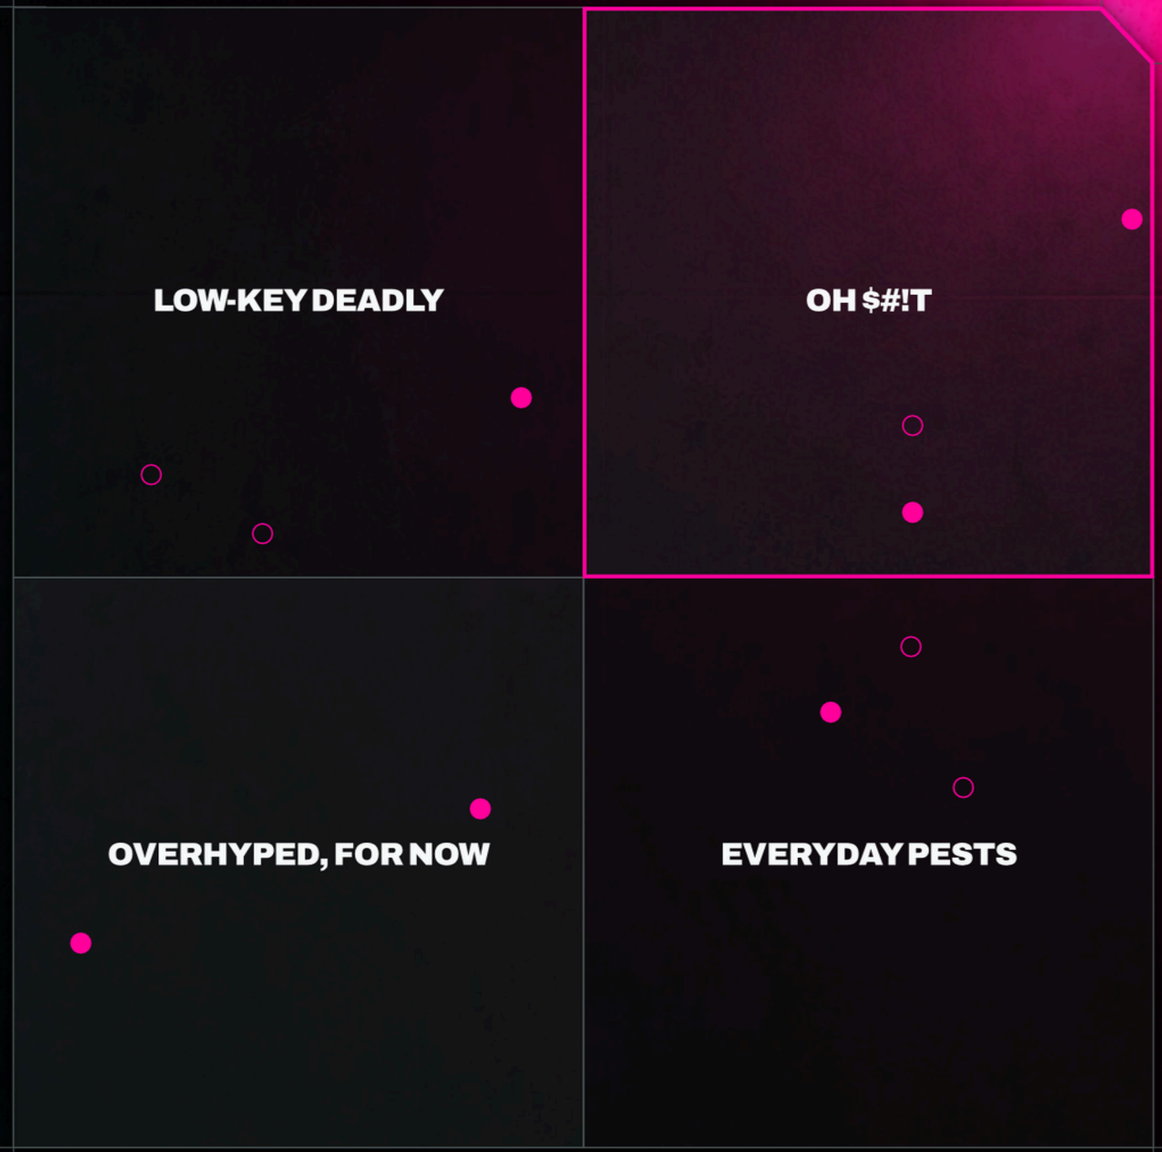
HUNTRESS

THE TRAGIC QUADRANT

Ranking the Cyber Tactics Wrecking Businesses



❯ [pucker_factor] risk, ranked
[status] hype filtered

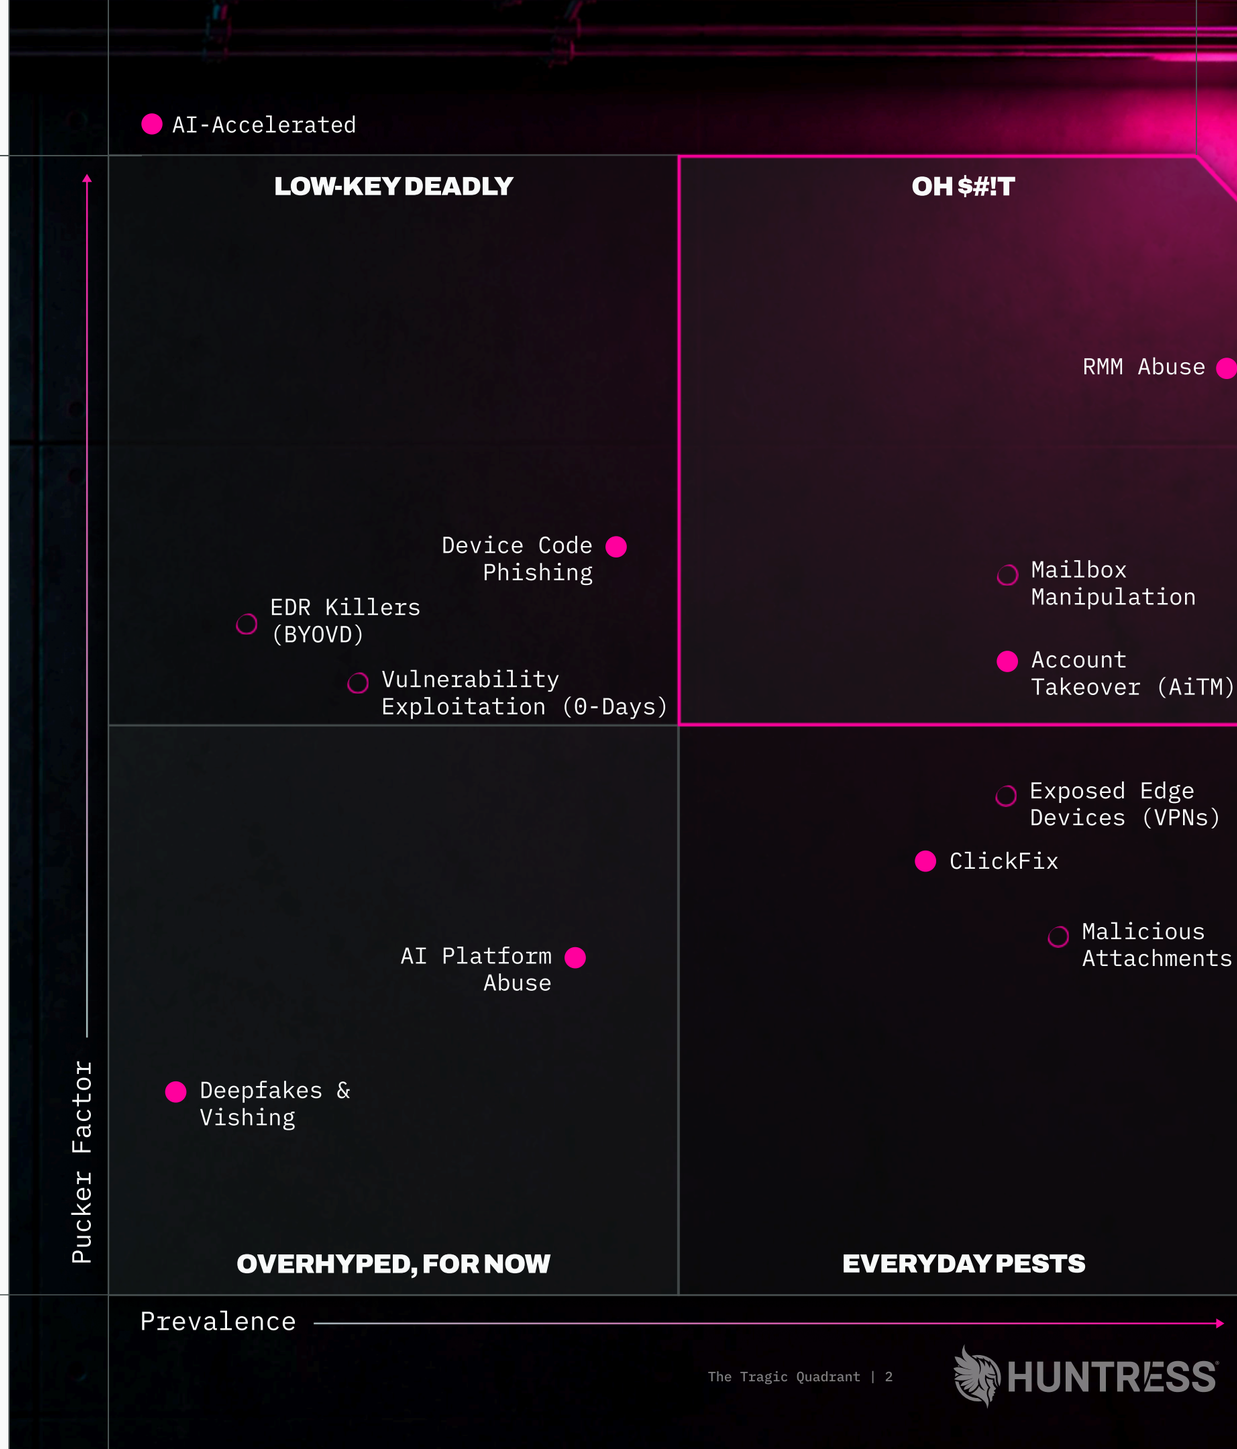
THE HUNTRESS TRAGIC QUADRANT

If you're running a business without a Fortune 500 budget, you probably don't have the time, resources, or risk tolerance to chase every cyber threat that the industry hypes up.

That’s why we created the Huntress Tragic Quadrant to help you beat the hype. It shows you the threats we’re catching and shutting down most often, not what's trending in the news cycle. This is *the* list, the tactics that really deserve your attention, and where you can close the gap before they become major business disruptions.

Each tactic on this quadrant is ranked based on how often we see it, and how damaging that tactic could be. The farther to the right means we commonly see that tactic in play across the environments we monitor. And the higher up is what we call “pucker factor,” or the closer you are to major damage when you see that tactic. It’s worth noting that none of these tactics work alone. A foothold from one tactic often opens the door to two or three more, and every step through that door pushes the potential damage higher.

If you only have time to tackle one thing, start in the top right. Those are the tactics that show up most often and have the highest chance of impacting your business immediately. Everything else on this chart still deserves a plan, but that corner is where detection and prevention pay off fastest.

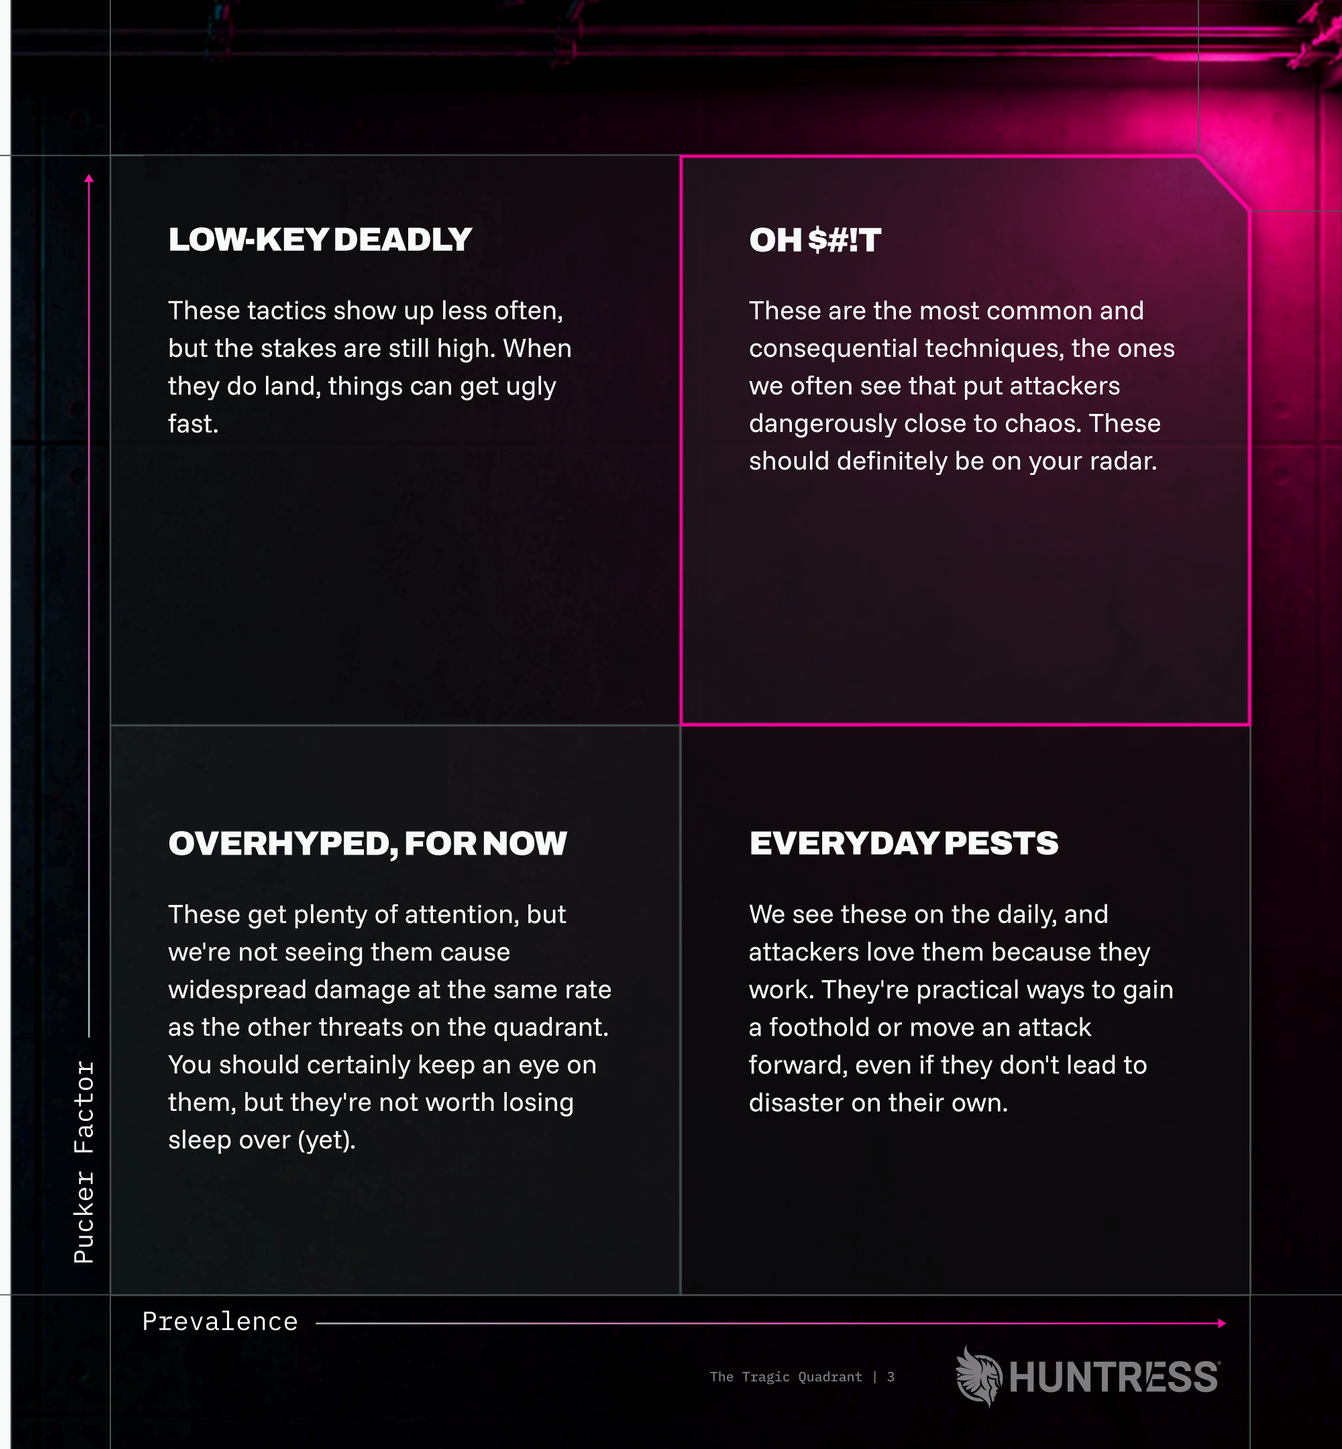


HUNTRESS

HOW THE QUADRANT WORKS

Not every threat deserves the same degree of attention. The Tragic Quadrant plots what we're seeing based on 1) how prevalent a technique is the farther right you go 2) how close it can put your organization to danger the higher you go. Wherever you see a solid dot, those are the tactics that we've seen are accelerated by AI.

This report breaks down some, but not all, of the threats we're seeing across our partner and customer environments. You'll learn why these specific threats landed on the Tragic Quadrant. From there, we'll share practical steps you can start taking to protect your business against them.



WHERE THE DATA COMES FROM

The Tragic Quadrant is built on telemetry across more than 5 million endpoints and 15 million identities, spanning nearly 300,000 protected organizations. The threats on this quadrant draw on data from detections and incidents we've observed in 2026, the Huntress 2026 Cyber Threat Report, and in-the-wild examples our Security Operations Center (SOC) has investigated across partner and customer environments.

About the Researchers



BEN NAHORNEY

Staff Technical Community Engagement Writer

Ben has been at this long enough to fondly remember a time when ransomware just asked for gift cards. Ben is an amphibian of sorts. He's often seen roaming among the land mammals who write about cybersecurity, but at other times he swims among aquatic life, diving into data lakes to surface trends hidden in telemetry—as he did for this report.



LINDSEY O'DONNELL-WELCH

Principal Technical Community Engagement Writer

Lindsey has more than a decade of experience translating complex security topics into compelling stories. Before Huntress, she was rooted in journalism, covering the intersection of technology and business. She's reported for top cybersecurity publications, delivering trusted coverage on everything from the SolarWinds supply chain breach to the SEC's cybersecurity disclosure rules.



BETH ROBINSON

Senior Technical Copywriter, Content

Beth's cybersecurity journey started at the NSA, where she spent four years immersed in computer network exploitation (CNE). Since then, she's stacked up over a decade of experience in security, including roles at the CIA, Mandiant, Bishop Fox, and runZero.



JAMIE LEVY

Senior Director, Adversary Tactics

Jamie has worked over 15 years in the digital forensics industry, conducting investigations as well as building out software solutions. She is a researcher, developer, board member of the Volatility Foundation, and co-author of The Art of Memory Forensics.



AARON DEAL

PR Manager

Aaron connects the technical side of Huntress' story to the media, partners, and security community who need to hear it. With 5+ years of experience in cyber, his background spans communications, social media, and copywriting. As an honorary member of Adversary Tactics, he partners extensively with the security team to turn complex topics into clear, compelling stories for a range of audiences.



DRAY AGHA

Senior Manager, Tactical Response

Dray specializes in digital forensics and incident response and is interested in defensive and offensive information security. He holds an OSCP certification (Offensive Security Certified Professional), among many other certifications, and is an active contributor to the security community through research, public speaking, and open-source tooling.

OH $#!T

LOW-KEY DEADLY

OH $#!T

OVERHYPED, FOR NOW

EVERYDAY PESTS

```
❯ [quadrant.status] prevalence: HIGH | pucker_factor: HIGH
  priority: close_the_gap before routine becomes disruption
recommendation: start here, preferably before the "oh $#!T" part
```

HUNTRESS

RMM ABUSE

Remote monitoring and management (RMM) tools are used by organizations of all types and sizes to manage computers from anywhere. Because these are trusted tools, spotting malicious activity can be a challenge. They also give an attacker the same keys your IT team already trusts: persistent access, remote command execution, and a way in that looks like business as usual.

Why it Lands Here

RMM abuse jumped 277% year over year in 2025, and in Q1 2026, 45% of endpoint-related incidents involved RMM abuse. It's the single most common threat category we see, and just one hop from ransomware or data theft. Threat actors are already using AI to generate fake document-share and service-agreement lures that trick users into downloading unwanted RMMs on their devices.

45% of endpoint-related incidents involved RMM abuse in Q1 2026

The Bottom Line

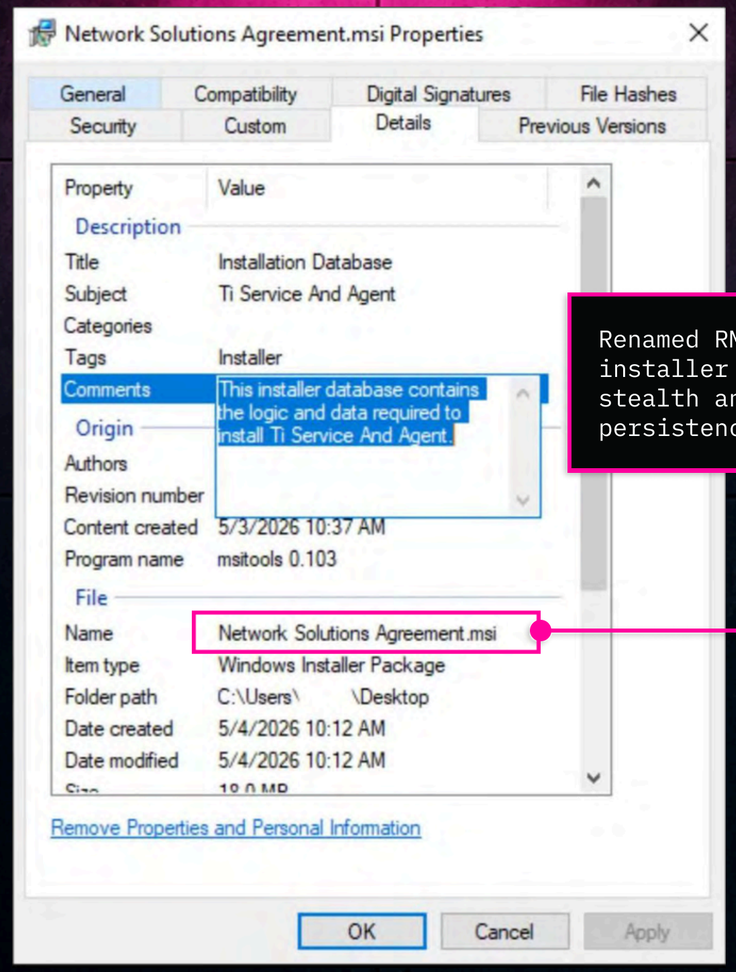
A legitimate RMM tool can give an attacker persistent access that looks like ordinary administrator behavior. They don't need their own tools when they can just borrow yours. Jamie Levy, Senior Director, Adversary Tactics, sums it up best:

“Why would you spend the cycles to develop or build from scratch when you can use a legitimate tool that you can just pull off the shelf?”

Ask your security team which remote monitoring tools they’ve approved, and what tips them off when an unapproved one shows up.

Close the gap with **Managed ESPM** to block rogue RMMs and **Managed EDR** to shut down RMM-driven attacks.

From the Wild



Renamed RMM installer for stealth and persistence

A **fake service agreement installed Tiflux**, then quietly stacked UltraVNC, Splashtop, and ScreenConnect on the same device. One phishing click gave the attacker multiple ways back in.

HUNTRESS

MAILBOX MANIPULATION

Once an attacker gets into an inbox, malware is optional. They can quietly alter mailbox settings to hide their activity and set up fraud. A hidden inbox rule can reroute a vendor's replies to a folder nobody checks, like RSS Feeds/Subscriptions or Archive. Then a swapped invoice redirects payment before the threat ever looks suspicious.

Why it Lands Here

Mailbox manipulation and persistence made up 19% of all identity-based threats Huntress tracked in 2025, and 24.6% of identity threat signals we've seen already in 2026. It skips malware and abuses trust in existing vendor relationships to commit business email compromise (BEC) fraud, gain access to steal sensitive information, or gain intel on a company.

24.6% of identity threat signals in 2026 point to mailbox manipulation

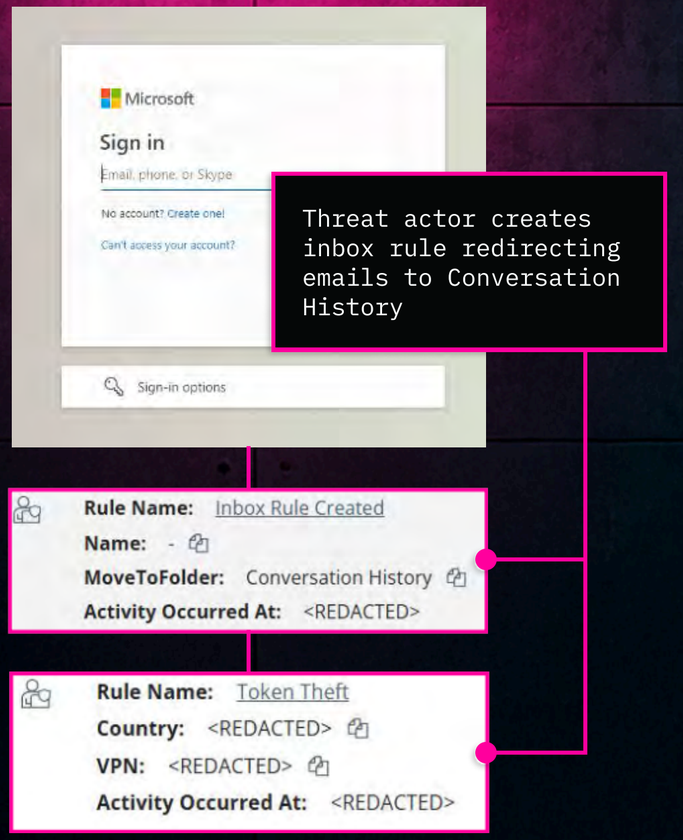
The Bottom Line

Mailbox manipulation is the quiet on-ramp to BEC fraud. It doesn't always trigger an alarm; sometimes the first sign is a vendor asking where the payment went.

Find out from IT which inbox rules are normal for your organization, and look twice at anything unfamiliar in a folder like RSS Feeds or Archive before writing it off as clutter.

Close the gap with Managed ISPM to harden mailbox rules and permissions, Managed SAT to train users to spot BEC tactics, and Managed ITDR to shut down active attacks.

From the Wild



After compromising credentials at a manufacturing organization, an attacker created an inbox rule in the Conversation History folder to reroute and hide suspicious email activity.

HUNTRESS

ACCOUNT TAKEOVER

AKA ADVERSARY-IN-THE-MIDDLE (AITM)

Account takeover comes in a lot of forms, and AiTM attacks are one of them. Here, the attacker slips in between a victim and a real Microsoft 365 login page. They grab the session token in transit and then forward everything on to Microsoft, so nothing looks off. As long as that session stays valid, the attacker can access the account, no password needed, no MFA prompt to stop them.

Why it Lands Here

AiTM attacks accounted for 18.9% of all identity-based threats Huntress tracked in 2025. This bypasses the MFA controls that most companies are told to rely on for identity protection, and drops the attacker straight into an authenticated account. The FBI has flagged AI-enabled phishing kits like Kali365, which help attackers run AiTM attacks faster at scale.

18.9% of all identity-based threats in 2025 were AiTM attacks

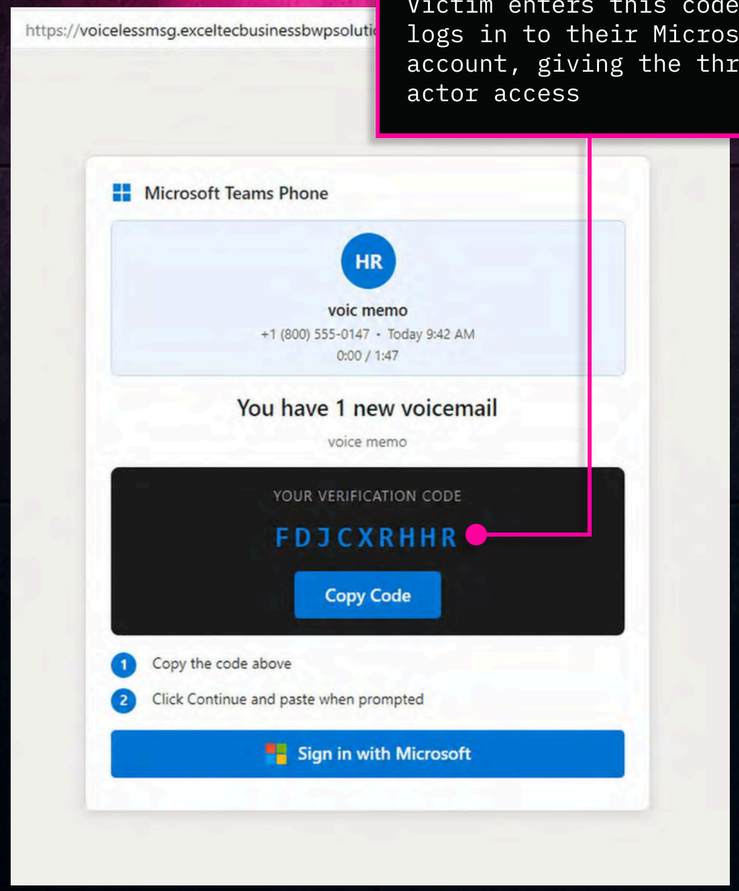
The Bottom Line

This tactic doesn't need your password or your MFA code. The stolen token works like a backstage pass, good for repeat visits until someone kills the session.

Ask your IT folks how long sessions stay active, and if they require a fresh login when someone signs in from a new device or location.

Close the gap with Managed ISPM to harden the controls attackers try to bypass, Managed SAT to train learners to spot AiTM tactics, and Managed ITDR to detect and respond to suspicious session persistence and hijacking.

From the Wild



A **phishing page** tricks a user into copying an authentication code and pasting it into Microsoft's real login screen, then completing the normal sign-in and MFA prompt. Once authenticated, the attacker silently steals the resulting session token, giving them full account access without ever needing the victim's password.

HUNTRESS

LOW-KEY DEADLY

LOW-KEY DEADLY

OH $#!T

OVERHYPED, FOR NOW

EVERYDAY PESTS

```
❯ [quadrant.status] prevalence: LOWER | pucker_factor: HIGH
                    signal: quiet doesn't mean harmless
recommendation: keep it on the list before it moves up the list
                    ---
```

HUNTRESS

DEVICE CODE PHISHING

A fake workflow prompt sends someone to Microsoft's real device code login. They enter the code, and the attacker gets a valid access token that can survive a password reset. Because the login screen is Microsoft's own device code page, there's no proxy, spoofed domain, or fake certificate for security tools to flag. The lure is often disguised as a routine workflow prompt, so it feels trustworthy.

Why it Lands Here

In early 2026, EvilTokens, a single Phishing-as-a-Service (PhaaS) kit, hit 344 organizations across five countries in just 16 days, using nothing but legitimate infrastructure and personalized lures to capture tokens at speed. Attackers used AI to vibe code the platform itself, along with AI tools to automate analysis of compromised inbox information with suggested follow-up targeting prompts.

1,380% increase in device code phishing attacks year over year

July - December 2025 vs. January - April 2026

The Bottom Line

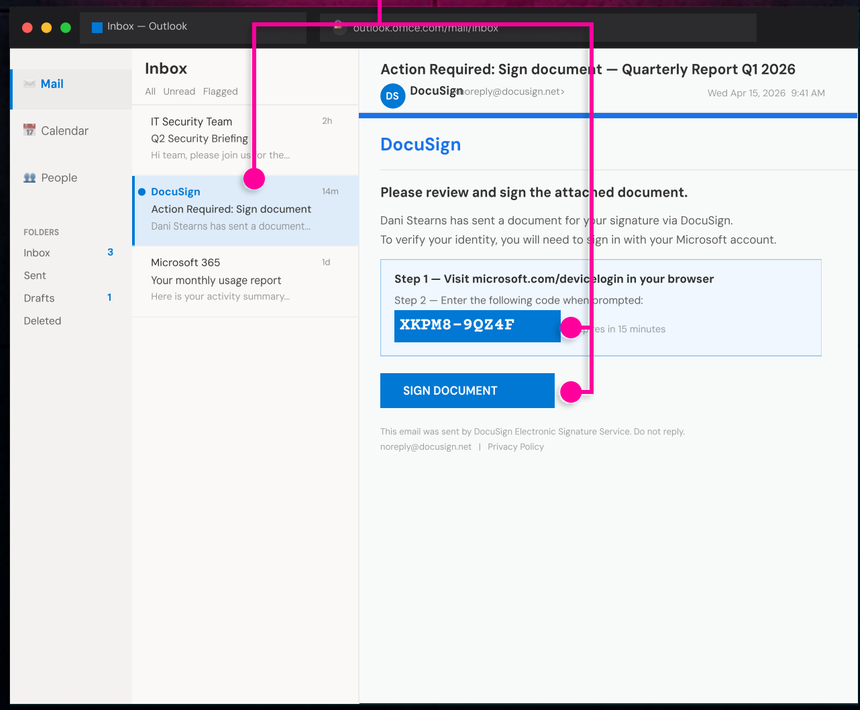
We've all entered a code to log into Netflix. But that same type of familiar prompt can hand over an entire Microsoft 365 account to an attacker, outliving password resets and bypassing MFA.

Check in with your IT team or security provider to block or restrict device code flow entirely if your team doesn't need it. If you see a login prompt show up out of nowhere, call or message the person or team who supposedly sent it before entering any code.

Close the gap with **Managed ITDR** to detect token abuse and revoke attacker access, **Managed ISPM** to harden the access controls against device code phishing exploits, and **Managed SAT** to catch and train users on these behaviors before they see it in the wild.

From the Wild

An EvilTokens device code phishing lure



Attackers quietly **hosted their token-harvesting infrastructure** on Railway, a legitimate developer platform most security teams never flagged, and wrapped phishing links in legitimate Cisco, Trend Micro, and Mimecast redirect URLs so the emails sailed through filters untouched.

HUNTRESS

EDR KILLERS

AKA BRING YOUR OWN VULNERABLE DRIVER (BYOVD)

Once attackers already have a foothold, some go beyond just killing the antivirus process. They land a signed-but-vulnerable driver for kernel-level access, then quietly weaken real-time protection, exclusions, and logging instead of shutting anything off outright. Some go a step further and unload the driver that feeds telemetry to a security tool, erasing the evidence at the source instead of fighting the tool's self-defense logic.

Why it Lands Here

This doesn't show up in most incidents, but when it does, it's the step that decides whether every other control you've invested in actually does its job. Kill the EDR first, and the ransomware, the credential theft, and the lateral movement that follow may happen without triggering alerts.

23.8% of driver abuse in 2025 traced back to one tool, Throttlestop/RWDriver

The Bottom Line

Once your endpoint protection is dead, everything downstream, encryption, exfiltration, and lateral movement runs dark, and nobody's getting an alert about any of it.

Get curious about how attackers try to defeat security tools. Start by checking with IT on whether your security software can block risky drivers, and whether a sudden crash would raise any flags.

Close the gap with **Managed EDR** to detect, contain, and remediate threats on every endpoint.

From the Wild

Threat actor registers a malicious driver as a legit disguised service

OemHwUpd

Key	Value
Detector(s)	<redacted>
Created At	<redacted> 12:37:12 UTC ( )
Updated At	<redacted> 12:37:12 UTC ( )
From Survey At	<redacted> 12:37:11 UTC ( )
Huntress Category	N/A
Persistence	Service (Kernel Driver) - OEM Hardware HAL Service
Binary	3111f4d7d4fac55103453c4c8adb742def007b96b7c8ed265347df97137fbee0
Virus Total	View Details
File Path	c:\programdata\oem\firmware\oemhwupd.sys
Name	OemHwUpd
Path	c:\programdata\oem\firmware\oemhwupd.sys
State	RUNNING
Status	OK
Command	\??\C:\ProgramData\OEM\Firmware\OemHwUpd.sys
Display Name	OEM Hardware HAL Service
Binary Mod Time	2026-01-22 00:06:09 EST
Binary Create Time	2026-01-22 00:06:09 EST
Event Summary	At <redacted> 12:37:12 UTC, a Windows Service with the name "OEM Hardware HAL Service" was detected on the endpoint.

A threat actor walked in through stolen SonicWall VPN credentials, then loaded a 15-year-old, revoked EnCase forensic driver to **kill 59 different security processes** from kernel mode.

HUNTRESS

VULNERABILITY EXPLOITATION

AKA 0-DAYS, N-DAYS

Attackers exploit and weaponize a known flaw, often within days of it going public or sometimes even before it goes public, to break into internet-facing systems. On-prem Exchange, WSUS, and file-transfer appliances are frequent targets, and a successful exploit usually opens up a foothold, like a webshell or remote code execution, with follow-on access.

Why it Lands Here

Most incidents still start with a stolen password or a phishing click. But when a serious vulnerability drops in a widely used platform, it hands attackers a ripe opportunity for access. Speed often decides the outcome here, not stealth.

“The remediation window that defenders rely on, the gap between ‘patch available’ and ‘actively exploited at scale,’ is already measured in hours and days for the most critical vulnerabilities, and that gap will continue to narrow.”

Muhammad Yahya Patel, Cybersecurity Advisor

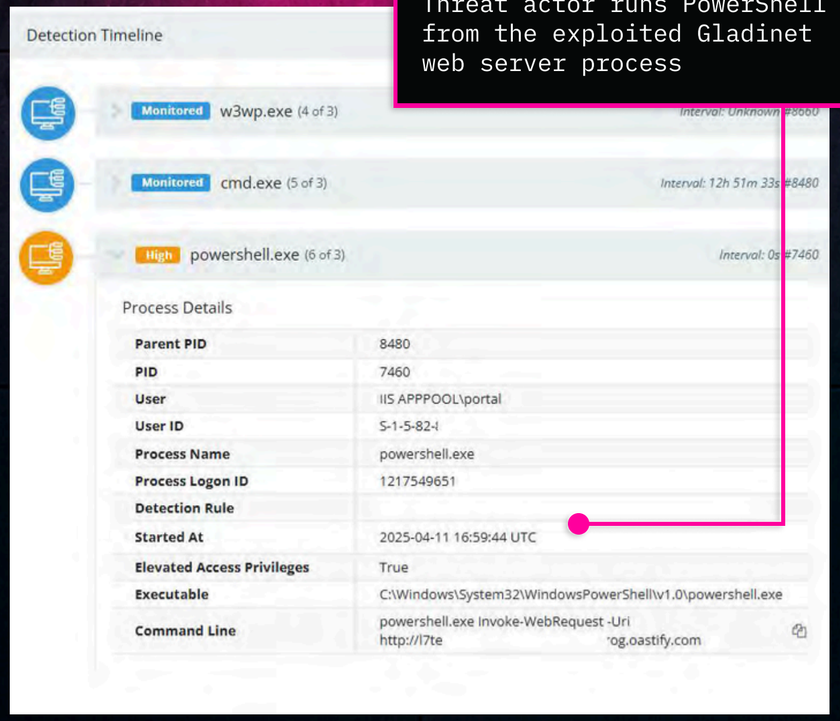
The Bottom Line

When a new vulnerability drops in a widely deployed platform, attackers move first, while security teams scramble to apply the latest patches.

Look into what your org’s patching process looks like, and how fast a critical fix moves from alert to actually being applied.

Close the gap with **Managed EDR** and **Managed ITDR** to detect and respond to active attacks.

From the Wild



Attackers exploited a **Wing FTP** Server flaw just one day after it went public, running commands and planting backdoor accounts before most teams had even patched. The same pattern played out with the **WSUS remote code execution flaw** and the Gladinet CentreStack and Triofox vulnerability.

HUNTRESS


❯ [quadrant.status] prevalence: HIGH | pucker_factor: MODERATE
signal: attackers use what works, often and efficiently
recommendation: treat the nuisance like it knows where the spare key is
---
EVERYDAY PESTS
LOW-KEY DEADLY
OH $#!T
OVERHYPED, FOR NOW
EVERYDAY PESTS
HUNTRESS

EXPOSED EDGE DEVICES

AKA VPNS & FIREWALLS

Edge devices like VPNs, firewalls, and RDP get exposed to the internet in ways that make them vulnerable to intrusions. Attackers can scan the internet for those gaps at scale, then use tactics like brute forcing, credential stuffing, password spraying, or CVE exploits to break into organizations running misconfigured setups.

Why it Lands Here

Roughly 70% of the active intrusions caught by our SOC start with adversaries authenticating through VPN access.

“Anything you expose to the internet will get hammered. RDP sits at the top of that list.”

Dray Agha, Senior Manager, Tactical Response

The Bottom Line

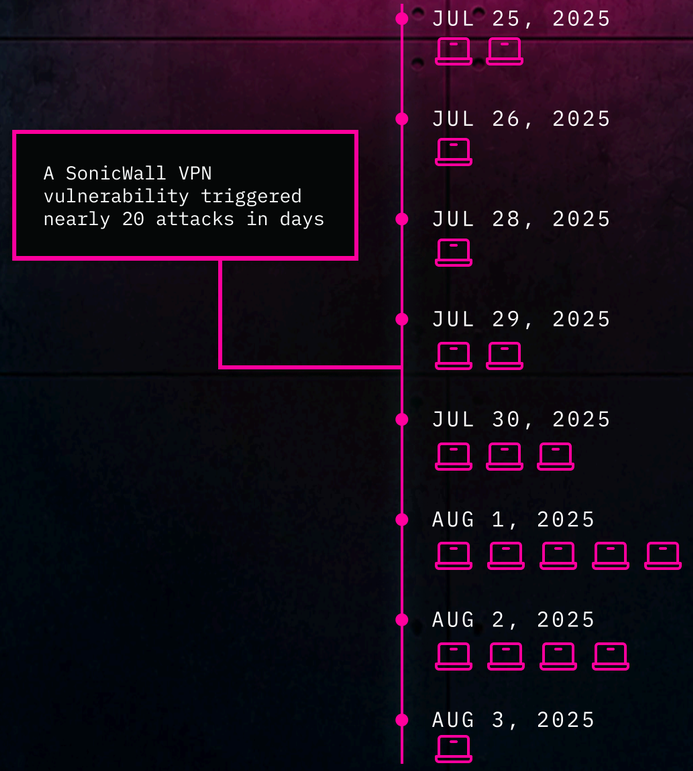
A stolen password with no MFA behind it turns VPNs or exposed RDP into a welcome mat.

Bring this question to whoever manages your company's network: Can someone log in remotely with just a password, or does it take a second step too?

Close the gap with **Managed SIEM** to spot suspicious activity early and **Managed EDR** to help contain threats before they spread.

From the Wild

SEPARATE COMPROMISED ACCOUNTS



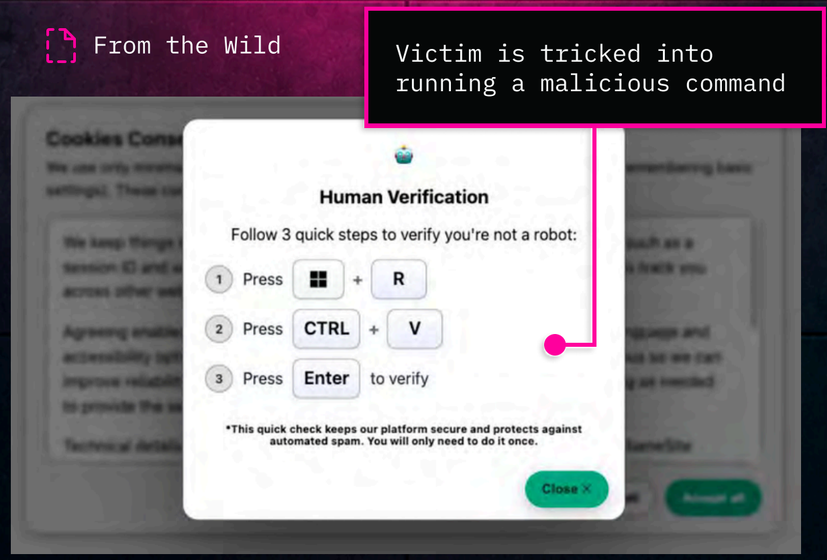
Attackers exploited a **SonicWall SSLVPN** flaw to gain administrative access, move laterally, steal credentials, disable security tools, and deploy Akira ransomware.

HUNTRESS

CLICKFIX

A fake CAPTCHA or a "verification failed" prompt walks someone through copying a command straight into the Windows Run box or a Mac Terminal, and a Control+V paste shortcut becomes the compromise lever. There's no exploit and no vulnerability, just a convincing enough moment of friction that people paste first and think later.

From the Wild



"Human Verification" prompted a victim to press Win+R, then paste one command into the Windows Run box. That **launched a multi-stage infection** ending in a LummaC2 infostealer compromise.

Why it Lands Here

ClickFix makes up 2.2% of EDR detection signals so far in 2026, but nearly 99% of those signals were high severity. A single paste opens the door to infostealers, RATs, and ransomware, though the real damage still depends on whatever gets dropped through it. AI is already speeding up how fast these fake CAPTCHA and verification prompts get customized, making them harder to spot as fake.

53.2% of all malware loader activity stemmed from ClickFix in 2025

The Bottom Line

ClickFix makes the victim do the attacker's dirty work with a shady copy-paste prompt.

Get IT to lock down the Windows Run box across the company. No Run box, no way to paste in the attacker's command.

"If I did not see [ClickFix] in a given day, I would be incredibly surprised and think something was wrong."

Nick Roddy, Security Operations Analyst

Close the gap with **Managed EDR** to disrupt active ClickFix attacks and **Managed SAT** to train users to spot the signs of an attack.

HUNTRESS

MALICIOUS ATTACHMENTS

A PDF feels safer than a link, especially when it fits the workflow. That’s the bet attackers make. They hide malicious links or payloads in familiar contracts, invoices, and signature requests, counting on the reflex that documents are harmless.

Why it Lands Here

57.7% of all phishing attacks in 2025 used malicious PDF attachments, and "e-signature request" was the single most common lure theme at 14.2%. Attackers only need one click to turn routine work into an open door.

57.7% of all phishing attacks in 2025 used malicious PDF attachments

The Bottom Line

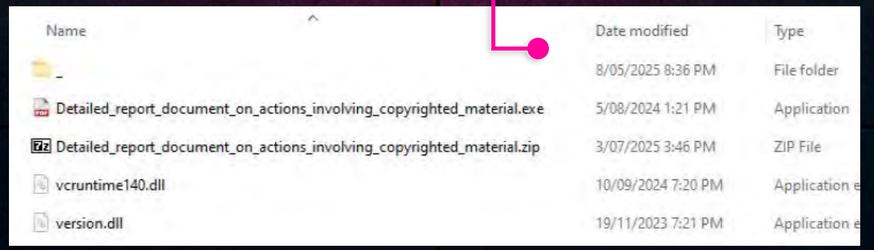
A malicious PDF blends in because it looks like the work everyone expects in their inbox.

You don't have to work in security to shut this down. If a signature or invoice shows up out of nowhere, confirm it through a second channel before you click and report it to your security team.

Close the gap with **Managed SAT** to test and train users in a controlled environment through phishing simulations, and **Managed EDR** to detect, disrupt, contain, and remediate malicious activity that makes it past the inbox.

From the Wild

Threat actor plants a malicious file next to a legitimate PDF reader app



A phishing email disguised as a copyright infringement notice **carried a ZIP attachment**. Inside, a real PDF reader sat alongside a malicious file that quietly launched a remote-access trojan on the victim’s computer.

HUNTRESS



OVERHYPED, FOR NOW

AI PLATFORM ABUSE

Attackers are abusing the trust people place in real AI platforms. They host, promote, or disguise malicious content behind a familiar AI tool or workflow. When the lure lives on a real platform and fits a familiar task, suspicion drops before the attack even starts.

Why it Lands Here

This is still an emerging pattern. Over the past nine months, hackers have weaponized shareable AI content, public mini-apps, and sponsored search placement across Claude, ChatGPT, and Grok. By putting malicious content on familiar domains, they borrow trust before anyone has a reason to be suspicious.

“AI security isn't limited to the models themselves. Attackers use legitimate AI platforms, shared conversations, public artifacts, and search results to make malware and malicious instructions look trustworthy.”

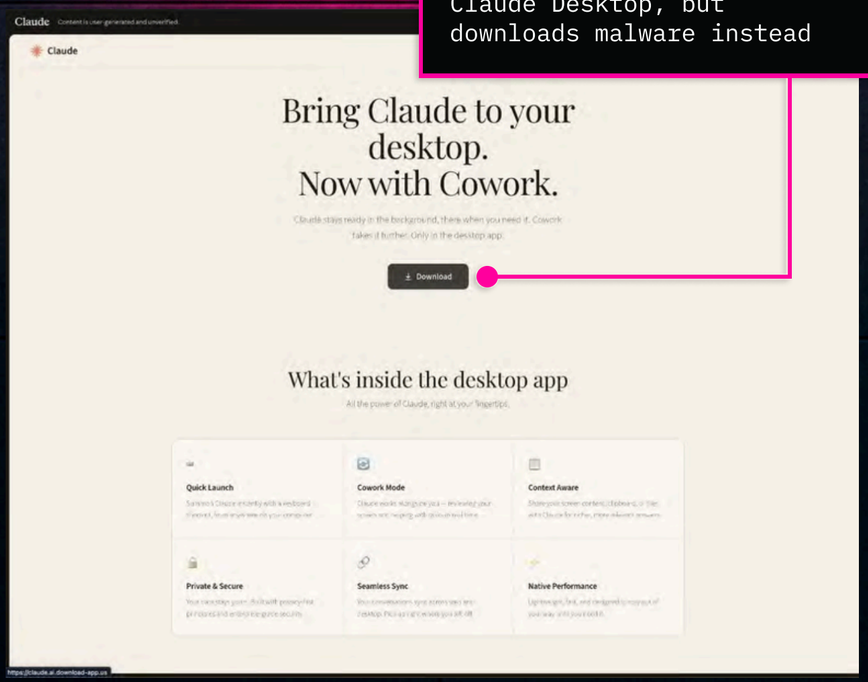
Lindsey O'Donnell-Welch, Principal Technical Community Engagement Writer, Huntress

The Bottom Line

A real AI platform can make malicious content feel safe. That trust is the attack.

Get AI tools directly from the vendor's official site or app store, never from a shared link, search ad, or "artifact" promising a desktop version. If an install prompt shows up after clicking a link inside an AI platform itself, that's the moment to stop and verify, not download.

From the Wild



FakeAgent used a malicious Claude Artifact hosted on the real claude.ai domain to send people looking for Claude Desktop to SectopRAT, hitting 29 organizations in two days.

Close the gap with **Managed SAT** to educate users on how to use AI tools safely and protect corporate data and **Managed EDR** to catch what happens after a suspicious download.

HUNTRESS

DEEPFAKES & VISHING

Deepfakes and voice cloning turn the thing employees trust most, a familiar face or voice on a call, into the attack surface. A fabricated video call from "the CEO" requesting an urgent wire, or a cloned voice confirming a fraudulent request over the phone, sidesteps every technical control because it never touches the network.

Why it Lands Here

This isn't common yet, and today it's usually a setup for the BEC or fraud attempt that follows rather than the payout itself. But AI has made it easier to turn a voice recording into a convincing voice clone.

“When you get a phone call asking you to provide sensitive information or perform an action, it’s more important than ever to slow down and take a few seconds to think: is this legit? No matter how important the call appears to be, there’s no harm in hanging up and calling back a verified number to confirm.”

Truman Kain, Principal Product Researcher

The Bottom Line

This turns a trusted face or voice into a weapon against your employees.

Agree on a callback number or a code word ahead of time, and use it to confirm anything urgent, like a wire transfer or password reset, before acting on it.

Close the gap with **Managed SAT** to prepare users to spot impersonation and **Managed ITDR** to detect suspicious account activity when deepfake or voice impersonation leads to identity compromise.

From the Wild

Threat actor fails a deepfake video authenticity test



Someone tried a **deepfake video scam** on Jim Browning during a Zoom call. Jim knew the test: hold up three fingers over his face. The caller refused and disconnected because he knew he wouldn’t pass.

HUNTRESS

YOUR NEXT MOVE

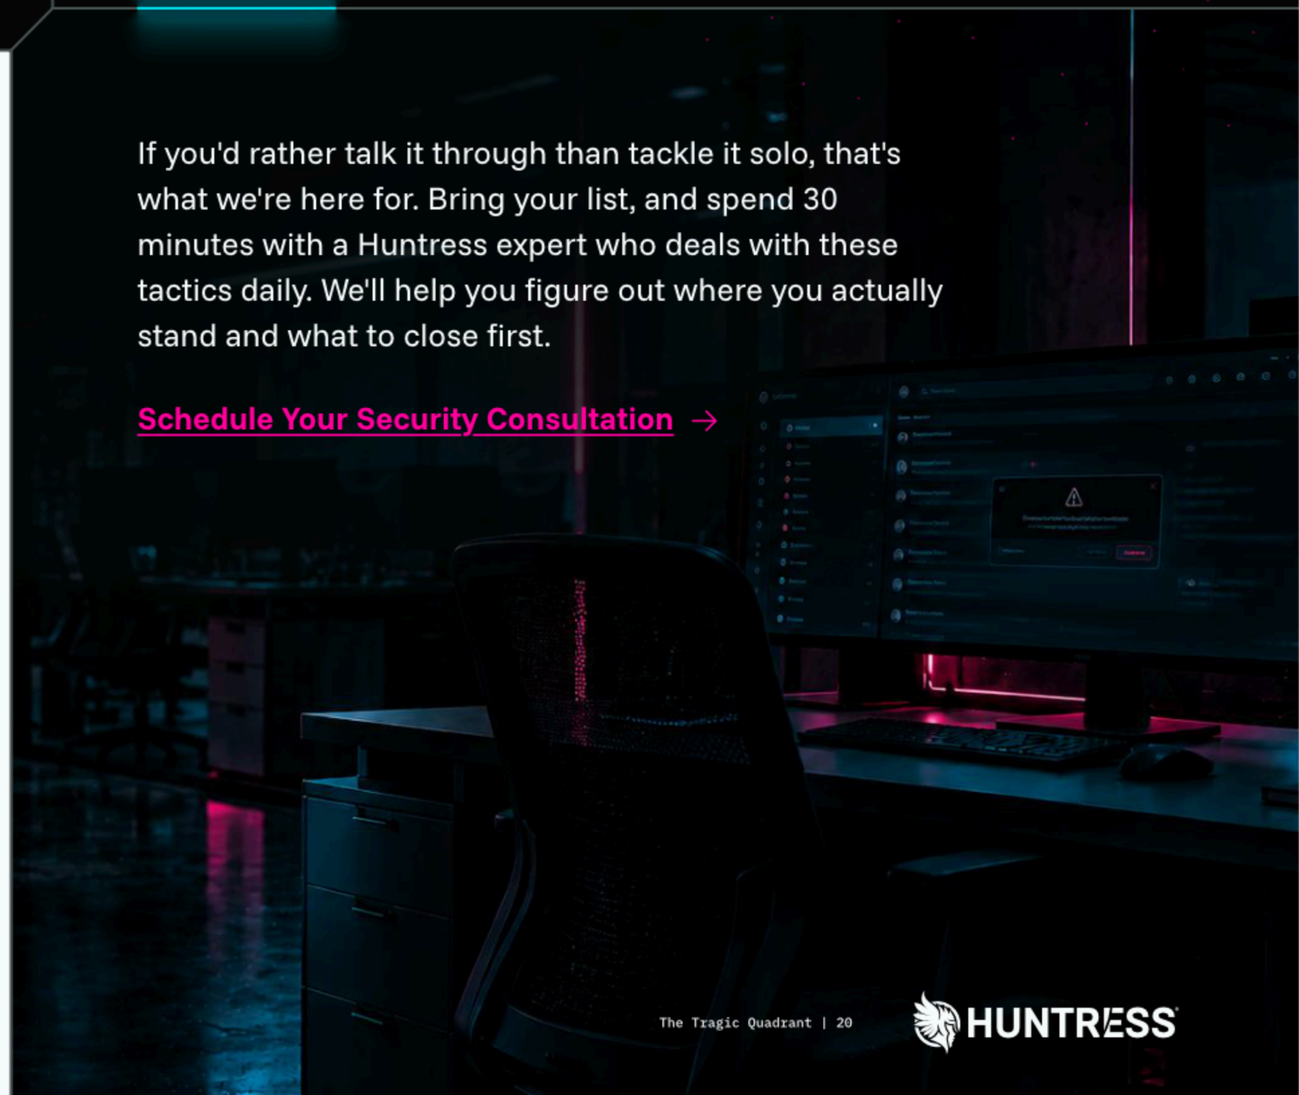
Every tactic on this quadrant comes with a defensive counterpoint and a clear next step to act on. Pick one tactic that sits high on both axes and close that gap first.

If you're unsure where to start, start in the top right, because those are the tactics that show up most often and hit hardest. For a lot of teams, that's RMM abuse. It looked like a niche concern a few years ago, and it's now driving nearly half of the endpoint-related incidents we investigate. Mailbox manipulation and account takeover sit right next to it for the same reason: high prevalence, high pucker factor, and a short hop to real damage. If you only close one gap this quarter, close one of these three.

While the other tactics may not be as pressing, they still deserve a plan. Know your baseline and what's normal for your business. Treat anything unexpected as something to verify through a second channel before you act on it. Lock down the basics that attackers are counting on to stay open. And most importantly, use this quadrant to get proactive, whether that's running a tabletop exercise with your team or bringing your questions to a security provider who can help you close the gaps that worry you most.

If you'd rather talk it through than tackle it solo, that's what we're here for. Bring your list, and spend 30 minutes with a Huntress expert who deals with these tactics daily. We'll help you figure out where you actually stand and what to close first.

Schedule Your Security Consultation →



HUNTRESS

ABOUT HUNTRESS

Huntress is a global cybersecurity company on a mission to make enterprise-grade products accessible to all businesses. Purpose-built from the ground up, the Huntress Agentic Security Platform unleashes end-to-end protection in the age of AI. From Endpoint Detection and Response (EDR) and Identity Threat Detection and Response (ITDR) to Security Information and Event Management (SIEM), Security Awareness Training (SAT), and Security Posture Management, the Huntress Agentic Security Platform is targeted protection for endpoints, identities, data, and employees, delivering trusted outcomes and valuable peace of mind.

Its 24/7 AI-centric Security Operations Center (SOC) is powered by a team of world-renowned engineers, researchers, and security analysts dedicated to stopping cyber threats before they cause harm. Huntress is often the first to respond to major hacks and incidents, with its expert security team sharing real-time tradecraft analysis and actionable advisories with the community.

Currently safeguarding more than 5 million endpoints and 15 million identities, Huntress empowers internal security and IT teams and Managed Service Providers (MSPs) worldwide to protect their businesses with enterprise-grade, accessible security products.



Contact us to learn more.

Jedi Services

hello@jedi-services.com

www.jedi-services.com/

JEDI SERVICES

In partnership with HUNTRESS