CRXfiltrate — Malicious Browser Extension Campaign

CLASSIFICATION **TLP:AMBER**
PREPARED BY 7AI Threat Research Team
DATE 2026-05-12

ABOUT THE NAME

CRXfiltrate is 7AI's name for this campaign. The underlying extension cluster was first publicly documented by Wladimir Palant in January 2025 as **Phoenix Invicta**; that name is preserved where it ties to his original research. Customers who received the V3 threat brief under the Phoenix Invicta name are looking at the same campaign — this is the V4 update under the public name. The public research is available at `https://7ai.com/crxfiltrate`.

TL;DR: A professionally operated campaign uses browser extensions disguised as productivity tools to silently execute operator-controlled JavaScript inside the browser, harvest identity data, and surveil browsing activity. We document 85,000+ installs across 22 cataloged cluster extensions, with a parallel delivery path that reaches victims through legitimate websites with no extension installed.

Why This Is Urgent

- **Extensions are still live on browser stores.** At paper validation, at least three cluster extensions remain installable: `Easy Dark Mode` on the Chrome Web Store (unlisted but live), plus two Microsoft Edge Add-ons listings (`1-Click Color Picker: Instant Eyedropper` and `AdBlock for Youtube: SkipAds`). New infections can occur today.
- **The campaign is actively maintained.** Reverse engineering of the production payload reveals a Webpack build with versioned release directories captured live across multiple iterations, a feature-flag system with server-side toggles, and source-code comments referencing internal project-tracker ticket numbers above 390. This is software product engineering, not a fire-and-forget malware drop.
- **Removed extensions keep running.** Extensions taken down from browser stores are not uninstalled from endpoints where they are already present. We have observed infections persisting 16+ months since public disclosure — only enterprise policy enforcement (Group Policy, Intune, Chrome Browser Cloud Management) reliably removes them from existing installations.
- **An extension audit alone is not enough.** The same operator infrastructure is also reached through third-party tracker chains and ad-chain sub-resources on legitimate websites — with no extension installed on the host. Endpoints with a clean extension inventory can still be exfiltrating data through the page-served path. DNS and proxy-log hunting is non-optional.

How to Check Your Exposure

1. Run the IOC sweep first. Use DNS logs, proxy logs, and network connection logs against the IOC block list. The highest-fidelity detection pattern is any request to `cdn23602612.ahacdn.me` containing the path `/500b-bench.jpg` — an operator-specific connectivity probe that is the highest-signal CDN-path indicator in the paper. This step finds both the extension-served and the page-served exposure paths. **Do not skip this step even if your extension inventory is clean.**

2. Search for malicious extension IDs on managed endpoints. The attached IOC block list contains 20 confirmed Chrome cluster extension IDs and 2 Edge extension IDs (plus 1 lower-confidence Chrome ledger-only entry). Check your endpoint fleet (via Intune, Chrome Browser Cloud Management, or EDR) for any installation of these IDs.

3. Check browser localStorage on suspect endpoints. The presence of `localStorage.gotAriaLabel` on any domain confirms the identity-harvester variant executed and the signed-in Google user's name and email were exfiltrated. The presence of `localStorage.zLastRedHer` timestamps the JavaScript execution backdoor's last fire. The identity-harvester capability is variant-gated — its absence does not rule out other variant infections.

Recommended Response Actions

Immediate (within 24 hours)

- Run the IOC sweep against DNS, proxy, and network logs to surface both extension-served and page-served exposure
- Force-remove all listed extension IDs from managed endpoints via Group Policy, Intune, or Chrome Browser Cloud Management
- Block all domains and IPs from the attached IOC block list at the proxy and firewall — observing the "DO NOT block" exceptions for shared CDN and shared-hosting IPs
- For users on confirmed compromised endpoints, reset credentials for SSO, email, and financial platforms

Near-term (within 1 week)

- Audit all installed browser extensions across the endpoint fleet against the extension ID list
- Deploy extension allowlisting to prevent installation of unapproved extensions
- Review browser sessions on compromised endpoints for unauthorized access to sensitive applications
- For endpoints showing DNS hits without matching extension presence, treat as page-served exposure and remediate at the network tier rather than the endpoint

Long-term

- Enforce browser extension governance — block all extensions not on an explicit allowlist, with particular scrutiny for any extension requesting `<all_urls>` host permissions or `nativeMessaging`
- Monitor for campaign expansion — new extensions and infrastructure are actively being developed; several reserved domains return HTTP 403 indicating staged but not-yet-active operations

• Treat browser-context attacks as a distinct detection problem; the standard EDR + SSL proxy + DNS stack is not architected to detect malicious activity originating from inside the browser's JavaScript execution environment

About This Threat

The campaign operates beneath the visibility of most endpoint and network security tools. Traffic to attacker infrastructure is consistently miscategorized as benign by web proxies, and payloads are uniquely customized per victim via server-side template substitution to defeat signature-based detection. Malicious behavior is activated server-side after extensions pass store review, meaning they appear clean during the review process and only begin malicious activity once installed at scale.

Documented scope

• **22 cataloged cluster extensions** (20 Chrome + 2 Edge; plus 1 lower-confidence Chrome ledger-only entry)

• **85,000+ documented installs** across the catalog (conservative floor; actual count exceeds this when "not listed" Edge entries and unverified secondary listings are included)

• **60+ active and reserved cluster domains** mapped across HZ-Hosting Bulgaria infrastructure (`5.149.x` range) plus secondary Gen 2 hosting

Capabilities present in source

• **JavaScript execution backdoor.** Operator-controlled JavaScript is delivered by the C2 server and executed inside the visited page's own JavaScript realm. The execution channel is not architecturally limited to ad fraud — the same channel can deliver credential harvesting, session token theft, or targeted attacks against authenticated banking, SSO, or admin console pages. The current shipped payload monetizes through ad injection and SERP hijacking, but the payload is the variable; the architecture is the constant.

• **Identity harvesting (variant-gated).** One captured payload variant scrapes the signed-in Google user's real name and email address from the Chrome sign-out element and exfiltrates the data to `doublestat.info`. We also observed a variant where this block was removed, so treat identity harvesting as variant-gated rather than universally present. When served, the captured identity links to detailed browsing activity indexed by a persistent install UUID.

• **Browsing surveillance via URL parameter exfiltration.** Every page the infected user loads triggers a C2 round-trip where the request URL itself carries the visited domain, complete page title, install UUID, and extension ID as query parameters. The response body returns the payload. Both happen in a single HTTP transaction — proxies that block based on response evaluation have already let the request URL through. Blocking the response prevents monetization; it does not prevent surveillance. DNS-tier or pre-connection blocking is required to stop exfiltration.

• **Demographic-targeting architecture (currently disabled).** The deployed JavaScript contains a function `generateChanelTargeting()` that builds a nine-bucket demographic code combining gender with an age band, including a bucket spanning the COPPA-protected under-13 range. The function is short-circuited by a `return;` statement at the top, but it is not deleted — activation requires only a server-side configuration change with no

extension update or store re-review. We do not have evidence this targeting is currently active in production, but the architecture exists.

• **Content-Security-Policy stripping.** Extensions use `declarativeNetRequest` rules to strip Content-Security-Policy and X-Frame-Options headers from every HTTP response. CSP is the browser's primary defense against injected JavaScript; with it stripped, anything the extension chooses to inject runs without restriction.

• **Manifest V3 prohibition bypass.** The extension's content script fetches operator-controlled JavaScript from the cluster's server and injects it into the page by calling `document.createElement('script')` and setting the response body as the new element's text. No HTML parsing, no remote script source, no bootstrapper file — the fetched text is treated as data until the moment the injected element is appended to the DOM, at which point the page executes it. The `declarativeNetRequest` CSP-stripping rule documented above is what makes this injection succeed on pages that would otherwise refuse inline script. Because no remotely-hosted script source is referenced, the technique falls outside the literal scope of MV3's remote-code prohibition.

Why default defenses miss this

The malicious activity occurs inside the browser's JavaScript execution environment. No file is written to disk by the payload. No process is spawned outside the browser. No registry key is created. The network activity originates from a normal Chrome or Edge process making normal-looking HTTPS requests to recently registered or low-reputation domains. EDR platforms are architected to detect process behavior, file changes, and network anomalies at the host level — browser-internal injection using the browser's own network stack is invisible to that telemetry model. Closing this detection gap requires extension-aware controls, infrastructure-level visibility into browser-originated network calls, or browser-resident security telemetry — none of which are standard equipment in most enterprise environments today.

This is an evolving situation and an ongoing investigation. The scope of the threat actor's activity, related indicators, and possible presence in your environment are based on information available to 7AI at the time of this report and may change.

Attached: IOC Block List (v4, 2026-05-12) — domains, IPs, extension IDs, and detection patterns.

For the latest version, contact your account team. The public research is available at `https://7ai.com/crxfiltrate`.