carahsoft

CROWDSTRIKE

2026 Threat Hunting Report

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through GSA Multiple Award Schedule (MAS), NASPO ValuePoint, The Quilt and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with CrowdStrike, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.

To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com

Explore More Resources:
carah.io/crowdstrikeresources

Join Events & Webinars:
carah.io/crowdstrikeevents

Discover Technology Solutions:
carah.io/crowdstrike

Learn About Procurement:
carah.io/crowdstrikecontracts

Connect With Our Team:
CrowdStrike@carahsoft.com
888-662-2724

CROWDSTRIKE

2026 THREAT HUNTING REPORT

Table of Contents

Introduction 3

Frontline Snapshot 7

- Intrusion Trends by Adversary 8
- Sector Targeting 11
- MITRE ATT&CK and MITRE ATLAS 14

Observations from the Front Lines 20

- Threat Hunting in the AI Era: Faster Vulnerability Weaponization and AI-Assisted Vulnerability Discovery 20
- Software Supply Chain Attacks: Defending Across a Cross-Domain Threat Landscape 24
- Vishing Surges as an Identity-Based Initial Access Vector 31
- Cloud Beyond the Perimeter: Cloud-Conscious Adversaries and Financial Targeting 36
- OVERCAST PANDA: Close Access Operations and the Endpoint Threat Beyond Network Defenses 43

Conclusion 47

Recommendations 48

CrowdStrike Falcon Platform, Products, and Services 50

Introduction

In 2025, adversaries prioritized evasion, as highlighted in the CrowdStrike 2026 Global Threat Report. Instead of forcing their way through the front door, they systematically undermined organizational trust by targeting supply chain partners, upstream developer ecosystems, legitimate software, and employees to gain initial access and move undetected. This tactical pivot firmly established 2025 as the year of the evasive adversary.

Throughout 2026, we have seen these techniques across trusted access paths organizations rely on every day. Adversaries understand the gaps between fragmented security controls and increasingly exploit areas where defenders often lack full visibility: valid credentials, cloud authentication flows, SaaS applications, edge devices, and unmanaged systems. They are also using AI to accelerate phishing, reconnaissance, and technical operations and targeting the AI systems now embedded across the enterprise. Frontier AI is collapsing the window between vulnerability discovery and exploitation, helping adversaries identify vulnerabilities, generate proof-of-concept (PoC) exploits, and map attack paths at increasing speed and scale. By chaining activity across identity, cloud, SaaS, endpoint, and perimeter infrastructure, adversaries can establish footholds, escalate privileges, exfiltrate data, or launch disruptive attacks before defenders can connect the signals and respond.

To keep pace with the evasive adversary, defenders need every advantage: faster analysis, richer context, and the ability to act before an intrusion escalates. Innovation remains a critical cornerstone to outmaneuver and disrupt rapidly evolving threats. AI-powered technologies, threat intelligence, and novel threat hunting techniques are required to anticipate the adversary's next moves, understand their evolving methodologies, and adapt defenses to stay ahead.

At the center of the fight is CrowdStrike Counter Adversary Operations, which unifies threat intelligence, managed threat hunting, and trillions of telemetry events daily from the AI-powered CrowdStrike Falcon® platform to detect, disrupt, and stop modern adversaries. Counter Adversary Operations comprises two closely integrated teams. The CrowdStrike Intelligence team identifies new adversaries, tracks malicious activity, and captures emerging cyber threat developments in real time. The CrowdStrike OverWatch team applies this intelligence through proactive threat hunting across customer telemetry to detect and address malicious activity, leveraging AI to hunt at a scale manual investigation cannot reach. Together, these teams help protect organizations from sophisticated adversaries by delivering intelligence and threat hunting capabilities that most organizations cannot replicate internally.

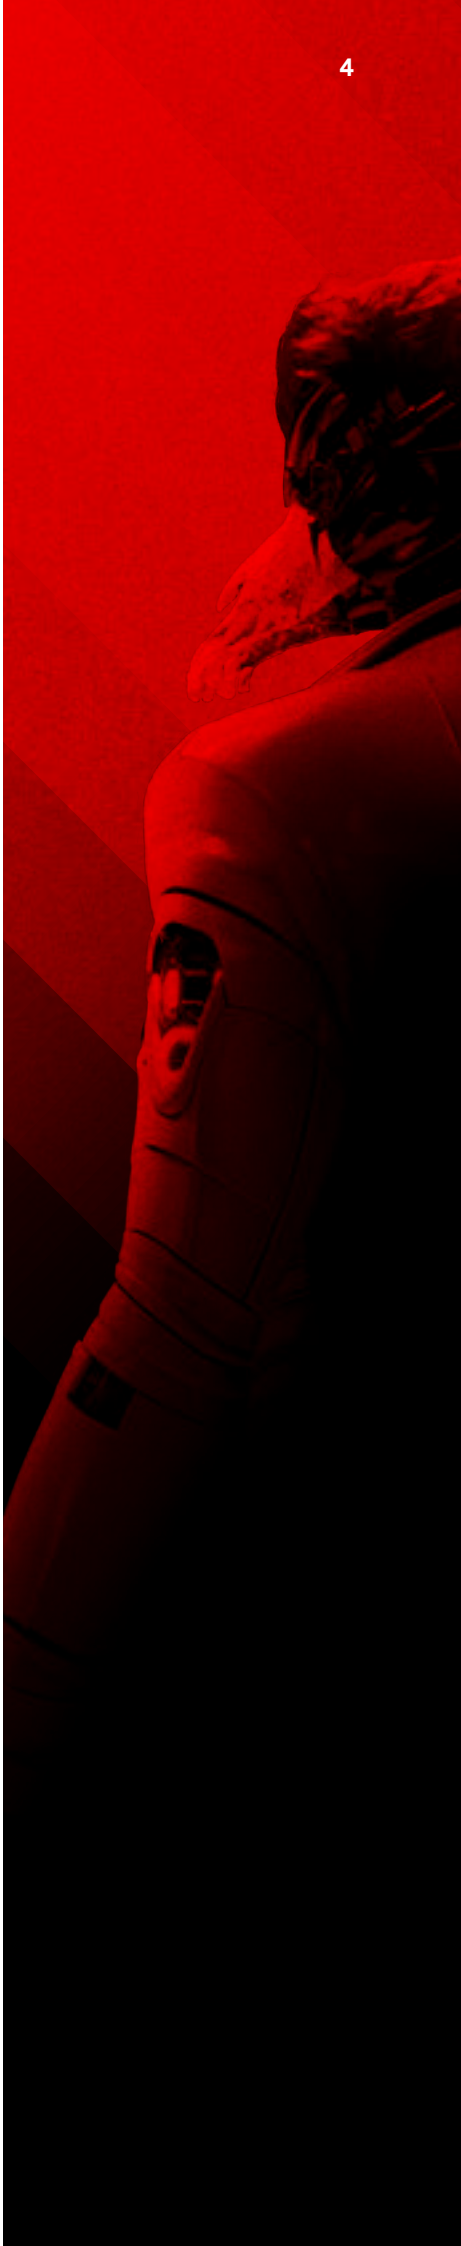

The CrowdStrike 2026 Threat Hunting Report builds on this global threat landscape with frontline hunting observations from CrowdStrike OverWatch. Drawing on activity observed over the past 12 months, from July 1, 2025, to June 30, 2026, the report shows how adversaries are evolving their tradecraft to exploit trust, target visibility gaps between disconnected security controls, and move at speeds that leave defenders little time to respond. Disrupting these attacks requires looking past isolated events to execute continuous, intelligence-driven threat hunting across every domain.

AI is now a tool, a target, and a force multiplier for adversaries. The CrowdStrike 2026 Global Threat Report showed that AI-enabled adversary activity surged **89%** in 2025 as attackers used AI to scale operations, accelerate tradecraft, and directly target AI infrastructure. CrowdStrike OverWatch has observed significant growth in AI agent-triggered detection leads, now tracking at **2.5x** the rate of human-triggered leads, underscoring how AI-driven activity is increasing the volume and velocity of signals threat hunters must assess. Our threat hunters have seen this shift firsthand. FAMOUS CHOLLIMA, associated with the Democratic People's Republic of Korea (DPRK), weaponized trusted AI-centric development environments to compromise cryptocurrency and blockchain companies. Financially motivated eCrime threat actors abused corporate large language model (LLM) access for LLMJacking and cost harvesting, including one campaign that sent nearly **200,000** API requests in an initial two-minute flood. Other adversaries exploited vulnerabilities in AI-related server software to mine cryptocurrency and harvest sensitive configuration information. These incidents reveal a challenging reality: The same AI tools driving modern businesses are creating underdefended attack surfaces that adversaries are already exploiting.

From January 2026 to June 2026, **88%** of CrowdStrike-observed exploitation of vulnerabilities with a public PoC was conducted within **48 hours** of the PoC's release. China-nexus adversaries moved even faster, with VAULT PANDA and GENESIS PANDA launching deliberate attacks within **24 hours** of public disclosure of a critical web application vulnerability. Frontier AI systems are likely to further compress these timelines by accelerating vulnerability discovery and exploit development, increasing pressure on defenders already struggling to keep pace.

Adversaries are also exploiting trust across the developer ecosystem. They are moving into CI/CD pipelines, container registries, package registries, and integrated development environment (IDE) extensions, where one compromised dependency or trusted component can rapidly propagate across downstream environments. The Node Package Manager (npm)[1] package ecosystem remains the primary vector for supply chain attacks, accounting for **87%** of all malicious software registry threats in the first half of 2026. Over the past year, STARDUST CHOLLIMA and ALTERED SPIDER drove some of the most consequential software supply chain attacks, showing how nation-state and financially motivated threat actors are targeting this attack surface. ALTERED SPIDER demonstrated the scale of the threat by compromising more than **300** software dependencies in a single day, harvesting credentials, and pivoting into cloud environments. As software development and AI workflows converge, the AI ecosystem is becoming the next software supply chain battleground.

1 npm is a package manager for the JavaScript runtime environment Node.js, with its public registry operated by GitHub.

Trust is also being exploited at the human layer. Voice phishing (vishing) has become a fast-moving identity-based initial access vector. CrowdStrike OverWatch detected a **2x** increase in vishing intrusions in the first half of 2026 compared with the second half of 2025. Human-to-human interaction makes vishing especially compelling as a social engineering technique, enabling adversaries to impersonate IT staff and manipulate users into granting access. CORDIAL SPIDER and SNARKY SPIDER used vishing to compromise single sign-on (SSO) accounts and rapidly exfiltrate data from SaaS applications, with SNARKY SPIDER moving from account takeover to data theft in under **five minutes**. By compromising identities, these adversaries can rapidly access sensitive data without moving laterally or escalating privileges.

Once adversaries gain trusted access, cloud environments offer a direct path to monetization. Within the reporting period, eCrime cloud-conscious activity surged **171%** as financially motivated threat actors targeted cloud environments for credential theft, cryptomining, LLM abuse, and access to digital financial assets. Traditional perimeter defenses do not apply when adversaries operate with valid credentials or tokens inside the cloud trust boundary. Over the last six months, there has been a **15x** increase in monthly device code phishing[2] attempts, as eCrime threat actors have rapidly scaled this technique to compromise cloud identities by abusing trusted authentication flows. SLIM SPIDER showed the sophistication of this shift in a multi-stage intrusion targeting cryptocurrency assets and instant payment infrastructure, while other eCrime threat actors hijacked cloud resources[3] for cryptomining at the victim's expense. As organizations move critical infrastructure and financial operations to the cloud, adversaries are following the money and refining cloud-specific tradecraft to stay hidden.

This report showcases the Counter Adversary Operations team's relentless pursuit to disrupt the adversary. In CrowdStrike customer environments, adversaries face a unified security solution that empowers every CrowdStrike threat hunter with extensive security telemetry spanning endpoint, identity, cloud, and next-gen security information and event management (SIEM), along with integrated intelligence, to disrupt modern threats.

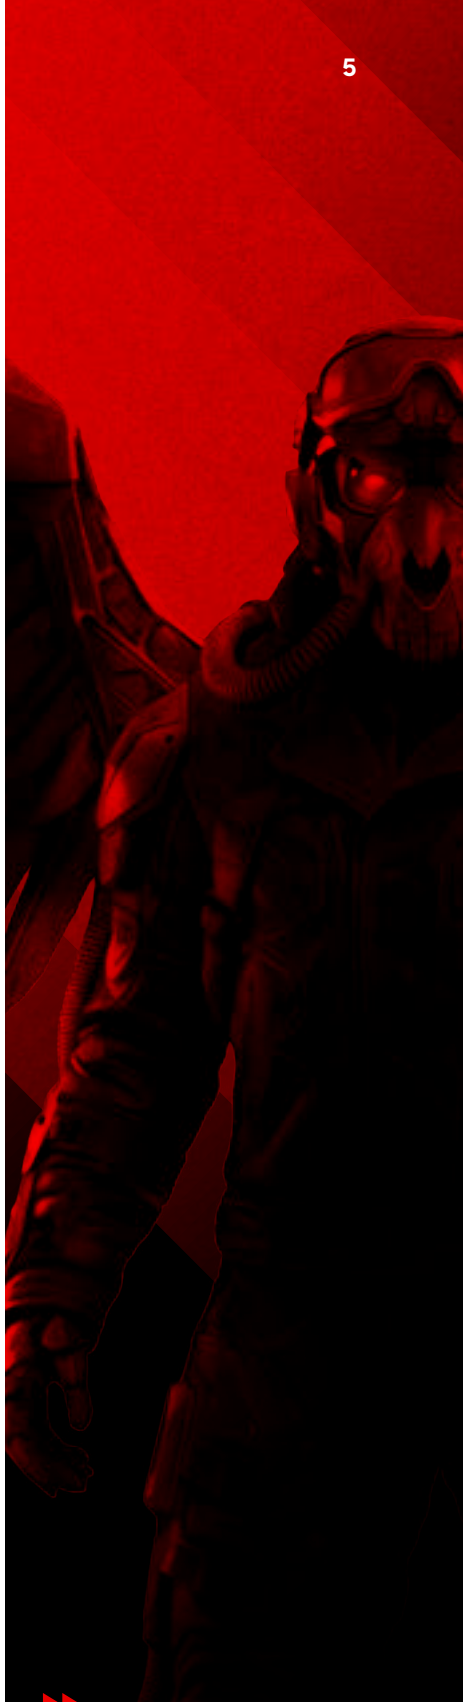

» EXPLORE THE CROWDSTRIKE ADVERSARY HUB FOR THE LATEST INSIGHTS ON ADVERSARIES, TRADECRAFT, AND ACTIVITY.

2 Device code phishing is an attack technique where adversaries abuse the OAuth 2.0 device authorization flow by tricking victims into entering attacker-controlled device codes on legitimate authentication portals, effectively granting the attacker access tokens to cloud services and applications. Unlike traditional phishing, it bypasses MFA since the victim authenticates legitimately; the attacker simply hijacks the resulting token.

3 Cloud resource hijacking refers to the unauthorized use of a victim's cloud computing resources by an adversary following account compromise. Rather than targeting data, adversaries exploit the victim's cloud infrastructure to run computationally intensive operations, most commonly cryptocurrency mining or LLM abuse, at the victim's expense.

ADVERSARY NAMING CONVENTIONS

BEAR
RUSSIA

BISON
BELARUS

BUFFALO
VIETNAM

CHOLLIMA
DPRK (NORTH KOREA)

CRANE
ROK (REPUBLIC OF KOREA)

HAWK
SYRIA

JACKAL
HACKTIVIST

KITTEN
IRAN

LEOPARD
PAKISTAN

LYNX
GEORGIA

OCELOT
COLOMBIA

PANDA
PEOPLE'S REPUBLIC OF CHINA

SAIGA
KAZAKHSTAN

SPHINX
EGYPT

SPIDER
eCRIME

TIGER
INDIA

WOLF
TÜRKIYE

Frontline Snapshot

The CrowdStrike Threat Hunting Report has historically focused on a single category of threat: interactive, hands-on-keyboard intrusions, where adversaries establish an active presence within a target environment and manually execute actions to achieve their objectives.

Adversaries are now exploiting vulnerabilities at unprecedented speed, weaponizing automated techniques to launch large-scale software supply chain attacks, and using AI to sharpen their precision and efficiency. CrowdStrike OverWatch has evolved this report in response. We intentionally expanded the scope of the analysis herein to reflect the modern threat in its entirety: interactive intrusions and automated attacks executing in tandem. Examining these elements together as a single, cohesive threat gives defenders a comprehensive picture of the modern threat landscape.

» SOPHISTICATED THREAT ACTORS ARE CHANGING THE GAME BY BLENDING SCALABLE AUTOMATION WITH TRADITIONAL HANDS-ON TACTICS TO MAXIMIZE THEIR IMPACT.

Adversaries no longer need to rely purely on hands-on-keyboard activity to achieve large-scale impact while evading detection. Sophisticated threat actors are changing the game by blending scalable automation with traditional hands-on tactics to maximize their impact. This trend spans both nation-state groups and financially motivated eCrime operations.

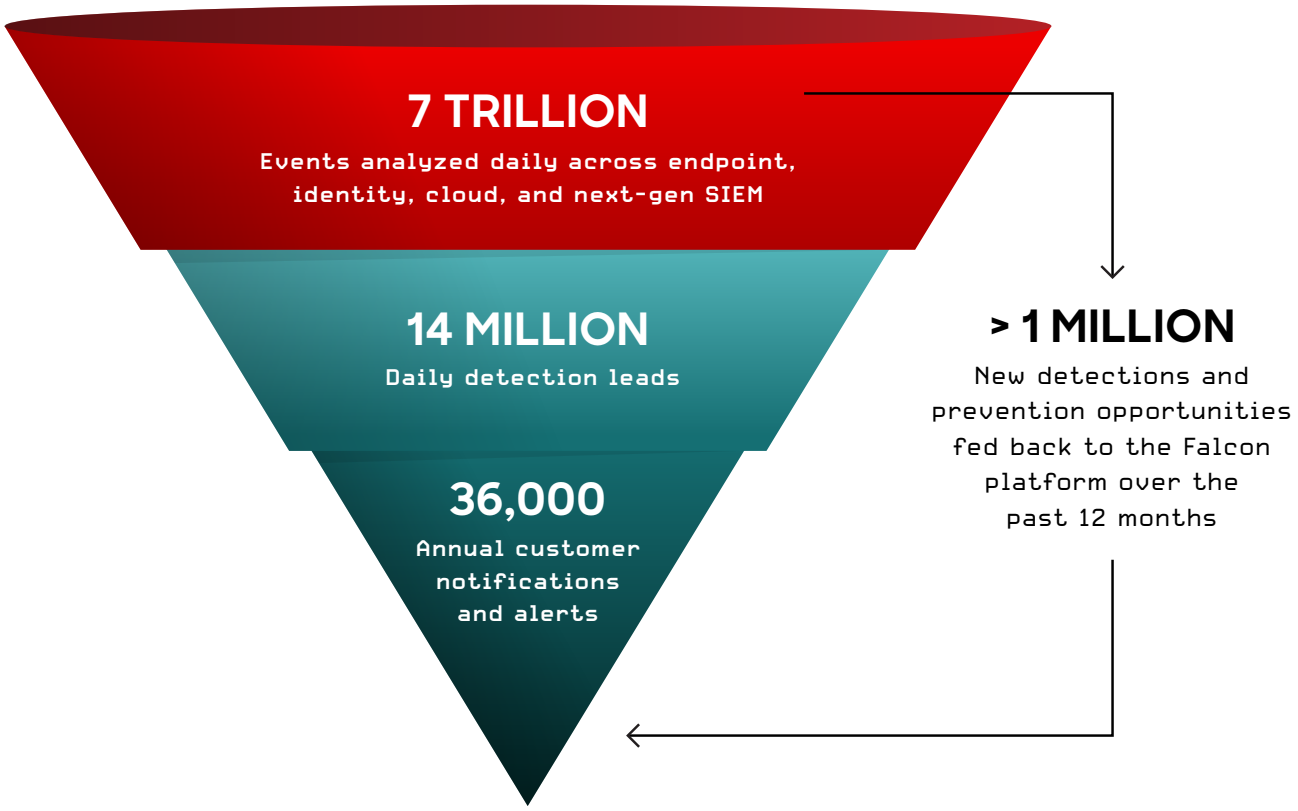

Figure 1. CrowdStrike OverWatch's threat hunting model

CrowdStrike OverWatch serves as a cornerstone of the CrowdStrike security ecosystem. By analyzing 7 trillion events daily, the CrowdStrike OverWatch team sifts through massive data volumes to surface more than 14 million daily detection leads for analysis. This continuous hunting annually produces more than 36,000 high-fidelity customer notifications and alerts. Each alert represents a pivotal moment where hunters identified potential malicious activity, notified the customer, and delivered actionable guidance to contain the threat quickly.

AI provides the scale necessary to hunt across this vast volume of activity, which is at a level that renders manual human-driven review unfeasible. These automated systems operate continuously, aiding in the categorization, identification, and prioritization of notable data streams, helping threat hunters to quickly triage activity and determine if additional investigation is warranted. This approach creates an operational synergy and effectively unifies the rapid processing power of machine-scale analytics with the nuance of expert decision-making.

Threat hunters drive a powerful protective cycle. CrowdStrike OverWatch findings feed directly back into the Falcon platform as new detections and preventions, creating a continuous feedback loop that evolves in real time and is designed to slow or stop the malicious activity hunters encounter. Over the past 12 months, this loop produced more than 1 million new detections and potential preventions, continuously hardening the Falcon platform and freeing threat hunters to focus on uncovering new threats.

Intrusion Trends by Adversary

Threat hunting is most effective when it is fueled by deep visibility into adversary behavior. Counter Adversary Operations combines professional managed threat hunting with best-in-class threat intelligence to ensure hunters always understand the adversary's motivations and behavior.

CrowdStrike maintains detailed intelligence on more than 290 attributed eCrime, nation-state, and hacktivist adversaries, alongside more than 150 active clusters of malicious activity that have not yet met CrowdStrike's standards to be tracked as a named adversary. This unmatched visibility gives threat hunters the valuable context that enables them to anticipate adversarial movements and keep pace with evolving tradecraft. In a landscape where automation increasingly obscures the human behind the attack, knowing the adversary is everything.

Figure 2 on the following page highlights the most prevalent adversary operations identified over the past 12 months.

10+

NEW ADVERSARIES NAMED IN 2026,
INCLUDING eCRIME ADVERSARY ALTERED SPIDER

290+

TOTAL ADVERSARIES NOW TRACKED BY CROWDSTRIKE

150+

ACTIVE MALICIOUS ACTIVITY CLUSTERS AND
EMERGING THREAT GROUPS TRACKED

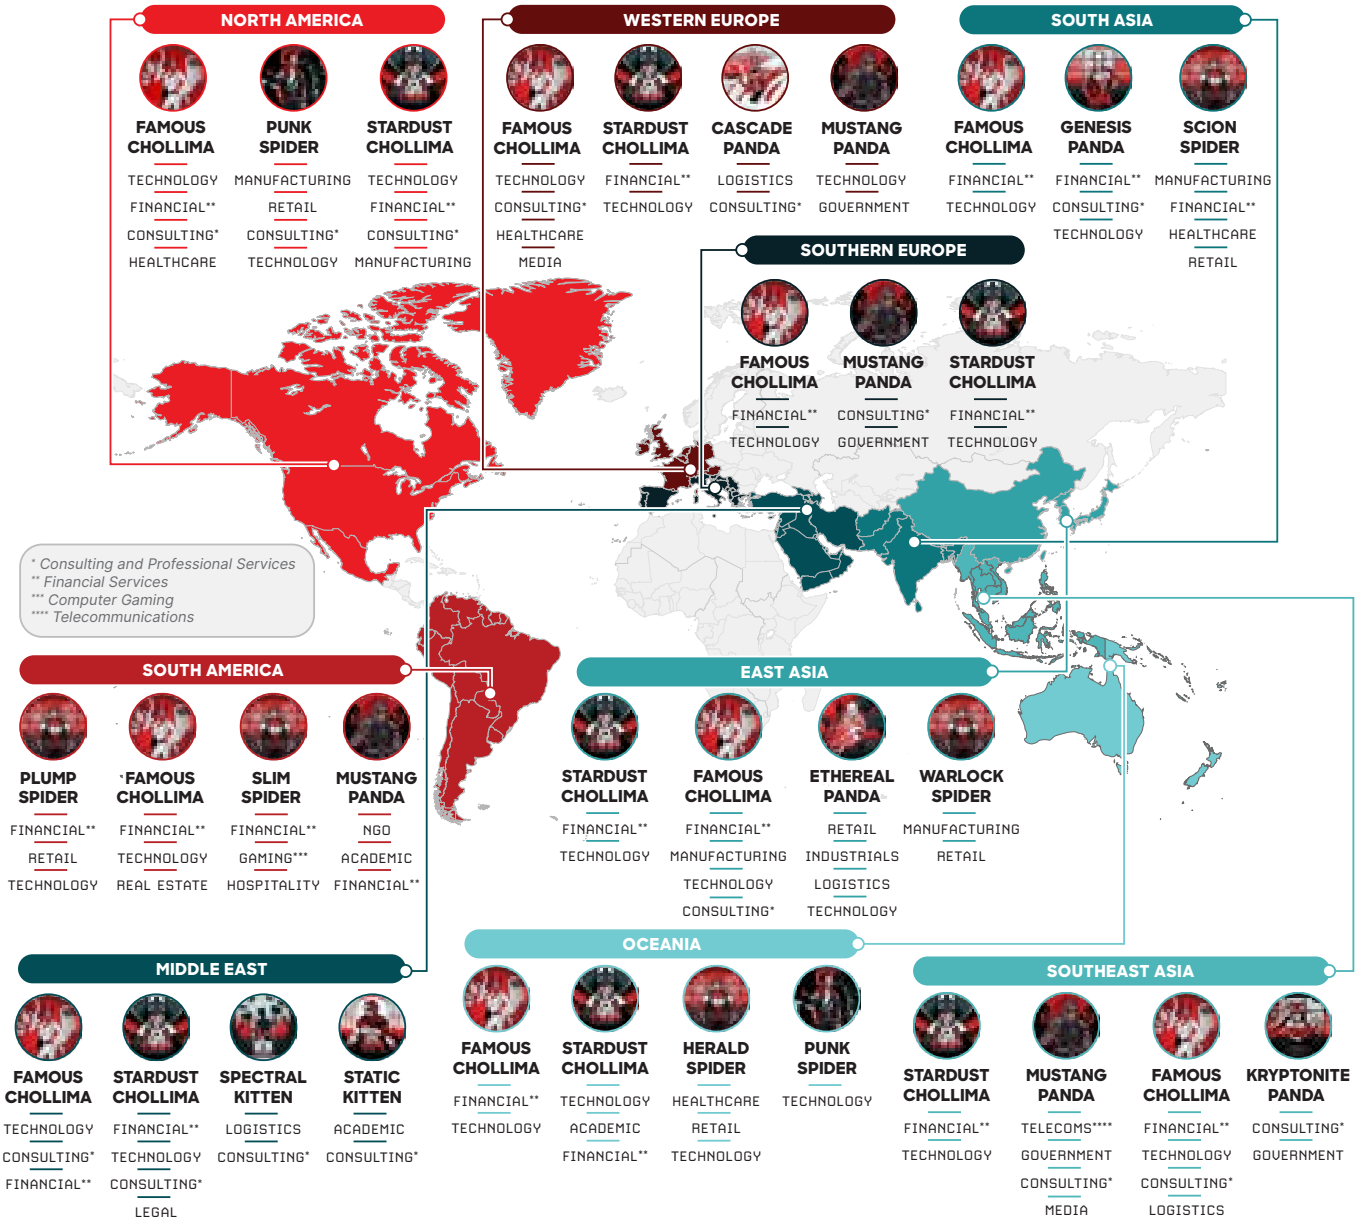

Figure 2. Adversary disruptions across the world, July 1, 2025, to June 30, 2026

This year's report expands its analytical scope beyond purely interactive intrusions to include automated attacks. This methodology change captures a more accurate and comprehensive view of evolving adversary behavior and the modern threat landscape that defenders face.

This year's data reveals an approximate 4% increase in overall intrusion activity, a stark contrast from the 27% year-over-year surge observed in last year's report. This plateau is expected and does not indicate a drop in adversary intent or a reduction in operational risk.

Instead, the data reflects a maturing threat landscape. Adversaries are investing greater time and resources into developing complex campaigns with increasingly innovative initial access points rather than the opportunistic, high-volume brute-force campaigns that characterized earlier years.

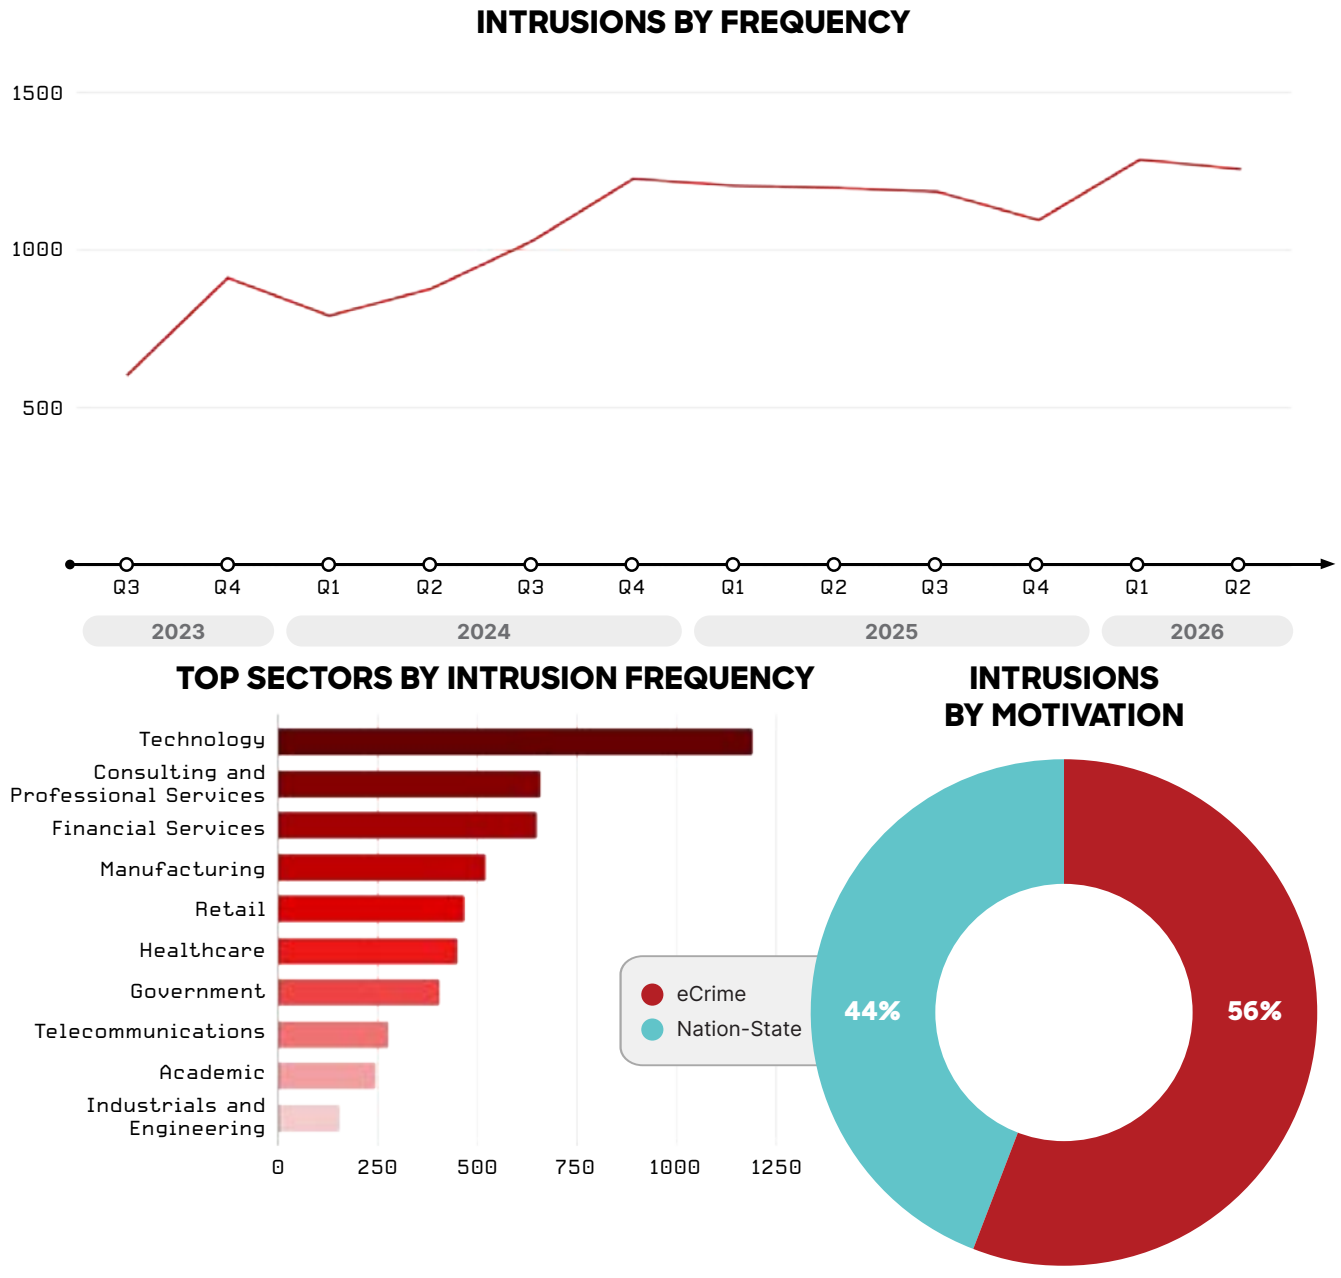

Figure 3. Intrusion breakdown, July 1, 2025, to June 30, 2026

Sector Targeting

Across industries, intrusion activity largely mirrors the broader trend established earlier in this report, marking a stabilization in overall volume that reflects adversary maturation rather than a reduced threat environment.

Technology remains the most targeted sector for the ninth consecutive year, a distinction that speaks to the high-value data and critical infrastructure it represents. The top five sectors shift notably when broken down by adversary type. Nation-state threat actors prioritize technology, financial services, and government entities, while eCrime adversaries concentrate on technology, consulting and professional services, and manufacturing entities. This divergence reflects their differing motivations, whether it's extracting sensitive intelligence or targeting industries rich in financial data and consumer information.

Sector rankings remained largely stable year-over-year, with most industries seeing only modest fluctuations in intrusion activity. The two most notable exceptions were the financial services and academic sectors, which recorded the largest increases of any sector at 11% and 17%, respectively, signaling that adversaries are increasingly targeting institutions where financial assets and sensitive research data converge.

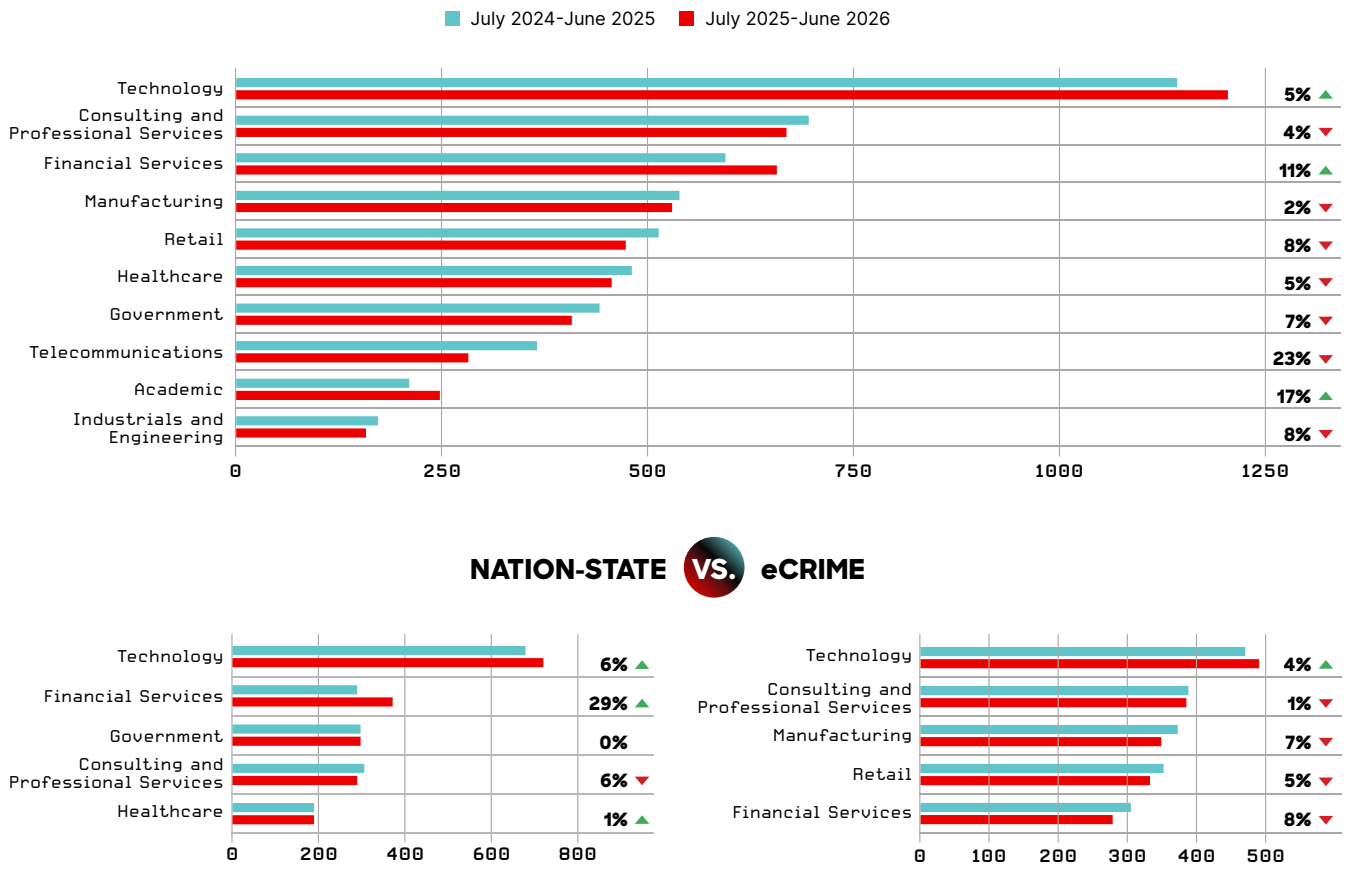

Figure 4. Targeted sectors by intrusion frequency, July 2024-June 2025 vs. July 2025-June 2026

Sector Spotlights

Technology

The technology sector retained its position as the most targeted industry for the ninth consecutive year, recording a 5% increase in intrusion activity. This represents modest growth consistent with the broader stabilization in overall intrusion activity observed throughout this reporting period.

Technology organizations hold vast repositories of intellectual property and proprietary data, and they serve as gateways to downstream software supply chain compromises. The sector also sits at the center of trusted digital ecosystems (spanning identities, software, cloud services, and developer workflows) that adversaries exploit to scale access and impact. This makes technology entities prime targets for adversaries of all motivations. They were targeted by both eCrime and nation-state adversaries in near-equal measure, with 59% of intrusions attributed to nation-state threat actors and 41% to eCrime threat actors.

DPRK-nexus adversaries posed a substantial nation-state threat, with FAMOUS CHOLLIMA conducting extensive IT worker infiltration operations across technology entities in North America, Europe, and Asia. FAMOUS CHOLLIMA's operations alone accounted for 55% of all nation-state intrusions and 44% of all intrusions targeting the technology sector during the reporting period, making them the single most active adversary targeting the sector. STARDUST CHOLLIMA contributed further to this DPRK-nexus activity.

The reliance on continuous availability also makes these organizations highly attractive targets for eCrime adversaries conducting ransomware operations, where disruption carries immediate operational consequences. PUNK SPIDER, MUTANT SPIDER, and HERALD SPIDER rounded out the top five adversaries targeting technology entities, driven by the clear financial incentives inherent to a sector rich in valuable data and assets.

Almost 54% of all technology sector targeting impacted North American institutions, reflecting the region's position as a global hub for technology development with a dense concentration of high-value technology firms, intellectual property, and digital infrastructure.

Financial Services

Financial institutions saw an 11% increase in intrusion activity this year compared to last year. The financial sector represents a dual-value target because these institutions hold significant financial assets, including cryptocurrency holdings, alongside high-value data such as business intelligence and customers' personally identifiable information (PII). For eCrime adversaries, the high availability requirements of financial operations make the sector a target for ransomware attacks, where operational disruption carries immediate and severe consequences.

FAMOUS CHOLLIMA'S OPERATIONS ALONE ACCOUNTED FOR 55% OF ALL NATION-STATE INTRUSIONS AND 44% OF ALL INTRUSIONS TARGETING THE TECHNOLOGY SECTOR DURING THE REPORTING PERIOD, MAKING THEM THE SINGLE MOST ACTIVE ADVERSARY TARGETING THE SECTOR.

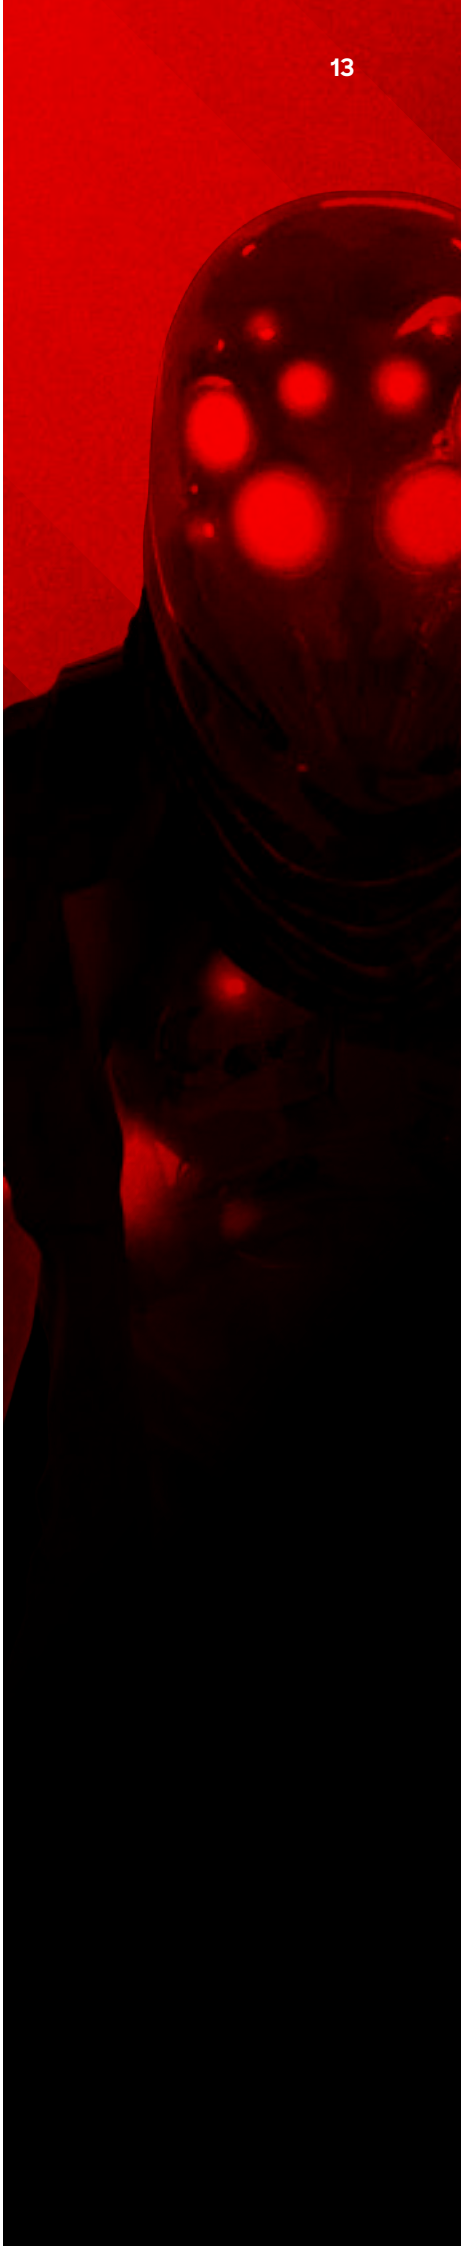

This duality is reflected directly in adversary motivations, with 57% of intrusions targeting financial services organizations originating from eCrime adversaries and 43% from nation-state threat actors. This split is broadly consistent with targeting patterns seen across other sectors in this report. Notably, nation-state intrusion volume targeting the financial services sector surged by 29%, marking the largest increase observed across all sectors in this report. The top five adversaries targeting the sector reflect these parallel interests: DPRK-nexus groups FAMOUS CHOLLIMA and STARDUST CHOLLIMA continue their well-documented pursuit of financial technology (fintech) entities and cryptocurrency exchanges, while eCrime adversaries PUNK SPIDER, PLUMP SPIDER, and MUTANT SPIDER are driven by direct financial motivation.

48% of all financial sector targeting impacted North American institutions, a trend consistent with the historical targeting patterns and the concentration of high-value financial infrastructure in the region.

Academic

The academic sector recorded the largest year-over-year increase in intrusion activity of any sector during this reporting period, marked by a 17% surge. The threat type distribution closely mirrors broader trends, with 55% of intrusions attributed to eCrime adversaries and 45% to nation-state threat actors.

Academic institutions present a compelling target for both adversary types. For eCrime adversaries, the sector holds large volumes of sensitive personal data, including Social Security numbers, health records, and financial information. This wealth of data makes the vertical a highly attractive target for opportunistic, financially motivated operations. For nation-state adversaries, academic institutions represent repositories of sensitive research, emerging technology, and intellectual property aligned with national intelligence collection priorities.

This breadth of motivation is reflected in the top five adversaries targeting the sector. FAMOUS CHOLLIMA, MUSTANG PANDA, STATIC KITTEN, and VAULT PANDA all represent nation-state activity, pursuing research and intellectual property consistent with DPRK, People's Republic of China (PRC), and Iran-nexus strategic collection priorities. VICE SPIDER was the sole eCrime adversary in the top five, yet it represented a significant share of all financially motivated targeting in the sector. When excluding FAMOUS CHOLLIMA activity due to its opportunistic and sector-agnostic nature, VICE SPIDER accounted for more than one-fifth of all remaining intrusion activity targeting academic institutions.

60% of all academic sector targeting impacted North American institutions, followed by Southeast Asian (11%) and Middle Eastern (8%) entities. This geographic concentration aligns with the density of high-value research institutions and universities in the region.

MITRE ATT&CK and MITRE ATLAS

CrowdStrike Counter Adversary Operations tracks intrusion activity against the MITRE ATT&CK® Enterprise Matrix, mapping observed adversary behaviors to a standardized framework that enables consistent analysis across intrusions. Figure 5 illustrates the top MITRE ATT&CK tactics and techniques CrowdStrike OverWatch identified across intrusion activity over the past 12 months. Because CrowdStrike OverWatch hunts for post-exploitation behavior regardless of how an adversary gains initial access, the most frequently observed activity clusters around the Discovery, Execution, and Stealth tactics, a pattern consistent with what CrowdStrike OverWatch has observed in prior years. This consistency is telling: Once inside an environment, adversaries reliably prioritize understanding where they are, executing malicious code through native and trusted means, and disguising their activity to avoid detection.

The top three techniques, which include Command and Scripting Interpreter, System Owner/User Discovery, and Masquerading, reflect this same ideology. Scripting interpreters offer adversaries a flexible, native means of executing malicious code without introducing external tooling; user and system discovery techniques allow them to orient themselves within the environment and identify paths to their objectives; and masquerading enables them to disguise malicious activity as legitimate processes, buying time against detection.

This tracking also reflects MITRE's April 2026 framework evolution, which retired the broad Defense Evasion tactic in favor of two more precise categories: Stealth and Defense Impairment. This distinction is meaningful for defenders. Stealth captures techniques where adversaries blend malicious activity into expected behavior, requiring teams to identify subtle anomalies within legitimate activity. Conversely, Defense Impairment covers the active sabotage or disabling of security controls and monitoring tools, representing a deliberate attempt to blind defenders entirely. Because of this structural shift, this report no longer lists Defense Evasion as a standalone tactic area in the MITRE ATT&CK framework.

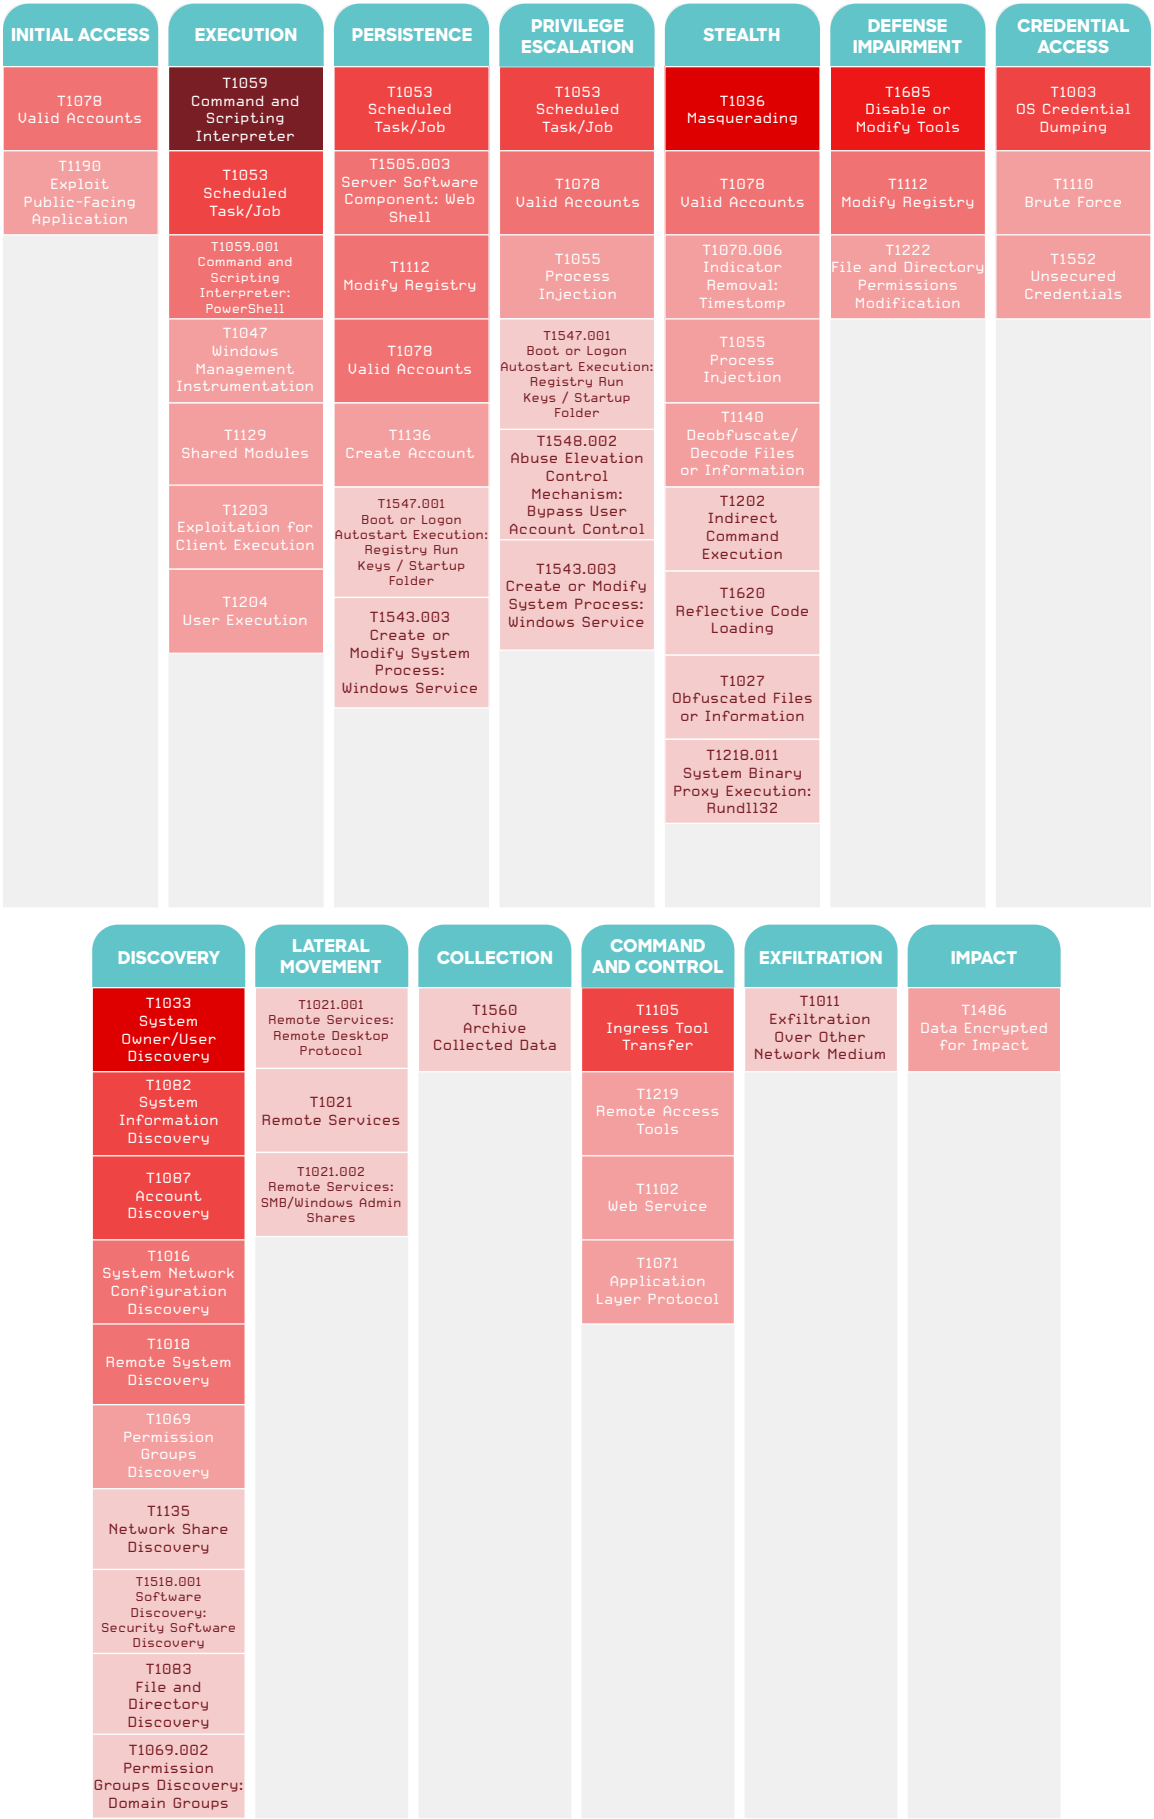

Figure 5. MITRE ATT&CK heat map highlighting the top techniques CrowdStrike OverWatch observed adversaries employ in each tactic area, July 2025-June 2026

As adversaries increasingly incorporate AI in their operations and target AI-enabled infrastructure, CrowdStrike Counter Adversary Operations simultaneously maps these specialized threat vectors to the MITRE ATLAS™ framework.[4] This unified mapping surfaces where threat actors are focusing their optimization efforts. Over the past 12 months, the top three identified MITRE ATLAS tactics were Resource Development, Initial Access, and AI Model Access. Figure 6 represents the precise concentration of the top MITRE ATLAS techniques across both AI-enabled and AI-targeted operations.

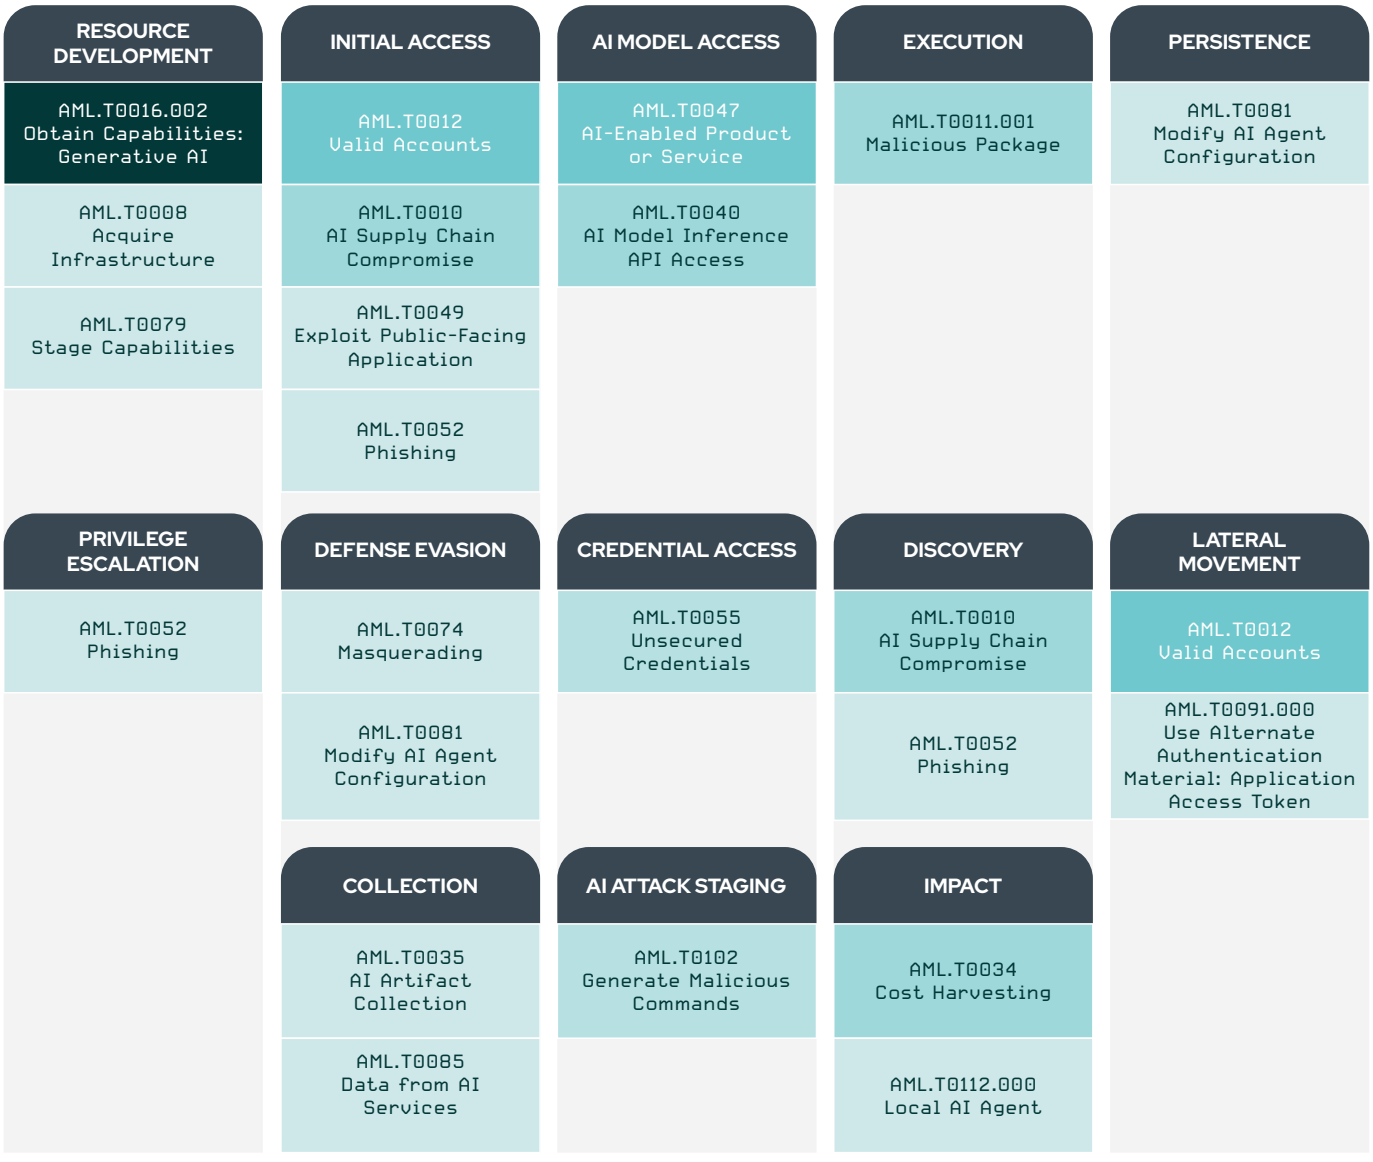

Figure 6. MITRE ATLAS heat map highlighting the top AI-enabled and AI-targeted techniques CrowdStrike OverWatch observed adversaries employ in each tactic area, July 2025-June 2026

4 To learn more about MITRE ATLAS, visit https://atlas.mitre.org/.

Resource Development

The most common method of AI adoption among threat actors is generating payloads or commands (AML.T0016.002: Obtain Capabilities: Generative AI). eCrime and nation-state threat actors routinely use generative AI, particularly for utility shell commands after gaining access.

FAMOUS CHOLLIMA demonstrated the most advanced AI usage, creating entire fake companies with AI-generated websites, GitHub accounts, and email infrastructure to support insider threat operations. Other nation-state threat actors employed LLM-generated reverse shells with professional commenting and resilient auto-reconnect logic or used LLM-generated commands for reconnaissance and data collection. Meanwhile, eCrime threat actors have utilized utility scripts, reverse shells, and credential harvesters likely written by LLMs.

LLM adoption elevates the capability floor: Lower-skilled adversaries can now produce relatively robust scripts with resilient error handling and sometimes evasion techniques. These allow some threat actors to achieve objectives faster than they could before.

Initial Access

Some threat actors are abusing AI tools as an intrusion vector. Similar to what CrowdStrike OverWatch observes overall in intrusions, Valid Accounts (AML.T0012) was the most commonly used Initial Access technique among adversaries. Notably unique to MITRE ATLAS, AI Supply Chain Compromise (AML.T0010) was the second most common MITRE ATLAS Initial Access technique. FAMOUS CHOLLIMA's campaign against AI-centric development environments was one of the most sophisticated examples of this technique in practice.

CASE STUDY:

FAMOUS CHOLLIMA Software Supply Chain Attack

Capitalizing on interest in AI-driven development tooling, in January and February 2026, DPRK-nexus adversary FAMOUS CHOLLIMA executed a targeted supply chain campaign against cryptocurrency and blockchain companies by weaponizing an AI-centric IDE.

FAMOUS CHOLLIMA constructed outwardly legitimate project repositories, primarily hosted on GitHub, while sharing at least one repository directly via Telegram. The trojanized repositories contained legitimate-looking project files alongside hidden malicious scripts in `package.json`, `.vscode/settings.json`, or post-install hooks. When developers opened these projects, the IDE's built-in terminal or task runner automatically executed the malicious commands, requiring no further interaction from the victim and allowing FAMOUS CHOLLIMA entry into the victim's environment.

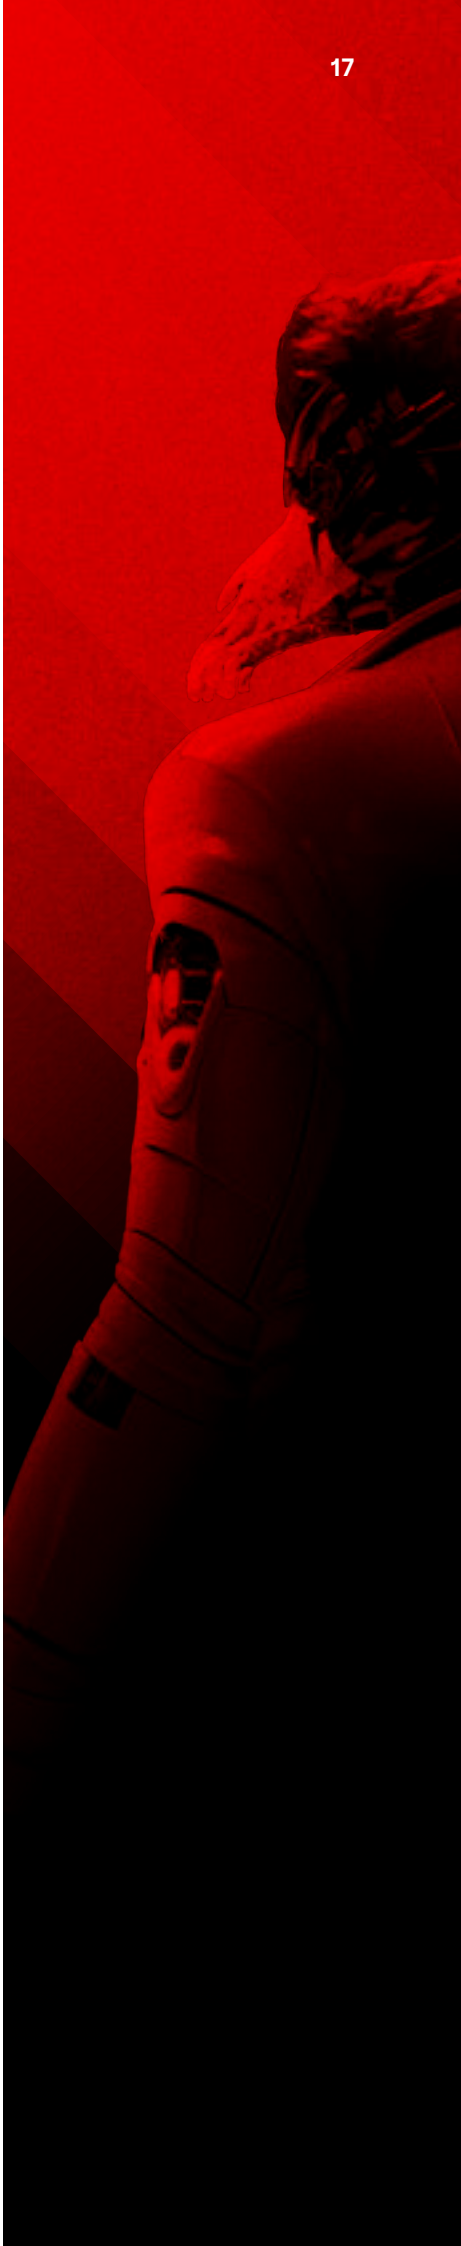

Like most software supply chain campaigns, whether AI-related or otherwise, this FAMOUS CHOLLIMA campaign resulted in observable behaviors relating to processes, files, and network connections. Endpoint security tools such as the Falcon platform provide threat hunting teams like CrowdStrike OverWatch with host-based telemetry regarding process-level behaviors, enriched with intelligence data, to enable quick detection and response.

AI Model Access

AI Model Access techniques, which enable the adversary to gain some level of access to an AI model, accounted for 16% of total MITRE ATLAS techniques observed during the past 12 months. Throughout 2026, threat actors have attempted to exploit vulnerabilities in AI-related server software. For instance, CrowdStrike's honeypot infrastructure detected threat actors attempting to exploit vulnerabilities that enable remote code execution (RCE). One exploit payload contained a malicious Model Context Protocol (MCP) server configuration that reads the parent process's environment variables, allowing the adversary to harvest and exfiltrate sensitive configuration information to an external webhook.

A likely separate threat actor exploited another vulnerability in AI-related server software at a Falcon platform customer to deliver a cryptomining payload via POST requests from a VPN exit on port 4000. CrowdStrike subsequently detected the affected process spawning Python subprocesses to download, extract, and execute a cryptocurrency mining binary named `gmon` from a remote URL. The `gmon` payload was written to a staging directory prior to execution, after which the staging directory was then deleted to conceal the activity.

Using AI for Impact

Impact covers the techniques adversaries use to manipulate, interrupt, or otherwise sabotage organizational AI systems and data. This activity translates into real costs for victim organizations. Impact techniques accounted for 8% of total MITRE ATLAS techniques observed during the past 12 months. The most common MITRE ATLAS Impact technique was Cost Harvesting (AML.T0034), a technique where adversaries deliberately drive a victim's AI services beyond normal operating capacity with the intent of increasing the cost of services and intentionally causing financial harm to the victim. This reflects a growing trend of threat actors seeking to compromise victims' AI platforms themselves. Several financially motivated adversaries have sought to obtain access to victims' corporate LLM API access, typically referred to as LLMJacking.

CASE STUDY:

LLMJacking

In May 2026, a likely financially motivated threat actor leveraging a cloud identity with a long-term access key conducted a large-scale LLMJacking campaign targeting a cloud computing service offering access to foundation models.

The threat actor first called `GetCallerIdentity` to enumerate account details and then used the `DryRun` flag to call `CreateSecurityGroup` and `RunInstances` to test account permissions, a technique commonly used to verify account permissions without executing the underlying operations. The threat actor then attempted to call `InvokeModel`, likely to access LLMs, but received error messages, including one indicating that a required use case form had not been completed.

Likely in response to these error messages, the threat actor called `GetFederationToken` to generate temporary credentials for a federated session. The threat actor used the session name `console-session` to masquerade the federated session as a legitimate interactive login. They then attached the `AdministratorAccess` policy to this session, creating a new set of temporary credentials carrying the same permissions as the compromised user.

With elevated access established, the threat actor likely logged in to the cloud administrative web console to submit the use case required to access LLMs. Within the `console-session` session, the threat actor called `PutUseCaseForModelAccess` to address the errors encountered during the initial `InvokeModel` attempts. The threat actor then used the compromised cloud identity account to call the `InvokeModelWithResponseStream` and `InvokeModel` APIs hundreds of thousands of times, likely as part of the LLMJacking campaign.

The impact was significant. During an initial two-minute flood, the threat actor sent nearly 200,000 API requests before throttling kicked in, after which only a small number of requests were logged. Most requests consisted of likely proxied customer requests, featuring variable input and output token sizes, high model diversity, and a wide range of `ValidationException` messages. The threat actor also sent likely health check messages consisting of single-token or few-token inputs directed primarily at a large foundation model in a single region. This stood in sharp contrast to the broad range of regions and models used by the proxied customer requests, suggesting the threat actor was actively monitoring service availability.

Observations from the Front Lines

Threat Hunting in the AI Era: Faster Vulnerability Weaponization and AI-Assisted Vulnerability Discovery

CrowdStrike Intelligence recorded a 42% year-over-year increase in zero-day exploitation from 2024 to 2025, reflecting adversaries' growing ability to discover, develop, and weaponize vulnerabilities before defenders have the opportunity to patch them. Now, the window between vulnerability disclosure and mass exploitation is measured in hours rather than days.

From January 2026 to June 2026, 88% of CrowdStrike-observed exploitation of vulnerabilities with a public PoC was conducted within 48 hours of the PoC's release. Exploitation timelines will almost certainly continue to compress across near-, mid-, and long-term horizons.

Though this pattern predates the widespread integration of frontier AI into vulnerability research, the ongoing implementation of such systems is poised to further compress the timeline between disclosure and active exploitation. Frontier models are likely contributing to the rising volume of disclosed vulnerabilities, exacerbating the challenges faced by network defenders as they attempt to cope with ever-shrinking patch windows.

Integrating CrowdStrike OverWatch's threat hunting capabilities with exposure management tools and solutions provides a vital backstop, mitigating damage during two crucial periods: before a patch is released and while organizations patch and restart affected instances.

The following case studies illustrate instances where adversaries such as UMBRAL BISON and GENESIS PANDA exploited zero-day or n-day vulnerabilities within hours of disclosure. CrowdStrike OverWatch's ability to identify and hunt for post-exploitation malicious behaviors acts as a critical fail-safe in cases like these, ensuring rapid and effective coverage against subsequent widespread exploitation by opportunistic adversaries. This approach ensures hunting coverage across multiple exploitation vectors, whether through network-based unauthenticated RCE or local privilege escalation (LPE) via binaries delivered to sensor-managed hosts.

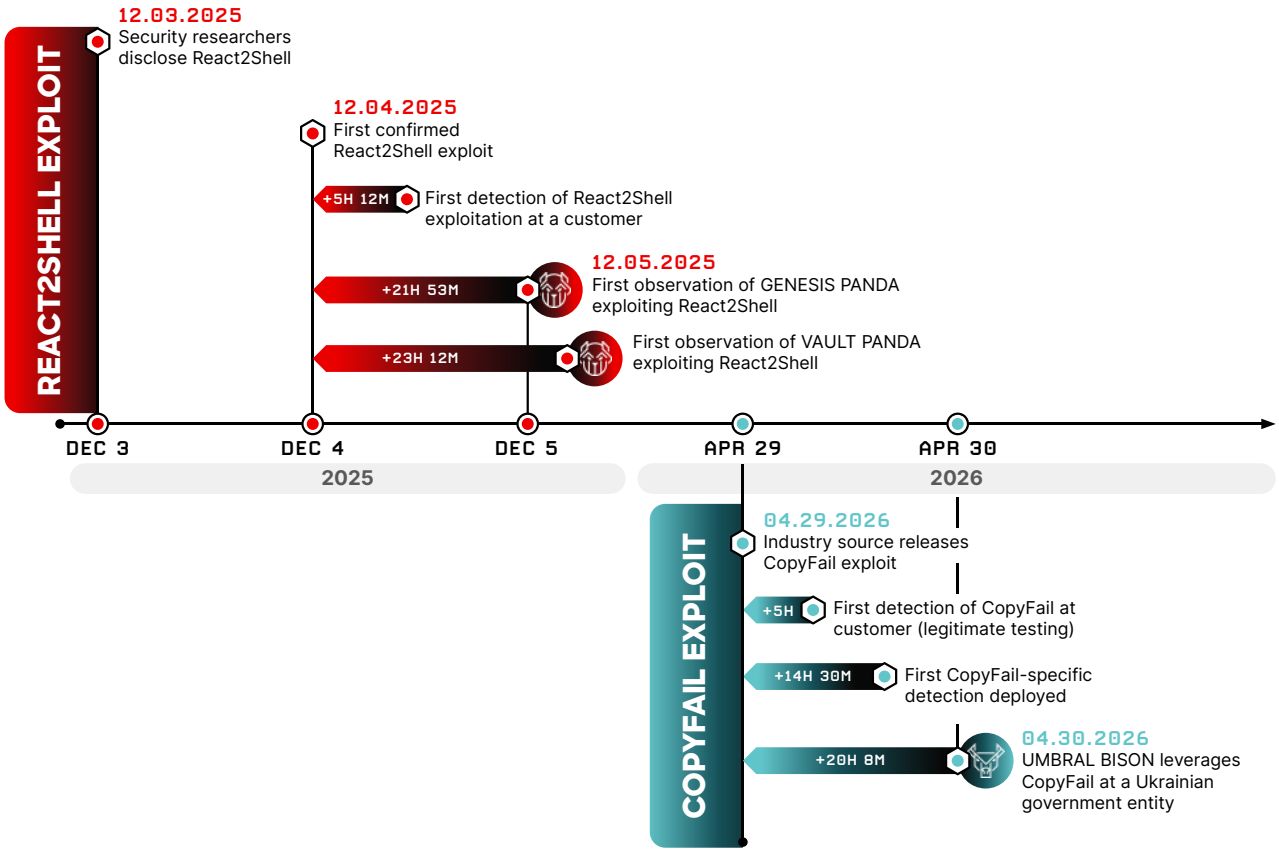

Figure 7. Timelines for UMBRAL BISON LPE and GENESIS PANDA RCE, including CrowdStrike OverWatch coverage and sensor detections

RCE CASE STUDY:

React2Shell Surge and Multiple PANDAs

React2Shell (CVE-2025-55182) is a critical insecure deserialization vulnerability that enables unauthenticated RCE in React Server Components and Next.js applications. Rapid exploitation of this flaw clearly illustrates the compressed window between vulnerability disclosure and weaponized exploitation.

The vulnerability was privately disclosed to React and Next.js administrators and publicly released alongside patches on December 3, 2025. Within 24 hours, multiple working PoC exploits were circulating openly, and multiple eCrime, nation-state, and unknown threat actors were leveraging React2Shell in the wild. **In the first four days following public disclosure, CrowdStrike OverWatch responded to more than 800 hunting leads related to suspected React2Shell exploitation at more than 80 different victims.**

Many of these hunting leads automatically generated detections based on existing post-exploitation behaviors, such as using the `curl` command to download tooling or malware from external infrastructure. This automated coverage enabled CrowdStrike OverWatch and customers to redirect focus toward threats not covered by existing detections that required immediate triage. Post-exploitation activity ranged from opportunistic deployment of commodity malware, such as cryptocurrency miners, to the use of low-prevalence remote access tools (RATs) by nation-state threat actors.

» IN THE FIRST FOUR DAYS FOLLOWING PUBLIC DISCLOSURE, CROWDSTRIKE OVERWATCH RESPONDED TO MORE THAN 800 HUNTING LEADS RELATED TO SUSPECTED REACT2SHELL EXPLOITATION AT MORE THAN 80 DIFFERENT VICTIMS.

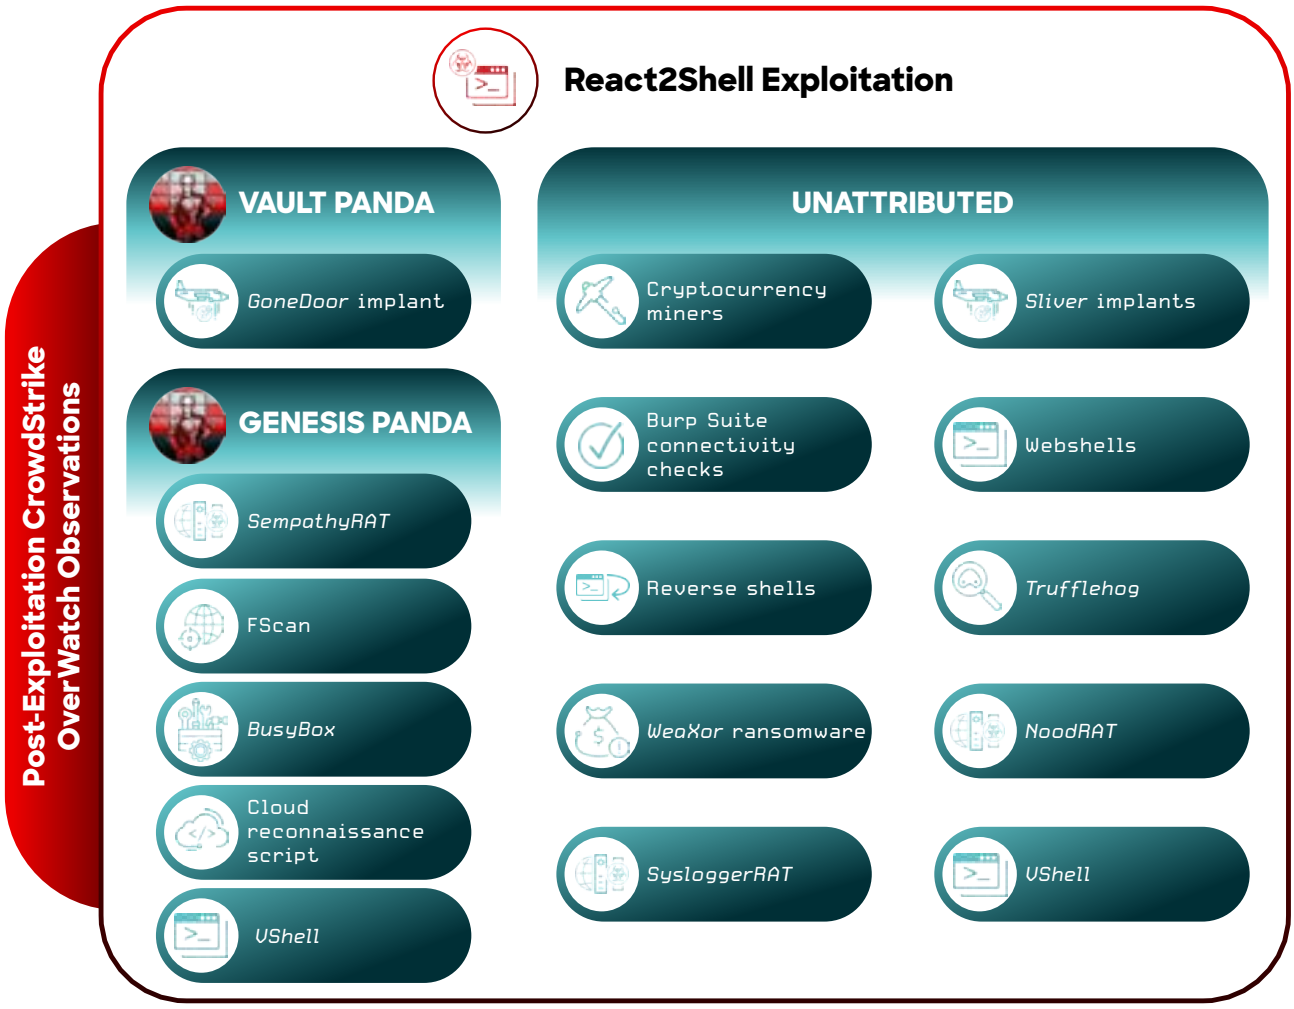

Figure 8. Post-exploitation observables and threat actors associated with React2Shell exploit activity from December 4 to 8, 2025

Within 24 hours of its public disclosure, China-nexus adversaries VAULT PANDA and GENESIS PANDA conducted deliberate, systematic attacks leveraging the React2Shell vulnerability. The speed of this response highlights their posture as adversaries who actively monitor vulnerability disclosures, rapidly validate exploitability, and pre-stage tooling in anticipation of a constantly changing attack surface.

VAULT PANDA leveraged React2Shell to deploy its *GoneDoor* implant, disguised as a legitimate Linux utility. CrowdStrike OverWatch observed the threat actor attempt to evade detection by deleting files from disk. In this case, VAULT PANDA used the `wget` command to download the *GoneDoor* implant, which the Falcon platform detected and blocked automatically. Where the Falcon platform did not automatically block the command, CrowdStrike OverWatch threat hunters promptly triggered UI alerts and communicated directly with local defenders, ensuring awareness of the malicious activity.

Concurrently, GENESIS PANDA exploited React2Shell to deploy other tooling alongside two RATs, *VShell* and *SempathyRAT*, at multiple victims across several sectors and regions. At least one intrusion was intended to extend beyond initial access to the cloud; specifically, GENESIS PANDA downloaded FScan, a publicly available scanning tool, for network enumeration. The adversary then executed a custom reconnaissance script to harvest cloud credentials and probe cloud services to identify lateral movement paths.

LPE CASE STUDY:

UMBRAL BISON's Rapid Adoption of CopyFail Exploit

On April 29, 2026, an industry researcher released a PoC exploit and technical details for a Linux LPE vulnerability publicly disclosed that same day, dubbed CopyFail (CVE-2026-31431). At the time of disclosure, CopyFail was addressed on the Linux mainline kernel but remained a zero-day vulnerability, as it had not been patched across many of the major Linux distributions, such as Ubuntu.

By April 30, 2026, CrowdStrike OverWatch detected widespread deployment of the CopyFail exploit, with approximately 94% of events in the first 24 hours related to testing behavior based on public PoC code. Threat hunters quickly triaged the incidents to distinguish malicious exploitation from legitimate testing, eventually uncovering Belarus-nexus activity just over 20 hours after public disclosure.

On April 30, 2026, CrowdStrike OverWatch identified UMBRAL BISON leveraging CopyFail at a Ukrainian government entity to perform LPE to `root`. UMBRAL BISON slightly modified the public exploit's formatting and execution, which did not alter exploit functionality but was very likely intended to evade static detection of the Python code. Using the escalated access, UMBRAL BISON deployed a `Mythic C2 Poseidon` command-and-control (C2) agent, which the Falcon platform terminated upon execution.

Combining the Falcon platform with CrowdStrike OverWatch threat hunting capabilities allows even novel exploits to be disrupted quickly. This is particularly important for organizations with Linux servers hosting critical assets that often run older Linux versions due to concerns regarding the complications and risks associated with patching and restarting the Linux kernel. In these instances especially, professional managed threat hunting helps minimize an intrusion's impact and provides insights to prevent future compromises.

OUTLOOK:

AI-Assisted Research Is Shrinking the Vulnerability Management Window

Speed is the defining characteristic of modern intrusions. The window between vulnerability disclosure and active exploitation continues to shrink, and the incorporation of agentic AI into vulnerability discovery workflows has the potential to render traditional patching cycles insufficient, leaving defenders forced to adopt a continuous emergency patching posture or accept significantly elevated risk.

This scenario is no longer theoretical. Two of the three recently disclosed highly stable Linux LPE exploits, CopyFail and Fragnesia (CVE-2026-46300), were allegedly discovered using AI or AI-assisted research, and threat actors wasted no time incorporating them into active operations.

Addressing this challenge requires a fundamental rethink of how organizations approach vulnerability management and threat hunting. It demands a strategic pivot toward defense-in-depth architectures and proactive, continuous asset discovery and threat hunting capabilities that can identify post-exploitation behaviors before adversaries achieve their objectives. Combining CrowdStrike OverWatch's threat hunting capabilities with CrowdStrike Falcon® Exposure Management provides organizations with a multifaceted approach to navigate these complexities.

Falcon Exposure Management addresses the front end of this challenge, delivering real-time visibility across the attack surface and integrating threat intelligence to prioritize the vulnerabilities most likely to be exploited. CrowdStrike OverWatch acts as a critical fail-safe, ensuring rapid and effective coverage and protecting against emerging threats.

While AI may increase the number of discovered vulnerabilities or provide threat actors with an alternative attack vector, the fundamental patterns of post-exploitation activity by adversaries remain largely unchanged. This enables CrowdStrike OverWatch and the Falcon platform to effectively identify and disrupt malicious actions before adversaries achieve their objectives. When telemetry indicates active escalation and attempted environment control, CrowdStrike OverWatch triggers rapid, high-fidelity alerts paired with direct notifications to drive immediate intervention. Isolated observations are processed with equal precision, ensuring defenders can cut through noise and focus on the threats that matter most.

Software Supply Chain Attacks: Defending Across a Cross-Domain Threat Landscape

In the CrowdStrike 2026 Global Threat Report, CrowdStrike identified a shift in adversary initial access techniques toward methods designed to evade traditional security controls. In the first half of 2026, that trend solidified, with software supply chain attacks increasingly leveraging malicious software packages uploaded to public software registries.

Because these software registries are used by developers and organizations across all industries and geographies, the scale of exposure is significant. A single malicious software package can reach millions of users a week, highlighting the risk of a widespread compromise from a single entry point. While adversaries have targeted a variety of public software registries, malicious npm packages dominated the activity, comprising 87% of all identified malicious packages over the six-month period. The npm ecosystem is an attractive target due to the established norm of projects having deep dependency trees of many small packages, the ability for packages to automatically run scripts upon installation, and the general popularity of JavaScript for full-stack software development. Millions of users rely on packages distributed through the npm registry.

Responding to software supply chain attacks is not straightforward. Malicious packages can reach hundreds of thousands of systems within hours of publication, representing a scale of compromise that is nearly impossible for network defenders to fully assess in real time. When multiple adversaries target different packages simultaneously, the problem compounds: Organizations can find themselves facing several overlapping compromises at once, each demanding its own triage, containment, and remediation.

A SINGLE MALICIOUS SOFTWARE PACKAGE CAN REACH MILLIONS OF USERS A WEEK, HIGHLIGHTING THE RISK OF A WIDESPREAD COMPROMISE FROM A SINGLE ENTRY POINT.

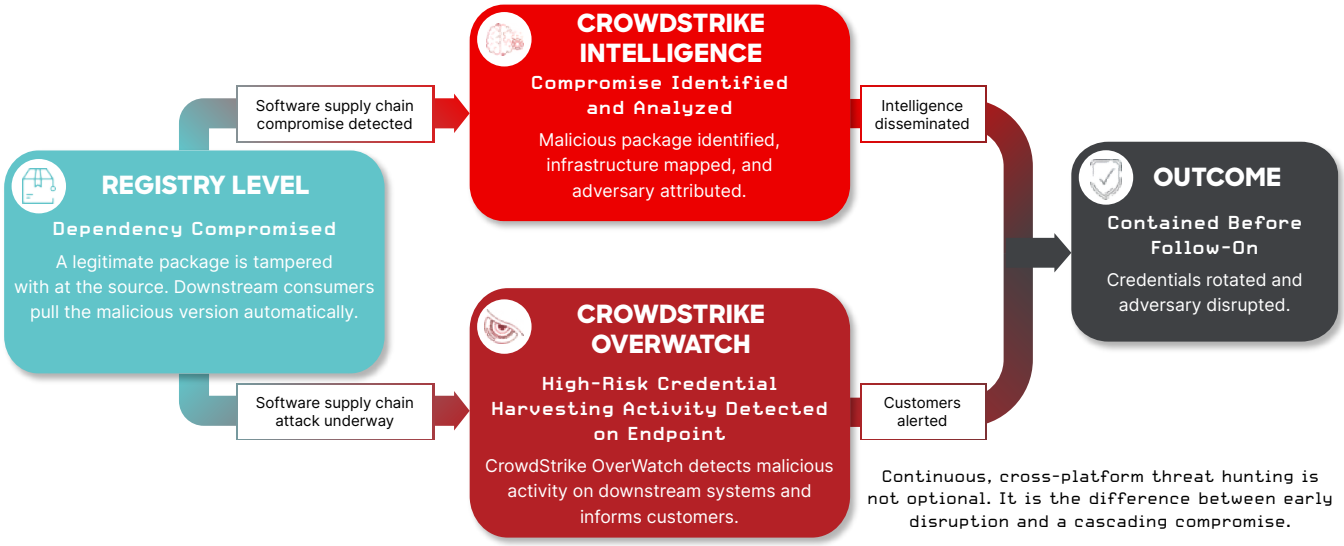

Figure 9. CrowdStrike Counter Adversary Operations' software supply chain compromise response

At their core, software supply chain attacks exploit trust: trusted access, inherited permissions, and automated propagation across interconnected systems. This year, adversaries moved deliberately into the developer ecosystem, targeting the tooling that underpins modern software development: CI/CD pipelines, container registries, and package registries. By compromising the tools developers trust implicitly, adversaries inherit that trust and everything that comes with it.

Over the past year, two adversaries have driven the most consequential software supply chain attacks: STARDUST CHOLLIMA, a DPRK-nexus adversary with a sustained focus on financial services and cryptocurrency, and ALTERED SPIDER, a financially motivated eCrime adversary operating a sprawling and rapidly evolving software supply chain attack operation. Together, their campaigns account for some of the most impactful software supply chain compromises in 2026 and illustrate the full spectrum of adversary sophistication targeting this attack surface.

Software Supply Chain: Compromise vs. Attack

A supply chain compromise occurs when a malicious actor successfully breaches an upstream vendor and gains the ability to modify that vendor's published software artifacts. A supply chain attack is the deliberate use of those modified artifacts and downstream trust relationships to compromise a downstream target.

In short: A supply chain compromise happens to the upstream vendor; a supply chain attack is the use of that compromise to target the ultimate victim.

2026 Software Supply Chain Trends

Software supply chain attacks are not new, but this year marks a meaningful inflection point in how they are conducted, the target base, and their downstream impact. CrowdStrike has identified five trends defining this evolution, which centers on the trusted developer ecosystem at the core of AI integration.

Trend 1: Developer Ecosystems Are the New Targets

As open-source dependency adoption has accelerated, adversaries have increasingly targeted open-source software (OSS) package registries, including npm and PyPI, as force-multiplying attack vectors. This technique can enable software supply chain attacks against thousands of dependent projects through a single malicious package. In 2026, however, attacks on the developer ecosystem went beyond compromising the package dependencies used in software products, extending directly to the tools and infrastructure used for software development. Modern software development tooling relies heavily on public registries and a large number of upstream trust relationships. As a result, CI/CD infrastructure like GitHub Actions, container registries, and IDE extensions such as Visual Studio Code and Open VSX plugins have all become delivery mechanisms for malicious code.

Adversaries such as ALTERED SPIDER and STARDUST CHOLLIMA have identified this trust as a systematic vulnerability and are shifting focus from software package dependencies to the tools, registries, and infrastructure that developers interact with every day. Compromising developer machines and developer tooling provides access to production environments, cloud credentials, signing certificates, and downstream customer environments that would be substantially more difficult to reach through conventional intrusion paths.

Trend 2: Attacks Are More Automated and Scalable

Earlier generations of software supply chain attacks, including the SolarWinds compromise, were highly targeted, operationally intensive, and focused on precision access to specific organizations. Modern adversaries are shifting toward industrialized, automated attack operations, conducting self-propagating campaigns that leverage programmatic package publication, typosquatting, dependency confusion, malicious update propagation, and credential harvesting at scale.

ALTERED SPIDER deployed the self-propagating malware *TeamPCPCloudStealer*, which can use stolen credentials to autonomously publish infected versions of additional packages. This process spreads the infection laterally across the npm and PyPI ecosystems without further operator involvement. A single set of stolen maintainer credentials became the seed for a cascading infection chain spanning hundreds of packages across multiple registries. Operating at a very high operational tempo during their May 2026 campaigns, ALTERED SPIDER compromised more than 300 software dependencies in a single day.

Trend 3: Identity Remains the Critical Entry Point

Abusing developer identities and privileged credentials remains the primary entry point for software supply chain compromises. What begins as a compromised developer identity at a single organization can quickly become the foundation for attacks targeting thousands of downstream systems.

STARDUST CHOLLIMA heavily relies on social engineering of individuals in developer roles to gain privileged access. In January 2026, the adversary published multiple malicious libraries to both the Python and Node.js clients of the dYdX cryptocurrency exchange to deliver the *PricePilot* malware, likely using credentials stolen from the dYdX development team to achieve the initial software supply chain compromise. In March 2026, STARDUST CHOLLIMA used stolen maintainer credentials to compromise the Axios npm package,[5] delivering platform-specific variants of their *ZshBucket* malware. In June 2026, STARDUST CHOLLIMA injected a malicious npm package as a dependency into at least 131 Mastra AI framework packages. This breach of the publishing environment was enabled by a social engineering operation in which the threat actor contacted a Mastra employee via LinkedIn and tricked them into clicking a malicious link over a video call. The compromised framework packages were then used to deliver the *TeaBundle* malware.

5 https://www.crowdstrike.com/en-us/blog/stardust-chollima-likely-compromises-axios-npm-package/

Trend 4: CI/CD Pipelines Are Under Attack

Upstream vendor CI/CD pipelines have emerged as one of the highest-value targets in the modern software supply chain attack surface. These pipelines possess production deployment permissions, cloud environment access, signing authority, artifact control, and infrastructure automation rights. A single pipeline compromise can become the launchpad for a sprawling attack at scale. Once compromised, they provide adversaries with trusted distribution channels and downstream propagation opportunities, allowing them to establish stealthy persistence that can be difficult to detect and even harder to fully remediate. These long-term implications compound the immediate credential theft risk. CI/CD environments are frequently configured with long-lived access credentials and minimal runtime monitoring, enabling unauthorized access to persist undetected well beyond the initial compromise window.

ALTERED SPIDER's compromise of Trivy's `trivy-action` repository in March 2026 demonstrates the severe operational leverage a pipeline compromise provides. Trivy's `trivy-action` repository is a publicly available GitHub Action: a predefined CI/CD component that other GitHub users can include in their own CI/CD pipelines. After gaining access to the `trivy-action` repository, the adversary used Git tag poisoning to force-push malicious commits across multiple release tags, silently replacing the CI/CD code included in each of those `trivy-action` releases. The result was that any organization pulling the affected `trivy-action` releases as part of an automated build or scan workflow unknowingly executed credential-stealing malware within its CI/CD environment. The compromise of a single trusted CI/CD pipeline component became a simultaneous attack against every organization that trusted it.

Trend 5: Adversaries Are Pivoting into the Cloud

Adversaries are increasingly leveraging software supply chain compromises to gain access to cloud infrastructure. In the majority of cases, harvesting developer credentials through malicious packages and trojanized tools was a means to an end: gaining entry to the cloud environments, data stores, and infrastructure those authentication factors controlled.

TeamPCPCloudStealer, ALTERED SPIDER's primary credential harvesting tool, was engineered specifically for this pivot. The malware systematically targeted cloud access keys, Azure credentials, Google Cloud tokens, Kubernetes service account tokens, SSH private keys, CI/CD pipeline secrets, container registry credentials, and cryptocurrency wallet keys. Every major cloud provider's credential format was explicitly targeted. Following the March 2026 Trivy compromise, CrowdStrike OverWatch identified multiple cases where adversaries leveraged credentials stolen during the exposure window to conduct immediate operations in victim cloud environments within hours of the initial software supply chain compromise.

Hunting ALTERED SPIDER Across Domains

In March 2026, CrowdStrike identified malicious commits across a substantial number of `aquasecurity/trivy-action` repository tags, indicating a software supply chain compromise of a CI/CD pipeline component trusted by development teams globally. This activity was attributed to ALTERED SPIDER, who performed Git tag poisoning at scale, force-pushed malicious commits across affected release tags, and silently replaced the code in each of those releases.

The software supply chain attack followed immediately: Every downstream pull of the now-compromised CI/CD pipeline component executed a credential-stealing payload, weaponizing a trusted tool against the very developers and pipelines that relied on it. This activity demonstrated how rapidly a software supply chain compromise can be converted into a software supply chain attack at scale.

Developers and automated pipelines consuming the affected dependency did not realize they were a part of a supply chain attack, unless they were monitoring for specific indicators of compromise, until CrowdStrike OverWatch detected evidence of *TeamPCPCloudStealer* attempting to execute on managed endpoints. Once active, the malware moved quickly, attempting to enumerate credentials, establish persistence, and exfiltrate harvested data to threat actor-controlled infrastructure. It also utilized a fallback mechanism that weaponized the victims' own GitHub accounts to stage stolen data.

The Trivy software supply chain compromise was just the beginning. ALTERED SPIDER had spent the exposure window stockpiling a significant number of credentials. Victims unaware of ALTERED SPIDER-infected packages executing in their environments were subject to follow-on operations from ALTERED SPIDER and their collaborators. These operations included *CipherForce* ransomware deployment, data theft, and abuse of stolen credentials to facilitate further software supply chain attacks against downstream targets. The speed and scale of these follow-on attacks became evident in the days after the initial compromise, as CrowdStrike OverWatch identified multiple cases of threat actors leveraging credentials likely stolen through *TeamPCPCloudStealer* to pivot into cloud environments, matching the cloud-pivoting pattern described above.

Tracking ALTERED SPIDER across successive campaigns requires drawing on extensive visibility spanning repository-level intelligence and endpoint telemetry to maintain a coherent picture of an adversary simultaneously operating across multiple platforms.

Figure 10 illustrates a theoretical ALTERED SPIDER software supply chain compromise, highlighting the adversary pivoting across domains and the visibility CrowdStrike OverWatch threat hunters have into cross-domain intrusion activity. It is constructed from CrowdStrike OverWatch's analysis of multiple observed intrusions. It does not represent a single end-to-end incident but rather a composite model of how these operations are designed to unfold, giving defenders a clearer picture of the full attack chain and where opportunities to detect and disrupt exist.

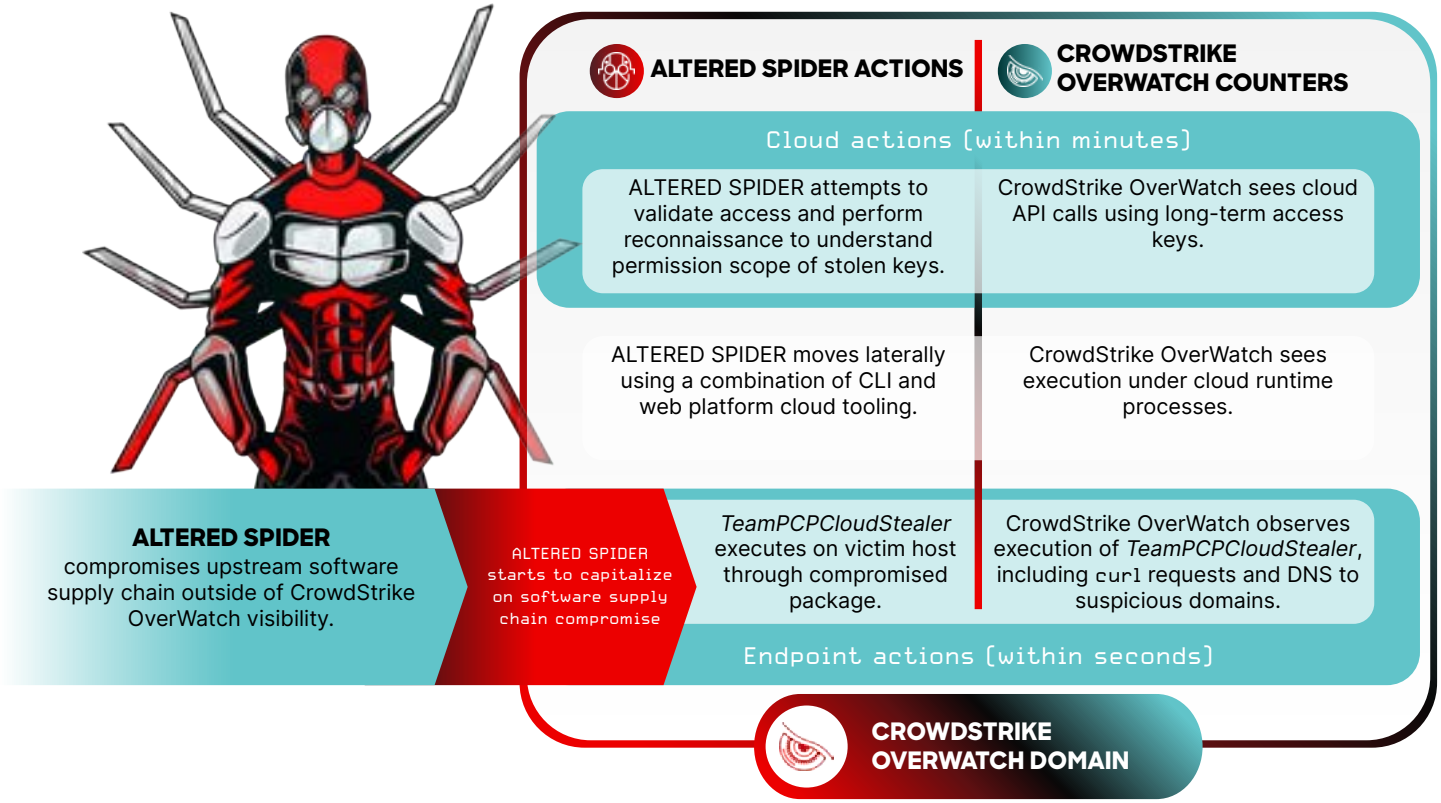

Figure 10. CrowdStrike OverWatch tracks ALTERED SPIDER across domains

This cross-domain visibility, which combines CrowdStrike Intelligence's analytical depth with CrowdStrike OverWatch's continuous endpoint monitoring, was essential in following ALTERED SPIDER across an attack surface no single data source could cover alone. Together, CrowdStrike Intelligence and CrowdStrike OverWatch provided customers with timely intelligence and warnings, enabling them to contain compromises and rotate credentials before ALTERED SPIDER could leverage stolen tokens for follow-on software supply chain attacks.

OUTLOOK:

The AI Ecosystem Is Becoming the Next Software Supply Chain Battleground

Emerging shifts in adversary tactics illustrate a fundamental challenge in defending the modern software supply chain. The attack surface is an entire ecosystem of tools, packages, repositories, pipelines, cloud environments, and, increasingly, AI frameworks that developers interact with every day. These same developer ecosystems are now being compromised. PyPI, npm, GitHub Actions, container registries, and open-source tooling are increasingly intersecting with AI development workflows. AI SDKs, machine learning (ML) dependencies, model tooling, and AI developer environments frequently contain cloud credentials, GPU infrastructure access, API tokens, and automation frameworks, making them high-value targets for software supply chain compromise.

The targeting of AI tooling by STARDUST CHOLLIMA and ALTERED SPIDER illustrates that AI ecosystems are already in the crosshairs. As agentic systems, MCP integrations, and dependency managers for AI agents proliferate, this attack surface will only expand. Continuous, cross-platform threat hunting is the difference between early disruption and a cascading software supply chain attack that no single defensive control is likely to stop.

Vishing Surges as an Identity-Based Initial Access Vector

Vishing as an initial access vector continues to accelerate. Between 2024 and 2025, CrowdStrike OverWatch detected a 134% increase in vishing intrusions. The trend shows no signs of slowing: The first half of 2026 has already seen twice the number of intrusions involving vishing as the likely initial access vector as the last half of 2025.

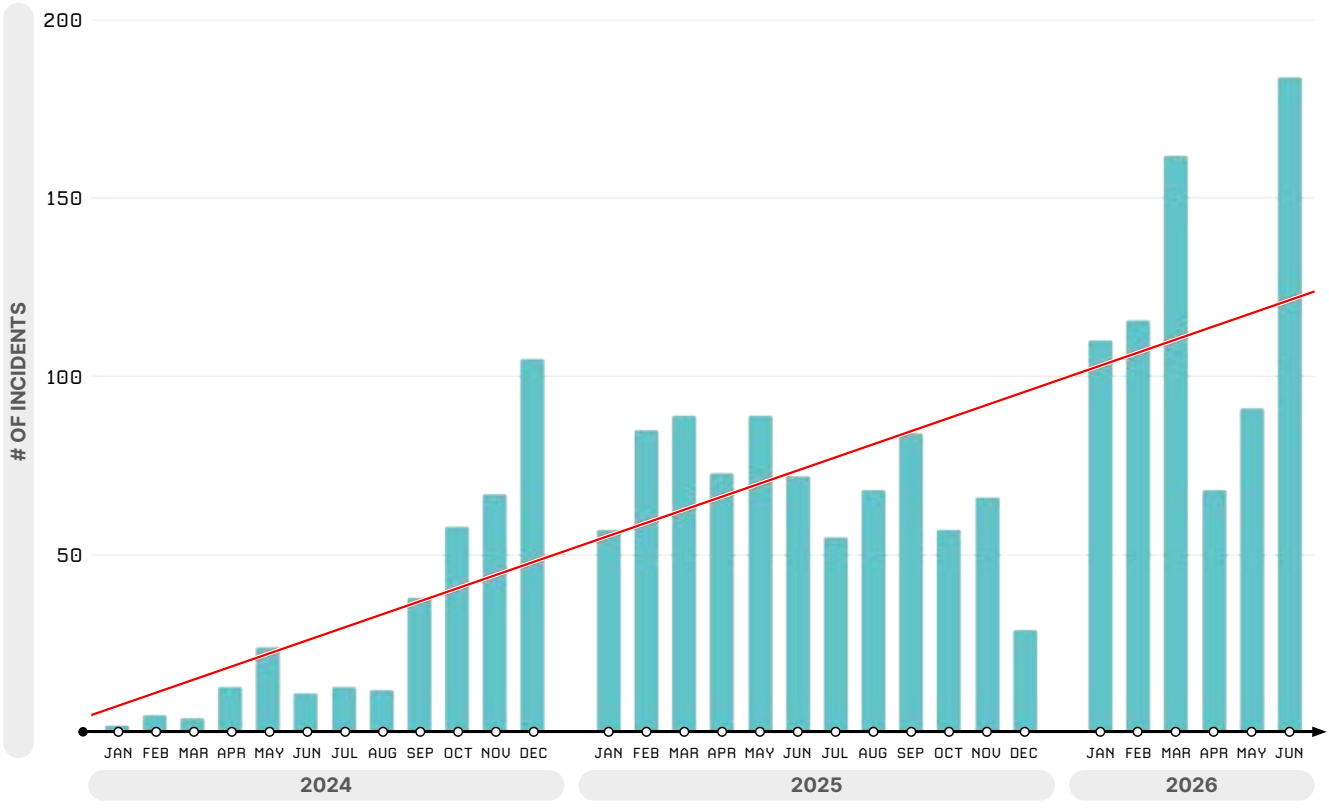

Figure 11. CrowdStrike OverWatch-detected intrusions in which vishing was the likely initial access vector per month, January 2024-June 2026

Vishing has emerged as a key technique among eCrime adversaries for several reasons:

- **It bypasses traditional security controls.** By convincing targeted users to enter credentials into adversary-in-the-middle (AiTM) pages or enable remote desktop sessions using common corporate tools like Microsoft Quick Assist, adversaries achieve access without executing malicious code on monitored hosts, sidestepping controls designed to catch malicious activity.
- **It exploits the human factor more than other social engineering techniques.** Direct human-to-human interaction enables a more compelling means of socially engineering targeted users than, for example, phishing emails. Threat actors often impersonate IT staff, citing the need to resolve a trivial support issue as a pretext for the call to manipulate targets into granting access.
- **It leaves defenders little to work with post-exploitation.** Compared to other social engineering techniques, such as ClickFix, SMS phishing (smishing), and email phishing, vishing provides incident responders with relatively few forensic artifacts, making post-incident investigations more challenging.

Vishing calls in isolation cannot be detected by enterprise security teams. Network defenders must instead understand the contextual activity that follows the vishing call to get ahead of these intrusions.

CASE STUDY:

CORDIAL SPIDER and SNARKY SPIDER Leverage Vishing to Steal Data from SaaS Applications

In 2026, CORDIAL SPIDER and SNARKY SPIDER emerged as two of the most prolific data theft and extortion adversaries. Both adversaries leverage vishing and similar identity-oriented tradecraft to exfiltrate data from SSO-integrated SaaS applications. In most observed cases, neither adversary interacts with managed endpoints, making centralizing logs from identity provider (IdP) and SaaS applications integral to detecting these fast-paced intrusions.

Both CORDIAL SPIDER and SNARKY SPIDER conduct vishing calls to direct targeted users to navigate to SSO-themed AiTM pages on the users' personal mobile devices. These devices are typically unmanaged, obscuring this stage of the intrusion from security teams' visibility. The AiTM pages are dynamic and highly configurable, allowing the adversaries to present several different prompt pages to the user according to the flow of the vishing call.

Once targeted users have entered SSO credentials and multifactor authentication

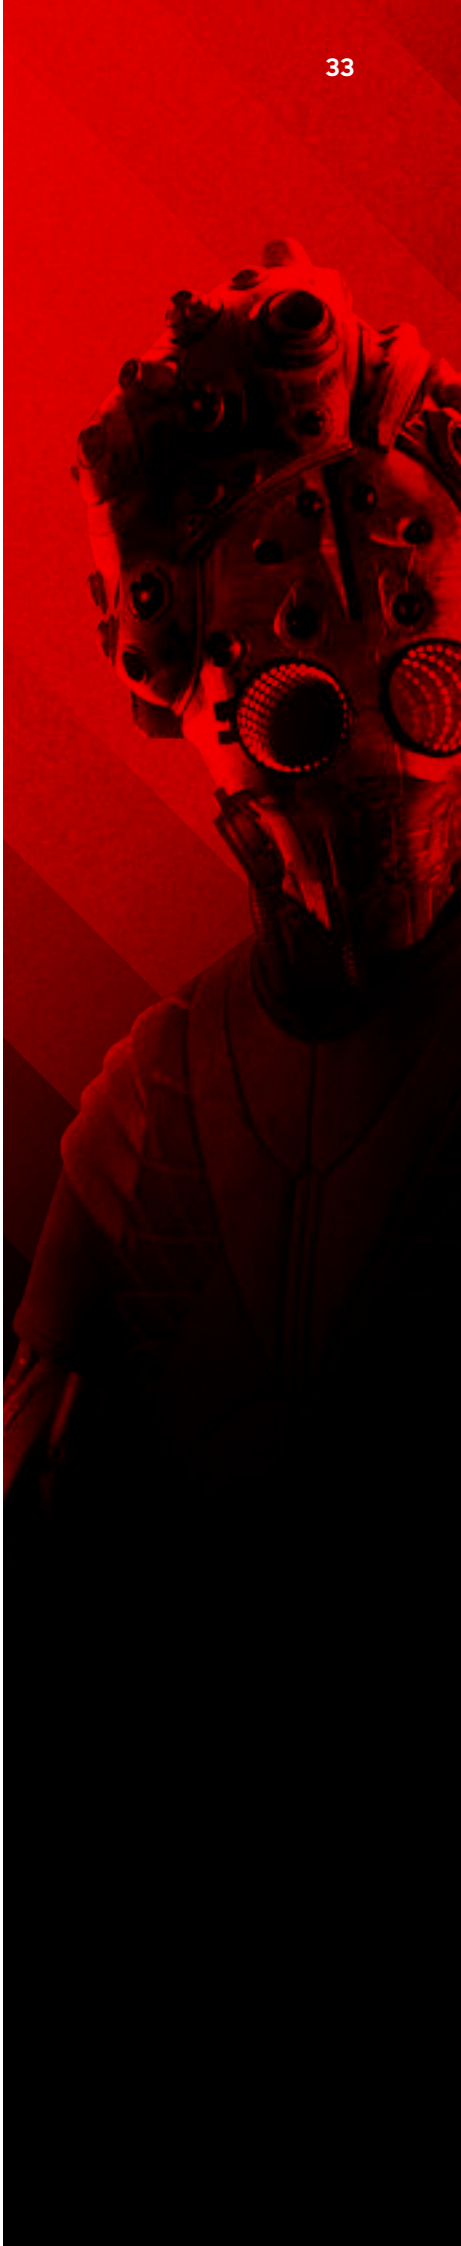

(MFA) codes, the adversaries typically log in to the compromised account from IP addresses associated with commercial residential proxy and VPN providers, such as Mullvad and NSOCKS. Unlike SCATTERED SPIDER, who often configured these services for IP addresses by matching the approximate geographic location of the legitimate account owner, CORDIAL SPIDER and SNARKY SPIDER typically do not configure their residential proxies and VPNs in any meaningful way. This often allows for geolocation- or other network-based anomaly detection.

To enable persistent access to compromised accounts, the adversaries configure mobile devices for MFA. CORDIAL SPIDER typically enrolls popular consumer mobile devices, whereas SNARKY SPIDER leverages the Genymobile Android emulator, a publicly available tool that allows users to operate connected Android devices from Linux, Windows, and macOS devices.

CrowdStrike OverWatch's visibility across CrowdStrike's vast telemetry puts threat hunters in a strong position to detect this activity. In addition to the Genymobile Android emulator, CrowdStrike OverWatch has observed threat actors install an SSO MFA application on VMware virtual machines (VMs) and Windows Quick Emulator (QEMU) devices, a low-prevalence behavior that serves as a high-confidence indicator of malicious activity and represents a valuable detection opportunity.

Having established persistent access to compromised IdP accounts, the adversaries quickly pivot to integrated SaaS applications, such as Microsoft 365 and Google Workspace. While SNARKY SPIDER manually downloads files from these platforms, CORDIAL SPIDER programmatically exfiltrates data using a likely custom Python-based tool. This often results in high-volume simultaneous `FileAccessed` or `SearchQueryPerformed` events associated with Python user-agent strings, which is likely to be anomalous in most environments.

Figure 12 shows CORDIAL SPIDER and SNARKY SPIDER techniques that can inform identity-oriented anomaly detection. Rather than relying solely on known attack signatures, detection content can flag unusual user behavior that falls outside what would normally be expected for that individual.

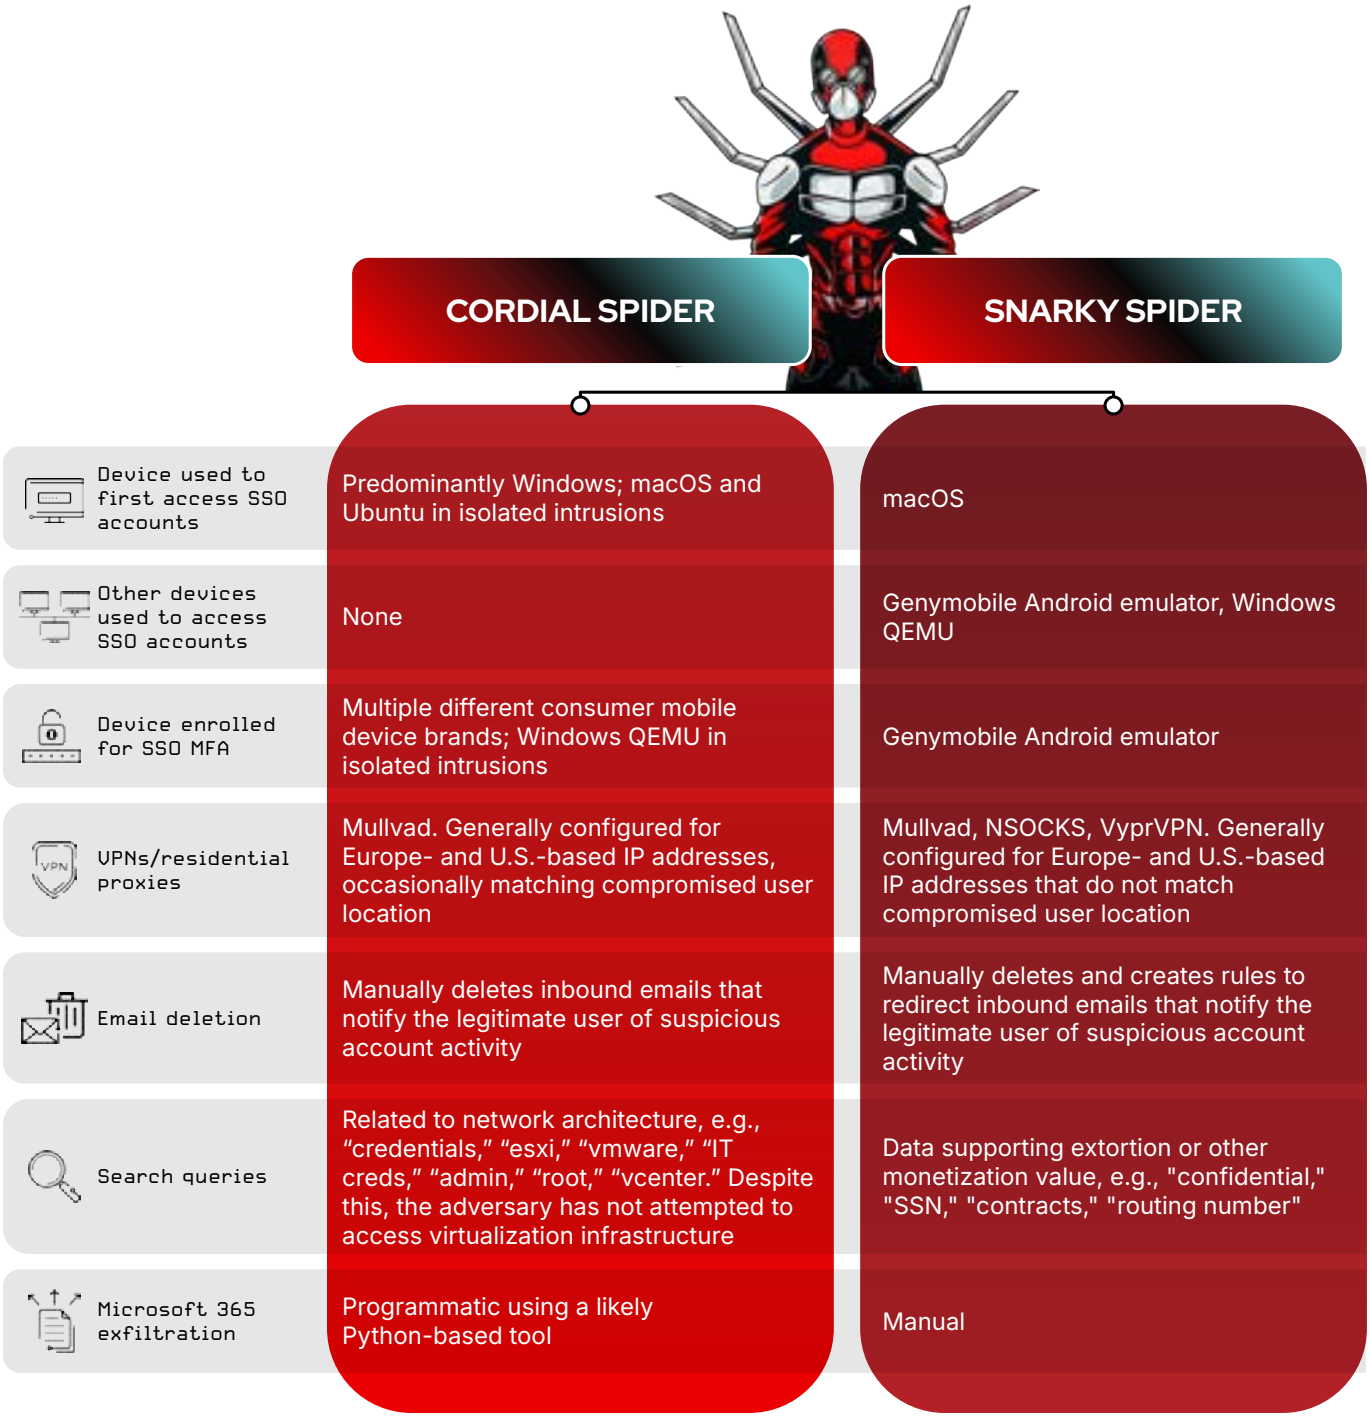

Figure 12. CORDIAL SPIDER and SNARKY SPIDER comparison

One challenge facing network defenders in CORDIAL SPIDER and SNARKY SPIDER intrusions is the speed at which adversaries can progress from initial access to data theft. **In one incident, SNARKY SPIDER moved from account takeover to data exfiltration in under five minutes.** These intrusions exploit the very feature that makes IdP and SaaS applications attractive technologies in enterprise environments. By compromising the credentials to a single account, eCrime adversaries like CORDIAL SPIDER and SNARKY SPIDER can rapidly access and exfiltrate sensitive data without moving laterally or escalating privileges.

How CrowdStrike OverWatch Detects CORDIAL SPIDER Vishing Activity

To track vishing, CrowdStrike OverWatch hunters focus on what the adversary needs to do next. For an adversary like CORDIAL SPIDER, the ultimate objective is exfiltrating data from SSO-integrated SaaS applications. Achieving this requires a series of prerequisite steps, including enrolling new devices for MFA to establish persistent access to compromised identities. This is where CrowdStrike OverWatch hunts.

In February, this strategic focus proved decisive. Despite successfully vishing a user into authenticating via a spoofed IdP, CORDIAL SPIDER could not outpace CrowdStrike OverWatch threat hunters, who detected the malicious activity in four minutes and alerted the customer to respond prior to data exfiltration.

» CORDIAL SPIDER COULD NOT OUTPACE CROWDSTRIKE OVERWATCH THREAT HUNTERS, WHO DETECTED THE MALICIOUS ACTIVITY IN FOUR MINUTES AND ALERTED THE CUSTOMER TO RESPOND PRIOR TO DATA EXFILTRATION.

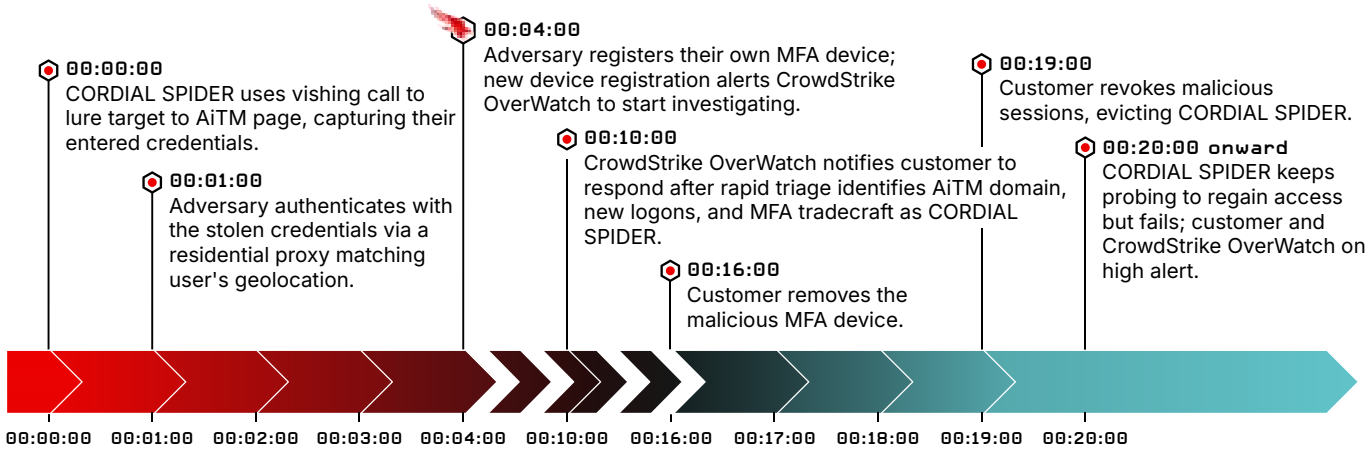

Figure 13. CORDIAL SPIDER detection timeline

OUTLOOK:

Vishing Is Becoming a Preferred Gateway for eCrime Activity

Vishing has established itself as one of the most popular initial access techniques among eCrime adversaries. The threat is accelerating. CORDIAL SPIDER and SNARKY SPIDER have clearly demonstrated the ability to pivot to data exfiltration within minutes of an account takeover.

Vishing calls in isolation cannot be detected. CrowdStrike OverWatch must focus on contextual activity immediately before and after the call. For some adversaries, this means detecting malicious activity on the host. For example, CURLY SPIDER leverages vishing calls to persuade the targeted user to establish a Microsoft Quick Assist session that the adversary uses to download further tooling. In the cases of CORDIAL SPIDER and SNARKY SPIDER, this entails detecting artifacts associated with the IdP compromise and follow-on access to SaaS applications.

Recommendations

- Instruct employees to ignore calls to their personal devices from individuals claiming to be IT staff who want to configure SSO passkeys, unless this practice is explicitly required for business operations.
- Enforce phishing-resistant MFA methods, such as FIDO2.
- Conduct regular threat hunting operations focused on identification of anomalous user behavior. This could include:
  - Detecting sign-ins from unexpected locations
  - Monitoring for enrollment of atypical MFA devices, such as the Genymobile Android emulator and Windows QEMU
  - Alerting on mass download of files from cloud-based storage suites originating from unusual IP addresses

Cloud Beyond the Perimeter: Cloud-Conscious Adversaries and Financial Targeting

The rapid adoption of cloud infrastructure has fundamentally transformed the cybersecurity landscape, creating new attack surfaces and challenging traditional security paradigms that relied heavily on network perimeters and endpoint controls. **Traditional perimeter defenses do not apply in cloud environments where adversaries operate with valid credentials or tokens within the trust boundary. CrowdStrike OverWatch threat hunting provides critical visibility into cloud control plane activity (including Entra ID sign-ins, device registrations, and application access patterns) which many organizations lack in their native cloud monitoring.**

The cases documented in this section illustrate how both financially motivated cybercriminals and nation-state threat actors have developed sophisticated techniques to abuse legitimate cloud services, from exploiting OAuth 2.0 authentication flows and hijacking computational resources at scale to deploying custom tooling designed specifically for cloud-native environments.

These evolving threats underscore the critical importance of implementing cloud-specific security strategies that go beyond traditional approaches. They require organizations to develop new capabilities for monitoring cloud control plane activity, detecting anomalous API patterns, and understanding the behavioral indicators that reveal malicious activity within the legitimate boundaries of cloud infrastructure.

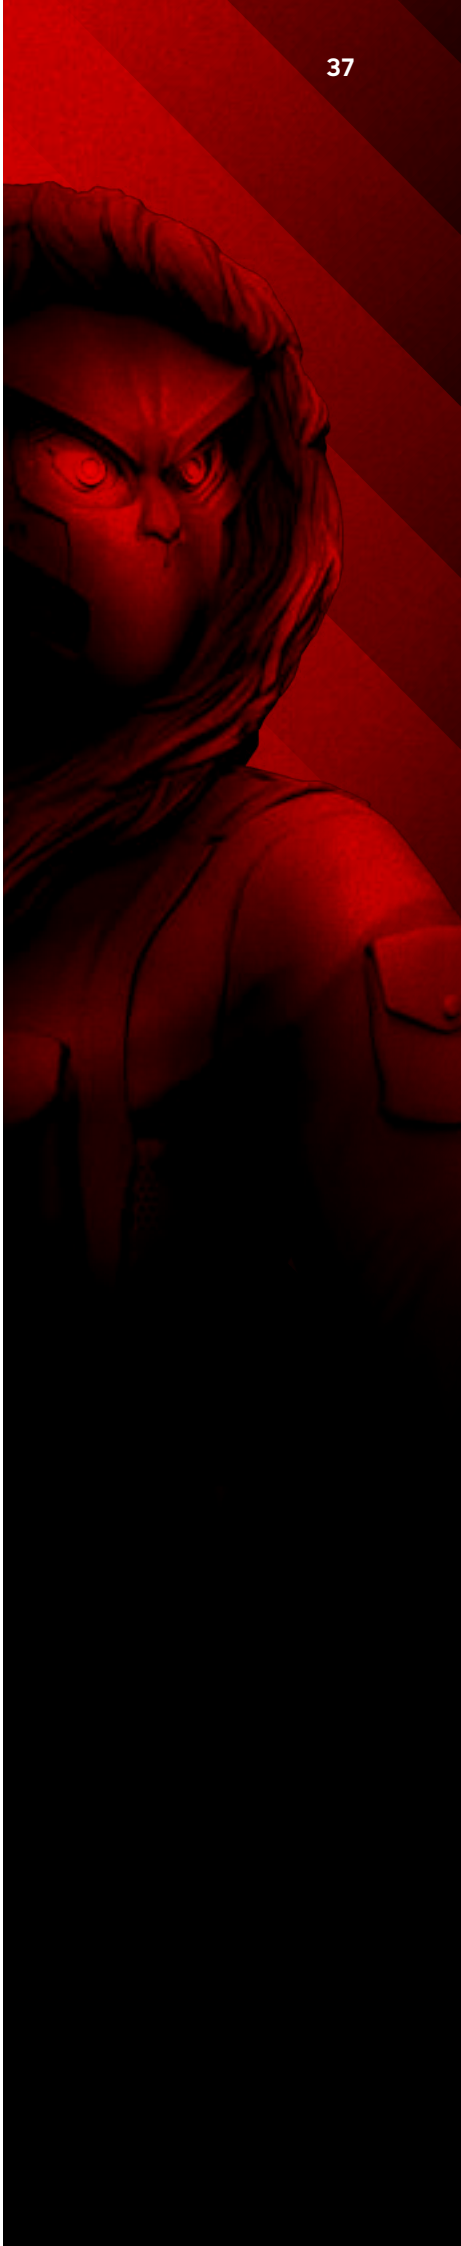

2026 OAuth 2.0 Phishing Operation Trends

Since at least February 2025, CrowdStrike has observed state-nexus adversaries abusing OAuth 2.0 flows for phishing to gain access to Entra ID accounts. This has primarily come in the form of the OAuth 2.0 device authorization grant flow,[6] also known as the device code flow, in which threat actors provide the victim user a valid device code with which to log in to their account.

In these instances, the threat actor can collect session tokens for access to the account, similar to most home entertainment applications allowing users to log in from their TV. State-nexus threat actors, including Russia-nexus adversary COZY BEAR, have used various methods to deliver these device codes, including HTML pages, redirections through a malicious Entra ID application, and email communication.

A diverse set of eCrime threat actors is now using newly established commodity phishing-as-a-service (PhaaS) offerings, most notably *EvilTokens* and *Kali365*, to conduct these phishing campaigns targeting device code flows. **Over the last six months, there has been a 15x increase in monthly device code phishing attempts.** Device code phishing is attractive because it allows the threat actor to obtain access to a target user's account via session tokens without requiring the user to input their credentials into a suspicious phishing page. Instead, it relies on legitimate Microsoft authentication services and flows that users implicitly trust. The dramatic increase aligns with the broader adoption of these PhaaS platforms.

Open-source infrastructure analysis identifies additional threat actors likely standing up similar infrastructure to support opportunistic device code phishing campaigns. Many of these financially motivated threat actors have adopted COZY BEAR's techniques, such as deploying dedicated device code and session management infrastructure, leveraging legitimate cloud-based hosting services, and delivering device code phishing pages via Entra ID application OAuth redirection at scale.

In addition to device code phishing, COZY BEAR has also abused the OAuth 2.0 authorization code flow on multiple occasions, targeting Entra ID accounts for email collection. This flow primarily exists for applications that require automated access to Entra ID accounts. When utilizing this flow, the user is provided a link to initiate the flow via Entra ID sign-in that redirects to a URL containing the authorization code. The target user is provided instructions to send back the resulting URL, including the authorization code, which the threat actor can then utilize to complete the flow and authenticate to the target user's account.

Despite the targeted and stealthy nature of this phishing technique, CrowdStrike OverWatch has been able to identify multiple such intrusion attempts, all attributed to COZY BEAR.

---

6 https://learn.microsoft.com/en-us/entra/identity-platform/v2-oauth2-device-code

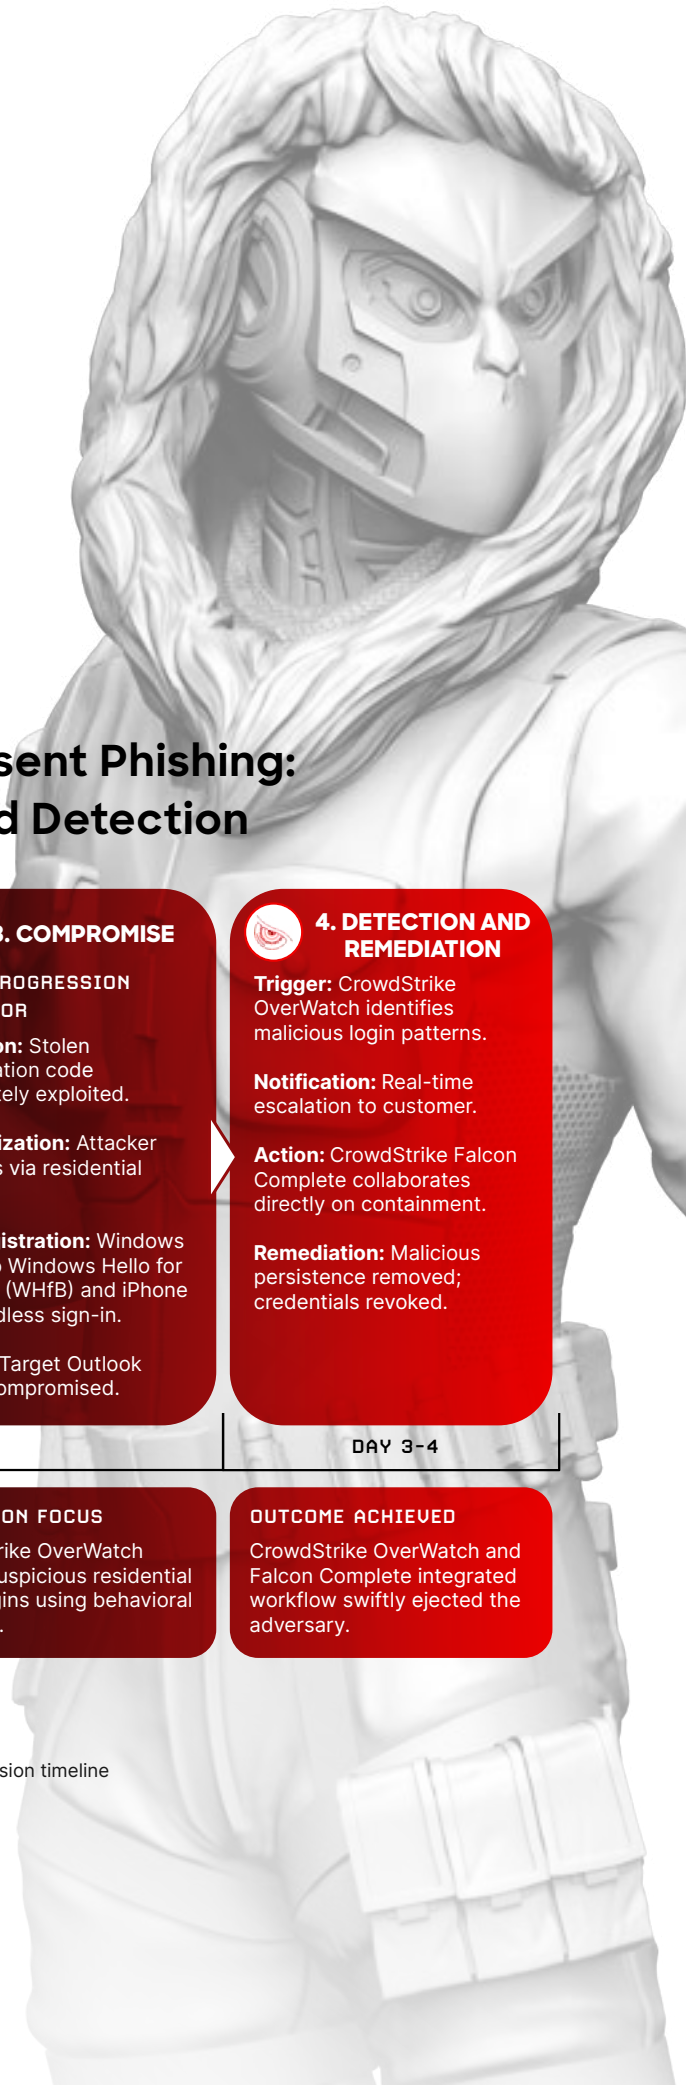

CASE STUDY:

COZY BEAR's OAuth 2.0 Techniques in Action

In mid-April 2026, CrowdStrike OverWatch and CrowdStrike Falcon® Complete identified COZY BEAR targeting a U.K.-based think tank through a combination of WhatsApp and OAuth 2.0 authorization code phishing, the same methodology described previously, executed against a high-value geopolitical target. This intrusion is representative of the sophisticated, targeted OAuth 2.0 operations COZY BEAR has refined over the past year, and it reflects the broader techniques that eCrime threat actors are now adopting at scale. The full details of this campaign are illustrated in Figure 14.

COZY BEAR OAuth Consent Phishing: Intrusion Timeline and Detection

1. INITIAL CONTACT

T+0 hrs: WhatsApp contact initiated by adversary.

Impersonation: German ambassador staff credentials.

T+24 hrs: Legitimate-looking meeting scheduled.

Pretext: Diplomatic engagement prior to CrowdStrike OverWatch visibility.

2. OAUTH EXPLOIT

Delivery: Phishing instructions sent directly via WhatsApp.

Execution: OAuth 2.0 authorization code flow initiated.

Auth: Victim authenticates, redirecting code to attacker infrastructure.

Cover: Fallback email set; invalid Teams link sent.

3. COMPROMISE

RAPID PROGRESSION INDICATOR

Escalation: Stolen authorization code immediately exploited.

Anonymization: Attacker accesses via residential proxy.

Dual Registration: Windows device to Windows Hello for Business (WHfB) and iPhone passwordless sign-in.

Access: Target Outlook Mobile compromised.

4. DETECTION AND REMEDIATION

Trigger: CrowdStrike OverWatch identifies malicious login patterns.

Notification: Real-time escalation to customer.

Action: CrowdStrike Falcon Complete collaborates directly on containment.

Remediation: Malicious persistence removed; credentials revoked.

DAY 1-2 | DAY 3 | DAY 3-4

WHY THIS WORKED

OAuth authorization codes bypass standard MFA protocols.

SPEED MATTERS

The critical transition from Phase 2 to Phase 3 took place in less than four hours.

DETECTION FOCUS

CrowdStrike OverWatch caught suspicious residential proxy logins using behavioral analytics.

OUTCOME ACHIEVED

CrowdStrike OverWatch and Falcon Complete integrated workflow swiftly ejected the adversary.

Figure 14. COZY BEAR phishing intrusion timeline

Cloud Resource Hijacking and Credential Theft for Financial Gain

Financially motivated adversaries are increasingly targeting cloud environments to access digital financial assets and hijack computational resources. **Over the last 12 months, cloud-conscious eCrime activity increased 171%.**

Unauthorized cryptomining operations can rapidly increase cloud compute costs, particularly when adversaries leverage automated tooling to provision resources across multiple regions within minutes.

Recovery is not always straightforward. Threat actors frequently employ techniques such as disabling instance termination via API modification to prevent organizations from reclaiming hijacked resources. They also establish persistence through cloud-native mechanisms such as backdoored cloud identity users, cross-account trust policy modifications, and atypical execution paths like the life cycle configurations of managed ML services. Many organizations lack the visibility to detect these techniques without dedicated behavioral hunting across the cloud control plane.

These operations represent substantial financial risk to organizations, with potential impacts including:

- **Resource hijacking:** $10,000-$100,000+ in unauthorized compute costs
- **Payment system compromise:** Direct financial theft in the millions
- **Data exfiltration:** Regulatory fines and intellectual property theft
- **Operational disruption:** Business continuity impacts and recovery costs

In addition, the proliferation of cloud infrastructure usage has caused various threat actors to target these environments as a stepping stone for access to other systems both inside and outside the cloud ecosystem. This is primarily conducted through credential-stealing operations from both secured and unsecured sources. SLIM SPIDER in particular has leveraged cloud access to steal credentials from cloud secrets manager solutions, which could then be leveraged to target digital financial assets.

CASE STUDY:

SLIM SPIDER's Cloud Intrusion Targeting Pix and Cryptocurrency Infrastructure

The following intrusion illustrates how financially motivated threat actors are increasingly using sophisticated, cloud-conscious tradecraft against high-value targets, underscoring the growing cost of that trend when left undetected. In late March 2026, CrowdStrike OverWatch observed SLIM SPIDER conducting a multi-stage intrusion at a Brazil-based financial institution, targeting the organization's cryptocurrency assets and instant payment accounts. SLIM SPIDER is a Brazil-based eCrime adversary who demonstrates an in-depth knowledge of Brazilian banking infrastructure, including cloud infrastructure for various providers and services.

To steal cloud credentials, SLIM SPIDER developed custom Bash scripts that queried the cloud instance metadata service, a built-in cloud service that provides temporary access credentials to running instances, using raw socket connections rather than conventional tooling. This approach allowed them to harvest temporary cloud credentials while minimizing their footprint. Once access was established, SLIM SPIDER enumerated all available secrets stored in the cloud credential manager and used `sed` to clone and modify secret-extracting scripts, swapping out target secret IDs to extract at least six secrets. The targeted names indicated a clear focus on credentials tied to digital financial assets.

Following exfiltration of digital asset custody secrets, SLIM SPIDER invoked `cast`, a component of the Foundry Ethereum developer toolkit, to derive the Ethereum wallet address associated with a stolen private key. Rather than relying on third-party libraries that could introduce detection risk, the threat actor implemented cloud-native cryptographic signing directly via OpenSSL within their Bash scripts. This deliberate choice reflected sophisticated operational security awareness and a nuanced understanding of cloud environments. SLIM SPIDER established access to nodes running in a cloud container service cluster, deploying implants masquerading as infrastructure-related binaries to blend with legitimate tooling. They then pivoted to Azure DevOps, likely using compromised credentials to execute malicious pipelines that deployed additional implants across a managed Kubernetes cluster. One such implant was named “spi,” likely referring to the Central Bank of Brazil's Sistema de Pagamentos Instantâneos, which serves as centralized infrastructure for Pix payments.

Infrastructure and Tooling

Analysis revealed that SLIM SPIDER operated various web-based panels to automate and streamline tasks in the intrusion chain. The exposed C2 panel displayed several compromised hosts from multiple Brazil-based banks and fintech organizations and likely exfiltrated archive files, including:

- **NEXUS // Scanner:** An API endpoint-scanning panel that uses Ollama AI to sort endpoints into 16 categories (including fintech, banking, payment, and cryptocurrency) and rank them based on availability and authentication options
- **Painel de Emails Entra ID:** An email reconnaissance panel that searches compromised Microsoft 365 mailboxes sorted into finance, admin, and Brazil categories
- **Painel Pix:** A transaction panel specifically designed to execute bulk unauthorized Pix transfers from compromised accounts, supporting all Pix key types

As organizations are hosting more critical services and their associated credentials in cloud services, intrusions such as this one underscore the significant potential impact financially motivated and cloud-conscious threat actors such as SLIM SPIDER can have. SLIM SPIDER's knowledge of the cloud attack surface allows them to target credentials associated with an organization's valuable digital currency assets, including custody credentials that control cryptocurrency wallets. Access to such assets can result in devastating financial loss for victims. This intrusion reflects a broader shift CrowdStrike OverWatch has observed over the past year: eCrime threat actors are demonstrating increasingly sophisticated cloud awareness, deliberately targeting the infrastructure and credentials that sit closest to high-value financial assets.

Based on this intrusion and ongoing tracking of SLIM SPIDER, CrowdStrike OverWatch has been able to develop a variety of behavioral detection indicators specifically designed to identify SLIM SPIDER's unique tradecraft for future intrusions, including their unique mechanisms for collecting secure credentials via raw socket connections and unsecured credentials from compromised hosts. These detection capabilities are now deployed across the CrowdStrike customer base to proactively identify similar cloud-focused eCrime operations.

CASE STUDY:

eCrime Threat Actor Conducts Cloud Resource Hijacking via Multi-Vector Deployment

On January 8, 2026, CrowdStrike OverWatch detected an eCrime threat actor hijacking cloud resources at a U.S. technology company to conduct a cryptomining operation, a clear example of the financially motivated cloud targeting CrowdStrike has observed increasing over the past year. After gaining initial access, the eCrime threat actor moved quickly and methodically, pursuing three parallel attack vectors to maximize the operation's reach while simultaneously establishing cross-account persistence to ensure their hijacking was not easily disrupted. The technical details of this campaign are illustrated in Figure 15.

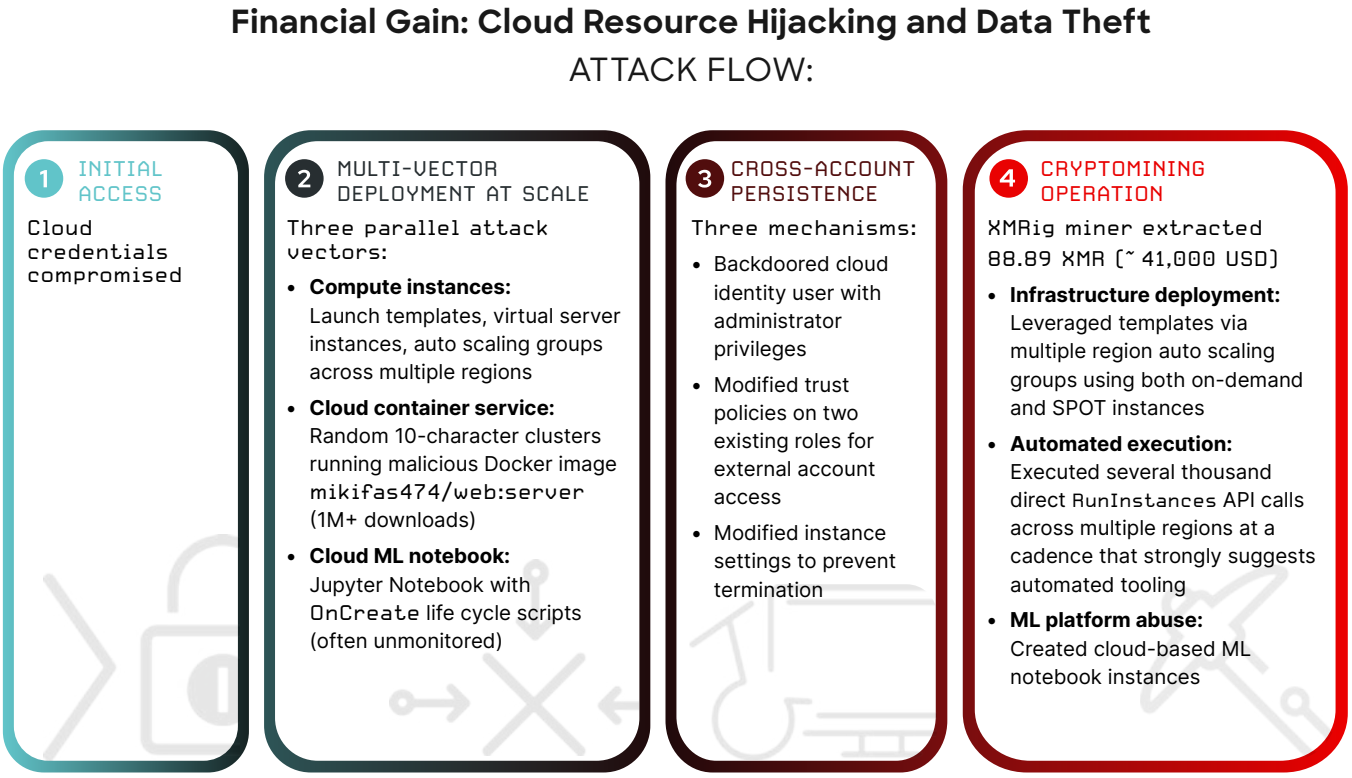

Figure 15. Cloud resource hijacking deployment flow

At the center of this attack was a malicious container configured to run XMRig, a popular cryptocurrency miner, using a container image that was downloaded more than one million times. To scale the mining capacity, the eCrime threat actor deployed container clusters with randomly generated 10-character names and registered task definitions, a blending technique designed to evade detection. The configured wallet had mined 88.89 XMR, worth approximately 41,000 USD, underscoring the real financial damage cloud resourcing hijacking can inflict in a short period of time.

Detecting this type of intrusion requires more than monitoring individual API calls to the cloud credential manager, which can appear as legitimate administrative activity. CrowdStrike OverWatch's behavioral analysis identified the anomalous pattern of rapid resource deployment, secret enumeration, and custom credential harvesting scripts. The team successfully connected activity across the environment that no single alert alone would have surfaced. This multilayered approach, combining cryptomining, privilege escalation, cross-account pivoting, and potential data exfiltration, required CrowdStrike OverWatch's advanced cross-domain correlation to map the comprehensive attack path across API call sequences that individually appeared benign but collectively revealed sophisticated eCrime operations.

OUTLOOK:

Behavioral Hunting Is Becoming Essential as Adversaries Refine Stealthy Cloud Attack Techniques

Detecting post-authentication abuse requires behavioral hunting for anomalous cloud API patterns, not just credential compromise detection. The registration of multiple devices within short time frames from unusual geographic locations, the use of nonstandard user-agent strings for passwordless credential enrollment, and authentication to mail applications from residential proxy services are all behavioral indicators that require active threat hunting to surface. Without CrowdStrike OverWatch's continuous monitoring, COZY BEAR's stealthy mail collection operations would have persisted undetected for months or years, enabling sustained intelligence collection against organizations critical to Western policy and strategy.

The same principle applies to financially motivated cloud threats. Organizations without behavioral hunting across their cloud control plane will struggle to detect LLMJacking and resource hijacking activity. Each individual API call appears legitimate in isolation. Only the pattern, velocity, and context reveal malicious intent, and identifying that requires continuous, expert-led monitoring.

As cloud adoption accelerates and organizations migrate critical infrastructure and financial operations to cloud platforms, both financially motivated and nation-state adversaries will continue refining their cloud-specific attack techniques. Attack surfaces will expand alongside the growing complexity of multi-cloud environments, and adversaries will look to exploit serverless computing platforms, container orchestration, and ML services as new vectors. Organizations that rely solely on cloud-native security controls without continuous behavioral hunting will find themselves increasingly exposed.

Recommendations

Organizations must implement comprehensive cloud security strategies that extend beyond traditional endpoint protection to include continuous behavioral monitoring of cloud control plane activity. This includes:

- Deploying advanced threat hunting capabilities specifically designed to detect anomalous cloud API patterns, unusual authentication flows, and suspicious resource provisioning activity that may indicate compromise
- Establishing robust identity and access management policies with conditional access controls, implementing least-privilege principles for cloud service accounts, and regularly auditing cross-account trust relationships and device registrations
- Investing in cloud-native security tools that provide visibility into container environments, serverless functions, and DevOps pipelines while maintaining detailed logging and monitoring of all cloud infrastructure changes and secret access patterns to enable rapid detection of and response to sophisticated cloud-based threats

OVERCAST PANDA: Close Access Operations and the Endpoint Threat Beyond Network Defenses

China-nexus adversary OVERCAST PANDA represents one of the most operationally sophisticated and persistent threats facing organizations with personnel traveling to China. Unlike the vast majority of cyber threat actors who operate at a distance, deploying phishing lures, exploiting vulnerabilities, or harvesting credentials, OVERCAST PANDA employs a fundamentally different doctrine: physical proximity to the target device. This approach allows the adversary to bypass conventional network-based security controls entirely, rendering perimeter defenses, endpoint detection, and user awareness training ineffective against the initial compromise vector.

Between March and May 2026, CrowdStrike OverWatch identified and disrupted a series of close access intrusions in which OVERCAST PANDA deployed their proprietary *FlowCloud* backdoor on victims' devices through direct physical access while those individuals were in China. CrowdStrike OverWatch's ability to surface intrusions that left minimal forensic evidence and produced no conventional network-layer indicators at the point of compromise reflects the depth of continuous behavioral monitoring that distinguishes managed threat hunting from automated detection alone. In each case, threat hunters identified anomalous endpoint activity consistent with *FlowCloud*'s distinctive execution chain, enabling immediate customer notification and rapid containment before the adversary achieved their collection objectives.

Figure 16. Summary of OVERCAST PANDA's target scope and attributed malware

A Methodical Tradecraft Designed to Exploit Business Travel

OVERCAST PANDA's close access methodology reflects the operational discipline and patience of a mature, well-resourced state intelligence program. The adversary does not compromise targets haphazardly. They likely exploit the predictable rhythms of international business travel, identifying precise windows during which high-value individuals are separated from their devices. Conference schedules, hotel accommodations, and meal hours all create recurring opportunities for an adversary with physical presence in the operating environment.

In March 2026, CrowdStrike OverWatch detected OVERCAST PANDA deploying *FlowCloud* on multiple devices belonging to employees of a U.S.-based agricultural biotechnology company attending a major industry conference in the Hainan province of China. Several compromises occurred at 19:52 and 21:57 China Standard Time (CST). These times align with dinner service at a conference of this nature, when attendees are most likely to leave devices unattended in hotel rooms. In at least one incident, the targeted laptop was confirmed to be inside a hotel room at the time of compromise, with no network-based exploitation activity observed. On each occasion, OVERCAST PANDA also executed an unexpected system restart immediately before *FlowCloud* began executing, a strong indicator of a physical device reboot consistent with the use of bootable media to directly access the host's storage outside of the running operating system.

This access methodology is deliberate and effective. By booting the target device into a secondary environment via removable media, OVERCAST PANDA bypasses the host's running operating system and all associated security controls. The adversary then directly accesses the device's storage, installs the *FlowCloud* implant components, and configures them for execution upon the device's normal restart. The target then powers the device back on, resumes their activities, and remains unaware that a persistent backdoor has been established. In a mid-2026 intrusion targeting an individual associated with a U.S.-based media entity traveling in China, CrowdStrike OverWatch identified the insertion of a bootable USB device shortly before the attempted compromise. This further corroborates the physical access tradecraft.

FlowCloud itself is a sophisticated, purpose-built tool uniquely associated with OVERCAST PANDA. Once deployed, *FlowCloud* provides the adversary with persistent, covert remote access to the target device, enabling keylogging, screen capture, file collection, and credential harvesting. The implant's components are engineered to blend into legitimate system file structures, and recent variants employ commercial-grade software protection to complicate forensic analysis. C2 communications are routed through legitimate cloud infrastructure, further obscuring the adversary's operational presence. The level of tooling investment reflected in *FlowCloud*'s design and continued development underscores OVERCAST PANDA's status as a capable, professionally operated state intelligence collection adversary.

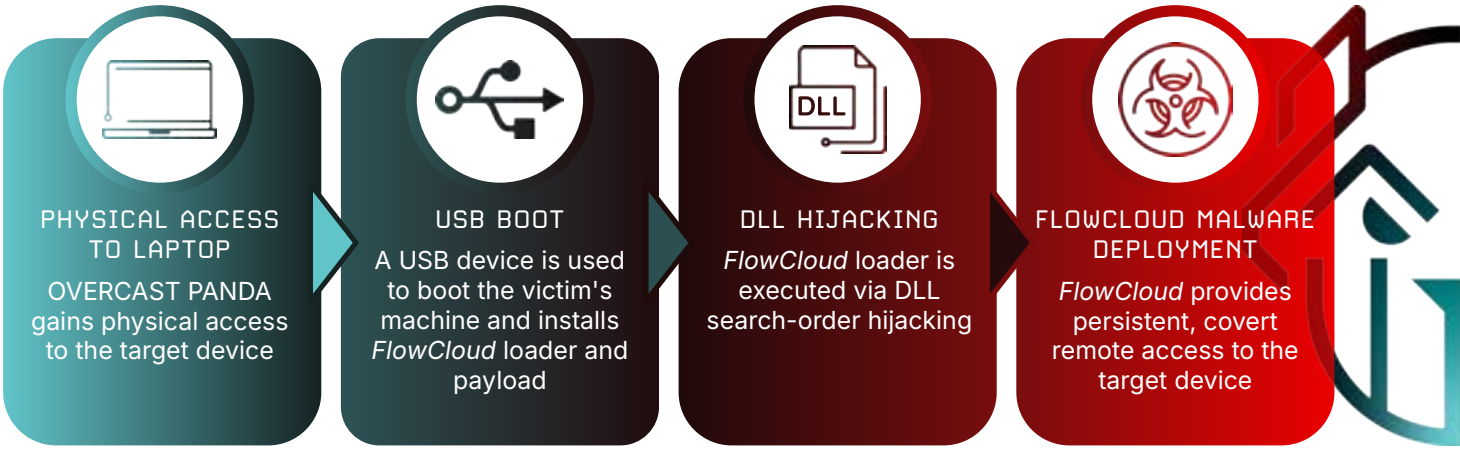

Figure 17. Summary of OVERCAST PANDA's tactics, techniques, and procedures (TTPs) in close access campaigns

Deliberate Victim Selection Aligned with Chinese State Priorities

The targeting pattern documented across OVERCAST PANDA's 2026 close access operations is not opportunistic; it is, however, expansive. Each intrusion very likely reflects a deliberate selection of individuals whose professional roles provide access to information directly relevant to Chinese state intelligence priorities. Any individual traveling within China is a potential target if they have access to such information.

The March 2026 targeting of agricultural biotechnology personnel is illustrative. China's 14th and 15th Five-Year Plans explicitly identify agricultural modernization and self-sufficiency as national strategic objectives, with particular emphasis on advancing seed technology and crop protection research. By targeting specialists attending an international agricultural conference, OVERCAST PANDA pursues collection objectives that map precisely to those stated national priorities. These individuals often travel with proprietary research data, collaboration agreements, and strategic intellectual property on the very devices they carry.

The adversary's sustained focus on foreign media organizations is likely equally deliberate. From at least July 2025 through May 2026, CrowdStrike OverWatch documented a persistent pattern of OVERCAST PANDA targeting individuals associated with foreign media entities operating within or visiting China. This targeting likely reflects an enduring Chinese state interest in understanding and monitoring how foreign narratives about China are constructed from within the country. The individuals targeted are not random journalists; they are the people who shape how China is perceived and reported upon in their home nations.

The Hainan province has emerged as a recurring operational environment across these intrusions. Multiple distinct close access deployments have been documented against conference attendees and hotel guests in this region, suggesting OVERCAST PANDA has established reliable operational infrastructure there and likely benefits from consistent physical access to venues frequented by visiting foreign professionals. The concentration of activity in this specific geography indicates a level of pre-operational planning and in-country presence that goes well beyond reactive targeting.

CrowdStrike OverWatch's detection of these intrusions is particularly significant. Because the initial compromise occurs outside the normal network environment, before the device is returned to the organization's infrastructure and before any network-based malicious activity is generated, conventional security tools relying on network traffic analysis or cloud log monitoring would not surface this activity. Only continuous, expert-led monitoring of endpoint behavioral telemetry, applied at the moment of and immediately following the device's return to normal operation, enables the identification of *FlowCloud*'s execution chain. CrowdStrike OverWatch delivered this capability across each of these intrusions, demonstrating the clear edge that professional managed threat hunting delivers continuously.

OUTLOOK:

Close Access Operations Expose Gaps in Standard Corporate Security Programs

The operational security implications of OVERCAST PANDA's close access methodology are serious and frequently underestimated by organizations operating standard corporate security programs. Phishing awareness training, strong password policies, and MFA do not defend against an adversary that bypasses the operating system entirely to install a backdoor on an unattended device. The time required to execute this methodology means that even brief, routine separations from a device can create the opportunity OVERCAST PANDA needs to compromise the host.

OVERCAST PANDA will almost certainly continue deploying *FlowCloud* against individuals traveling to China over the next six months. The adversary's consistent operational tempo, their continued refinement of close access tradecraft, and the sustained alignment between their targeting and Chinese state intelligence priorities all indicate this campaign will persist. Executive-level personnel and technical specialists affiliated with strategically important sectors, including media, agriculture, biotechnology, manufacturing, and international organizations, should be considered at elevated risk.

Organizations should treat travel to China by personnel with access to sensitive intellectual property, proprietary research, or information of strategic value as a material security risk. It is imperative to implement the controls below before the next trip, not after an incident. While OVERCAST PANDA's operations specifically target individuals traveling to China, these types of operations can occur in any travel situation. Organizations should implement the following recommendations regardless of travel destination, particularly for travel within China.

Recommendations

Implement the following enhanced protocols for employees traveling for business or with business assets:

- **Require** full-disk encryption with pre-boot authentication for all devices
- **Disable** Wi-Fi auto-connect features
- **Implement** BIOS/UEFI passwords to prevent boot-level tampering
- **Prohibit** employees from traveling with devices containing proprietary research data or strategic business information
- **Maintain** continuous physical custody of devices, and never leave laptops unattended
- **Brief** executive and technical personnel, particularly those with access to novel research or strategic intellectual property, on the elevated targeting risks posed during travel
- **Enable** USB blocking via CrowdStrike Falcon® Device Control policies for devices entering targeted regions, even if they are designated as burner hosts

Conclusion

The past year marked a defining chapter in the evolution of threat hunting. From the acceleration of vulnerability exploitation to the fast-spreading software supply chain attack, threat hunters face an unprecedented need for speed and precision. Whether motivated by financial gain, espionage, or long-term access, evasive adversaries exploit trusted relationships and move beyond traditional attack surfaces to evade detection.

CrowdStrike OverWatch threat hunters have revealed how adversaries no longer operate in silos. They execute cross-domain attacks across AI tools and identity, endpoint, and cloud environments, using both hands-on-keyboard tradecraft and automated attacks that challenge traditional security tools. Defenders are rising to meet the challenge with faster detection, deeper context, and coordinated defense rooted in intelligence.

This report underscores why proactive, intelligence-driven hunting is essential. Security teams must integrate telemetry across the enterprise, operationalize threat intelligence, and leverage automation to scale human capability. CrowdStrike OverWatch does exactly this, using AI to help threat hunters continuously identify and prioritize the data most closely linked with adversarial activity, and ensuring human-led expertise is always the final determinant.

It is not enough to respond; defenders must anticipate, pivot, and relentlessly pursue the adversary.

As adversaries sharpen their capabilities, the CrowdStrike Counter Adversary Operations team remains resolute in detecting and disrupting the world's most sophisticated threat actors. This commitment ensures that wherever adversaries go, the team is already there.

Recommendations

1 Secure AI to reduce emerging business and operational risk

AI is increasingly becoming embedded across business, development, and cloud environments, creating a high-value target for adversaries. Threat actors are abusing corporate LLM access, exploiting AI-related server software, targeting AI-centric development environments, and using AI to accelerate phishing, reconnaissance, and vulnerability exploitation. Organizations should secure AI systems like critical infrastructure. This starts with inventorying AI applications, model endpoints, API keys, service accounts, AI developer tools, and cloud/GPU resources. Security teams should enforce least-privilege access to models and APIs, protect AI credentials, monitor abnormal LLM usage and cost spikes, consider segmenting LLM infrastructure from internal networks as appropriate, assess AI dependencies for supply chain risk, and harden AI development environments. They should also incorporate AI-targeted and AI-enabled threats into incident response plans and threat hunting workflows.

2 Treat identity and SaaS as primary attack surfaces

Identity and SaaS platforms sit at the center of enterprise access, data, and business operations, making them prime targets. Adversaries are using vishing, OAuth abuse, device code phishing, and stolen session tokens to compromise accounts and rapidly access SaaS data. Organizations should enforce phishing-resistant MFA, monitor anomalous sign-ins and MFA device enrollment, apply least-privilege access for human and non-human accounts, and detect unusual SaaS activity such as mass file access, abnormal user agents, and suspicious token use.

3 Eliminate cross-domain blind spots to stop high-impact attacks

Today's most disruptive intrusions succeed by exploiting gaps between security domains, tools, and teams rather than weaknesses in any single control. Adversaries chain activity across endpoints, cloud environments, SaaS applications, identity systems, developer workflows, and unmanaged hosts to evade detection. Organizations should consolidate telemetry, apply cross-domain correlation through extended detection and response (XDR) and next-gen SIEM workflows, and automate enrichment with threat intelligence to see full attack paths and accelerate response.

4 Secure the software supply chain and developer workflows

Trust in software updates, open-source dependencies, and development pipelines has become a critical business dependency and a prime target for adversaries. Malicious packages, stolen maintainer credentials, compromised CI/CD pipelines, and trusted developer tools are enabling adversaries to steal credentials and pivot into cloud environments. Organizations should harden developer environments, enforce code signing and dependency validation, scan repositories and packages for anomalies, protect developer identities, and assess third-party risk to reduce the likelihood that trusted software becomes a vehicle for compromise.

5 Proactively reduce the attack surface across the enterprise

Adversaries are moving faster than traditional vulnerability management cycles can support. They exploit public vulnerabilities within days, abuse cloud and identity misconfigurations, target exposed services, and turn trusted developer and AI workflows into paths for access and credential theft. Organizations should continuously identify and prioritize the exposures most likely to create real attack paths, including critical vulnerabilities, overprivileged accounts, risky cloud configurations, unmanaged assets, exposed internet-facing systems, and weaknesses in SaaS, AI, and developer environments. By hardening systems across the board, security teams can reduce the attack surface, contain the blast radius, and stop intrusions before they gain momentum.

6 Prioritize proactive threat intelligence and hunting

When attacks unfold in minutes or seconds, reactive defense is no longer enough. An intelligence-driven approach helps organizations understand which adversaries are targeting them, how they operate, and where they are likely to strike next. By applying threat intelligence and adversary tradecraft analysis through proactive hunting, teams can identify stealthy footholds (including AiTM activity, suspicious MFA enrollment, cloud abuse, and supply chain anomalies) before attacks escalate. Specialized AI agents can further accelerate intelligence analysis, hunting, triage, and response, helping defenders turn insight into action earlier in the attack life cycle.

7 Strengthen human resilience against social engineering and rapid intrusions

As adversaries increasingly rely on phishing, vishing, and trust abuse to bypass technical controls, human decision-making remains critical to preventing breaches. Organizations should reinforce awareness programs with real-world adversary tactics, including IT impersonation, AiTM phishing, suspicious MFA requests, and social engineering designed to capture credentials or enroll unauthorized devices.

For security teams, preparedness under pressure is essential. Regular tabletop exercises and red/blue team operations help organizations identify gaps in detection, decision-making, and response, ensuring teams can act quickly and effectively when attacks unfold. Continuous rehearsal strengthens organizational resilience and reduces the likelihood that minor failures escalate into major incidents.

CrowdStrike Falcon Platform

AI and Cloud-Native

Leverages the network effect of crowdsourced security data while eliminating the management burden of cumbersome on-premises solutions

Single Lightweight Sensor

Provides frictionless and scalable deployment and stops all types of attacks while eliminating sensor bloat and scheduled scans

Charlotte AI

Powers the CrowdStrike portfolio of agentic and generative AI capabilities across the Falcon platform, including natural-language chat, Charlotte AI AgentWorks for custom agent development, and specialized turnkey agents

Charlotte Agentic SOAR

Provides native security orchestration, automation, and response (SOAR) capabilities within the Falcon platform to automate security workflows, reduce manual effort, and enable faster, more consistent threat response

CrowdStrike Enterprise Graph

Unifies and contextualizes security telemetry across domains, connecting users, assets, behaviors, and adversary activity into a single shared view of the enterprise to give humans and AI agents the context to see attack paths, reason over complexity, and act faster; specialized graphs include:

- **Asset Graph:** Solves one of the most complex customer problems today: identifying assets, identities, and configurations accurately across all systems (including cloud, on-premises, mobile, internet of things, and more) and connecting them together in a graph form, including cloud, on-premises, mobile, internet of things (IoT), and more
- **Intel Graph:** Enables security teams to proactively defend against emerging threats with intelligence-driven insights by mapping relationships between threat actors, tactics, vulnerabilities, and real-world attacks
- **Threat Graph:** Uses cloud-scale AI to correlate trillions of data points from multiple telemetry sources to identify shifts in adversarial tactics and map tradecraft to automatically predict and prevent threats in real time across CrowdStrike's global customer base

Falcon Foundry

Allows customers and partners to easily build custom, low-code applications that harness the data, automation, and cloud-scale infrastructure of the Falcon platform to solve your toughest cybersecurity challenges

CrowdStrike Marketplace

Offers an enterprise marketplace of technology partners where you can discover, try, buy, and deploy trusted CrowdStrike and partner applications that extend the CrowdStrike Falcon platform, without adding agents or increasing complexity

CrowdStrike Products

Endpoint Security

FALCON PREVENT | NEXT-GENERATION ANTIVIRUS

Protects against all types of threats, from malware and ransomware to sophisticated attacks, and deploys in minutes, immediately protecting your endpoints

FALCON INSIGHT XDR | DETECTION AND RESPONSE FOR ENDPOINT AND BEYOND

Offers industry-leading, unified endpoint detection and response (EDR) and XDR with enterprise-wide visibility to automatically detect adversary activity and respond across endpoints and key attack surfaces

FALCON FIREWALL MANAGEMENT | HOST-BASED FIREWALL

Delivers simple, centralized host firewall management, making it easy to manage and control host firewall policies

FALCON DEVICE CONTROL | USB SECURITY

Provides the visibility and precise control required to enable safe usage of USB devices across your organization

FALCON FOR MOBILE | MOBILE THREAT DETECTION

Protects against threats to iOS and Android devices, extending endpoint security to your mobile devices, with advanced threat protection and real-time visibility into app and network activity

FALCON FORENSICS | FORENSIC CYBERSECURITY

Allows you to quickly respond and recover with automated forensic data collection, enrichment, and correlation

FALCON INSIGHT FOR ChromeOS | PROTECTION FOR ChromeOS DEVICES

Delivers industry-first native detection and response for ChromeOS devices without requiring additional agents or mobile device management (MDM) solutions, providing unified visibility through the Falcon console

FALCON FOR LEGACY SYSTEMS | PROTECTION FOR LEGACY OPERATING SYSTEMS

Delivers anti-malware protection for Windows XP, Server 2003, Vista, and more while minimizing impact on resource-constrained systems and integrating with the Falcon platform

FALCON ADVERSARY OVERWATCH: ENDPOINT | THREAT HUNTING

Provides 24/7 managed endpoint threat hunting, proactively monitoring your environment to identify novel attacks, misuse of legitimate tools, credential compromise, insider threats, and adversary pivots from endpoint activity into other domains

Counter Adversary Operations

FALCON ADVERSARY OVERWATCH | INTELLIGENCE-LED THREAT HUNTING

Provides 24/7 protection across endpoints, identities, cloud workloads, and next-gen SIEM, delivered by expert threat hunters powered by AI and threat intelligence to detect advanced intrusions and expose adversary tradecraft, credential misuse, and vulnerability exploitation

FALCON ADVERSARY INTELLIGENCE | SOC AUTOMATION

Cuts investigation time from days to minutes across the SOC with personalized, real-time threat intelligence, automating malware analysis and response through workflows and integrations while continuously monitoring the open, deep, and dark web for fraud and emerging threats

FALCON ADVERSARY INTELLIGENCE PREMIUM | ADVERSARY INTELLIGENCE

Delivers industry-leading intelligence reporting on 290+ adversaries globally and enables you to defend against AI-powered adversaries with threat AI agents built to reason across data, hunt for threats, and act decisively to automate and accelerate complex analyst workflows

FALCON COUNTER ADVERSARY OPERATIONS ELITE | ON-DEMAND ANALYST

Provides an assigned analyst who leverages AI-powered investigative and threat hunting tools, enhanced by deep adversary intelligence, to detect and disrupt adversaries across your IT environment and beyond

Cloud Security

FALCON CLOUD SECURITY: PROACTIVE SECURITY

Provides unified cloud security posture management (CSPM), including data security posture management (DSPM), application security posture management (ASPM), AI security posture management (AI-SPM), cloud infrastructure entitlement management (CIEM), AI model scanning, infrastructure as code (IaC) scanning, and cloud compliance posture insights

FALCON CLOUD SECURITY: CLOUD RUNTIME PROTECTION

Cloud Detection and Response: Breach protection including threat intelligence, agent-based and agentless cloud detection and response (CDR), and multi-cloud support

Cloud Detection and Response with Containers: Includes the features and capabilities of Cloud Detection and Response and adds container and Kubernetes protection, which can be deployed across on-premises, hybrid, and multi-cloud environments

Cloud Detection and Response with Managed Containers: Container security and runtime protection for cloud service provider-managed containers, including threat intelligence, CDR, container image security, and Kubernetes protection

FALCON CLOUD SECURITY: CNAPP

Includes the features and capabilities of both Proactive Security and Cloud Runtime Protection for Falcon Cloud Security

FALCON ADVERSARY OVERWATCH: CLOUD | THREAT HUNTING

Provides 24/7 managed cloud threat hunting, proactively monitoring your cloud control plane, host operating system, and data plane to detect advanced intrusions and expose adversary tradecraft, credential misuse, and vulnerability exploitation across cloud environments

SaaS Security

FALCON SHIELD | SaaS SECURITY

Unifies SaaS security posture management (SSPM), identity protection, threat detection, and AI governance in a single AI-native platform; continuously monitors SaaS applications, identities, connected apps, data, and AI agents to uncover risks and misconfigurations, detect active threats, and prioritize remediation; and combines prevention, detection, and automated response to close SaaS security gaps and stop breaches before they happen

AI Detection and Response

FALCON AI DETECTION AND RESPONSE | AIDR

Protects employee AI adoption and AI development at runtime by securing the prompt and agentic interaction layer with unified visibility, real-time threat detection, data protection, access controls, and automated response across endpoints, SaaS applications, and cloud environments

Next-Gen Identity Security

FALCON NEXT-GEN IDENTITY SECURITY

Secures human, non-human, and AI identities by combining initial access prevention, modern secure privileged access, identity threat detection and response (ITDR), SaaS identity security, and agentic identity protection to stop identity-driven breaches across domains

FALCON IDENTITY THREAT DETECTION

Provides unified visibility across hybrid identities and AI-driven threat detection to expose identity-based threats before they escalate

FALCON IDENTITY THREAT PROTECTION

Secures hybrid identities with AI-driven threat detection and response to stop identity-based attacks in real time

FALCON PRIVILEGED ACCESS

Eliminates standing privileges by enforcing just-in-time access based on real-time risk, removing the complexity of traditional privileged access management (PAM) solutions

FALCONID

Neutralizes credential theft and session hijacking with phishing-resistant MFA powered by real-time CrowdStrike Falcon platform telemetry

FALCON ADVERSARY OVERWATCH: IDENTITY | THREAT HUNTING

Provides 24/7 managed identity threat hunting, proactively detecting identity-based attacks, monitoring criminal forums for stolen credentials, and enforcing MFA challenges to prevent unauthorized access

Browser Security

FALCON SECURE ACCESS | BROWSER SECURITY AND CONNECTIVITY

Transforms any browser into a secure enterprise browser, delivering real-time protection against phishing, data loss, and identity theft on both managed and unmanaged devices, and provides browser-native data loss prevention (DLP) and governance, secure remote access to SaaS and private applications, and centralized policy enforcement that blocks web threats and governs generative AI use without disrupting user workflows

Next-Gen SIEM

FALCON NEXT-GEN SIEM | SIEM

Unifies first-party Falcon telemetry and third-party data on a single platform, enriching every signal with intelligence from 290+ tracked adversaries so analysts and AI agents can detect, investigate, and respond at machine speed

FALCON ONUM | HIGH-PERFORMANCE DATA PIPELINE

Delivers precise, real-time control over telemetry in motion, ensuring clean, high-quality signals reach security and analytic workflows in milliseconds, not minutes

FALCON ADVERSARY OVERWATCH: NEXT-GEN SIEM | THREAT HUNTING

Delivers end-to-end threat disruption by correlating first- and third-party Falcon Next-Gen SIEM data and proactively hunting advanced threats across network edge devices, SaaS applications, email security, operating systems, and more

Data Security

FALCON DATA SECURITY FOR ENDPOINT | REAL-TIME ENDPOINT DATA SECURITY

Delivers real-time visibility, encryption detection, and behavioral analysis to stop unauthorized data exfiltration across Windows and macOS devices

FALCON DATA SECURITY FOR CLOUD | RUNTIME CLOUD DATA SECURITY

Provides real-time monitoring and classification of sensitive data in motion across cloud environments using eBPF, enabling organizations to detect and respond to data risks without added complexity and with minimal performance impact

Security and IT/OT Operations

FALCON EXPOSURE MANAGEMENT | EXPOSURE MANAGEMENT

Provides continuous exposure and attack surface visibility, prioritizing true CVE and misconfiguration exploitability based on asset and network context to automate remediation to proactively reduce cyber risk and prevent breaches

FALCON EXPOSURE MANAGEMENT: CAASM

Allows you to discover and monitor managed and unmanaged assets in real time and visually map assets and their relationships, revealing deep host insights into applications, browsers, CVEs, and misconfigurations

FALCON FOR IT | IT AUTOMATION

Delivers endpoint visibility, operational control, and fleet-wide remediation to discover emerging risks, automate security and IT workflows, and reduce security and operational risk

FALCON FOR XIoT | XIoT ASSET VISIBILITY AND PROTECTION

Delivers real-time visibility, risk prioritization, and detection into unmanaged and adjacent extended internet of things (XIoT) assets, backed by XIoT-specific indicators of attack (IOAs) and indicators of compromise (IOCs) to reduce security, safety, and operational blind spots

FALCON FILEVANTAGE | FILE INTEGRITY MONITORING

Provides real-time, comprehensive, and centralized visibility that boosts compliance and offers relevant contextual data

Managed Services

FALCON COMPLETE | MANAGED DETECTION AND RESPONSE

Provides 24/7 expert-driven protection across endpoints, identities, cloud workloads, and third-party Falcon Next-Gen SIEM data, combining elite security expertise, AI-powered technology, and proactive threat hunting to detect, disrupt, and remediate sophisticated threats at machine speed

CrowdStrike Services

INCIDENT RESPONSE

Provides 24/7/365 elite incident response to contain threats, restore operations, and mitigate breach impact

Incident Response Services | Provides comprehensive breach response and recovery, from triage and investigation to cross-domain remediation and restoration, backed by world-class threat intelligence and delivered by a highly experienced incident response team

Services Retainer | Provides prearranged, on-demand access to CrowdStrike experts for rapid incident response and proactive consulting services that strengthen defenses over time

STRATEGIC ADVISORY SERVICES

Develops and matures the security program to improve defenses

Tabletop Exercise | Simulates incident response scenarios that expose process gaps and improve coordination across the full team, from hands-on-keyboard analysts to executive stakeholders

Maturity Assessment | Comprehensively evaluates your organization's security posture, identifying gaps, benchmarking capabilities, and providing a prioritized roadmap to strengthen defenses against evolving threats

Regulation Readiness and CXO Advisory | Helps you understand and prepare for cyber-related regulation mandates, including the evolving risk and governance responsibilities of the board of executives

Insider Risk Services | Strengthens your insider risk strategy by assessing and optimizing your current detection, prevention, and response capabilities

RED TEAM SERVICES

Tests and validates defenses through emulated attacks that expose weaknesses

Penetration Testing | Provides attack emulations that test the detection and response capabilities of your people, processes, and technology to identify vulnerabilities

Red Team/Blue Team Exercise | Increases response readiness under expert guidance, as a red team attacks systems in a simulated exercise and a blue team mounts the defense

Cloud Breach Emulation and Response Exercise | Helps organizations test and enhance their CDR capabilities through real-world adversary simulation

Adversary Emulation Exercise | Gauges readiness to defend against a sophisticated adversary infiltration that employs advanced tradecraft

AI SECURITY SERVICES

Secures the AI powering your organization and uses AI to defend with scale, precision, and speed

Frontier AI Readiness and Resilience Service | Provides continuous vulnerability scanning across applications and code bases, powered by frontier AI models; findings are prioritized by adversary risk and exploitability, with remediation directly through Falcon for IT and a trusted partner ecosystem

Shadow AI Visibility Service | Identifies sanctioned and unsanctioned AI systems, agents, and extensions across cloud, SaaS, and endpoint environments, uncovering unauthorized access to sensitive data and closing the inventory gap between documented and active production agents

AI Red Team Services | Exposes vulnerabilities in the generative AI stack that could be exploited by testing LLM integrations for sensitive data exposure and adversarial manipulation

AI Systems Security Assessment | Provides Falcon-powered discovery and threat-informed testing to uncover shadow AI, risky integrations, and governance gaps, delivering clear visibility and actionable guidance

AI for SecOps Readiness | Provides expert guidance on integrating AI into detection and response workflows with tailored use cases, architectural guidance, and a roadmap to increase response speed, precision, and scale

TECHNICAL ASSESSMENT SERVICES

Audits and addresses security gaps across endpoints, cloud, and SaaS applications to tangibly reduce risk

Technical Risk Assessment | Highlights security vulnerabilities, weaknesses, and gaps in the IT environment across endpoint devices, applications, and user identities

Identity Security Assessment | Audits identity security practices and defense posture for weaknesses, including Active Directory domain configuration, account configuration, privilege delegation, and potential attack paths

Cloud Security Assessment | Identifies misconfigurations and vulnerabilities in the cloud estate that could be exploited by adversaries

Compromise Assessment | Exposes and addresses undetected threat activity through a one-time threat hunt, available for endpoint, cloud, and SaaS applications

SaaS Security Assessment | Assesses SaaS environments for security gaps across configurations, access controls, data policies, and third-party integrations

PLATFORM PROFESSIONAL SERVICES

Helps ensure your CrowdStrike Falcon deployment is expertly configured, optimized, and aligned to your security needs; specialists provide best-practice implementation and deep module expertise to maximize protection, improve efficiency, and achieve security outcomes faster

TRAINING AND SECURITY UPSKILLING

Builds security acumen and closes the skills gap through CrowdStrike University, offering on-demand training, personalized learning paths, and eight certifications for deep Falcon module expertise

CROWDSTRIKE PULSE SERVICES

Provides continuous consulting engagements via focused sessions on a recurring cadence (biweekly, monthly, or every two months) tailored to your needs, aligned with your priorities, and adapted as needed, enabling consistent progress, improved resilience, and strategic maturity that evolves at the speed of the adversary

About CrowdStrike

CrowdStrike (NASDAQ: CRWD), a global cybersecurity leader, has redefined modern security with the world's most advanced cloud-native platform for protecting critical areas of enterprise risk – endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft, and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting, and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity, and immediate time-to-value.

CrowdStrike: We stop breaches.

Learn more: www.crowdstrike.com

Follow us: Blog | X | LinkedIn | Facebook | Instagram | YouTube

Start a free trial today: https://www.crowdstrike.com/trial