blackpoint

ANNUAL THREAT REPORT

2026

Table of Contents

Foreword

A valid username; a legitimate password; a reliable tool your team uses every day. Throughout 2025, these simple symbols of trust quickly became the adversary's welcome mat.

For IT and security teams, this shift raises the stakes. Between VPN sessions, remote monitoring tools, software installation, and more, organizations must rethink what "routine" really means in an era where attackers have learned to weaponize routine itself.

Our latest research illustrates attacker behaviors, where they failed, and strategic defense recommendations for the year ahead. Despite the increasing speed and scale of adversary operations, our 24/7 SOC continually disrupted intrusions before they could escalate, intervening in early stages, dismantling malicious workflows, and preventing payloads from ever taking hold. Their work demonstrates how human-led, real-time detection and response remains the most decisive advantage defenders have.

As 2026 unfolds, the tactics detailed here will continue shaping the threat landscape. But with stronger visibility, tighter governance of trusted workflows, and partnership-driven defense, organizations can stay ahead of those who expect to be underestimated.

On behalf of the entire Blackpoint Cyber team, we are proud to share this year's findings and look forward to supporting our clients, partners, and the channel community throughout 2026.

Sincerely,

Gagan Singh
Chief Executive Officer,
Blackpoint Cyber

Executive Summary

Threat actors didn't need to break in, they only needed to convince users, endpoints, and infrastructure to let them in. And the Blackpoint SOC spent 2025 closing those open doors faster than ever before.

The Year of 'Trusted' Compromise

If 2025 taught defenders anything, it is that the era of relying solely on complex, esoteric exploits to breach a perimeter is fading. It is being replaced by a more pragmatic, ruthless reality: Attackers no longer need to *break* in if they can simply *log* in.

Throughout the past year, the Blackpoint Security Operations Center (SOC) observed a decisive shift in tradecraft. Adversaries moved away from relying exclusively on zero-day vulnerabilities and bespoke malware, choosing instead to repurpose the very tools organizations trust to operate.

Whether it was SSL VPN gateways granting overly broad access, legitimate Remote Monitoring and Management (RMM) tools quietly doubling as Command and Control (C2) infrastructure, or Fake CAPTCHA prompts tricking users into executing code via the Windows Run dialog, the pattern was clear: **If the platform is trusted, the malicious activity is trusted, too.**

Blackpoint's Annual Threat Report details the operational realities of 2025, a year defined by the abuse of authorized workflows.

Initial access no longer hinges on breaking past hardened perimeter controls; it relies on blending into the workflow an organization has already built itself. A VPN credential becomes a golden ticket to accessing the crown jewels. A rogue RMM install hides in plain sight among legitimate IT tools. A trojanized installer impersonates the software users expect to download anyway.

However, this shift in attacker strategy has also revealed their weakness. By relying on recognizable behaviors rather than invisible exploits, they become detectable to those watching the context of an action, not just the code.

In 2025, the Blackpoint SOC disrupted 56% of all incidents before a payload could even be deployed. This proves that with the right visibility, we can close the doors attackers are trying to walk through.

Global Threat Trends

The Year in Review

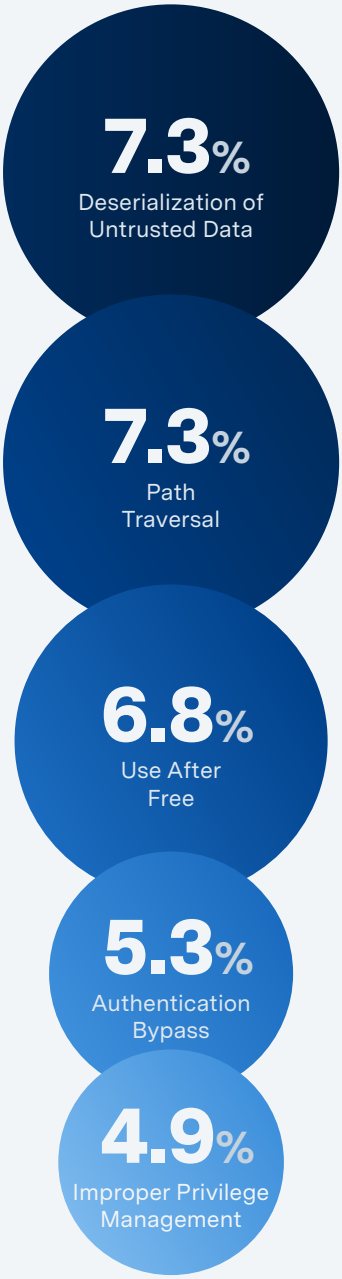

Top Exploited Vulnerabilities

The vulnerability landscape in 2025 was characterized not just by the sheer volume of disclosures — **over 45,000 vulnerabilities** were disclosed throughout the year — but by the specific types of flaws that attackers prioritized. Blackpoint's Adversary Pursuit Group (APG) focused its threat notifications on vulnerabilities with direct relevance to Managed Service Provider (MSP) managed environments, prioritizing technologies widely deployed across customer networks.

While not every notification reflected confirmed in-the-wild exploitation at the moment of publication, each was assessed as operationally meaningful based on exposure, impact, and likelihood of attacker interest. The year's notifications were dominated by specific Common Weakness Enumerations (CWEs) that highlight recurring fragility in software:

- **Deserialization of Untrusted Data (7.3%):** Allows attackers to manipulate data sent to an application, potentially leading to remote code execution.
- **Path Traversal (7.3%):** Enables access to files and directories stored outside the web root folder.
- **Use After Free (6.8%):** A memory corruption flaw that can be exploited to execute arbitrary code.
- **Authentication Bypass (5.3%):** Critical flaws that allow attackers to skirt login mechanisms entirely.
- **Improper Privilege Management (4.9%):** Flaws that fail to correctly restrict access permissions.

The Most Notable Vulnerabilities of 2025

Several specific vulnerabilities defined the year, catching the attention of both defenders and attackers due to their wide reach and high impact.

Gladinet CentreStack and TrioFox

CVE	CVSS	Disclosure Date	Type
CVE-2025-11371	6.1	October 2025	Local File Inclusion (LFI)
CVE-2025-14611	7.1	December 2025	Hard coded cryptographic keys

These vulnerabilities, disclosed in late 2025, highlighted the increasing use of exploit chaining against file-sharing infrastructure. CVE-2025-11371 enables authenticated file reads, allowing attackers to extract sensitive secrets such as machine keys. Threat actors can weaponize the stolen secrets via CVE-2025-14611 to bypass ViewState protections and trigger remote code execution (RCE), resulting in full server compromise. Blackpoint's SOC actively responded to alerts indicating likely exploitation, observing unusual HTTP requests from external IPs and non-standard user agents targeting CentreStack servers.

SonicWall

CVE	CVSS	Disclosure Date	Type
CVE-2025-40602	6.6	December 2025	Privilege escalation
CVE-2025-23006	9.8	January 2025	Deserialization of untrusted data

Disclosed in December 2025, this privilege escalation vulnerability underscored how attackers increasingly chain multiple vulnerabilities into a single exploitation path. SonicWall confirmed active exploitation at the time of disclosure, with CVE-2025-40602 reportedly being paired with CVE-2025-23006 to escalate privileges and gain administrative access to edge devices. SonicWall devices remained a consistent target throughout the year, accounting for 59% of SSL VPN incidents observed by the SOC. This activity reinforces that patching individual CVEs in isolation is no longer sufficient against modern attack chains.

Fortinet

CVE	CVSS	Disclosure Date	Type
CVE-2025-64446	9.8	November 2025	Path traversal
CVE-2025-59718	9.8	December 2025	Authentication bypass
CVE-2025-59719	9.8	December 2025	Authentication bypass

Scoring a critical 9.8 CVSS, CVE-2025-64446 allows unauthenticated threat actors to gain administrator-level access to the FortiWeb Manager panel. Multiple researchers reported active exploitation where attackers created admin accounts for persistence. Additionally, proof-of-concept (PoC) exploits were made publicly available, significantly lowering the barrier to entry for threat actors.

Citrix “CitrixBleed2”

CVE	CVSS	Disclosure Date	Type
CVE-2025-5777	9.3	June 2025	Out-of-bounds read

Dubbed “CitrixBleed2” due to its resemblance to previous massive Citrix exploits, this vulnerability allows threat actors to find valid cookies to hijack active NetScaler admin sessions. The availability of multiple PoCs shortly after its June disclosure made this a high-priority threat, as these devices often serve as centralized entry points into an organization’s network.

Top Three Attack Campaigns

In 2025, isolated events took a backseat to repeatable, scalable attack campaigns. The Blackpoint SOC tracked three dominant campaigns that drove the vast majority of malicious activity: Fake CAPTCHA/ClickFix, SSL VPN Abuse, and RMM Abuse.

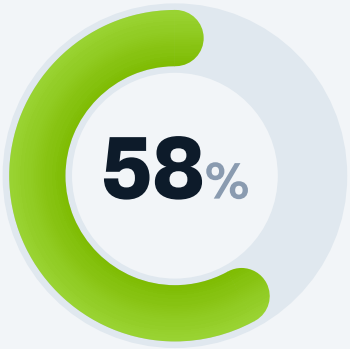

Fake CAPTCHA/ClickFix:

By far the most prevalent campaign of the year, this technique accounted for over half of all identifiable incidents. It leverages fake verification prompts (often mimicking Cloudflare or Google CAPTCHAs) to lure users into executing malicious payloads on their own systems.

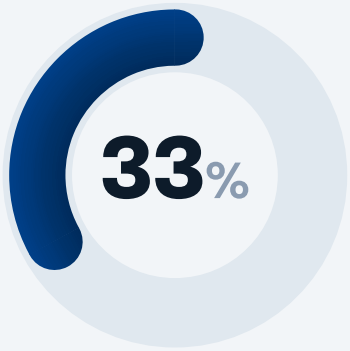

SSL VPN Abuse:

Network edge devices remain a critical bridge between internal networks and the internet. Their high value and frequent exposure made them a consistent target, accounting for nearly a third of identifiable incidents.

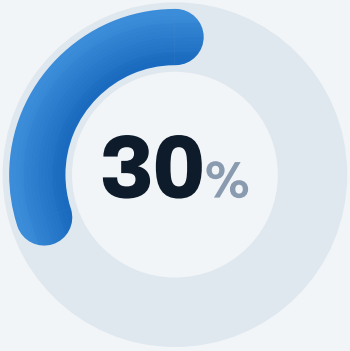

RMM Abuse:

The abuse of legitimate Remote Monitoring and Management tools highlights the shift toward "living off the land." By using trusted tools like ScreenConnect and AnyDesk, attackers blend in with normal IT activity, making detection significantly harder for traditional security controls.

Incidents may be assigned multiple tags based on observations made by the SOC during the course of the investigation.

Field Note

EDR tools frequently miss or delay alerts on RMM abuse, fake CAPTCHA chains, VPN-sourced lateral movement, and installer-based intrusions.

Having humans in the loop that are trained to identify the tradecraft rather than wait for the alert has become a critical differentiator.

From a SOC operations perspective, this trend validated Blackpoint's people-plus-telemetry model and solidified the SOC as the partner's most reliable early-warning system.

The Most Impacted Industries

In 2025, the Blackpoint SOC handled incidents spanning nearly all industries; however, Manufacturing, Healthcare, MSP, Construction & Engineering, and Financials were targeted most frequently.

Organizations in these industries are critical, and threats facing these industries range from rogue RMMs to fake CAPTCHA to suspicious initial access, persistence, or lateral movement. These incidents could have led to stolen credentials, theft of information, deployment of malware variants meant to provide persistent access or ransomware; or initial and persistent access that could be sold to other threat groups by Initial Access Brokers (IABs). IABs are financially motivated threat actors that profit through the sale of remote access to compromised networks. These accesses are often sold on cybercriminal forms like RAMP, XSS, Exploit, and BreachForums.

Fig. 01 Top Five Most Impacted Industries

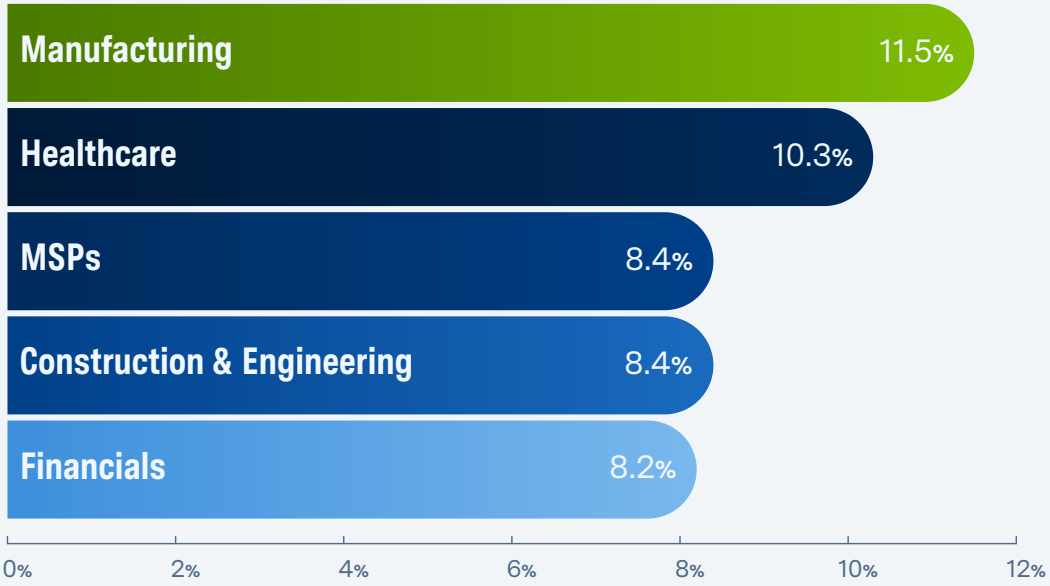

Percentage of total incidents observed by the Blackpoint SOC from Jan. 1, 2025 - Dec. 31, 2025.
Percentages may not total 100% due to rounding.

Manufacturing (11.5%)

Manufacturing organizations are attractive targets for threat actors for several reasons. Many rely on outdated systems that cannot be easily patched or segmented, creating a flat and enticing environment where ransomware operations can thrive. Beyond the technical vulnerabilities, the downstream impact of a manufacturer going offline extends well beyond the organization itself — delays cascade to customers, distributors, and partners, and can even affect the national or global economy. This widespread pressure accelerates negotiations and significantly increases the likelihood of a ransom payment.

Healthcare (10.3%)

Healthcare organizations remain a prime target due to the critical nature of their operations and the sensitivity of the data they hold. The Blackpoint SOC observed frequent targeting of this sector, where threats ranged from rogue RMM installations to suspicious lateral movement.

MSPs (8.4%)

Threats facing MSPs reflects a strategic shift, as attackers increasingly view MSPs not as the end target, but as a force multiplier. MSPs sit at the crossroads of trust and access, where compromising a single technician account or RMM instance can quietly unlock access to dozens or hundreds of downstream client environments.

Construction & Engineering (8.4%)

Construction and engineering organizations are heavily targeted by threat actors for several key reasons. These organizations operate on tight schedules with strict deadlines, meaning they cannot afford significant downtime Additionally, their IT environments tend to be highly decentralized, spanning multiple job sites, subcontractors, and shared credentials.

Financials (8.2%)

Financial organizations continue to be a target because they combine direct monetary value, sensitive data, and privileged access within always-on environments. This concentration of value and trust enables attackers to rapidly monetize intrusions while exploiting legitimate access paths that are difficult to distinguish from normal operations.

Attacker Behaviors

Tactics, Techniques, and Procedures (TTPs)

Initial access no longer hinges on breaking past hardened perimeter controls; it relies on blending into the workflow an organization has already built for itself.

If 2025 taught defenders anything, it's this: Attackers no longer need zero-days, bespoke malware, or elite tradecraft to compromise a business. They simply repurpose the very tools organizations already trust.

Across thousands of investigations, the Blackpoint SOC saw a clear pattern that modern intrusions don't begin with technical wizardry. They begin with normal behavior: a user installing an app, a remote session authenticating successfully, or a CAPTCHA prompt asking someone to 'verify' access.

Whether it was SSL VPN gateways granting overly broad internal access, legitimate RMM tools quietly doubling as Command and Control (C2) infrastructure, or the Windows Run dialog executing attacker-supplied commands, threat actors leaned heavily on one idea: If a platform is trusted, the activity will be trusted, too. And that trust gives them everything they need.

The shift is striking. Initial access no longer hinges on breaking past hardened perimeter controls; it relies on blending into the workflow an organization has already built for itself.

A VPN credential becomes a golden ticket to accessing the crown jewels. A rogue RMM install hides in plain sight among legitimate IT tools. A trojanized installer impersonates the software users expect to download anyway. A fake CAPTCHA turns a habitual click into remote code execution. None of this triggers early alarms because nothing looks overtly malicious.

2025 was the year attackers fully embraced this model at scale. They exploited human patterns, operational shortcuts, and architectural blind spots, and on top of this, they exploited unpatched CVEs. Unless defenders rethink what "routine" really means, including what should be trusted, what should be inspected, and what should be allowed to execute, 2026 will stay on the same trajectory.

The real problem isn't that attackers have become dramatically more sophisticated; it's that they no longer need to be.

Read the **blog** for more information.

1. SSL VPN Compromises:

The Front Door to Business Disruption

In 2025, the Blackpoint SOC consistently observed SSL VPN compromise as one of the most abused initial access vectors by threat actors.

This is largely due to how these appliances are designed and deployed. VPN gateways must be internet-facing to provide remote access into trusted network zones, and they often operate outside the visibility of traditional endpoint and network security tooling.

The combination of public exposure and limited telemetry makes SSL VPNs especially attractive targets, particularly for ransomware groups. Once access is gained, either through exploitation of a vulnerability or the use of stolen credentials, the threat actor enters the environment from a position that appears legitimate to many security controls.

From there, it is common to see a rapid transition into internal discovery, credential theft, data exfiltration, and ultimately large-scale encryption activity.

The Blackpoint SOC most frequently observed SSL VPN compromises originating from two primary paths: exploitation of vulnerabilities within the VPN appliance itself or abuse of valid but compromised credentials.

In vulnerability-driven intrusions, threat actors target the internet-facing VPN portal directly, then pivot from the appliance into internal systems, quickly moving toward the organization's crown jewels.

Case Study:

Inside the Front Door, Then Across the Network

In one significant incident involving compromised credentials, the Blackpoint SOC observed lateral movement originating directly from the SSL VPN address pool, followed by a rapid pivot into remote execution against the organization's crown jewels, most notably domain controllers.

This behavior is characteristic of an adversary who gains initial access through legitimate remote authentication, confirms internal reachability from the VPN subnet, and immediately begins operating against high-value systems—all within minutes.

Rather than deploying an interactive payload on a user workstation, the threat actor leveraged remote administration protocols to execute directly from the VPN-assigned address. This approach allows the attacker to bypass many endpoint-focused detections while maintaining an access pattern that appears routine to several security controls.

In this specific incident, the threat actor successfully authenticated to a FortiGate SSL VPN using compromised credentials and was assigned the IP address **10.0.0[.]4**. Instead of establishing a prolonged presence on an endpoint, the actor pivoted almost immediately into remote execution against a domain controller. This shift in targeting significantly compresses the intrusion timeline.

Fig. 02 Tools & Malware Observed with SonicWall SSL VPN Abuse

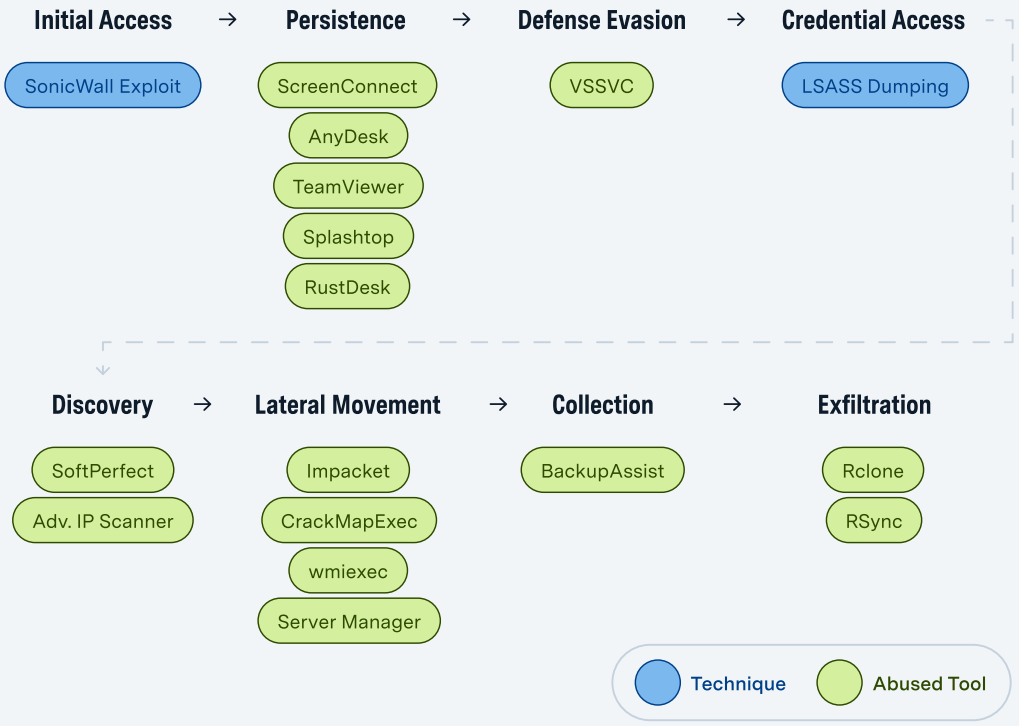

Once execution on a domain controller is achieved, the adversary gains control over the identity plane, enabling rapid privilege escalation and laying the groundwork for business wide impact.

After establishing the ability to execute remotely on the domain controller, the threat actor shifted to Windows Management Instrumentation (WMI) as the primary remote management mechanism. WMI is a legitimate and commonly used administrative tool, but in this case, its use stood out immediately because it did not align with the normal behavior or administrative patterns of the compromised account.

The WMI activity originated from the known rogue SSL VPN address, **10.0.0[.]4**, and was used to create Volume Shadow Copies on the

Field Note

Over **60%** of SSL VPN abuse incidents were very likely related to pre-ransomware activity. Blackpoint's Adversary Pursuit Group (APG) makes this determination based on observed behavioral indicators, tools deployed, and alignment with well-known ransomware tactics, techniques, and procedures (TTPs). There exists the potential that additional SOC-observed incidents have had the potential to escalate into ransomware.

domain controller. Shadow copy creation provides a point-in-time snapshot of the system, allowing access to files and system state that are typically locked during normal operation.

On a domain controller, this capability carries clear downstream value, as it can enable access to directory-related artifacts from a snapshot context. Even in cases where those artifacts are not immediately accessed within the same activity window, the combination of VPN-sourced remote execution and shadow copy manipulation is a strong indicator of preparation for follow-on actions.

This pattern commonly precedes credential access, persistence mechanisms, or broader domain-level staging as part of a larger intrusion lifecycle.

Why This Entry Vector Persists:

SSL VPN compromises remain consistently effective because they turn a single internet-facing system into trusted access to the internal network with minimal effort.

Once a VPN session is established, many environments treat that connection as if the user is physically on the corporate network, removing many friction points an attacker would otherwise have to overcome. This is especially valuable for ransomware operators because it dramatically shortens time to impact.

A valid VPN session provides a stable platform for internal discovery, credential access, lateral movement, and staging, often without the attacker needing to run noisy exploits or deploy malware on endpoints early in the intrusion.

Overprivileged network access is a major contributor. In many organizations, VPN users land in internal segments that are broadly routable and lightly segmented. Access controls are often permissive, meaning a single compromised VPN session can reach core services such as Active Directory, file servers, remote management interfaces, virtualization platforms, and backup infrastructure.

When routing and ACLs are too open, attackers can move quickly from initial access to privilege escalation and then to organization-wide ransomware deployment.

Abuse of legitimate RMM tools made up **30.3%** of identifiable incidents the Blackpoint SOC triaged throughout 2025.

2. RMM Abuse:
Access by Design

While not a new trend, the Blackpoint SOC observed a sharp increase in Remote Monitoring and Management (RMM) abuse as an Initial Access vector and persistence mechanism to retain access to environments after an initial compromise.

Abuse of legitimate RMM tools made up **30.3%** of identifiable incidents the Blackpoint SOC triaged throughout 2025. Several new campaigns were observed this year as users were targeted with phishing emails disguised as Social Security Statements, Tax Documents among others, to trick users into installing these tools.

While these are old tricks, threat actors can still rely on a steady percentage of users to fall for these phishing emails, allowing the threat actor the access they are after.

These RMM tools are attractive to threat actors as they already provide built-in remote access features, making Command and Control (C2) communications easily accessible for less advanced threat actors, as opposed to a custom-written implant or malware. Another reason for their attractiveness is that these RMM tools are often expected programs that are widely used by MSPs and IT staff for legitimate purposes.

The Blackpoint SOC often observes networks with multiple RMM/remote access tools installed on various devices. This can make a rogue installation of an additional tool by an attacker hard to spot and easy to miss.

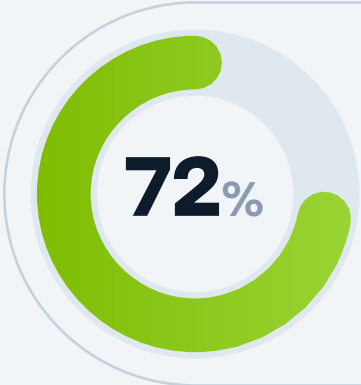

of the time, the SOC detected and responded before traditional EDR agents alerted.

Fig. 03 Rogue RMM Tool Incidents Observed by Blackpoint's SOC

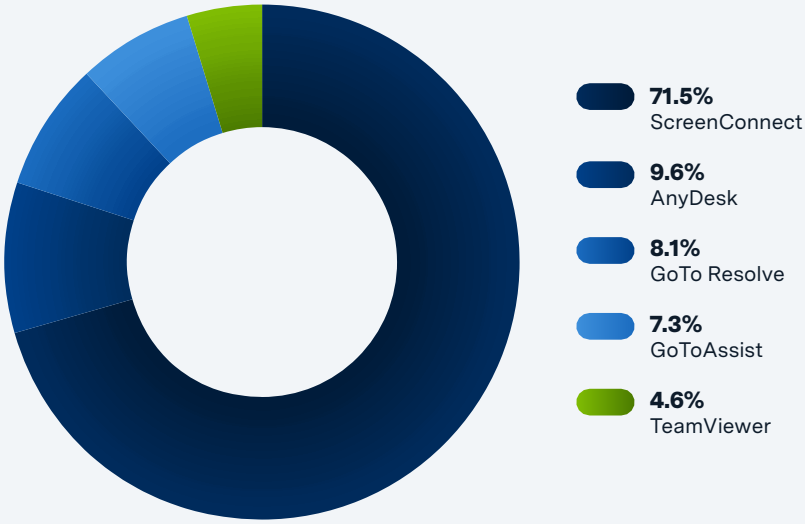

Percentage of total incidents observed by the Blackpoint SOC from Jan. 1, 2025 - Dec. 31, 2025. Percentages may not total 100% due to rounding.

Read the **blog** for more information.

Campaign Deep Dive:

Backdoor Your Backdoor

The Blackpoint SOC observed several specific campaigns throughout 2025 that abused a variety of RMM Tooling including ScreenConnect, GoTo, and FleetDeck. ScreenConnect was the most frequently observed tool, tied to an ongoing and widespread campaign dubbed Rogue ScreenConnect, appearing in more than **71%** of rogue RMM incidents *[Fig. 03]*.

As mentioned, the initial install often comes from executables in phishing emails disguised as Social Security and tax documents or installers of other legitimate applications. Once the user installs the initial payload, the tool in question is often programmed to run additional commands and payloads upon successful installation.

Often, these additional payloads include PowerShell or remote msiexec commands to download additional tools, including a second legitimate remote access tool. These commands are often web requests or remote installs that target a threat actor-controlled domain which is hosting the second stage payload, and the domains are typically recently registered, which allows them to bypass URL reputation-based security tools.

Case Study:
Phished for FleetDeck

In one campaign, the Blackpoint SOC observed threat actors utilizing malicious FleetDeck installers to allow attackers to take over a device.

While FleetDeck is far from an 'unknown' RMM tool, it's not commonly seen used for legitimate purposes. Campaigns involving FleetDeck also follow standard TTPs for rogue RMM tool abuse, as they start with a user being phished and tricked into allowing the initial application install to run.

In one analyzed incident, a user received a phishing email masquerading as an invoice. Not realizing the threat from this file, the user downloaded and ran the phishing payload **INVOICE#345634.exe**, allowing the initial FleetDeck installation to occur. Further research into the phishing payload confirmed that it was a binary for the FleetDeck installer.

Fig. 04 Tools & Malware Observed in Rogue ScreenConnect Campaign

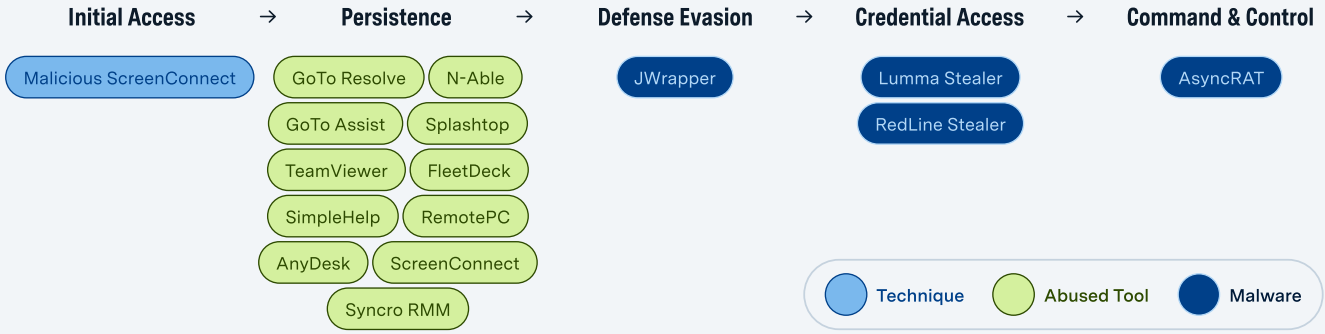

After the initial installation of the FleetDeck service, a remote installation command ran via msiexec. This command, launched by FleetDeck, called out to a suspicious domain: **hxxp://server[.]bnotuyiyed[.]online**. Within that domain, a **ScreenConnect.ClientSetup.msi** file was hosted and served as the target for the remote installation command.

Suspicious domains and top-level domains (TLDs) are often used by attackers, specifically for hosting ScreenConnect control infrastructure. These domains and

TLDs are an indicator that the Blackpoint SOC monitors heavily, as they can be a tell-tale sign that a given ScreenConnect install is illegitimate. The ScreenConnect install then served as a second foothold onto the device.

In many cases, the initial FleetDeck installation may be remediated and removed from the device. However, the additional ScreenConnect install may have a better chance of blending into the expected applications for that device.

Due to ScreenConnect being a common tool used by MSPs, duplicate installs of ScreenConnect may go unnoticed, as it can be an easy mistake to assume it is the legitimate install of the expected application.

Read the **blog** for more information.

3. Trojanized Installers:

When Installation Equals Compromise

Trojanized installers remained a reliable initial access mechanism throughout 2025 because they abuse normal user behavior and standard business workflows. Rather than defeating perimeter controls directly, the attacker convinces a user to download and run software that looks legitimate.

The installer runs in a trusted context, usually with user approval and, in many environments, with administrative elevation. After execution, it performs enough expected installation behavior to avoid immediate suspicion while also dropping a payload that establishes persistent access.

This technique is effective because it turns an ordinary business action into an initial access event. The user is not opening a suspicious attachment or running a random script. They are installing a tool they believe they need to do their job, often under time pressure and from the first convincing download page they find.

The process looks routine, the execution is user-driven, and the activity can blend into normal installation patterns long enough for the attacker to gain a foothold.

Case Study:

The Apps Users Trust Most (Microsoft Teams)

In one campaign observed by the Blackpoint SOC, a fake Microsoft Teams installer, **MSSetup.exe**, served as the complete initial access mechanism. The intrusion began with user-driven execution that appeared consistent with routine software installation, which delayed suspicion and reduced early friction.

The first stage activity occurred under a legitimate user context and blended into normal business behavior long enough to establish a foothold, and the user was directed to a malicious domain, **teams-install[.]top**, which hosted a page mimicking a legitimate download portal: “Download Microsoft Teams for desktop and mobile and get connected across devices.”

Fig. 05 Tools & Malware Observed in Trojanized Installers Campaign

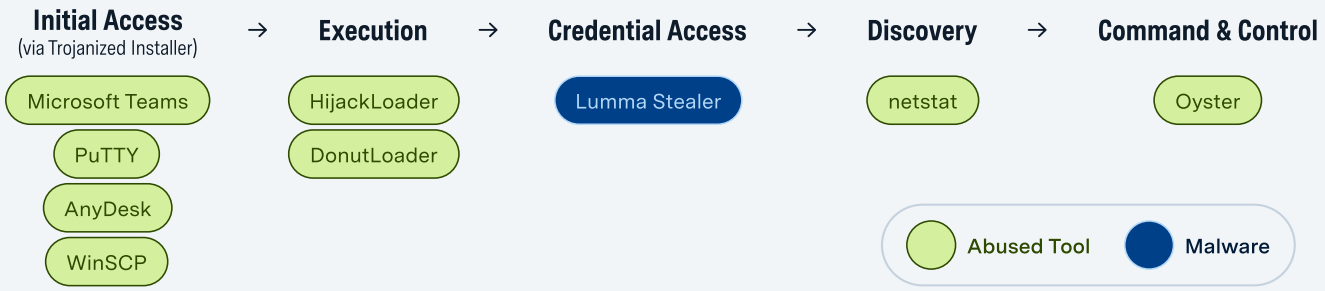

The kill chain then transitioned into payload activation using native Windows execution paths rather than an obvious standalone malware binary. The activity shifted into DLL-based execution through signed Windows utilities, with payload components staged in user-writable directories.

This approach reduced exposure to simple application controls and allowed the actor to run code through trusted host processes. During this phase, the Oyster backdoor was deployed and established as the access layer on the endpoint, providing a durable platform for follow-on operations.

After Oyster was in place, the operator moved into hands-on-keyboard activity focused on domain-aware discovery. The observed actions prioritized validating the compromised user context and privilege, identifying domain controllers and trust

The pattern is clear: If the platform is trusted, the malicious activity is trusted, too.

relationships, and enumerating server-class assets through directory queries. This discovery was then followed by reachability checks against internal systems to confirm where authentication and remote access were viable.

The activity was consistent with staging for lateral movement, rather than opportunistic reconnaissance, and indicated intent to expand beyond the initial host.

Specific command observations included:

- **rundll32.exe** executing malicious DLLs from AppData paths.
- **net.exe** and **nltest.exe** commands used to map domain admins, trust relationships, and local groups.
- PowerShell scripts leveraging **DirectoryServices** and **DirectorySearcher** to query for servers and user accounts.

This sequence reinforced why trojanized installers remained effective throughout 2025. Initial execution was user-approved and operationally plausible, payload activation relied on common Windows binaries and user-space staging, and post-compromise discovery used standard administrative tooling that can appear benign when viewed in isolation.

4. Fake CAPTCHA & ClickFix:

From Verification to Execution

Fake CAPTCHA campaigns gained strong traction this year, becoming one of the highest volume and most successful ways threat actors delivered malware. They worked because they leaned into something defenders rarely get to control: routine human behavior.

Everyone is used to clicking “verify,” solving a CAPTCHA, or confirming access to a file. Attackers took that muscle memory and paired it with one of Windows’ most powerful built-in features: the Run dialog box **Win+R**. By telling victims to paste a string into that box, attackers turned a normal, trusted shortcut into the start of a full compromise.

The Run dialog can execute anything the current user has permission to run, including cmd, PowerShell, mshta, or any other Living-off-the-Land (LotL) binary the threat actor wants to abuse. That makes it the perfect bridge between social engineering and real execution.

To amplify this threat, most of these lures appear on sites the user already trusts, often legitimate websites that have been compromised and quietly turned into delivery platforms. The pages are styled to look like Cloudflare checks, secure document portals, or normal access gates, so nothing feels out of place.

Throughout 2025, Blackpoint’s SOC tracked these campaigns deploying a diverse array of malware, including Lumma Stealer, RedLine, NetSupport RAT, AsyncRAT, and loaders like MintsLoader and Latrodectus. The tooling and payloads shifted constantly, but the core idea never changed: convince the user to run one command to deploy the payload.

Fig. 06 Malware Variants Observed in Fake CAPTCHA/ClickFix Campaign

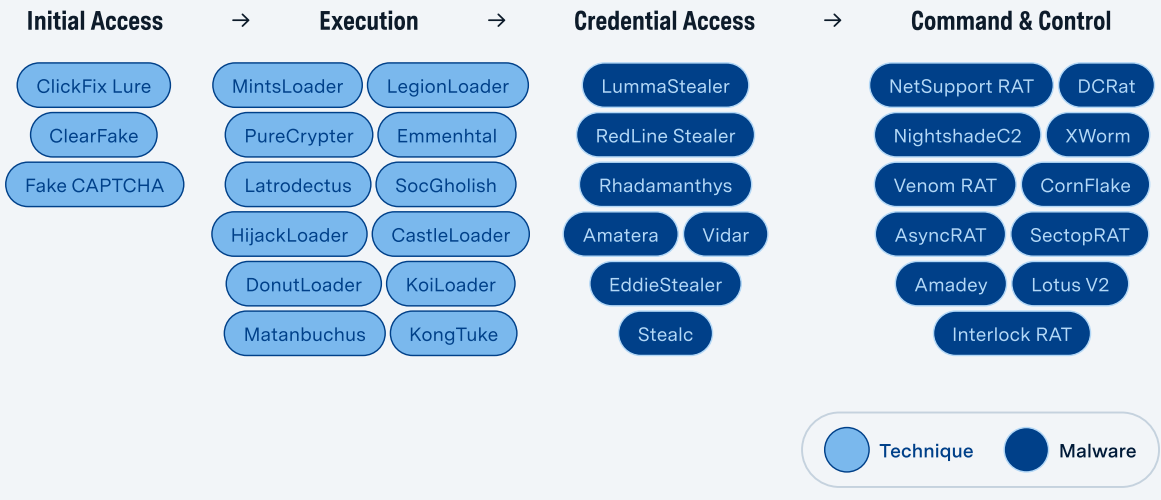

Field Note

Binance is one of the largest cryptocurrency exchanges in the world. Binance Smart Chain (BSC) is Binance's proprietary competitor to Ethereum, offering the ability to run decentralized apps and smart contracts.

Think of a smart contract as a digital agreement, which executes specific actions once the correct conditions are met. Etherhiding abuses BSC smart contracts to host malicious code within the Blockchain..

Technique Spotlight:
Etherhiding and Smart Contract Abuse

One of the biggest evolutions in fake CAPTCHA campaigns this year was how threat actors began managing their compromised websites at scale.

Instead of manually updating JavaScript on thousands of hacked WordPress sites every time infrastructure changed, they moved their payload logic into the blockchain, turning decentralized smart contracts into a centralized control layer for their campaigns.

In one incident the Blackpoint SOC investigated, users were browsing a legitimate company WordPress site that had been quietly injected with malicious JavaScript and a Cloudflare-style CAPTCHA telling them to press **Win + R** and paste a command to 'verify' themselves.

The injected JavaScript was not just another web-based loader pulling malware from a server. Instead, it reached out to the Binance Smart Chain (BSC) using a **JSON-RPC eth_call** to retrieve the next stage directly from a smart contract.

This technique, known as Etherhiding, stores base64 encoded JavaScript inside an immutable blockchain contract that cannot be taken down once deployed. By updating a single contract, the attacker can silently change payloads for every infected site and victim at once, without touching their compromised web infrastructure.

The code returned from the contract built the fake CAPTCHA overlay and copied a malicious mshta-based PowerShell command to the victim's clipboard. It also created a unique UUID for each visitor, stored it as a cookie, and polled the smart contract to see if that UUID appeared in the list of victims who executed the payload.

Case Study:

What Comes After the Click (NetSupport RAT)

One of the most common outcomes the Blackpoint SOC observed this year was the deployment of NetSupport RAT, a legitimate remote access tool abused by threat actors for initial access.

In one incident, the fake CAPTCHA chain led to the execution of a one-line curl-based command that downloaded and launched a malicious batch script, kicking off a quiet NetSupport RAT deployment using only native Windows utilities.

That initial batch script handled the entire staging and deployment process. It began by suppressing output and enabling delayed variable expansion to reduce visibility. It then defined four key variables: the URL of a second-stage ZIP archive, the local staging directory, the extraction directory, and the full path to the final NetSupport payload, **client32.exe**.

PowerShell was used to retrieve the ZIP archive from attacker-controlled infrastructure and save it locally with a hidden window. The archive was then extracted using the **.NET System.IO.Compression.ZipFile** class into a newly created directory under the user's AppData folder.

Immediately after extraction, the script launched **client32.exe**, which is a NetSupport Manager binary repurposed here as a Remote Access Trojan, giving the attacker live interactive control of the system.

To maintain access, the script created a Current User Run registry key named **Support11** that pointed to the NetSupport binary, ensuring it would be launched every time the user logged in.

Fig. 07 Fake CAPTCHA Prompt and the Windows "Run" Utility

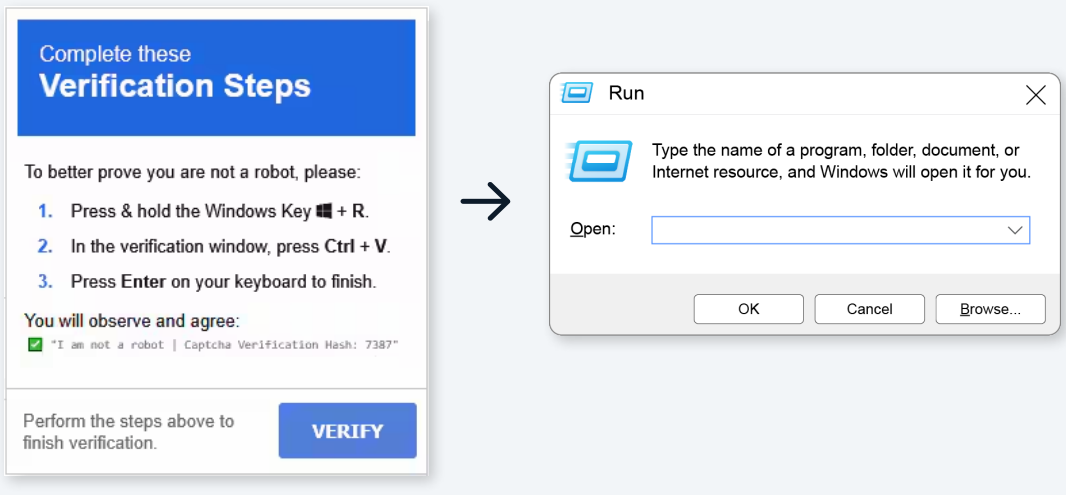

5. AiTM Phishing:

MFA Working as Designed

Adversary-in-the-Middle (AiTM) phishing remained one of the most effective cloud-focused attack techniques throughout 2025. Unlike traditional phishing, these campaigns did not rely on stealing credentials and hoping for the best. Instead, they targeted something far more valuable: the authenticated session created after a successful sign-in.

AiTM attacks work by placing the attacker directly in the middle of the authentication flow. When a victim enters their credentials and completes MFA, those responses are relayed in real-time to the legitimate identity provider. From the user's perspective, everything looks normal.

From the attacker's perspective, the goal is not the password or the MFA code, but the session cookies or access tokens issued at the end of the login process. Once captured, those session artifacts can be replayed to access cloud services as the victim, often without triggering MFA again until the session expires or risk conditions change.

This is why AiTM campaigns continued to succeed even in environments with MFA enabled. MFA was never bypassed or broken. It worked exactly as designed. The problem is that once MFA succeeds, trust shifts to the session itself, and AiTM attacks are built specifically to steal and reuse that trust. No malware is required, no exploit is needed, and no brute force activity occurs.

Fig. 08 Observed AiTM Phishing Attack Chain

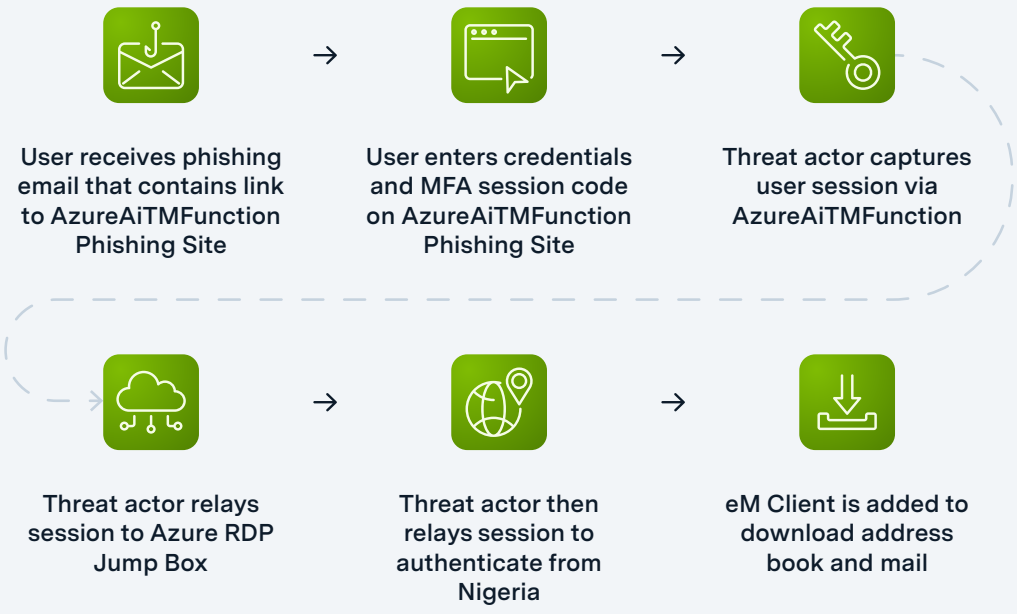

Case Study:

The Anatomy of an AiTM Phish

In one campaign the Blackpoint SOC observed in 2025, a threat actor leveraged an open-source framework known as AzureAiTMFunction to conduct an adversary-in-the-middle phishing attack against a Microsoft 365 user. Rather than hosting a traditional phishing site on rented infrastructure, the attacker deployed the phishing logic as an Azure Function, allowing them to host Microsoft look-alike login pages in a serverless, low-overhead way that blended in with legitimate cloud traffic.

When the user interacted with the phishing email and followed the link, their authentication traffic was silently proxied through the Azure Function. As the victim entered credentials and completed MFA, the framework relayed those requests in real-time to Microsoft while simultaneously capturing the authenticated session material.

In this case, Axios was used as the HTTP client to manage redirects, headers, cookies, and multi-step authentication, which is why many AiTM campaigns present an Axios user agent during authentication.

Once the session was captured, the attacker immediately relayed it to a remote RDP jump box hosted on a VPS. These jump boxes are commonly used to proxy cloud access and obscure attacker origin. During post-exploitation, the attacker used the stolen session to consent to a legitimate Azure application, eM Client, granting persistent access to the user's mailbox and contacts. This allowed the attacker to collect email data and harvest contacts for future phishing activity, extending the impact beyond the initial account.

Strategic Defense and Recommendations

Treat Remote Access as a High-Risk, High-Privilege Activity

The Annual Threat Report (ATR) repeatedly highlights SSL VPN abuse as one of the most common and damaging initial access vectors. MSPs should:

- Enforce MFA universally on VPNs without exception; eliminate legacy authorization and fallback paths.
- Aggressively patch or replace aging VPN appliances, especially SonicWall, Fortinet, and other devices shown to be targeted throughout the ATR.
- Strictly segment VPN address pools, so remote sessions cannot directly reach domain controllers, hypervisors, backups, or RMM servers.
- Block direct RDP/remote admin access from VPN ranges, forcing jump hosts or PAWs.
- Enable conditional access logic for VPN logins (geo anomalies, first-seen devices, TOR/proxy IPs).
- Centralize and monitor VPN authentication logs for credential spraying, impossible travel, and abnormal session durations.

Lock Down RMM Tools Like Your Business Depends on It (because it does)

RMM abuse made up **13%** of all SOC-triaged incidents in 2025, and the report shows attackers frequently install multiple RMMs to maintain persistence. Fortified defense for MSPs entails:

- Implement Managed Application Control (MAC) or equivalent application allowlisting to prevent any unapproved RMM installation.
- Create and enforce an authoritative RMM inventory and document exactly which RMM tools are approved and where they should exist.
- Block installation of executables for all non-admin users.
- Disable local admin rights for technicians except when elevation is formally requested and logged.

- Remove all legacy or unused RMM agents; attackers hide inside “old” agents left behind.
- Block uncommon **TLDs (.online, .top, .live)** to disrupt RMM-related malicious infrastructure seen in campaigns.

Make Software Installation a Controlled Event, Not an Attack Vector

The ATR shows that trojanized installers consistently delivered Oyster backdoors, RATs, stealers, and loaders—all because users installed something they believed was legitimate. To stay proactive, MSPs are advised to:

- Ban ad-hoc software downloads. Publish an approved software catalog and require all installations to flow through it.
- Block execution from Downloads, Temp, and AppData, where nearly all malicious installers staged payloads.
- Enforce allowlisting for installers and signed binaries from vetted publishers.
- Deploy browser isolation and DNS filtering to eliminate malvertising and SEO-poisoned download pages.
- Alert on installer behavior spawning **rundll32**, **PowerShell**, **cmd**, **msiexec**, **wscript**; all behaviors shown in the ATR’s kill chains.
- Review newly dropped DLLs and AppData folders for persistence mechanisms.

Counter Fake CAPTCHA & ClickFix Attacks with Behavioral Controls

Fake CAPTCHA campaigns were one of the highest-volume intrusions of 2025, often leading to NetSupport RAT, Lumma/RedLine/Vidar, and a variety of loaders. To stay ahead of this growing trend, MSPs should:

- Disable/limit Windows Run dialog **Win + R** via GPO for standard users.
- Restrict PowerShell and require constrained language mode or admin-only execution.
- Monitor LOLBAS chains such as:
  - **cmd** → **curl** → **powershell**
  - **mshta** → **powershell**
- Block outbound Web3 RPC endpoints (Binance Smart Chain), which enable Etherhiding delivery.
- Alert on clipboard-based execution patterns often used in Fake CAPTCHA payloads.
- Block low-reputation TLDs used to host encrypted loaders.

Strengthen Identity, MFA, and Privileged Access

Multiple sections of the ATR emphasize identity compromise, misuse of valid accounts, and abuse of routine admin behavior. It's up to MSPs to:

- Move to phishing-resistant MFA (FIDO2/WebAuthn).
- Monitor MFA fatigue, token theft, and session reuse (especially via AiTM).
- Enforce privileged separation: technicians should have separate accounts for admin and non-admin tasks.
- Implement Conditional Access with device, location, and session-risk evaluation.
- Audit and remove dormant privileged accounts, including stale VPN accounts.
- Deploy Privileged Access Workstations (PAWs) for high-risk admin tasks.

Harden the Cloud: M365 & Entra

The ATR's cloud section emphasizes identity-driven compromise and the volume of disabled accounts. MSPs can proactively prioritize this threat when they:

- Audit Conditional Access policies for gaps in location, device compliance, and sign-in risk scoring.
- Restrict external application consent to admins only.
- Monitor for suspicious inbox rules, OAuth grants, and mass MFA resets.
- Enable continuous access evaluation (CAE) where supported.
- Enforce least privilege across Azure roles, especially "Global Admin" and "Privileged Authenticator".
- Perform periodic reviews for shared mailboxes, service accounts, and non-MFA accounts.

Invest in Data Resiliency & Response Readiness

Attackers frequently inhibit recovery by manipulating shadow copies, backups, and hypervisor infrastructure. Staying vigilant requires MSPs to:

- Enable immutable backups and off-network replication.
- Perform regular bare-metal recovery tests monthly or quarterly.
- Segment backup infrastructure from production networks.
- Protect hypervisors with MFA and network segregation to prevent RMM-based or VPN-based pivots.

Conclusion

Across 2025, the most effective attacks weren't powered by stealth—they were rooted in trust. Threats actors advanced by blending into normal operations, leveraging software, remote management platforms, and built-in system utilities. This shift will accelerate in 2026, as adversaries exploit familiarity as their primary attack surface.

But familiarity cuts both ways. The more attackers rely on human habits and approved tools, the more defenders can rely on context, behavior, and early-stage detection. Fortunately, that is where the advantage shifts back. If 2025 was the year attackers weaponized trust, then 2026 must be the year we redefine it.

At Blackpoint Cyber, trust and dedicated support are at the core of everything we do. It's our mission to help clients and partners mature their security programs from reactive response to proactive risk reduction. Our 24/7 human led, AI-powered SOC is built to rapidly detect and stop cyberthreats in industry-leading response times. Quick action is essential when a single VPN session, a fake installer, or a rogue RMM tool can open the door to a much larger incident; the faster that activity is detected, the easier it is to contain before it becomes downtime or broader impact.

As 2026 progresses, the tactics described in this report will keep shaping how the threat landscape unfolds. Effective security will rely on strong collaboration, clear visibility, and defenders who understand how these intrusions take shape in real environments. We're here to ensure that you can navigate the year ahead with confidence, knowing that your business, environment, and operations are protected.