique C2 configurations, while Gh0st RAT produced 37,575 hashes and 201 configurations. Hash volume alone does not fully represent operational diversity; configuration diversity provides additional context on whether submissions reflect broad infrastructure variation, repeated repackaging, or concentrated reuse of C2 settings.

Configuration diversity reveals a more detailed picture among families. AsyncRAT showed high diversity at 73.18%, while Quasar RAT, SheetRAT, DarkComet, and XWorm also exhibited comparatively high diversity relative to the overall group, suggesting broader variation in submitted configurations and infrastructure. By contrast, Cobalt Strike, Gh0st RAT, and StealC exhibited notably low configuration diversity relative to their hash volumes. Cobalt Strike's 2.38% diversity and Gh0st RAT's 0.53% diversity likely reflect repeated submissions, repackaging, shared builders, or common configurations reused

across many samples. StealC showed a similar pattern at 11.08%, indicating that its H1 2026 sample volume likely reflected recurring use of a smaller set of configurations rather than broad C2 variation.

Insikt Group Reporting

In H1 2025, we assessed that RATs such as AsyncRAT, XWorm, and REMCOS RAT remained widely used because their flexible capabilities supported multiple intrusion objectives. H1 2026 reporting reinforced that assessment. AsyncRAT, Cobalt Strike, REMCOS, and XWorm appeared among the top malware families in both Recorded Future Malware Intelligence submissions and Insikt Group reporting, showing consistency across these two sources. Across Insikt Group reporting, RATs were the most prominent malware type, followed by stealware, loaders, backdoors, and ransomware, illustrating continued demand for tooling that supports persistent remote access, credential theft, staged payload delivery, and post-compromise monetization.

The following subsections highlight representative H1 2026 activity from some of the most frequently reported malware families as identified by Insikt Group, grouped by malware type. The examples illustrate how prominent families were delivered and used across distinct campaigns, with an emphasis on recurring initial access methods, execution chains, and functionality.

RATs

Adversaries used phishing lures paired with fake verification prompts as well as trojanized developer tooling as initial access vectors before deploying RATs.

AsyncRAT

- Distributed by Kimsuky through a multi-stage LNK-based infection chain hosted on GitHub and Dropbox
- Delivered via invoice-themed phishing targeting German manufacturing organizations that exploited CVE-2024-4345

ModeloRAT

- Deployed through the ClickFix-style CrashFix campaign abusing fake browser-crash prompts and native Windows utilities
- Delivered via a malicious browser extension in a KongTuke CrashFix campaign

XWorm

- Distributed through receipt-themed phishing targeting Latin American enterprises
- Concealed inside a malicious Node Packet Manager (NPM) package disguised as a game
- Delivered via social engineering lures impersonating the Indian Income Tax Department

REMCOS RAT

- Deployed using the ZPHP JavaScript loader, which was spread through ClickFix and fake CAPTCHA prompts targeting US state, local, tribal, and territorial (SLTT) government organizations
- Delivered through a malicious OpenClaw skill

PlugX

- Distributed via a trojanized Claude Code installer
- Delivered through a dynamic-link library (DLL) sideloading chain using a spoofed MSI installer
- Deployed by the China-nexus threat group Camaro Dragon against Qatari targets

DOGCALL (RokRAT)

- Delivered through a North Korean APT37 Facebook-based pretexting campaign that lured targets into installing a tampered PDFelement installer
- Delivered through spearphishing emails carrying malicious Hangul Word Processor (HWP) documents disguised as interview materials

Stealware

Stealware targeted browser credentials and cryptocurrency wallets, and, in some cases, developer credential stores.

MacSync

- Delivered through Google-sponsored results and ClickFix-style lures hosted on legitimate platforms, including claude.ai and Medium

Vidar

- Distributed through compromised WordPress sites
- Distributed through trojanized GitHub repositories posing as leaked Claude Code
- Distributed through short-form video tutorials on TikTok and Instagram Reels that promised free premium software

Amatera

- Delivered via cloned Claude Code installation pages using the InstallFix technique
- Delivered through a fake CAPTCHA chain that abused a signed Microsoft App-V script

GlassWorm

- Harvested GitHub tokens from developer credential stores and extension storage before pivoting to account takeover

Atomic macOS Stealer (AMOS)

- Delivered through malicious OpenClaw skills
- Delivered through Cursor AI agent sessions

Loaders

Loader activity involved multi-stage social engineering and script-based staging designed to smuggle payloads past defenses.

DonutLoader

- Delivered through a phishing email containing a malicious batch file that abused the legitimate `SyncAppvPublishingServer.vbs` component to stage an AutoIt script, which reconstructed loader shellcode in memory and injected it into the legitimate `colorcpl.exe` process to launch REMCOS RAT

CastleLoader

- Distributed through TAG-176's Finger protocol-based ClickFix infection chain, in which victims were lured to fake-verification pages that invoked the Finger protocol to retrieve a Python-based loader ahead of the CastleLoader payload

HijackLoader

- Delivered via phishing pages impersonating Telefónica's O2 service that triggered automatic executable downloads, which decrypted and loaded an embedded SnappyClient payload into memory
- Delivered separately via Viber-delivered decoy documents targeting Ukrainian government and military personnel and ultimately loaded REMCOS RAT in memory

Backdoors

State-sponsored threat actors used backdoor malware to support espionage efforts focused on long-term, covert access to government, diplomatic, and edge-device infrastructure.

Cobalt Strike

- Delivered through the exploitation of public-facing applications, including Exchange, Openfire, and GeoServer vulnerabilities used by StrikeShark to deploy SharkLoader ahead of the Cobalt Strike Beacon payload
- Delivered via spearphishing PDF lures impersonating a Ukrainian telecommunications provider that FrostyNeighbor used to deliver PicassoLoader before deploying Cobalt Strike

LOTUSLITE:

- Delivered via a politically themed ZIP archive that abused DLL sideloading to target US government and policy-related entities
- Delivered in an updated v1.1 campaign via a malicious Compiled HTML Help (CHM) file that targeted India's banking sector and South Korean government entities

BRICKSTORM:

- Deployed by VerdantBamboo after the threat actor obtained valid secure shell (SSH) credentials for a default account from a victim's managed service provider and exploited a sudo misconfiguration to gain root-level access to edge appliances, including an Egnyte Storage Sync virtual appliance, a pfSense firewall, and a Synology Network Attached storage (NAS)

Ransomware

Double extortion remained the dominant extortion model, while some ransomware threat actors expanded to triple extortion tactics by combining encryption and data theft with additional coercive measures. Common techniques used in ransomware campaigns throughout H1 2026 remained consistent with those observed in H1 2025, including ClickFix-based social engineering for initial access and the abuse of legitimate remote monitoring and management (RMM) tools to maintain access and move laterally. Ransomware operators continued to favor trusted user workflows and legitimate administrative tools because they can support multiple intrusion phases while blending into routine enterprise activity.

Agenda

- Building on 2025 activity that deployed a Linux encryptor against Windows hosts through Windows Subsystem for Linux (WSL), a new campaign automated WSL installation with a batch script to broaden its pool of encryptable hosts, exfiltrated data through the legitimate `s5cmd` utility to Amazon S3, sustained remote access with the commercial Remotely tool, randomized local security authority subsystem service (LSASS) dump file extensions to obscure credential theft, and removed MeshAgent components after use to eliminate remote-access artifacts.

Interlock

- Threat actors deploying Interlock ransomware exploited a Cisco Secure Firewall Management Center zero-day (CVE-2026-20131) to execute code as root on unpatched devices
- Hive0163 used a ClickFix lure that hijacked victims' clipboards to run a malicious PowerShell command installing the NodeSnake first-stage client of an AI-generated C2 framework; the campaign also deployed Interlock RAT, which communicated with hard-coded C2 infrastructure via web sockets

Medusa

- Deployed by the China-linked Storm-1175 group, which chained n-day and zero-day exploits against public-facing applications before moving laterally over remote desktop protocol (RDP) through Cloudflare tunnels using remote monitoring and management tools such as Atera, AnyDesk, and DWAgent

Rhysida

- Impersonated Microsoft Teams and Google Authenticator in malvertising campaigns to deliver CleanUpLoader, which was in turn used to drop the Vidar infostealer on victim systems

Reynolds:

- Embedded bring-your-own-vulnerable-driver (BYOVD) functionality directly into its ransomware payload to disable endpoint security controls before encryption; it dropped the vulnerable NsecSoft NSecKrnl driver, registered it as a service, and exploited CVE-2025-68947 to terminate processes associated with security products, including Microsoft Defender, CrowdStrike, Sophos, and Symantec; this payload-integrated BYOVD design is not novel, as earlier BlackByte encryptors had incorporated similar BYOVD functionality

Prinz Eugen

- Enhanced anti-forensic capabilities by prioritizing recently modified files and employing key wiping and self-deletion mechanisms

pe32s

- Adopted post-quantum cryptographic algorithms, making recovery of its encryption keys more difficult and reducing the likelihood that future cryptographic advances will help victims decrypt their files

LeakNet

- Added ClickFix lures delivered through compromised legitimate websites as an initial access method, and used PsExec for lateral movement after initial compromise

The Gentlemen:

- Used remote access software, specifically AnyDesk, for persistence
- Analysis of the hastalamuerte and The Gentlemen ransomware-as-a-service (RaaS) ecosystem identified shared references, scripts, and snippets related to extracting Veeam credentials, including references to CVE-2023-27532 affecting Veeam Backup & Replication. The same reporting described commands for creating accounts, assigning database permissions in Veeam environments, and interacting with Veeam Backup databases; the threat actors also stopped Veeam and other backup services and deleted shadow copies to inhibit recovery

- In May 2026, The Gentlemen's internal chats were leaked following the breach of virtual private server (VPS) provider 4VPS, exposing 22 chat rooms from the group's Rocket.Chat server that revealed its RaaS organizational structure, tradecraft, and operations

Mobile Malware

The mobile threat landscape continued to evolve during H1 2026, with publicly reported activity centered on Android banking trojans, Near Field Communication (NFC)-enabled malware, spyware, and other financially motivated mobile malware. Consistent with trends observed throughout H1 and H2 2025, newly reported mobile malware families, variants, and campaigns continued to prioritize credential theft, account takeover, payment fraud, remote device control, transaction manipulation, and payment-card abuse. Industry reporting published in 2026 found that 34 active mobile malware families targeted more than 1,200 financial applications across 90 countries, underscoring the global scale of mobile banking malware activity entering H1 2026.

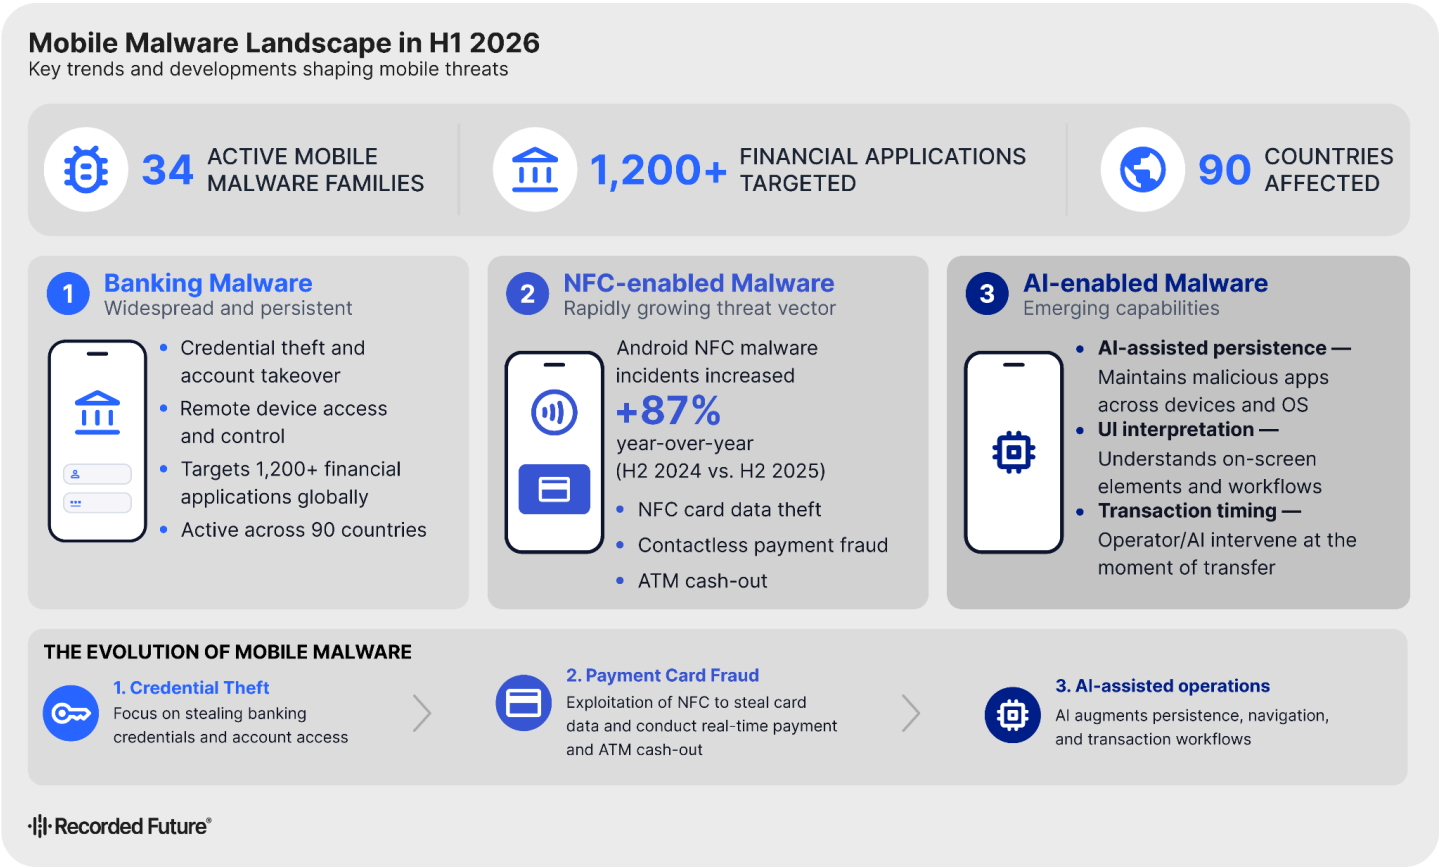

Figure 8: *Mobile malware landscape in H1 2026 (Source: Recorded Future)*

One of the most notable developments observed between H2 2025 and H1 2026 was the continued growth of NFC-enabled malware. Android NFC malware incidents increased by around 87% year-over-year (YoY) in H2 2025. This trend continued to accelerate in H1 2026, with public reporting identifying a 188% increase in NFC-based Android attacks between January and April 2026. These

attacks primarily sought to facilitate payment fraud, steal payment card data, and conduct ATM cash-out schemes. Examples include:

- In January 2026, NFCShare, an Android banking trojan, was distributed through a Deutsche Bank-themed phishing campaign. Victims were instructed to install a malicious Android Package (APK), place a payment card near the device, and enter the card PIN. The malware then read NFC payment card data and exfiltrated it to a remote WebSocket endpoint.
- In April 2026, a new NGate variant targeting Android users in Brazil trojanized the legitimate HandyPay NFC relay application. The malware relayed payment card NFC data to attacker-controlled devices to facilitate unauthorized contactless payments or ATM cash-out schemes while also stealing card PINs.

These events illustrate that NFC abuse is not limited to credential theft; it increasingly supports real-time payment-card fraud workflows that combine social engineering, device permissions, and payment-card proximity.

Early examples of AI-linked mobile malware emerged in H2 2025, while H1 2026 introduced additional AI-enabled capabilities, including AI-assisted persistence mechanisms and operational workflows. Examples include:

- In February 2026, PromptSpy was identified as a likely proof-of-concept (PoC) Android threat that used Google's Gemini to interpret user interface (UI) elements and provide step-by-step instructions for keeping the malicious app pinned in the recent apps list, improving persistence across different device layouts and Android variants.
- In March 2026, PixRevolution was identified as an Android banking trojan targeting Brazil's PIX payment system through an operator-in-the-loop model, in which a human or AI agent could monitor a victim's screen in near real time and intervene at the moment of transfer.

These events suggest that AI-enabled mobile threats are beginning to augment persistence, UI navigation, and transaction timing rather than replace established mobile banking trojan tradecraft.

TTP Trends

Key Takeaways

- The most common malware tactics, techniques, and procedures (TTPs) observed in H1 2026 TTP Instances reflected continued adversary reliance on command execution, obfuscation, user execution following social engineering, host discovery, and payload transfer.
- Supply-chain compromises in H1 2026 increasingly targeted trusted developer infrastructure, including package ecosystems and AI-enabled tooling, where privileged access to credentials, repositories, continuous integration/continuous delivery (CI/CD) systems, and cloud services allowed attackers to propagate malicious code and expand compromise across downstream environments.
- Magecart activity continued to evolve in H1 2026, with threat actors abusing trusted third-party services, alternative browser execution mechanisms, fake payment interfaces, and administrator-aware evasion to conceal payment-card theft within legitimate e-commerce workflows.

Insikt Group identified the most prominent MITRE ATT&CK techniques observed in Recorded Future TTP Instances published in H1 2026. Prominence is based on the volume of ATT&CK technique references in those published TTP Instances. Across H1 2026 TTP Instances, Command and Scripting Interpreter (T1059) led with 320 references, followed by Obfuscated Files or Information (T1027) with 201, User Execution (T1204) with 187, System Information Discovery (T1082) with 180, and Ingress Tool Transfer (T1105) with 178. These findings indicate that H1 2026 malware activity commonly relied on script- and command-based execution, evasive packaging or encoding, user execution following social engineering, and payload staging to achieve its objectives.

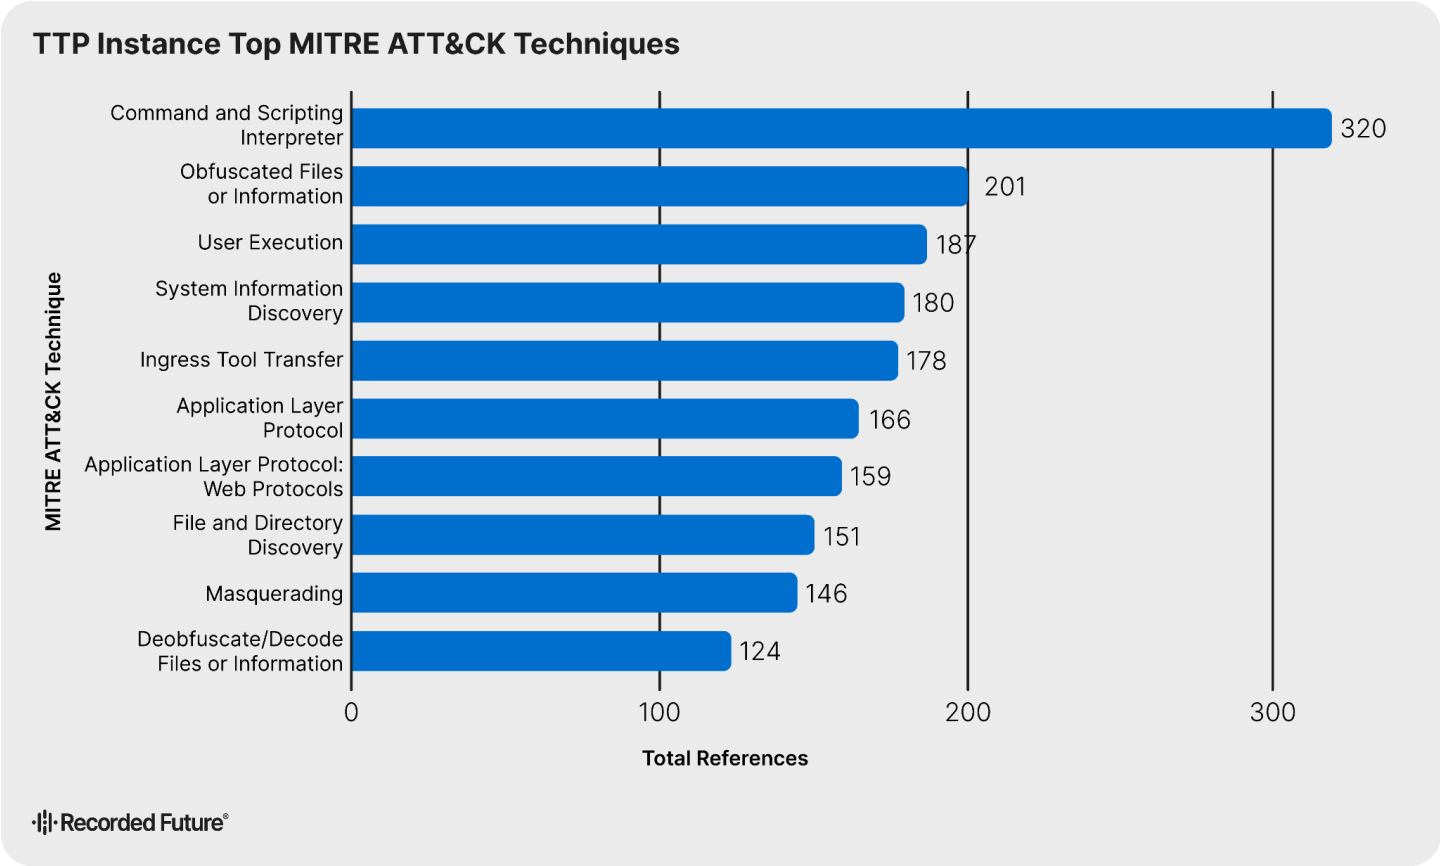

Figure 9: *Top MITRE ATT&CK techniques observed in H1 2026 TTP Instances (Source: Recorded Future)*

The top ten ATT&CK techniques appeared in activity linked to cybercriminal, financially motivated, and state-sponsored threat actors. RATs, stealware, loaders, backdoors, and ransomware appeared most frequently in campaigns using these techniques, consistent with the findings discussed in the **Prominent Malware** section. The techniques supported credential theft, data theft, remote access, payload delivery, espionage collection, and ransomware deployment. Together, they show how threat actors used execution, obfuscation, discovery, C2, and tool transfer after initial compromise.

Examples of TTP Instances demonstrating the top five ATT&CK techniques included the following:

- Lazarus used Command and Scripting Interpreter to execute commands after exploiting React2Shell to deploy COPPERHEDGE against financial and blockchain-related organizations.
- Mini Shai-Hulud operators used Obfuscated Files or Information to complicate analysis and conceal malicious behavior within compromised developer workflows.
- Threat actors used User Execution in a GlassWorm campaign, relying on developers to install trojanized Open Visual Studio Extensions (VSXs) or extension packs that pulled malicious dependencies through the marketplace's normal extension-installation process.
- A Brazilian phishing campaign deploying agenteV2 banking stealer against financial users used System Information Discovery to profile infected systems after execution.
- A REMCOS RAT campaign deploying DonutLoader relied on Ingress Tool Transfer to stage payload delivery and enable follow-on malware execution after initial compromise.

Supply-Chain Compromises

In H1 2026, supply chain compromises were marked by threat actors abusing trusted access paths, software distribution channels, and intermediary platforms to reach downstream victims at scale. Per ReversingLabs' 2026 Software Supply Chain Security Report, the cybersecurity firm identified a 73% increase in malicious open-source packages during 2025, with similar threats continuing to trend into H1 2026. The available evidence from H1 2026 shows repeated compromise of systems that aggregate sensitive customer, patient, financial, or operational data, including medical sales platforms, digital health services, insurance environments, financial institutions, and blockchain governance controls. This activity reflected a broader shift toward exploiting relationships of trust, where compromised credentials, administrative privileges, multi-sig authority, and supplier-managed systems gave attackers access to high-value data or assets without requiring direct compromise of every downstream organization.

Abuse of package managers and developer tooling was prominent in H1 2026 supply-chain activity, with threat actors using compromised maintainer accounts, package lifecycle scripts, and developer platforms to steal credentials, propagate malicious packages, and sustain compromise across software ecosystems. On July 9, 2026, Sonatype published a blog post discussing Q2 open-source malware trends, which mentioned that the security firm had identified more than 464,000 malicious open-source packages during the quarter, and that 96.6% of the packages were discovered on npm, demonstrating that the threat of package manager abuse continues to persist.

In April 2026, Socket reported that compromised SAP Cloud Application Programming Model (CAP) and Cloud MTA-related npm packages, including mbt and multiple `@cap-js/*` packages, introduced malicious preinstall behavior that executed automatically during npm install, downloaded a Bun runtime, and ran an obfuscated payload in developer and CI/CD environments. The payload harvested GitHub, npm, cloud, Kubernetes, Docker, messaging, cryptocurrency wallet, browser, SSH, and developer-tooling credentials; probed cloud metadata services; extracted CI secrets from runner memory; and used stolen access to create GitHub repositories for encrypted exfiltration and to publish additional compromised npm packages. Socket assessed with medium confidence that the activity was linked to TeamPCP and related Shai-Hulud-style supply-chain campaigns, highlighting how package lifecycle scripts can turn routine dependency installation into credential theft, propagation, and persistence.

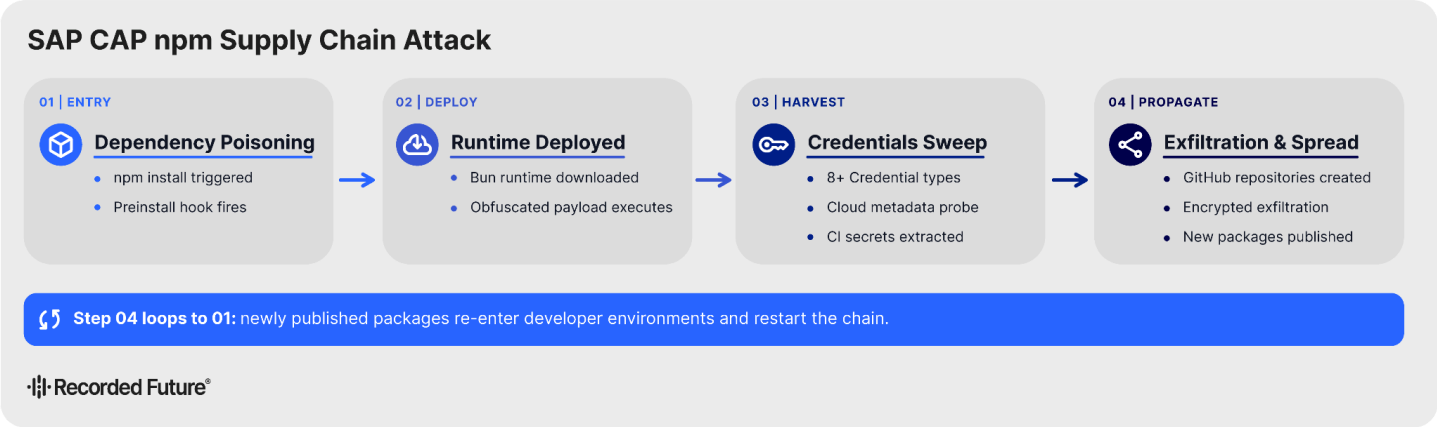

Figure 10: Visual representation of TeamPCP's compromise of SAP CAP's npm packages (Source: Recorded Future)

Activity attributed to or closely resembling TeamPCP's "CanisterWorm" campaign, reported throughout H1 2026, followed a similar pattern. Threat actors compromised legitimate npm publisher accounts and republished trojanized package versions that ran postinstall scripts to harvest developer credentials during routine dependency installation. Stolen npm publishing tokens then let the attacker enumerate a victim's own packages, inject hooks, and republish the packages automatically, carrying the compromise from one maintainer to the next without operator involvement. Exfiltration ran partly through Internet Computer Protocol (ICP) blockchain canisters that resist conventional takedown;the worm also seeded malicious packages in PyPI. These cases show how a single compromised maintainer account can sustain compromise across package ecosystems.

Additionally, in at least one instance, CanisterWorm-linked activity demonstrated destructive capabilities by wiping Kubernetes clusters belonging to Iranian targets, while related variants established persistence and enabled lateral movement via stolen SSH keys and exposed Docker application programming interfaces (APIs). This behavior showed how the compromise of third-party or supply-chain assets could expand beyond developer infrastructure into cloud and containerized environments.

The Shai-Hulud campaigns also showed how quickly supply-chain tradecraft can spread once tooling becomes open-source. In May 2026, OX Security reported that TeamPCP appeared to leak Shai-Hulud malware source code to GitHub through compromised accounts, with repositories containing deployment instructions and copycat modifications. OX Security observed patterns consistent with earlier Shai-Hulud activity, including uploading stolen credentials to newly created GitHub repositories, communicating with predefined C2 infrastructure, exfiltrating secrets and wallets, and targeting Claude Code configurations by adding hooks that execute malware when the tool starts. These examples underscore that developer ecosystems are not only distribution channels but also credential stores, automation surfaces, and propagation infrastructure.

Late 2025 and H1 2026 reporting also showed how AI-enabled developer tools are becoming part of the software supply-chain attack surface because of the access and automation privileges they hold within development environments. Microsoft reported in December 2025 that Shai-Hulud modified Claude

Code configuration files to add session hooks that re-executed malware when the tool started, turning a trusted developer workflow into a persistence mechanism. In March 2026, Trend Micro reported that TeamPCP compromised LiteLLM, an AI gateway used to connect applications with multiple large language model providers, and used the access to harvest API keys, cloud credentials, SSH keys, Kubernetes secrets, and other developer credentials before moving into downstream cloud and container environments. The emerging risk around AI-enabled development tools is not limited to the exploitation of the underlying AI models. Their integrations with repositories, CI/CD systems, cloud platforms, package ecosystems, and other privileged services can give attackers broader access once a trusted component is compromised.

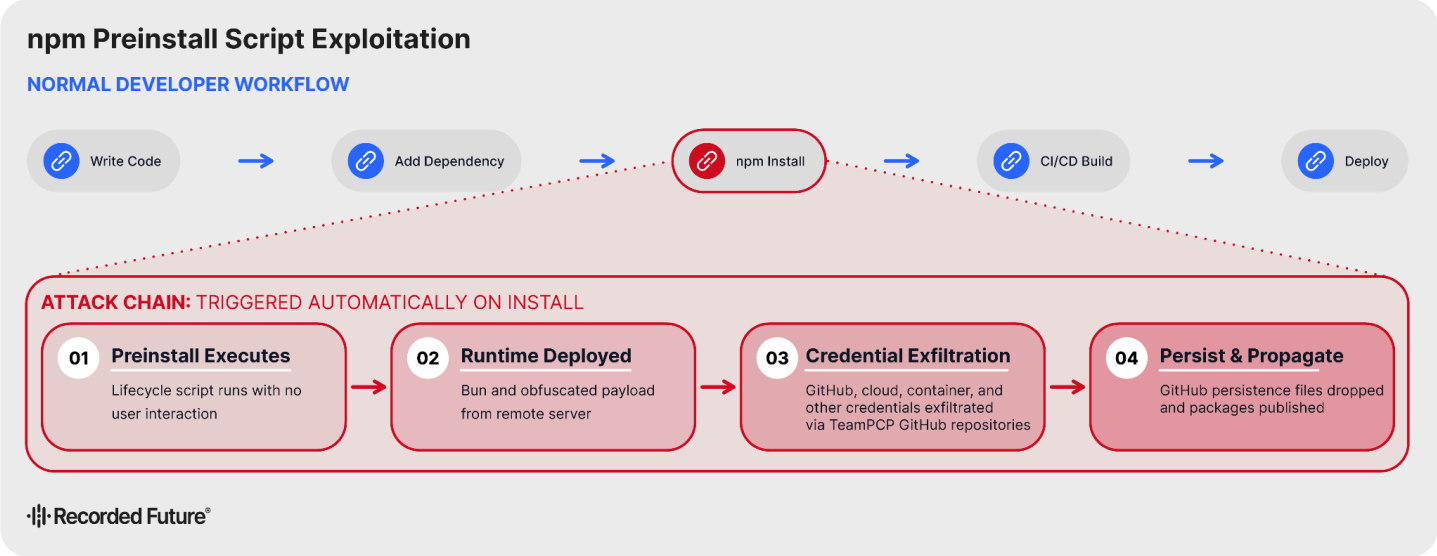

Figure 11: *Visual representation of TeamPCP's poisoning of GitHub repositories with Shai-Hulud-like malware (Source: Recorded Future)*

Magecart Infections

Compared with H1 2025 and H2 2025, Magecart activity observed during H1 2026 showed greater emphasis on the abuse of trusted third-party services for skimming payload delivery and data exfiltration. While H1 2025 and H2 2025 campaigns emphasized techniques such as Content Security Policy (CSP) evasion, blockchain-hosted payload delivery, WebSocket-based payload retrieval, analytics tag impersonation, and fake payment forms, H1 2026 public reporting highlighted the use of legitimate services, including Google Tag Manager and Stripe, as part of the skimming infrastructure.

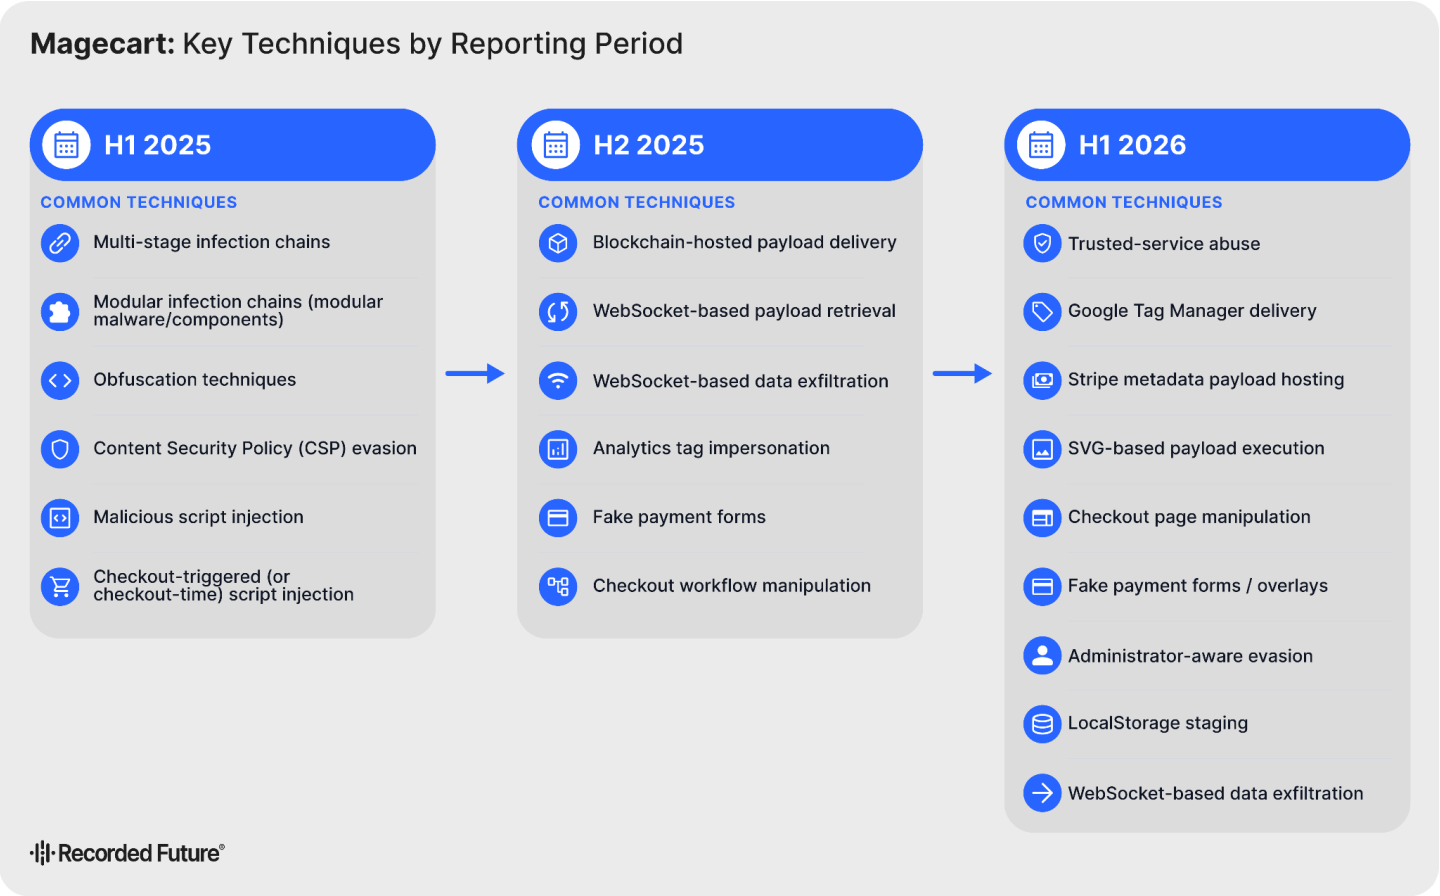

Figure 12: *Evolution of Magecart techniques based on public reporting from H1 2025 through H1 2026 (Source: Recorded Future)*

Magecart activity reported during H1 2026 demonstrated continued evolution in client-side web-skimming tradecraft, with threat actors abusing trusted third-party services, adopting alternative browser execution mechanisms, and manipulating checkout workflows to harvest payment card data:

- WooCommerce: A long-running campaign targeted WooCommerce websites by hiding the legitimate Stripe payment form, injecting a counterfeit iframe-based payment form, harvesting payment and personal data in real time, using WordPress Admin Bar detection for administrator-aware evasion, and exfiltrating the collected data through HTTP POST requests.
- Magento / Adobe Commerce: A campaign targeting Magento and Adobe Commerce websites abused Google Tag Manager to load skimming code, retrieved malicious payloads from Stripe customer metadata, harvested payment card information from checkout pages, and exfiltrated the collected data through Stripe APIs.
- Magento / Adobe Commerce SVG Campaign: A separate campaign concealed skimming code within hidden SVG elements, executed the payload through scalable vector graphics (SVG) onload events, displayed a counterfeit Secure Checkout overlay to collect payment card information, and exfiltrated the harvested data to attacker-controlled infrastructure.

Mitigations

Organizations can use the following measures to mitigate the threats discussed in this report.

Vulnerability Exploitation

- **Recorded Future® Vulnerability Intelligence**: Use Vulnerability Intelligence to gain timely, comprehensive insights into publicly and privately known vulnerabilities, tailored to specific security needs, and to help prioritize remediation measures. With Vulnerability Intelligence, Recorded Future customers also gain access to Nuclei templates created by Insikt Group, enabling proactive scanning and rapid detection of emerging vulnerabilities.
- **Recorded Future Attack Surface Intelligence**: Use external Attack Surface Intelligence to discover unknown and unmanaged internet-exposed assets, including forgotten infrastructure, shadow IT, subsidiary assets, third-party infrastructure, and cloud-connected services. Teams can use Attack Surface Intelligence to prioritize externally visible exposures based on exploitability and active threat activity, reduce public-facing misconfigurations and unnecessary exposures, and bring newly discovered assets into vulnerability management and remediation workflows.
- **Exposure and Asset Inventory:** Conduct regular inventories of all internet-facing hosts, including remote access appliances, edge devices, and third-party software. Prioritize monitoring gateway-layer products (such as Secure Sockets Layer [SSL]-VPNs, firewalls, and virtualization underlays), which were frequently exploited in H1 2026. Run scheduled scans to identify shadow IT assets and newly exposed services to fill critical security gaps.
- **Rapid Patch Prioritization**: Develop and enforce rapid patch deployment workflows for vulnerabilities with known public PoCs or active exploitation. Prioritize vulnerabilities that are remotely exploitable, require no authentication, or affect widely deployed software (such as Microsoft, Ivanti, or Fortinet). When immediate patching is not possible, apply mitigation controls such as web application firewall (WAF) rules, intrusion prevention system (IPS) signatures, access restrictions, and segmentation.
- **Harden Internet-Facing and Security Appliances**: Enforce multi-factor authentication on all administrative interfaces and disable unused or legacy services (such as insecure SSH or Telnet). Segregate edge-device management networks from production networks, and restrict the internal subnets a compromised VPN or firewall can reach.

Malware Intrusions

- **Recorded Future Hunting Packages**: Implement Hunting Packages developed by Insikt Group to monitor intrusions associated with malware families, as well as their delivery, execution, discovery, credential access, C2, and data collection behaviors.
- **EDR, Heuristic, and Behavior-Based Analysis**: Tune endpoint detection and response (EDR) and security information and event management (SIEM) detections to correlate suspicious sequences of activity rather than alerting on individual behaviors in isolation. Prioritize

combinations such as user-driven execution followed by script or installer activity, system discovery, ingress tool transfer, credential access, unusual C2, lateral movement, or post-exploitation tooling.

- **Application Allowlisting and Script Control**: Restrict unnecessary PowerShell, command shell, `msiexec`, HTA, JavaScript, and package manager execution. Alert on unusual script execution from user directories, browser download paths, temporary folders, and developer workstations.
- **Supply-Chain and Developer Environment Hardening:** Audit npm, GitHub, CI/CD, cloud, container, SSH, messaging, and developer-tooling credentials. Enforce least privilege for package publishing, protect CI secrets, review package lifecycle scripts, and monitor for unauthorized repository creation or dependency changes.
- **Ransomware and Remote Access Tool Monitoring:** Inventory and restrict remote monitoring and management (RMM) tools such as AnyDesk, SimpleHelp, Net Monitor, and similar remote access software. Monitor for ClickFix-style command execution (for example, `msiexec`, `rundll32`, or `regsvr32`), PsExec use, scheduled task creation, log clearing, suspicious PowerShell, attempts to disable endpoint security, backup service termination, and unusual data staging or exfiltration.
- **Mobile Malware and Fraud Monitoring:** Require users to install mobile apps only from trusted app stores and, where possible, block application sideloading. Enforce mobile device management controls, and monitor for suspicious APK installation, accessibility-service abuse, overlay behavior, NFC relay indicators, unusual device enrollment, and transaction timing consistent with operator-assisted payment fraud.

Magecart Attacks

- **Recorded Future Payment Fraud Intelligence**: Use Payment Fraud Intelligence to monitor ongoing Magecart e-skimmer infections, stay current with the latest Magecart TTPs, enhance prevention strategies, and take action on compromised cards before fraud occurs.
- **Regular Security Audits and Vulnerability Scanning:** Conduct regular audits and automated scans of Magento, Adobe Commerce, WordPress, WooCommerce, plugins, themes, tag managers, analytics tools, and payment integrations that support e-commerce workflows.
- **Content Security Policy (CSP):** Deploy and maintain a strict CSP to control resources loaded on payment pages and reduce unauthorized script execution. Review CSP rules regularly so trusted third-party services do not become broad allowlist paths for skimming infrastructure.
- **Third-Party Integrations Monitoring:** Maintain an inventory of payment-page scripts and third-party integrations, require explicit business justification for each script, and monitor services such as Google Tag Manager, Stripe integrations, analytics tools, and payment widgets for unauthorized changes.
- **Payment-Page Integrity Monitoring:** Monitor checkout pages for unauthorized JavaScript, SVG-based execution, fake payment forms, unexpected redirects, administrator-aware evasion, and script changes that appear only for non-administrator users.

Outlook

Across H2 2026, the most persistent risk will come from threat actors using legitimate tools, trusted platforms, and routine workflows to reduce detection opportunities while preserving operational flexibility. Organizations should expect malicious activity to increasingly blend into normal endpoint, identity, cloud, developer, mobile, and payment activity, making behavioral correlation and defense in depth more important than reliance on single indicators or malware-specific IoCs.

Threat actors will likely continue to favor vulnerabilities that are remotely accessible, require little or no authentication, and offer clear operational value. AI is unlikely to change those fundamentals, but it may shorten the path from vulnerability discovery to exploit development, reducing the time defenders have to validate exposure and deploy mitigations. Rising vulnerability discovery and submission volumes associated with AI-assisted research suggest that this pressure is already emerging.

Malware operators will likely continue to rely on RATs, stealware, loaders, and backdoors because these tools provide reliable capabilities for remote access, credential theft, staged payload delivery, and monetization. Ransomware operators will likely continue pairing incremental payload improvements with established methods such as ClickFix lures, remote access tooling, backup targeting, and anti-recovery techniques.

Supply-chain compromises will likewise remain a major concern as package managers, developer tooling, CI/CD systems, and code repositories provide access to sensitive credentials and downstream propagation paths. As AI-enabled coding tools gain broader permissions across these environments, threat actors will likely increasingly target their credentials, configurations, and trusted execution paths.

AI-enabled cyber threats will likely develop along two tracks. AI-enabled malware will primarily augment existing tradecraft through functions such as persistence, reconnaissance, UI interpretation, payload adaptation, and post-compromise tasking rather than replace it with fully autonomous attack execution. Agentic AI poses a separate emerging risk, with early H2 2026 reporting on the Hugging Face incident representing an AIM3 Level 5 event in which an autonomous agent coordinated actions, identified vulnerabilities, sought internet access, and operated across systems in a testing environment. However, reliable use in in-the-wild attacks will likely remain constrained by model access, reliability, and infrastructure over the next six to twelve months. Given those constraints, state-sponsored threat actors may be best positioned to experiment with autonomous attack chains. Improvements in open-source models could gradually lower those barriers for less-resourced adversaries.

Defenders should prioritize high-risk exposures, protect developer and CI/CD secrets, constrain privileges granted to AI-enabled tools, and maintain isolated, regularly tested recovery systems.

Recorded Future reporting contains expressions of likelihood or probability consistent with US Intelligence Community Directive (ICD) 203: Analytic Standards (published January 2, 2015). Recorded Future reporting also uses confidence level standards employed by the US Intelligence Community to assess the quality and quantity of the source information supporting our analytic judgments.

About Insikt Group®

Recorded Future's Insikt Group, the company's threat research division, comprises analysts and security researchers with deep government, law enforcement, military, and intelligence agency experience. Their mission is to produce intelligence that reduces risk for customers, enables tangible outcomes, and prevents business disruption.

About Recorded Future®

Recorded Future is the world's largest intelligence company. The Recorded Future Intelligence Operations Platform provides the most complete coverage across adversaries, infrastructure, and targets. By combining precise, AI-driven analytics with the Intelligence Graph® populated by specialized threat data, Recorded Future enables cyber teams to see the complete picture, act with confidence, and get ahead of threats that matter before they impact your business. Headquartered in Boston with offices around the world, Recorded Future works with more than 1,900 businesses and government organizations across 80 countries.

Learn more at recordedfuture.com