vega

2026

Quarterly Threat Report

Vega Threat Research

Table Of Contents

Executive Summary 3

Threat Landscape 4

Notes From The Trenches 5

Detections We Loved This Quarter 7

Summary 9

Executive Summary

Over the past quarter, we observed threat actors continuing to favor techniques that are **fast to deploy, easy to adapt, and difficult for traditional prevention controls to consistently stop.** Rather than relying on sophisticated malware or complex intrusion chains, attackers increasingly combined **social engineering, native system utilities, rapid exploitation of one-day vulnerabilities, security control tampering, and evasion infrastructure** to gain execution and move quickly after initial access.

A common theme across these observations is the **operational efficiency of modern attacks.** Many of the techniques we observed were not highly complex, but they were effective because they exploited **gaps between user behavior, endpoint visibility, vulnerability exposure, cloud identity controls, and automated analysis.** This reflects a threat environment where attackers increasingly advance by abusing what organizations already trust: workflows, credentials, exposed systems, and administrative tools.

Threat Landscape

In Q2, the threat landscape was shaped by five recurring themes:

- Exploitation of newly disclosed vulnerabilities
- User-assisted initial access
- Ransomware and data extortion
- Cloud and identity abuse
- Software supply chain compromise

Across these areas, attackers relied on exposed software, trusted identities, third-party services, user trust, and visibility gaps across cloud, SaaS, identity, and CI/CD environments.

Vulnerability exploitation

Vega Research observed continued exploitation of exposed enterprise applications soon after disclosure. Industry reporting also highlighted exploitation affecting browsers, operating system components, and other internet-facing infrastructure.

Initial Access Through ClickFix

ClickFix remained a common initial-access technique, observed across more than half of Vega customer environments. Attackers use fake CAPTCHAs, software updates, rogue AI Skills to trick users into running commands directly, often abusing trusted utilities to bypass controls focused on malicious downloads and email attachments.

Cloud and identity abuse

Reported activity included misuse of long-lived tokens, OAuth grants, session cookies, SaaS credentials, CI/CD secrets, and cloud permissions. With valid access, attacker activity can resemble legitimate behavior and evade endpoint-focused detections.

Ransomware and disruption

Groups including Qilin, The Gentlemen, and Akira combined data theft, encryption, extortion, EDR tampering, and attacks on backups, identity systems, and administrative tooling. Ransomware remained a threat to both data protection and business continuity.

Network edge targeting

Firewalls, VPNs, and remote-access platforms continued to attract threat actors because they are internet-facing, trusted, and often lack endpoint-style telemetry. Activity involving PAN-OS and Fortinet flaws highlighted continued focus on security appliances.

Software supply chain compromise

Campaigns across npm, GitHub, and CI/CD environments showed how malicious dependencies can steal developer secrets, tokens, environment variables, cryptographic keys, crypto wallets. Asurion-impersonating npm packages, Megalodon and Mini Shai-Hulud showed how compromise can begin inside normal developer workflows and spread through trusted automation

Explore Vega Threat Hub to stay up to date on the latest incidents, adversary tools and techniques, detections, and IOCs.

Notes From The Trenches

During Q2, Vega observed user-assisted malware execution, developer workflow compromise, cloud and identity-plane abuse, rapid vulnerability exploitation, security-control tampering, evasion infrastructure, and anti-forensic activity around sensitive data. Across these incidents, attackers repeatedly abused trusted workflows, valid access paths, and visibility gaps between endpoint, cloud, identity, developer, and data environments.

1

ClickFix and user-assisted execution

ClickFix was one of the most common patterns Vega observed, appearing in more than 50% of customer environments. Attackers used lures such as fake CAPTCHAs, fake macOS updates, fake Claude-related content, browser-fix prompts, and software-installation themes to convince users to run commands in trusted local tools such as Terminal, Windows Run, Script Editor, or shell sessions.

The technique appeared across macOS and Windows environments and used native tools including shell interpreters, scripting utilities, command-line downloaders, finger.exe, and wscript.exe. In several cases, one of which was documented by Vega Research, the compromise path shifted from conventional browser or email download activity into local command execution performed by the user.

2

Evasion infrastructure and selective lure delivery

Several social-engineering-driven cases involved Traffic Distribution Systems used to control when, where, and to whom malicious content was delivered. This made some lures harder for automated analysis and sandboxing tools to capture, with malicious infrastructure only becoming visible through manual investigation.

3

Developer workflow and supply chain compromise

Vega observed activity linked to compromised npm packages, including suspicious infrastructure communication, payload retrieval, and execution on at least one endpoint. The activity began inside routine package-management workflows rather than through an obvious phishing lure or standalone exploit, making it harder to separate malicious behavior from normal developer activity.

The risk extended beyond the affected endpoint. Depending on user and system context, a compromised developer workflow can expose developer secrets, access tokens, repositories, package-publishing permissions, CI/CD credentials, and cloud credentials.

4

Cloud, identity, and control-plane abuse

Vega observed suspicious administrative activity across cloud and identity environments, including role assumption, object access, secret retrieval, decryption operations, identity creation, and privilege changes. Because the activity used valid APIs, SSO-linked roles, and privileged access paths, it required context to distinguish it from legitimate administration.

The strongest signals came from unusual access patterns, unfamiliar locations, abnormal privilege use, high-volume secret access, identity changes, and DNS resolutions to threat-intelligence-validated command-and-control infrastructure. These cases reinforced the need to connect identity behavior, DNS activity, workload telemetry, and threat intelligence enrichment with cloud control-plane logs.

5

Rapid vulnerability exploitation

Vega observed active exploitation of CVE-2026-22679, a critical unauthenticated RCE in Weaver E-cology, only days after the vendor patch was released. The case involved RCE validation, payload delivery attempts, PowerShell-based retrieval, discovery commands, and process-name evasion using a renamed PowerShell binary.

This reinforced the narrow response window defenders face once technical details become available, and showed how quickly exploitation can move from disclosure to customer impact.

6

Security-control tampering and credential-access enablement

Vega observed attacker behavior aimed at weakening security controls and enabling follow-on access. In one case, threat actors enabled WDigest, a configuration change that can expose cleartext credentials in memory and support credential collection, privilege escalation, and lateral movement without relying on a highly visible malware payload.

7

Suspicious deletion and evidence removal

Vega observed secure-deletion activity in a shared file environment containing sensitive verification-related data. Multiple accounts with little or no prior activity on the affected file server generated unusually high file-operation volume, including overwrite and delete patterns consistent with anti-forensic tooling.

This activity was notable because the impact was not limited to possible data access. Overwrite and deletion activity can remove evidence, disrupt investigation, affect recovery, and limit the organization's ability to reconstruct what happened.

Detections We Loved This Quarter

Q2 highlighted sophisticated attack patterns across cloud, identity, and developer environments. Here are the most valuable behavioral signals we tracked this quarter:

Endpoint & Ransomware

ClickFix Win+R Registry Entries → Captures victims running malicious commands via the `Win+R` prompt. Because Windows logs entries in the `RunMRU` registry key before execution, it provides an ultra-early warning before a process even spawns.

SHub Stealer C2 Beacon → Spotlights macOS AMOS/SHub infostealers by flagging `curl` processes that transmit specific telemetry event names (e.g., `password_obtained`, `payload_started`) alongside a campaign build hash.

Forensic Artifact Deletion → Targets attackers trying to cover their tracks by detecting the forced, quiet deletion (`del /f /q`) of high-value incident response targets: Windows Prefetch files, Defender logs, and RDP session histories.

Volume Shadow Copy Deletion → A dead-giveaway ransomware signal. It flags `vssadmin.exe` clearing out local backups (`/delete /shadows /all /quiet`) to prevent easy recovery before the encryptor runs.

Cloud & Infrastructure

EC2 Instance Creating an IAM User → EC2 instances have no legitimate reason to create IAM users. This detects a compromised instance using its profile role to establish a persistent IAM identity that survives instance termination.

High-Volume Secrets Harvesting → Triggers when a single AWS identity requests 20 or more secrets within 15 minutes—a clear pattern of automated credential harvesting rather than normal application behavior.

Identity & Session Hijacking

Okta Session Cookie Reuse → Catches session hijacking by identifying a single Okta device token (`dthash`) appearing across multiple distinct IPs, ASNs, and user agents within a tight 5-minute window.

PIM Activation + Sensitive Operation → Flags an identity that activates an elevated Entra ID role and immediately performs a high-risk action (like weakening Conditional Access or adding service principal credentials) within 15 minutes.

Dev Environment & Supply Chain

Bot Identity Publishing GitHub Releases → Automated bot tokens rarely have a valid reason to publish releases. This alerts on compromised CI/CD pipelines where an attacker uses the inherited `GITHUB_TOKEN` to push malicious code.

Token Revocation Directory Wipe → A destructive supply chain "dead man's switch." Implants monitor if a stolen token has been revoked, and if so, immediately execute an `rm -rf` on the developer's home directory to destroy evidence and maximize damage.

Summary

Across Q2 investigations, **incidents observed by Vega manifested across multiple data sources,** combining social engineering, command execution, developer workflow abuse, cloud access, identity compromise, DNS activity, and sensitive data operations. Modern incidents are therefore best understood as **connected attack paths,** not standalone endpoint, cloud, identity, network, or data-access events.

A recurring pattern was the **abuse of legitimate tools and trusted workflows.** ClickFix relied on users executing commands locally, the npm supply-chain case abused standard developer workflows, and cloud activity used valid SSO, IAM, API, and administrative mechanisms. Individual actions often appeared normal in isolation; **sequence, timing, user context, and downstream behavior** exposed the malicious activity.

Identity was a consistent thread across these investigations, including credential access, developer secret exposure, cloud role abuse, privileged changes, and suspicious use of legitimate accounts. Identity was often the link **between endpoint activity, developer workflows, cloud access, and sensitive data operations.**

Cross-domain correlation was central to understanding these investigations. Decisive evidence often came from relationships between endpoint, cloud, identity, DNS, developer, and file activity rather than from a single indicator. This context helped **separate compromise from benign activity, identify evidence gaps, and clarify risk and response priorities.**