CONNECTWISE

2026 MSP THREAT REPORT

How Modern Attacks Abuse Trust and Identities

Contents

Foreword

The most damaging incidents in 2025 didn't rely on novel exploits or advanced malware. Instead, **attackers succeeded by abusing trust.**

Attackers consistently exploited identities that were assumed to be legitimate, software that was assumed to be safe, automation that was assumed to be reliable, and users who believed they were following normal instructions. These attacks did not break systems; they blended into them.

For managed service providers (MSPs), this represents a fundamental shift in the nature of risk. MSPs are no longer defending isolated endpoints or individual networks. They are defending interconnected ecosystems where access, updates, tooling, and users are deeply interdependent. When trust fails at any point in that ecosystem, the consequences extend quickly across clients, vendors, and shared infrastructure.

The MSPs that succeed in the years ahead will be those that move beyond reactive security models and rethink how trust is granted, monitored, and enforced across the environments they manage.

Patrick Beggs
Chief Information Security Officer, ConnectWise

Introduction

The 2026 MSP Threat Report analyzes the most significant threats observed throughout 2025, drawing from real-world incident response investigations, ConnectWise partner intelligence, and industry research. The findings focus specifically on attack patterns that materially impacted MSP-managed environments and the small and midsized businesses (SMBs) they support.

Rather than cataloging every emerging tactic or malware family, this report concentrates on how attackers consistently gained access, why defenses failed, and what those failures reveal about modern MSP environments. Across investigations, adversaries favored techniques that were reliable and repeatable.

Each section of this report examines a major threat category observed in 2025 and the growing influence of artificial intelligence on threats and attacks. We analyze real incidents and translate those findings into actionable guidance to strengthen defenses in the year ahead.

This report is a blueprint for enhancing security posture throughout the year. Share its insights with your teams, integrate its findings into risk assessments, and prioritize recommended actions that best fit your operational environment.

Key Findings

Initial access was rarely sophisticated, but it was highly effective

Most successful intrusions began with credential abuse, misconfigured VPN infrastructure, or user-initiated execution. Attackers did not need zero-day exploits when valid access paths already existed.

Trust was the primary attack surface

Adversaries repeatedly abused trusted identities, software updates, system utilities, and user behavior. Security controls that assumed trust rather than verifying it were routinely bypassed.

Ransomware surged to record levels with faster, more disruptive attacks

Ransomware activity intensified in Q4 2025, making it the most dangerous period of the year. Groups such as Akira moved rapidly from access to impact, targeting backup infrastructure, harvesting credentials, and exfiltrating data before defenders could respond.

AI increased attacker scale, not visibility

While AI's role is often invisible in incident telemetry (sensor data collected from security incidents), its impact was evident in phishing quality, fraud realism, malware iteration speed, and operational efficiency.

Software supply chains amplified blast radius

Compromises of upstream packages, installers, and update mechanisms allowed single intrusions to cascade across thousands of downstream environments, often without immediate detection.

User-mediated execution became a dominant technique

Techniques such as ClickFix demonstrated that convincing users to manually execute commands is one of the most reliable ways to bypass traditional endpoint and email defenses.

Reactive security models consistently failed MSP environments

Detection after execution was often too late. Environments with limited identity monitoring, weak application controls, or poor visibility into execution context suffered the greatest impact.

The 2025 Threat Landscape

The cyberthreat landscape facing MSPs and their clients intensified throughout 2025, driven by a sharp increase in ransomware activity and a broader shift in how attackers achieve scale and impact. Rather than relying on single exploits or isolated malware campaigns, adversaries combined multiple access paths, including credential abuse, social engineering, software supply chain compromise, and AI-assisted techniques, to move quickly from intrusion to monetization.

Ransomware **was the most visible and damaging result of this convergence.** Victim counts increased significantly year over year, reversing prior declines and reinforcing ransomware as the preferred endgame for a wide range of threat actors. This growth was not fueled by new encryption techniques, but by more reliable access methods. Attackers focused on stealing credentials, abusing remote-access infrastructure, and exploiting trusted workflows to deploy ransomware faster and at greater scale.

Several factors contributed directly to the rise of ransomware:

- **Supply chain attacks accelerated** across public package registries, targeting maintainer accounts, CI/CD pipelines, and upstream distribution channels. These incidents reinforced a critical reality: MSPs inherently inherit the security posture of their vendors and software ecosystems. Even trusted packages, updates, and automation tools can become delivery mechanisms for malicious code when upstream maintainers are compromised.
- **User-driven attacks such as ClickFix surged in frequency and sophistication.** These campaigns persuade users to run commands manually under the guise of browser verification, file repair, or system troubleshooting. By shifting execution onto the user, attackers bypass traditional security controls and leverage built-in system utilities to initiate infection chains. The increasing reliability of this technique has made it a preferred initial access method across multiple malware families.
- **VPN abuse remained one of the most consequential intrusion vectors.** Weak credential hygiene, legacy MFA implementations, and misconfigurations enabled attackers—including ransomware groups such as Akira—to obtain privileged access with little resistance. Once inside, adversaries moved quickly, disabling defenses, stealing data, and deploying ransomware with minimal dwell time. These incidents highlight the elevated risk MSPs face when remote access infrastructure is not actively monitored and maintained.
- **2025 marked a turning point in the practical use of AI within cybercrime.** While AI often leaves no direct artifacts in incident telemetry, its influence is unmistakable: more convincing phishing lures, rapid malware iteration, deepfake-enabled fraud, and automated obfuscation techniques. AI has lowered the barrier to entry for attackers while increasing operational scale and precision.

MSPs must assume that attackers will use multiple avenues to gain access: some technical, some human, some leveraging trusted software itself. Building resilience requires a layered approach, including stronger identity security, hardened remote access, application control, behavioral monitoring, and user education that reflects modern attack patterns.

While the landscape continues to evolve rapidly, MSPs can stay ahead by adopting informed strategies and proactive controls that protect their clients with confidence.

Ransomware in 2025: A Shift in Frequency and Strategy

The 2025 ransomware environment underscored two realities: **threats are increasing in quantity, and attackers are adapting strategies that exploit existing trust relationships, human behavior, and pervasive access mechanisms.**

According to recent industry reporting, overall ransomware victim counts surged sharply. One security firm observed a 58% year-over-year increase in the number of victims added to data leak sites—the highest level on record. Other telemetry echoed this trend, with incident counts rising significantly across geographies and sectors and fragmented extortion groups filling the void left by high-profile law enforcement takedowns.

Volume and Reach

Ransomware activity accelerated in both scale and visibility, reshaping the overall threat landscape.

- Ransomware incidents increased dramatically, outpacing prior years and reversing earlier declines.
- The number of active extortion groups grew, even as some traditional ransomware-as-a-service (RaaS) gangs dissolved, resulting in a highly fragmented but resilient threat ecosystem.

Tactics and Techniques

As ransomware volumes surged, attackers also refined how they operated.

- Double extortion remained the dominant model, combining file encryption with data theft and extortion via public leak sites.
- Attackers increasingly leveraged AI-assisted tooling and automation to scale phishing, credential abuse, and lateral movement activity, making initial access faster and more widespread.
- Established groups splintered into new variants or were replaced by smaller, agile operators, complicating attribution and response.

Impact and Resilience

The growing sophistication and reach of ransomware translated directly into real-world consequences.

- High-impact ransomware events affected organizations of all sizes, from local governments and hospitals to major enterprises, often forcing extended downtime and costly recovery.
- Although some defensive measures gained traction, such as immutable backups and incident response planning, the sheer volume and speed of attacks strained detection and containment capabilities.

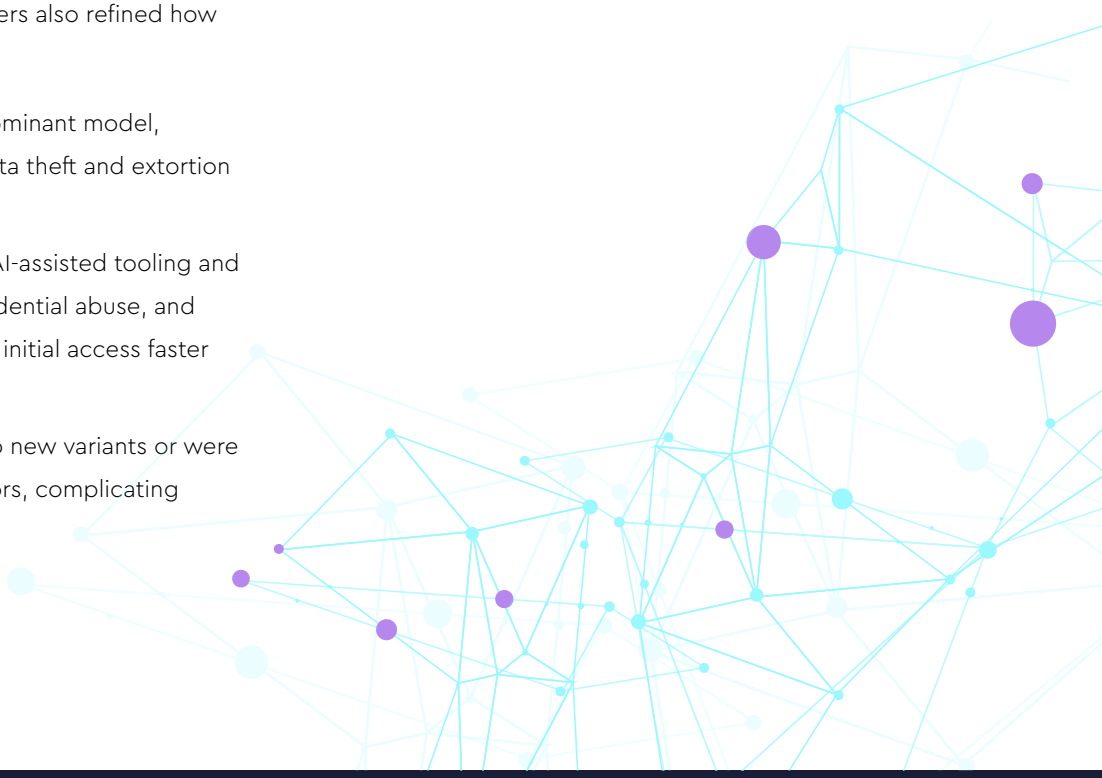

Akira Ransomware

Overview

Akira was one of the most disruptive threats we tracked this year, with activity particularly focused on midsized businesses across North America and Europe. The attack uses a double-extortion model of data theft followed by file encryption.

The group's affiliates leveraged stolen VPN credentials to gain access to their victims' networks. Once inside, Akira operations followed a consistent lifecycle: scan, steal, encrypt. There was little to no effort to dwell quietly. Operators moved swiftly, often succeeding in their overall goals.

For SMBs, the risk isn't that they're being singled out; it's that they're part of a much larger pool of opportunistic targets. They're not picking specific targets, they're following the easiest path in. For far too many SMBs, that path starts at the VPN.

Key characteristics of Akira activity in 2025 included:

- **Rapid escalation from access to impact** with minimal dwell time between initial compromise and ransomware deployment
- **Early disruption of backup infrastructure** aimed at preventing recovery and increasing operational pressure
- **Consistent use of double extortion** pairing encryption with aggressive data theft and leak-site pressure
- **A focus on SMB-centric environments** where limited visibility and weaker identity controls increased the likelihood of success

This combination of speed, opportunism, and operational discipline made Akira particularly effective throughout 2025 and reinforced the broader trend of ransomware groups prioritizing reliability over novelty.

Notable Incidents

Akira Operations

As a RaaS, Akira develops ransomware, maintains infrastructure, and recruits affiliates to help carry out attacks. A large number of Akira ransomware events this year began with VPN compromises, followed by rapid lateral movement and ransomware deployment. In many cases, the attackers used scanners such as **Advanced IP Scanner** or **SoftPerfect Network Scanner**; dumped Active Directory and credential information using tools such as **Mimikatz** and **Ntdsutil**; staged archives of data using tools such as **WinRar**; and encrypted files via the final **akira.exe** payload across multiple systems. There was also notable exfiltration activity across some Akira=related incidents where we observed **Rclone, WinSCP,** and **FileZilla** being utilized in exfiltration attempts.

MFA Bypass

Multiple incident reports and vendor analyses indicate that Akira operations were logging into SonicWall SSL VPN accounts even when OTP-based MFA was enabled. The cause was likely the use of previously stolen OTP seeds or appliance-resident secrets (credentials or seed material migrated from older device configurations), allowing attackers to generate or enroll OTP devices and bypass the second factor.

Defense Disabling Via Vulnerable/Legitimate Drivers (BYOVD)

Observed post-compromise activity often focused on interfering with backup platforms and disabling endpoint protections. Akira affiliates have been observed harvesting backup credentials, removing or corrupting restore points, and using a bring-your-own-vulnerable-driver technique (deploying a legitimate driver alongside a malicious helper driver) to disable endpoint protection services prior to exfiltration and encryption.

What This Means for MSPs

For MSPs and their clients, Akira's activity highlights several recurring weaknesses:

- Privileged VPN access is a consistent starting point
- Backup Infrastructure is being targeted directly
- RDP is enabled far too often on critical systems
- Data staging and exfiltration happen fast

Strategies for mitigation

- **Early detection and access control**
  MSPs must prioritize detecting attackers during early stages of an attack. If your clients rely on VPN based access and centralized backup platforms, these are the services that need network segmentation, MFA hardening, and continuous active monitoring. The encryption stage is too late. If you're not catching them at the VPN login or recon phase, you're likely catching them too late.
- **Privileged access management**
  Ransomware operators depend on elevated privileges to disable security controls, access backups, and deploy payloads at scale. Privileged access should be governed by the principle of least privilege, with administrative rights being removed wherever possible. Enforcing just-in-time elevation and maintaining full auditability of privileged sessions helps limit blast radius and reduces the likelihood that initial access escalates into widespread compromise.
- **Managed detection and response (MDR)**
  Modern ransomware campaigns intentionally blend into normal activity, making signature-based detections insufficient. MDR provides continuous, behavior-based monitoring to identify early indicators such as anomalous authentication, suspicious tool usage, and attempts to impair security controls. Rapid investigation and containment at this stage are critical to stopping attacks before encryption or extortion occurs.
- **Security information and event management (SIEM)**
  SIEM strengthens ransomware defense by centralizing and correlating logs across identity, endpoint, network, and backup systems. By analyzing activity holistically, SIEM helps surface coordinated attack behavior that may appear benign in isolation, such as privilege escalation combined with backup enumeration. Off-host log retention also preserves visibility if attackers attempt to delete or tamper with local logs.
- **Business continuity and disaster recovery (BCDR)**
  Even with strong prevention and detection, MSPs must assume some attacks will succeed. Immutable backups protect against ransomware operators who deliberately target backup infrastructure early in the attack lifecycle. Enforcing immutability, isolating backup access from production credentials, and routinely testing restoration processes ensure recovery remains possible even when encryption and extortion attempts occur.

Related content

How to detect ransomware >>
How to prevent ransomware >>

Software Supply Chain Attacks

Overview

Modern software development relies heavily on the reuse of trusted components from public repositories. These repositories, and the dependencies, packages, build tools, and distribution platforms that support them, collectively form the "software supply chain."

Threat actors frequently target package maintainers or build pipelines because compromising these upstream sources enables them to distribute malicious, signed updates that dependency managers automatically include in numerous downstream applications. By exploiting the inherited trust across the supply chain, attackers can bypass security measures, maintain persistence, and scale impact more efficiently than targeting individual entities directly.

In recent supply chain attacks, maintainers were compromised through phishing and credential theft techniques, such as token theft or MFA fatigue, reinforcing that identity security remains a critical weakness in supply chain defense.

Notable Incidents

npm

The default package manager and public registry for Node.js and JavaScript, npm, was hit by a worm-style campaign in September 2025 dubbed "Shai-Hulud."

Shai-Hulud: What happened

- Attackers used stolen maintainer credentials and tokens to publish trojanized updates to hundreds of packages.
- These updates executed install-time scripts that harvested additional secrets (including npm, GitHub, and cloud keys) and planted hidden GitHub Actions workflows to persist and spread to more projects.
- GitHub and npm maintainers removed hundreds of compromised packages from the registry and implemented several mitigation steps, including blocking uploads matching campaign IOCs, rolling out short-lived tokens, and enforcing phishing-resistant MFA.

The impact: While the total number of impacted downstream vendors remains unknown, these attacks highlight the severe implications for application security and public trust in open-source ecosystems.

Shai-Hulud 2.0: What happened

In late 2025, a second and even more aggressive campaign emerged, which was dubbed "Shai-Hulud 2.0." This variant was among the fastest spreading npm supply chain attacks ever observed.

- Unlike its predecessor, Shai-Hulud 2.0 leveraged automation extensively, enabling rapid propagation across thousands of packages in a short timeframe.
- The worm exploited compromised accounts and abused npm's ecosystem to inject malicious code at scale, focusing on persistence and lateral movement.
- It introduced enhanced evasion techniques, including dynamic payload delivery and more sophisticated credential harvesting, targeting not only developer secrets but also CI/CD environments.

The impact: The campaign's velocity and automation amplified its impact, forcing npm and GitHub to accelerate incident response measures and expand detection capabilities.

Shai-Hulud 2.0 underscores the growing trend of weaponized automation in supply chain attacks, raising urgent concerns about the resilience of package registries and the need for stronger identity and integrity controls.

PyPI

In another ongoing supply chain compromise campaign, PyPI, has continued to observe waves of phishing attempts where attackers send out "verify your account" emails to lead to similarly named but spoofed domains designed to trick users into entering their credentials.

The campaign began as early as May 2023, but there have been multiple reported large-scale attempts that persist to this day. Fortunately, PyPI has not had any confirmed cases of compromise where attackers published malicious packages via phished accounts.

Ruby

In mid-September 2025, the Ruby community was hit by internal strife among the GitHub repo maintainers, wherein Ruby Central asserted control over key RubyGems and Bundler GitHub repos, removing or restricting several long-term maintainers. Since the governance shakeup changed who can publish, the speed at which fixes get pushed out, and how keys and releases are managed, it has a potentially significant impact on future supply chain trust.

Rust

In September of 2025, the Rust team and researchers disclosed two malicious crates: faster_log and async_println. These crates impersonated a popular logger and scanned source trees to exfiltrate crypto wallet keys. They had been live since late May with thousands of downloads before it was resolved and the crates were removed, the accounts that published the crates were suspended, and the rotation of any exposed secrets were disabled.

NuGet

Since 2023, NuGet, the public package registry many .NET apps use, have seen recurring waves of supply chain abuse. Attackers routinely spin up new publisher accounts and upload malicious packages with names similar to legitimate and popular packages that are introduced into projects through typos, name confusion or misplaced trust in an unfamiliar publisher.

Some campaigns also hijack existing accounts or smuggle code into MSBuild .targets/.props so it runs during restore/build, turning a normal dependency update into an execution path. Microsoft and security vendors routinely flag and remove these packages and have tightened checks, but the pattern persists, highlighting how a single upstream package can quietly cascade into many downstream applications through ordinary updates.

Whether it's threat actors or disputes amongst maintainers, it's clear that the supply chain itself is delicate and at risk and requires the same safeguards and scrutiny that a company would place on its office environments, if not more so. Considering a successful compromise of a supply chain can turn one target into many, the importance of proper security controls surrounding the supply chain should not be taken lightly.

What This Means for MSPs

DevOps who use these supply chains generally have their own process in place, such as:

- ☑ Selecting reputable packages
- ☑ Reviewing maintainer reputation and release history
- ☑ Verifying signatures or checksums
- ☑ Locking versions
- ☑ Consuming through an approved internal repository,
- ☑ Generating software bill of materials (SBOMs) to record what is included

Combined, these are a solid layered approach to reducing risk before code enters the CI/CD pipeline, but they do not guarantee safety, and not every development team implements them consistently.

Strategies for mitigation

For MSPs, having a proper global application approval process in place is important, but without proper controls, it doesn't necessarily prevent end users from installing potentially unsafe apps on their work machines.

Industry best practices should be followed by those in charge of approving and/or pushing applications in commercial environments including:

- Implementing endpoint privilege management with application control
- Safelisting apps that can be installed, while also verifying the installers, using data such as file hash and signer identity, will help prevent the local installation of malicious applications
- Accepting installers from only official vendor channels, verifying hashes and publisher signatures for any application install, and testing in a controlled environment before rolling out the application globally
- Small test rollouts to groups of users with lesser access before pushing to higher-risk machines with access to important company or user information can greatly reduce risk
- Blocking suspicious download sites globally can contribute to preventing the inadvertent download of malicious programs

MSPs have very little control over upstream compromise, especially updates to applications that have already been installed. However, they can cut the risk, minimize impact, and catch bad releases early by sourcing only from trusted vendors, gating installs with application approval tools, and staging updates in a safe test and small canary rollouts with a ready rollback

Potentially Unwanted Programs Trojanized for Higher Impact

Overview

Not all threats carry the same weight. Potentially unwanted programs (PUPs) or potentially unwanted applications (PUAs) traditionally sat near the bottom of the threat list. They typically include "free" software that may serve a legitimate purpose for the victim, including PDF editing, file conversion, contact information lookups, or enabling dark mode on any website. Unlike other forms of free software (e.g., open source), PUPs don't have a monetary user cost; the "cost" comes in the form of ads, additional software, hijacked search results, crypto mining, collection of data about your contacts, or even victim machines being enrolled as residential proxies.

Historically, these behaviors were considered undesirable but tolerable. Excessive focus on PUPs risked distracting defenders from higher-impact threats. However, **that assumption no longer holds.**

In 2025, attackers increasingly weaponized PUPs as a delivery mechanism for persistence, credential theft, and remote control. What was once endpoint clutter has become a reliable entry point for more serious compromise.

Notable Incidents

In 2025, the ConnectWise Cyber Research Unit™ (CRU) and other threat researchers worldwide observed multiple campaigns distributing what would normally be classified as traditional PUPs. It was later revealed that PUPs had more malicious capabilities. Two prominent examples have been tracked: TamperedChef and EvilAI.

TamperedChef

TamperedChef has been distributed via apps for image searching, recipe searching, AI assistants, PDF editors, calorie counting, and product manual searching.

These trojanized apps are pushed via:

- Malvertising (ads purchased on platforms such as Google)
- SEO poisoning (gaming ranking algorithms to promote malicious sites).

To further attempt legitimacy, code-signing certificates are used to verify the installers, which are often promptly revoked once abuse is discovered.

Typical PUP behavior

On their own, the trojanized applications already exhibit undesirable behavior typical of many PUPs:

- Bundled adware
- Additional installers
- Enrolling the device as a residential proxy, often presented during installation as a tradeoff for free access

This enrolls them as an accessory to a service that makes traffic appear to come from a residential ISP ASN. It routes traffic from unknown users through that machine's IP address rather than using well-known VPNs that can more easily be tracked and filtered.

Affected application names observed include: AllManualsReader, ManualReaderPro, JustAskJacky, OpenMyManual, TotalUserManuals, EffortlessPDF, AppSuite PDF Editor, RecipeLister, OneStart Browser, SodukuFunZone, and multiple similarly branded variants, totaling dozens of signed installers.

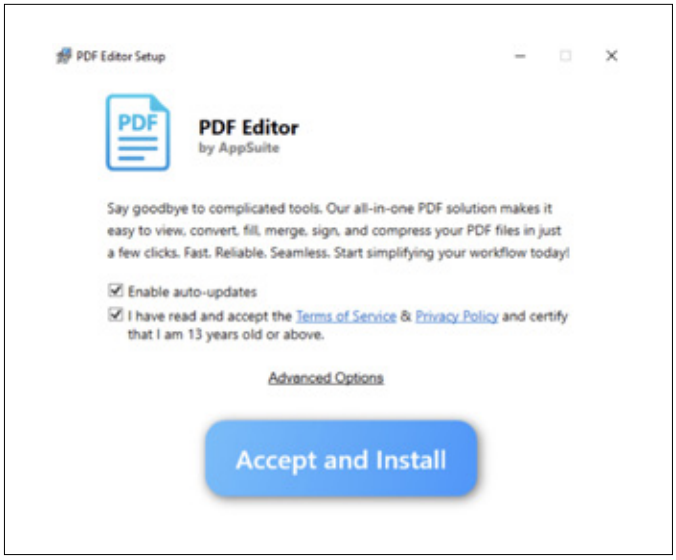

Figure 1:AppSuite PDF Editor was one of several trojanized apps used in the TamperedChef campaign

Escalation beyond traditional PUP behavior

TamperedChef, the malicious component of these applications, executes as a JavaScript file using NodeJS. It makes itself persistent through a scheduled task and in a registry Run key. It will stop browser processes to copy the "Web Data" and "Preferences" files containing cookies that may permit session hijacking as well as theft of other potentially sensitive data.

Trojan malware also operates as a backdoor enabling:

- Download of files on the victim machines
- Registry manipulation
- Execution of arbitrary processes

These capabilities allow a wider range of attacks with more potential impact than cookie theft or enrollment as a residential proxy.

What This Means for MSPs

The resurgence of trojanized PUPs in 2025 highlights a core challenge for MSPs: modern threat actors are increasingly using channels that appear legitimate to both users and security tools. These campaigns exploit gaps in application hygiene, endpoint privilege, and user behavior, which are areas that often receive less attention than perimeter defenses or ransomware prevention.

Preventing these threats requires stronger guardrails around:

- What software can be installed
- How elevation is granted
- How endpoints are monitored for unexpected persistence mechanisms or credential harvesting behavior.

Strategies for mitigation

- Application control, endpoint privilege management, and policy-driven software sourcing should be considered essential controls to prevent this class of threat.
- Behavioral monitoring, especially around script interpreters, browser data access, and persistent scheduling, plays a critical role in early detection. Since installers rely heavily on user consent and legitimate system utilities, traditional file-based detections may not trigger until long after initial infection.
- MSPs should also reassess their processes for approving utilities and ensure that consumer-grade applications, even when code-signed, undergo proper validation before being allowed into managed environments.

Ultimately, 2025 demonstrated that PUPs are no longer endpoint clutter. They have become a reliable entry point for attackers seeking persistence, credential access, and remote control, and MSP environments must treat them accordingly.

SSL VPN Compromises

Overview

Throughout 2025, we observed widespread abuse of SSL VPN infrastructure as a primary entry point into SMB and MSP client networks. Most successful compromises this year did not rely on new exploits; they stemmed from incomplete configuration hygiene, credential reuse, and exposure of legacy infrastructure.

In many cases, the initial intrusion vector was not particularly advanced. Publicly exposed SSL VPN interfaces, especially those from SonicWall and Fortinet, were routinely targeted by mass-scanning and credential-stuffing campaigns. Even environments with MFA in place were not immune, especially those using appliance-based OTP implementations. Our team responded to numerous incidents this year in which MFA was technically enabled, but attackers still managed to authenticate by enrolling their own device or using credentials retained from prior firmware generations.

This trend reflects an uncomfortable reality for the MSPs and SMBs: SSL VPNS are often treated as a "set it and forget it" infrastructure and rarely receive the same attention as endpoint or email security. However, this year made it clear, when these devices are left exposed and unmanaged, they become the single easiest way into the network.

Notable Incidents

Ivanti (Pulse Connect Secure VPN)

In early 2025, Ivanti Connect Secure SSL VPN appliances were widely compromised through exploitation of CVE-2025-0282, a critical pre-authentication remote code execution vulnerability that allowed attackers to gain unauthorized access to VPN gateways without valid credentials. At least 17 organizations were confirmed breached, with investigations attributing the campaign to a Chinese state-aligned threat actor focused on long term access rather than immediate monetization.

SonicWall VPN Intrusions

Multiple SMB organizations were compromised through SonicWall Gen7 VPNs that had been upgraded from older Gen6 hardware without properly resetting their inherited credentials. This vulnerability (later tied to CVE-2024-40766) allowed attackers to log in with valid credentials even when MFA was enabled. Several of these incidents resulted in full domain compromise in under two hours.

Fortinet Brute Force Campaign

There was a spike in brute-force login attempts across Fortinet VPN portals in early August. These attacks were confirmed to be part of a coordinated credential stuffing wave by threat actors. Several partners had login portals exposed without proper perimeter defense mechanisms in place, raising serious concern even in the absence of confirmed exploitation.

SonicWall Cloud Backup Breach

SonicWall disclosed that attackers had compromised its MySonicWall cloud portal and accessed firewall configuration backups for a subset of customers. These backups included hashed credentials, VPN group settings, and appliance secrets. For MSPs managing dozens of SonicWall clients through shared portals, this introduced a wider surface area for compromise, especially for any clients that had not rotated any credentials post migration.

Cisco ASA Zero Day Activity

While this campaign appeared to target larger orgs and government entities, we note it here because many legacy ASA appliances are still in use across SMB environments. Cisco confirmed two zero days (CVE-2025-20362 and CVE-2025-20333) were exploited in the wild to implant persistent malware directly on the firewall OS. This reinforces that even "trusted" edge devices should be monitored and routinely upgraded.

What This Means for MSPs

VPN appliances are one of the few ways SMBs allow persistent privileged access into their networks. While the intent is often for legitimate purposes, these devices become high-value targets for attackers and high-risk points of failure for MSPs. The activity observed in 2025 showed us:

- Default and migrated credentials remain dangerously common
- MFA is not enough if implemented poorly
- Monitoring VPN login telemetry must become standard
- Vendors aren't immune to supply chain breaches
- Strategies for mitigation

MSPs relying on VPNs for remote access, or managing clients who expose VPNs without proper hygiene, need to treat these services as critical infrastructure, not background noise. VPNs should be continuously monitored, regularly hardened, and protected with strong authentication and access controls.

The Continued Rise of ClickFix in 2025

Overview

Throughout 2025, ClickFix-style social engineering emerged as one of the most reliable and repeatable initial access methods observed across multiple intrusion campaigns. Rather than relying on exploitation or malicious attachments, ClickFix attacks manipulate users into manually executing attacker-provided commands under the pretense of routine verification steps, CAPTCHAs, browser prompts, or security-related actions. This execution model shifts the initial access burden onto the user, allowing attackers to bypass many traditional security controls by operating entirely through legitimate system utilities and trusted execution paths.

Typical ClickFix Attack Chain

This pattern has become prominent enough that it was formalized within the MITRE ATT&CK Framework as Malicious Copy and Paste (T1204.004), reflecting broader industry recognition that manual user execution represents a distinct and recurring adversary behavior. In practice, ClickFix does not represent a single lure or campaign style, but a general execution pattern that can be adapted to different themes, interfaces, and delivery contexts.

As the year progressed, attackers increasingly diversified how this pattern was presented to users. Variants such as FileFix reframed manual execution as a required step to open or repair a downloaded document, while ConsentFix extended the same copy-and-paste social engineering model into browser-based identity workflows by abusing legitimate OAuth consent and authorization flows within the browser rather than delivering endpoint malware. Other campaigns introduced more aggressive visual pressure, including fake system error and BSOD-themed ClickFix lures (tracked by researchers at Securonix as PHALT#BLYX), designed to create urgency and reduce hesitation by implying system instability or data loss. Despite differing themes, each of these variants relied on the same core mechanic of convincing the user that manually executing provided instructions was both safe and necessary.

Over time, ClickFix has evolved from relatively simple copy-and-paste lures into a flexible execution pattern used by multiple threat actors and malware families. As defenders and users became more familiar with earlier browser verification and CAPTCHA-style lures, attackers adapted by introducing new narratives that aligned with common user expectations around file handling, system errors, and consent workflows. The underlying execution model remained unchanged, but the surrounding context continued to evolve, reinforcing ClickFix's effectiveness as a durable and adaptable initial access method.

Notable Incidents

Early 2025: Establishing a Reliable Initial Access Pattern

At the start of the year, ClickFix activity was already present on a meaningful scale and was being actively operationalized by multiple threat actors. Early campaigns favored consistency and reliability over technical complexity.

Common characteristics of early ClickFix campaigns included:

- Fake browser verification messages, update prompts, and security warnings used to convince users to manually execute short commands
- Initiation of the infection chain through legitimate Windows utilities
- Simple download-and-execute chains that delivered commodity malware

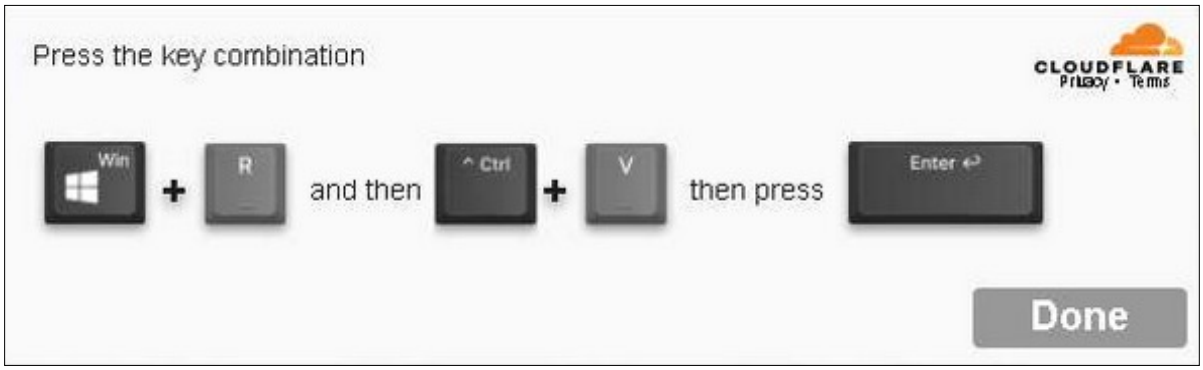

Figure 2: Clickfix lure pretending to be a Cloudflare CAPTCHA

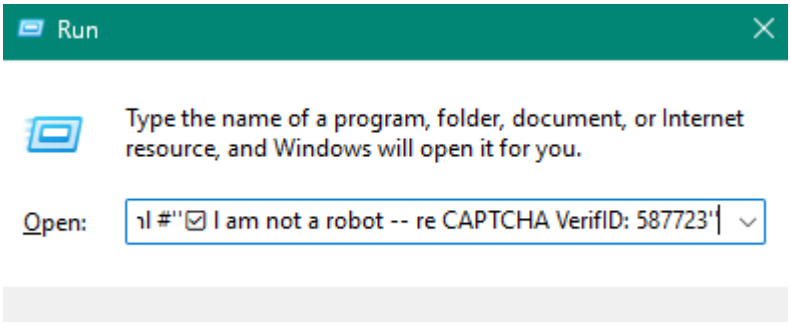

Figure 3: What it looks like to victims when they paste the malicious command into the Run box

SmartApeSG campaigns

Campaigns attributed to SmartApeSG exemplified this early operational model. Payloads observed during this period frequently included NetSupport Manager as a remote access tool, deployed with basic user-level persistence mechanisms. Infrastructure rotated frequently, often using compromised sites or short-lived domains, but the underlying lure mechanics and execution flow remained stable, indicating early confidence in ClickFix as a dependable initial access method.

Filch Stealer campaigns

Other early-year ClickFix campaigns demonstrated how easily this execution pattern could support different malware families without changing the user-facing lure. Campaigns delivering Filch Stealer used the same execution entry point to stage multi-step PowerShell loaders that ultimately deployed information-stealing payloads with limited remote access capability. While Filch Stealer showed more complexity than early SmartApeSG deliveries, the ClickFix component itself remained minimal, serving only as a reliable bootstrap mechanism.

Latrodectus campaigns

Similarly, Latrodectus campaigns observed during this period reinforced the trend toward loader-based delivery while retaining the same social engineering foundation. In these cases, the user-executed step initiated a chain involving script execution and installer-based staging, ultimately leading to reflective loading or sideloading techniques. Although the downstream tooling differed, the initial access method remained consistent, underscoring ClickFix's flexibility across both stealer- and loader-centric campaigns.

Moving beyond single actor campaigns at scale

Early 2025 ClickFix activity also overlapped with a broader ecosystem of commodity stealers and remote access tools, including Lumma, Vidar, and similar families frequently observed in opportunistic intrusion activity. Telemetry from multiple incidents showed ClickFix lures being reused to deliver these payloads with minimal modification, further demonstrating that attackers did not need bespoke infrastructure or advanced tooling to benefit from the technique. In many cases, the same execution flow could support a stealer, a RAT, or a loader with only minor backend changes. These campaigns established the technique as a scalable and low-friction initial access method.

Later 2025: Refinement and Operational Maturity

By the latter half of the year, ClickFix campaigns demonstrated clear signs of operational refinement rather than fundamental change. The core execution model of manual user execution through trusted system utilities remained intact, but attackers increasingly focused on reducing friction and improving operational resilience. Compared to early-year activity, later campaigns showed greater emphasis on infrastructure agility, payload flexibility, and minimizing observable artifacts.

Fake CAPTCHAs

One of the most visible shifts during this period was the growing use of fake CAPTCHA and browser verification overlays embedded into otherwise legitimate websites. These pages closely mimicked trusted services and prompted users to copy short commands into the Run dialog or a command shell, maintaining the same ClickFix interaction pattern while improving visual credibility. In several incidents, this initial step triggered PowerShell-based loaders that retrieved additional stages entirely in memory, leaving little to no conventional executable footprint on disk. These campaigns often culminated in stealer-style payloads or loader frameworks capable of deploying follow-on malware.

Multi-stage archive-based delivery chains

Another prominent late-year pattern involved multi-stage archive-based delivery chains, including the use of password-protected archives paired with locally staged extraction utilities. In these incidents, the ClickFix prompt initiated a short loader that fetched both an encrypted archive and the tooling required to unpack it. The extracted contents typically included lightweight loader scripts or executables that established persistence and launched final payloads. Data and indicators from these campaigns aligned with information stealers such as Vidar, Lumma, StealC, and related families, though the modular nature of the loaders allowed operators to change payloads without altering the user-facing lure.

Infrastructure abstraction

Several campaigns observed during this period also demonstrate deliberate infrastructure abstraction. Rather than embedding full staging URLs in the initial command, attackers relied on URL shorteners or redirect services to keep the visible execution step minimal and static while rotating backend hosts as needed. This approach reduced the operational cost of infrastructure churn and allowed the same ClickFix lure to remain effective even as payload locations changed.

MSHTA-based execution chains

While many later campaigns favored PowerShell-centric loaders, others revisited MSHTA-based execution chains associated with broader ClearFake activity. In these cases, ClickFix lures led users to execute commands that launched remote HTA or HTML content, which then staged additional script-based execution. Although MSHTA is a well-known abuse vector, its continued use in ClickFix campaigns underscores the emphasis on leveraging legitimate, signed binaries rather than introducing novel tooling. Some of these chains also showed overlaps with reflective loaders and HijackLoader-style behavior, suggesting opportunistic blending of techniques rather than tightly coupled malware families.

The end result: CickFix emerged as a proven and repeatable technique

Across these varied campaigns, the unifying theme was efficiency. Commands became shorter, staging logic moved off-host, and payload delivery became increasingly modular. The user-facing interaction changed little, even as backend execution chains grew more flexible. ClickFix matured into a delivery method that prioritizes reliability and adaptability over stealth through obscurity. Attackers did not abandon the technique as awareness grew; they optimized it instead. The result was a consistent initial access pattern capable of supporting a wide range of malware families and operational goals.

What This Means for MSPs

ClickFix campaigns succeed not because of novel malware or advanced exploitation, but because they exploit user trust and legitimate system functionality. For MSPs managing diverse client environments, this shifts the defensive challenge away from blocking specific payloads and toward identifying risky execution contexts and user-initiated process chains.

Strategies for mitigation

Traditional perimeter controls and file-based detections offer limited protection when users are convinced to manually execute commands using built-in tools. As a result, visibility into how processes are launched becomes more important than the identity of the final payload. Execution chains that originate from browsers, script hosts, or the Run dialog box warrant closer scrutiny regardless of the malware family involved, particularly when they rapidly transition into scripting engines or command interpreters.

User awareness also plays a critical role, but it must reflect modern attack patterns rather than legacy phishing models. ClickFix lures do not rely solely on malicious links or attachments. They depend on convincing users that manual code execution is a legitimate verification or recovery step. Effective training should explicitly address this behavior and reinforce that legitimate services do not require users to paste commands into the Run dialog box or a command shell.

Ultimately, ClickFix's persistence and refinement throughout 2025 demonstrate that social engineering-driven execution remains a powerful and adaptable intrusion method. MSPs that prioritize execution context, behavioral telemetry, and user-initiated risk, rather than relying solely on malware signatures or static indicators, will be best positioned to detect and disrupt these campaigns early in the attack chain.

AI Cybersecurity Threats

Overview

Artificial intelligence (AI) continued to reshape the cyber threat landscape throughout 2025, but not always in ways that are directly observable from incident data alone. Supply-chain compromises or VPN intrusions include technical indicators that provide clear evidence of how attackers gained access, but AI's role in modern attacks often remains hidden behind the scenes.

Why AI activity is hard to observe

Threat actors are increasingly using AI to accelerate development, improve lures, automate decision-making, or obfuscate their code, but these activities typically occur before the payload ever reaches a victim environment. By the time an attack surfaces within an MSP-managed environment, the components directly attributable to AI have usually been stripped away, leaving a significant visibility gap for incident responders.

As a result, it is difficult to determine whether generative AI was used to craft phishing lures, produce malicious scripts, automate payload delivery, or support operational planning. Our visibility into AI-enabled activity depends heavily on intelligence shared by external research organizations, cloud providers, and security vendors who monitor upstream development pipelines, large-scale phishing infrastructure, and AI-tool abuse patterns that MSP-level incident data simply cannot reveal.

Why AI's impact on the threat landscape still matters

We include these broader findings in this report because AI is exerting a massive influence across every stage of the threat landscape, even when that influence is not directly measurable in the incidents we respond to. The evolution of deepfake-enabled fraud, LLM-generated phishing content, prompt-manipulation attacks, and AI-assisted malware tooling demonstrates a rapid acceleration that cannot be ignored. While its fingerprints may not always be visible in the artifacts we analyze, AI is undeniably shaping attacker capabilities, lowering operational barriers, and expanding the scale and speed at which threat actors can operate.

For MSPs, understanding this invisible dimension of threat evolution is critical. The challenges ahead are not only about detecting malware or blocking suspicious scripts; they are about recognizing that the tools used to create these threats are becoming more powerful, accessible, and automated. AI is enabling attackers to move faster, disguise intention more effectively, and iterate through defenses with unprecedented agility. Even when we cannot point to a specific indicator that "AI was used here," the impact is evident in the increasing volume and adaptability of the threats we face.

Notable Incidents

Deepfake-enabled fraud and impersonation

One of the most alarming trends in 2025 was the rise of deepfake-enabled fraud. In July, a businessman in China lost over $622,000 in a Zoom call scam where attackers used real-time face-swapping technology to impersonate a trusted contact. The victim was convinced to authorize a wire transfer during the call, only to discover later that the conversation never occurred with the real individual. In the United States, AI voice scams dubbed "Grandparent Scam 2.0" targeted seniors by cloning voices of loved ones and requesting emergency funds. These scams often bypass traditional fraud detection due to their emotional manipulation and realism.

AI-generated phishing

Microsoft's Threat Intelligence team flagged a phishing campaign in August 2025 that used LLM-generated SVG files to bypass email security filters. These files, disguised as PDF attachments, contained obfuscated JavaScript payloads that redirected victims to credential-harvesting sites. The attack used business-related terminology and synthetic structures to evade detection, suggesting the payloads were crafted using generative AI. SVG-based phishing attacks have become increasingly popular due to their ability to embed dynamic content and evade static analysis. In multiple documented cases, these files passed undetected through VirusTotal and other scanners, highlighting a significant blind spot in current defenses.

Man-in-the-prompt attacks on AI tools

A newly identified threat vector, Man-in-the-Prompt (MitP) targeting browser-based AI tools like ChatGPT, Gemini, and Copilot. Researchers at LayerX demonstrated how malicious browser extensions can intercept and manipulate prompts in real time, injecting hidden instructions or exfiltrating sensitive data without user awareness. These attacks exploit the document object model (DOM) used by AI interfaces, allowing extensions with basic scripting access to read and write prompts. Proof-of-concept attacks showed how compromised extensions could silently alter user inputs, extract confidential information, and delete chat history to cover tracks.

AI-assisted malware development

In the fourth quarter, threat intelligence researchers at Google's Threat Intelligence Group (GTIG) observed increasingly sophisticated uses of AI tools by malicious actors, including the emergence of PROMPTFLUX, which leverages large language models during execution to dynamically generate malicious scripts, obfuscate code, and create malicious functions on demand. Additional threats include FRUITSHELL, a normal reverse shell that code comments to evade LLM-based security scanning, and PROMPTSTEAL (also known as LAMEHUG), a Hugging Face API utilization which enables the malware to use more than just public-facing chatbots.

Prompt injection

According to GTIG, threat actors continue to employ prompt injection techniques to bypass AI safety guardrails (e.g., claiming to be working on capture-the-flag exercises or cybersecurity research papers). Using LLMs in malware has also led to operational security failures on the part of operators, inadvertently exposing hard-coded information such as C2 domains and encryption keys to AI models. The quarter also saw a growing marketplace for AI-enabled malicious services, including deepfakes, malware generation, and phishing campaigns, that use several different models that specialize in different areas of development and exploitation (e.g., WormGPT and FraudGPT). Most of what is observed in the report is still experimental, but it highlights that we should expect to see more of over the next year.

What This Means for MSPs

For MSPs, the rapid evolution of AI-driven cyberthreats means that traditional playbooks focused on signature-based detection and user awareness are no longer enough. Attackers are using generative AI to automate phishing, create deepfake-based social engineering, and develop polymorphic malware that adapts faster than static defenses.

Strategies for mitigation

MSPs must now assume that every client organization could face attacks engineered or enhanced by AI, even from relatively low-skilled actors. Mitigation requires adapting controls, training, and processes to reflect this shift.

Strengthen verification against impersonation and fraud

- Implement stricter verification for financial and account access requests.
- Require multi-channel confirmation, such as verbal verification using known phone numbers.
- Expand training beyond email awareness to include AI-generated voice and video impersonation.

Harden software supply chain controls

- Treat software dependencies as potential intrusion vectors, especially with AI-assisted code obfuscation.
- Encourage the use of signed packages, software bills of materials (SBOMs), and integrity checks in CI/CD pipelines.
- Discourage use of AI utilities from unknown GitHub repositories, which have become common infection carriers.

Reduce risk introduced by AI tooling

- Enforce browser extension allowlists and monitor for unapproved data exfiltration.
- Disable auto-sync features that may expose chat history or model interactions.
- Monitor AI usage itself as a potential attack surface

Apply AI defensively with guardrails

- Use AI for log correlation, phishing detection, and anomaly detection under human oversight.
- Protect against prompt injection and data leakage through access controls, sanitization, and private or on-prem AI deployments where possible.

Translate AI risk into operational action

- Proactively brief clients on AI-enabled threats and their business impact.
- Demonstrate how controls such as dual verification and extension policies reduce risk.
- Integrate these updates into security awareness programs and incident response plans.

From Insights to Action: How ConnectWise Can Help

The findings presented in this report underscore a critical truth: the threat landscape is not only expanding but also evolving in ways that challenge traditional security models.

Defending against modern threats requires shifting security efforts earlier in the attack lifecycle. Identity must be protected as a primary control plane, remote access must be tightly governed and continuously monitored, and execution behavior must be visible, not just outcomes. And when prevention fails, recovery must be fast, reliable, and resistant to attacker interference.

ConnectWise delivers integrated Cybersecurity and Data Protection solutions to enforce the principle of least privilege, detect suspicious behavior early, and ensure rapid recovery even in worst-case scenarios.

Privileged Access Management

Identity and access failures were at the center of many of the most damaging incidents in 2025. ScreenConnect® Privileged Access Management helps MSPs reduce this risk by enforcing least-privileged access and controlling how administrative sessions are initiated, elevated, and audited. By removing standing privileges and tightly governing remote access, MSPs can limit the blast radius of credential compromise and reduce the likelihood that a single stolen account leads to full domain compromise.

Managed Detection and Response (MDR)

Many of the attacks outlined in this report progressed rapidly because early warning signs went unnoticed. ConnectWise MDR™ provides continuous monitoring and expert-led response across endpoints, identities, and network activity. By correlating behavioral signals rather than relying solely on signatures, MDR helps MSPs detect suspicious activity earlier in the intrusion lifecycle, including techniques that bypass traditional defenses.

Security Information and Event Management (SIEM)

Modern attacks rarely leave a single, obvious indicator. Instead, they generate weak signals spread across identity systems, endpoints, network devices, and cloud platforms. Without centralized visibility, these signals are easy to miss.

ConnectWise SIEM™ helps MSPs aggregate and correlate telemetry across client environments, providing the context needed to identify abnormal patterns and escalation paths. SIEM enables faster investigation, supports proactive threat hunting, and strengthens incident response by ensuring critical events are not analyzed in isolation.

For MSPs managing complex, multi-tenant environments, SIEM provides the foundation required to move from reactive alerting to informed, intelligence-driven defense.

Business Continuity and Disaster Recovery (BCDR)

Ransomware groups in 2025 increasingly targeted backup infrastructure early in attacks, attempting to delete, encrypt, or corrupt recovery points before deploying ransomware. In these scenarios, traditional backup strategies often failed when they were needed most.

Immutable backup is no longer optional. It is a core resilience requirement.

x360Recover from Axcient, a ConnectWise company, provides immutable backups designed to resist tampering, even when attackers gain administrative access. By protecting backup integrity and enabling rapid recovery, x360Recover ensures MSPs can restore systems confidently when preventative and detective controls are bypassed.

Methodology: How We Build This Report

The ConnectWise MSP Threat Report was first introduced in 2020 to provide MSPs with practical, real-world insight into the evolving cybersecurity landscape. The report highlights the threats, attack patterns, mitigations, and solutions most relevant to environments MSPs manage every day.

Our annual MSP Threat Report is made possible by the research and findings from the ConnectWise Cyber Research Unit™ (CRU). This elite team is composed of experienced threat hunters and cybersecurity professionals with deep expertise in engineering, IT admin, security operations, incident analysis, and incident response. The team gathers threat intelligence 24/7 from a wide range of sources, including real-world incident response investigations, telemetry from ConnectWise partner and SMB client environments, ransomware leak sites, and malicious infrastructure such as botnets and command-and-control networks. By correlating these data points, the CRU identifies emerging trends, common attack techniques, and recurring points of failure across MSP-managed environments.

The goal of this report is not simply to document threats, but to translate complex threat intelligence into actionable insight for MSPs.

About ConnectWise

ConnectWise powers IT businesses by simplifying operations, enhancing experiences, and driving growth. Trusted by IT solution providers worldwide, ConnectWise sets the standard for innovation and service delivery. For more than 40 years, ConnectWise has been committed to partner success, delivering software, services, and an open ecosystem of integrations. The ConnectWise Platform provides unmatched scale and AI-driven automation across PSA, RMM, cybersecurity, and data protection, helping our partners deliver and secure services more efficiently. Discover how ConnectWise is transforming the IT industry at connectwise.com.