

EXECUTIVE SUMMARY

2026 THREAT DETECTION REPORT

A security leader's guide to the
year's cyber threats and risks

EXECUTIVE SUMMARY

2026

THREAT DETECTION

REPORT

Securing identities, AI, and browsers should be top of mind as security leaders set their security strategy for 2026 and beyond. The 2026 Threat Detection Report has the tactical guidance your team needs to tune out the noise and focus on diminishing the risk your organization faces from cyberthreats.

Red Canary's global view of the threat landscape

The Threat Detection Report is Red Canary's annual macroanalysis of cyber threats, adversary techniques, and security trends gleaned from the extensive **security operations work** we performed in partnership with nearly 2,000 organizations in 2025. The report is designed to inform security leaders and practitioners about the current state of the threat landscape—and to provide them with actionable analysis to help guide their security operations strategy.

This executive summary synthesizes the most important information from the report that every CISO, security leader, and technical executive must know in order to protect their organization in 2026 and beyond. While the data in this report is derived from threats Red Canary detected in 2025, we know from more than a decade of similar analysis that many trends, threats, and techniques remain prevalent for years.

The scope of our analysis

We derived the intelligence that comprises the 2026 Threat Detection Report from more than **110,000 confirmed threats** detected across **4.5 million endpoints, networks, cloud systems, identities, and SaaS applications**. However, that only tells part of the story of what we detected. Just as important is the scope of our collection and the vast volume of signals, alerts, and potentially malicious events that we examine but suppress because they are benign or inconclusive—in other words, the noise security teams would otherwise have to scrutinize or willfully ignore.

To that point, **Red Canary processed 329 billion records per day** and **ingested 305 petabytes of telemetry in 2025**. We triaged **419 million alerts, investigated 8.5 million potentially malicious events**, and, in the end, **detected 110,000 threats**, including **34,000 higher-severity threats** that are the primary focus of this report.

Read our Methodology to understand the breadth of our data and how we count.

[Explore](#)


That this level of reduction is possible is a bug for our industry but a feature of Red Canary. We achieve it daily across our customers': EDR tools, cloud platforms, identity providers, email clients, SaaS applications, SIEMs, network applications, and other third-party security tools.

The key risks highlighted in this summary:




The critical importance of identities

Identity compromise attempts are ubiquitous and increasingly risky as identities are the gateways to everything in the cloud, including apps and AI tools.



The two main flavors of AI abuse

Adversaries are leveraging AI for threat development and also seek to compromise AI tools to gain fast and easy access to sensitive datastores.



Browsers are the endpoints that matter most

Adversaries are targeting browsers to steal credentials and deliver malware because the browser is the new endpoint where users authenticate and access apps.



The critical importance of identities

We detected an **850 percent increase in identity threats** in 2025, which accounted for 53 percent of overall detection volume, up from 20 percent in 2024. Likewise, **T1078.004: Cloud Accounts**, a technique that includes identity compromises within its scope, was the **most prevalent ATT&CK technique** we detected this year. While this increase may sound alarming and seem to suggest that adversaries have massively ramped up campaigns targeting identities, that's not exactly what's happening.

WE DETECTED AN

850%

increase in identity threats in 2025

WHICH ACCOUNTED FOR

53%

overall detection volume

UP FROM

20%

in 2024

TOP TECHNIQUES DETECTED IN 2025


1. Cloud Accounts




2. PowerShell




3. Windows Command Shell




4. Data from Cloud Storage




5. Ingress Tool Transfer




6. Email Forwarding Rule




7. Windows Management Instrumentation




8. Malicious Copy and Paste



9. Email Hiding Rules



10. Obfuscated Files or Information



Identities are the most critical security boundary in your organization.

The reality is that adversaries have always placed a premium on identity compromise. That premium has only increased as organizations have adopted identity providers that authorize access to a wide array of accounts once a user authenticates. The increase suggests a growing interest among adversaries to compromise identities, but it also reflects increased adoption of Red Canary's identity security products and improved detection coverage for identity products—particularly how we've been using **agentic AI** to more accurately identify anomalous logins.

The bottom line is this: Identities are the most critical security boundary in your organization, and adversaries are increasingly prioritizing identity compromise to access cloud systems, SaaS apps, corporate AI tools, and more.


Identity-related alerts are notoriously noisy, contextually confusing, and difficult to action. It's important for organizations to monitor suspicious logins, but it's equally important not to get mired in them.

As such, we recommend the following preventative steps:

- Invest heavily in identity security controls.
- Enforce the principle of least privilege.
- Implement **zero trust security controls**.
- Leverage conditional access policies to universally enforce MFA.
- Restrict access based on device types, user agent strings, and more.
- Limit the use of unsanctioned VPNs.

Learn how Red Canary uses agentic AI to detect malicious logins

Meet our agents






The two main flavors of AI abuse

We continue to think that AI offers more benefits to defenders than it does to adversaries. However, it's becoming clear that adversary abuse of AI systems falls into two distinct categories:

1. Attempts to compromise **corporate AI systems and tooling**
2. Using AI to develop or **execute threats or campaigns**

ADVERSARIES WILL USE AI AND COMPROMISE AI

There's nothing new about defending against both/either.



Defending against threats developed with AI

AI-powered threats represent an evolution in tooling, not a revolution in attack techniques. Adversaries leverage AI for the same reasons that everyone else does: as a force multiplier and for automation. Some of the specific benefits adversaries may gain by using AI include:

- writing malicious code faster or modifying existing code to evade detection
- writing more convincing phishing messages in unfamiliar languages
- processing stolen data
- performing reconnaissance
- intercepting sensitive information via malicious AI

While not exhaustive, this list highlights an important point that security leaders need to understand: AI isn't presently enabling adversaries to perform actions they couldn't perform previously, but it does afford them the ability to operate faster and in more volume.

Protecting environments from AI-powered threats relies on the same fundamentals as any existing threat:

- defense in depth by way of sophisticated **detection and response capabilities**
- **zero trust** to limit the blast radius if an intrusion occurs
- continuous monitoring to catch threats early and often

The one potential novelty in the above list—and a great segway that straddles the line between AI-powered threats and threats to AI infrastructure—is adversary use of AI to intercept sensitive information. As users become increasingly comfortable interacting with and granting access to agentic AI systems and AI helpers, adversaries will seek to exploit this comfort with malicious AI agents masquerading as legit AI tools that may request access to user accounts and information.

The prevalence of threats like **JustAskJacky**, malware masquerading as an AI helper app, and the emergence of **malicious skills in OpenClaw** demonstrate that adversaries are actively exploiting user interest in AI assistants to deliver malware.

In some ways, this is no different from traditional social engineering and phishing schemes that seek credential access, which requires employee vigilance, user awareness training, and a slew of other security controls to mitigate. However, the magnitude of the threat may be greater given AI's ability to move fast and process large quantities of data.

Securing AI tools and systems

The business world is diving head first into AI adoption, often without considering all of the security implications that come along with it. As businesses connect agents and other AI tools into their infrastructure and give them access to their applications and data, adversaries will inevitably attempt to compromise those AI tools.

Threats to AI infrastructure, including **prompt injection** and model hijacking, exploit deep interconnectivity with development environments, cloud resources, and data stores. Adversaries may try to compromise AI tools to gain access to the resources available to those tools.

More surreptitiously, there's a looming **supply chain risk**—very similar to what we've seen with **npm packages**—whereby adversaries will likely try to infiltrate the AI development supply chain by compromising publicly available **MCP servers** and other AI tools.

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

Securing the use of AI within an organization requires a layered approach. We recommend the following:

Treat AI tools as privileged infrastructure.

Lock down credentials.

Audit open source AI components and tools.

Understand and monitor your AI inventory.

Segment and contain AI tooling.

Establish governance and training.

Fight fire with fire

Speed is the primary benefit that AI affords adversaries, but defenders can leverage AI to move faster as well. Responsibly integrate **AI into your security operations center** to whatever extent is possible.

One of the best ways to counteract the benefits adversaries may gain from AI is to use it yourself. Leverage automation, via AI or otherwise, to match the speed and adaptability of AI-powered adversaries.

Check out our guide on building AI agents into your SOC.

Get the guide






Browsers are the endpoints that matter most

In a world of identity providers and cloud-everything, users primarily authenticate and access sensitive information through their browsers. Adversaries have noticed, and so the browser is where an increasing amount of the action is happening.


The majority of the **top 10 threats** we detected this year abuse the browser in some way. They deliver payloads via the browser or they target information stored in or transiting the browser. Common delivery mechanisms include malvertising, SEO poisoning, fake update schemes, and **paste and run**. **Stealers** and **malicious extensions** may target browser-based cookies or session tokens, and adversaries commonly leverage adversary-in-the-middle schemes and toolkits to steal credentials in real time with the help of spoofed login pages.

TOP 10 THREATS DETECTED IN 2025


1. Amber Albatross




2. JustAskJacky




3. Tampered Chef




4. NetSupport Manager




5. LummaC2




6. Scarlet Goldfinch




7. KongTuke




8. SocGhosh



9. MintsLoader



10. Rhadamanthys



Organizations must have optics into their browsers and implement security controls to detect and mitigate the risk posed by browser-borne threats. Beyond defense in depth and advanced monitoring, important browser security controls to consider include:

- Leverage ad-blocking solutions where possible.
- Manage browsers through MDM, group policy, or admin consoles.
- Implement allowlists to limit extension installation.
- Use short-term access tokens to limit the impact of stolen session tokens.
- Consider **token protection features** that bind token use to familiar devices.

Developing a security strategy for 2026

We encourage you to read the entire Threat Detection Report (or at least the parts of it that align with your security priorities) and to share it with your team and colleagues. The **report PDF** highlights what's new and novel, and the **full web version** includes extensive analysis of the most pressing trends, threats, and techniques, including mitigation, detection, and testing guidance.

Explore our new Threat Library

Dive into all of our analyses of
common threats and techniques
from previous years.



You can use this report to inform your own threat model by comparing what we detected with what you observe in their own environment. Additionally, use the analysis as a playbook to prioritize threats and techniques, focusing on the early-stage adversary activity that will empower your security team to detect threats early and often, before adversaries achieve their objectives.

The data we use to compile this report is not mere software signals but the result of hundreds of thousands of hours of expert-led investigation and threat confirmation. We hope this report helps you achieve better security outcomes in 2026 and beyond.



red canary

a **zscaler** company