

Malvertising delivered through an authenticated Carnival Cruise Line email via a lapsed brand domain (cclpromos.com)

Author: [Daniel Jones](#), Tuxxin LLC (tuxxin.com), contact@tuxxin.com Findings documented 2026-06-25. Resolution verified 2026-08-27. Published 2026-09-10.

When bad actors take over a known-good domain

In June 2026 a genuine Carnival Cruise Line casino email led its recipients to malware. The message was a real "Players Club" booking confirmation. It passed every authentication check and came from Carnival's own infrastructure.

One promotional link was the problem. It ran through Carnival's own Salesforce click-tracker and landed on cclpromos.com. Carnival used that domain for its Players Club promotion until about 2022, then let it lapse. A third party re-registered it on 2024-11-22.

The domain now runs a device-aware traffic distribution system (TDS). It shows automated scanners a blank, benign page and routes real people, by device, into malware and scams. Because the carrier email passed SPF, DKIM, and DMARC, the attack inherited full Carnival brand trust.

This report walks the delivery path, the cloaking that kept it hidden, every payload variant observed (several new to public threat intelligence), and the indicators of compromise. It closes with how the vector was disclosed and fixed. After coordinated disclosure, Carnival re-acquired the domain on 2026-08-26, and the malicious routing was verified dead across all vantages on 2026-08-27 (sections 10 and 12).

1. Executive summary

- A live, authenticated Carnival casino booking-confirmation email links, through Carnival's legitimate click-tracker (click.carnivalcruiselineemail.com), to cclpromos.com, a domain Carnival no longer controls.
- cclpromos.com runs a FingerprintJS-based cloaker that shows automated scanners a benign skeleton while routing real human browsers, by device, into a rotating advertiser pool that delivers malware and scams.
- Observed payloads include Windows fake-security MSIX malware (the publicly documented PhantomJack family), a PhantomJack variant with a forced auto-update channel that I track as PseudoJack, multiple fake-antivirus scareware pages, two fullscreen tech-support browser-lockers that each push a scam phone number, a Firefox search-hijacker, an iOS App Store scam, and browser push-notification hijacking.
- Several variants were absent from prior public threat intelligence and are documented here for the first time (section 6).
- The domain stayed live throughout the observation window (2026-06-13 to 2026-06-24). The operator rotated downstream payload domains continuously while the cclpromos.com entry point persisted.

- The vector is resolved as of 2026-08-27. After escalation via certified notice and a hand-off to Carnival's security team, Carnival re-acquired cclpromos.com, removed it from the third-party network, and repointed it to a legitimate Carnival property. The malicious routing was independently verified dead across datacenter, residential, and mobile vantages (sections 10 and 12).

2. Impact and severity

The message carries valid SPF, DKIM, and DMARC. Those are the three checks that prove an email really came from the domain it claims and was not altered in transit. All three passing means the message inherits full Carnival brand trust and passes the recipient's normal scam check.

The link leads to a device-aware traffic distribution system, or TDS: a routing layer that fingerprints each visitor and serves a different result by device and network. Sections 5 and 7 show how it works here. It serves a platform-appropriate attack. Windows and iOS are confirmed, and networks of this kind commonly also carry Android and geo-targeted variants.

The rotation includes fullscreen tech-support browser-lockers that require no installation. They trap the browser with a fullscreen takeover, a captured mouse cursor, and looping "your PC is locked" audio, then push a scam phone number. For a non-technical recipient this is harder to escape than a malware installer and more likely to succeed.

The audience is high value. A casino and Players Club booking flow reaches recipients with active bookings.

The vector was live and reachable from the email throughout the observation window. Carnival re-acquired the domain on 2026-08-27, which remediated it (section 12).

3. The delivery vector: the carrier email

The carrier email is genuine, fully authenticated Carnival mail sent from Carnival's own Salesforce infrastructure.

Field	Value
From	"Carnival Cruise Line" <funships@carnivalcruiselineemail.com>
To	(recipient redacted)
Subject	"Game on Daniel, thanks for booking" (Players Club booking-confirmation)
Date	Sat, 13 Jun 2026 15:31:52 -0600
Booking reference	(redacted)
Sending infrastructure	Salesforce Marketing Cloud / ExactTarget (mta7.carnivalcruiselineemail.com, 136.147.186.150)
Attachments	None (multipart/alternative, text and HTML)

Authentication: SPF pass; DKIM pass (two signatures, @carnivalcruiselineemail.com and @s7.y.mc.salesforce.com); DMARC pass (the domain publishes p=REJECT). The promotional call-to-action in the message body is

a click.carnivalcruiselineemail.com tracking link, which is Carnival's legitimate Salesforce click-tracker, and its decoded destination resolves to cclpromos.com.

4. Domain background and history

cclpromos.com was Carnival's own "Carnival Players Club" casino-promotion site. The 2021 Wayback Machine snapshot shows the page title "Carnival Players Club," hundreds of "Carnival" references, links to www.carnival.com/casino-signup and help.carnival.com, and the same `/carnivalplayersclub/` path the malicious email links to today (source: web.archive.org/web/20210613081457/).

Site content ceased around August 2022. The current registration is dated 2024-11-22 (registrar SnapNames, per WHOIS). The domain is now parked with a domain-monetization provider, and its nameservers and MX point to a parking platform. The root cause is non-renewal of an in-use brand domain, which auto-renew plus registry-lock under an enterprise registrar would have prevented.

The exposure is limited to the lapsed link target. Carnival's email-sending domain (carnivalcruiselineemail.com), which DKIM-signs the mail and hosts the click-tracker, is properly secured: enterprise registrar (MarkMonitor), registrar locks in place, expiry 2027-10-01, and Salesforce/ExactTarget nameservers.

5. Attack chain, cloaking, and anti-analysis




Figure 1. The delivery path: an authenticated Carnival email, through the legitimate click-tracker, to the lapsed domain and its device-aware cloaker, which serves a benign page to scanners and platform-specific malware to real visitors.

5.1 De-cloaked redirect chain

A live capture from a residential US Windows browser (2026-06-15) recorded the complete path a real recipient traverses (defanged):

```
hxxps://click.carnivalcruiselineemail.com/?qs=... (Carnival's legitimate click-tracker; destination id 4908) 302
-> hxxp://www.cclpromos.com/carnivalplayersclub/?utm_source=bgmarketing&utm_medium=em&utm_campaign=... 200 (lapsed brand domain; FingerprintJS cloaker fires here)
-> hxxp://jubaaa[.]com/int?dat=... 200 (TDS feed)
-> hxxps://tratobid[.]com/xfc/sfclick?u=426470a5-... 302 (click broker / affiliate)
-> hxxps://cinga.ngapp[.]online/?clickid=...&cid=9961&dkw=cclpromos.com&pid=2496 200 (fake "security check" lander)
-> hxxps://bacon.ngapp[.]online/impression?ext_name=NetGuard&cid=9961 200 (impression beacon)
-> hxxps://der.ngapp[.]online/downloadproxy/.../?ext_name=NetGuard&cid=9961 302
-> hxxps://file.ngapp[.]online/prvcytls/NetGuard.Msidx (forced download; 147,462,271 bytes)
```

The broker, feed, and payload hops rotate per visit. An earlier capture routed cclpromos.com to ww38.cclpromos.com, a secondary monetization endpoint. cclpromos.com is the constant entry point. The

lander receives `dkw=cc1promos.com` , which confirms that monetization is keyed off the cclpromos referrer.




Exhibit 01. The fake "click here to install" download page a real recipient reaches.

5.2 Cloaking (evasion)

The landing page loads a FingerprintJS script (`/js/fingerprint/iife.min.js`) and routes on a fingerprint (`fp=`) parameter plus a per-visit `tr_uuid` . Automated scanners, curl, and headless tools fail the check and receive a benign skeleton, which is why reputation engines initially rate the domain clean. Only real human browsers reach a live payload. Public scanning (urlscan.io) flagged the page as one of "10,000+ similar pages," which points to a templated, mass-scale malicious-monetization network rather than a one-off.



Exhibit 06. The decoy page the cloaker serves to scanners: a header row of repeated "word" filler and a few random tokens such as randomWL100000k1, with the rest of the page blank. Reputation engines see this near-empty page and rate the domain clean.

5.3 Anti-analysis and anti-debugging (confirmed 2026-06-24)

Advertiser landers in this network ship a tamper and automation-detection layer that does three things:

- It hooks DOM APIs (`querySelector` , `querySelectorAll` , and others) and inspects the caller's stack trace to detect DevTools-console evaluation (it tests for the `@debugger eval code` stack frame) and injected or extension scripts (`<anonymous>`), verifies caller identity with function-hash checks, and flags such access as `pup-qs` or `pup-qla` .

- It triggers a runtime-constructed `debugger` trap with no literal `debugger;` token in the source, which defeats static scanning. It was observed live as a browser stalling on "Debugger paused."
- It loads an `adblockdetect.js` check and reports to the operator's own Sentry telemetry endpoint (`js.sentry-cdn.com`).

Together these freeze or blank the page whenever developer tools are open, which explains why the payloads are difficult to reproduce under analysis and why automated scanners only ever capture the benign skeleton.




Exhibit 13. A browser frozen on "Debugger paused", the lander anti-analysis layer detecting DevTools.

6. Payload variants

The cloaker serves different attacks by device and rotates downstream domains as they are taken down. These variants were observed between 2026-06-13 and 2026-06-24.

6.1 Windows fake-security MSIX malware (PhantomJack family)

Forced `.msix` downloads of about 148 MB each arrive as fake security and privacy apps: SafeWatch, NetGuard, LeakGuard, and PrivacyKeeper. These are the publicly documented PhantomJack browser-hijacker family, distributed through the PseudoTDS ad-tech-abuse campaign (PhantomJack and PseudoTDS are both Trinity Cyber's names). ReversingLabs classifies the SafeWatch sample as Win32.Trojan.PhantomJack, which is a vendor's independent read on the family rather than mine. SafeWatch.msix (SHA-256 D66895...923B3) is an exact match to a published PhantomJack indicator set (Trinity Cyber, "Blurred Lines: AdTech Abuse Delivers Browser Hijackers Through the Microsoft Store"). Delivery domains observed: ngapp[.]online and swatchapp[.]online.

Each package is signed with a certificate whose organization field is a bare GUID, issued by a Microsoft Marketplace CA and valid for exactly three days. SafeWatch's certificate (organization 0FE13B24-D232-46E4-B668-5D9597B36A78, issuer Microsoft Marketplace CA G 026) was valid 2025-04-01 to 2025-04-04. PrivacyKeeper's (organization 22C37D45-CD8D-464F-85CA-4517AE16D585, issuer Microsoft Marketplace CA G 027) was valid 2025-10-26 to 2025-10-29. Both had already expired before the packages were served in June 2026. These are the MSIX signing certificates, a separate artifact from the .appinstaller publisher CNs in section 6.2.




Exhibit 04. A forced SafeWatch.msix download (recipient details redacted).




Exhibit 05. Three .msix payloads in the browser download history.

Together they are a standing channel to push an arbitrary package at any time. The channel is present in the manifests; I did not observe it delivering a second-stage payload.

Trinity Cyber's report notes only that these apps are MSIX files, so the forced auto-update configuration is documented here for the first time. All four sample hashes were unknown to MalwareBazaar and ThreatFox before I filed them on 2026-09-10. SafeWatch is 147,724,942 bytes and PrivacyKeeper is 148,652,394 bytes, which a browser rounds to 141 MB in its download bar. That size is consistent with padding to exceed antivirus and sandbox size limits (I did not test any scanner's ceiling).

For context, Microsoft disabled the ms-appinstaller URI handler by default in December 2023 after abuse by Storm-1113 and others, reporting that multiple cybercriminals were selling a malware kit as a service that abuses the MSIX file format and the ms-appinstaller protocol handler. These samples arrive as direct .Msix downloads rather than through an ms-appinstaller URI, which is consistent with that handler being off by default.

6.3 Fake-antivirus scareware

Several fake-AV landers rotate through the chain:

- A fake McAfee "Safety warning / Run Quick Scan" via hackshielded[.]com.
- A "Protect Your PC from Viruses" lander via gum-promos[.]club.
- A fake McAfee "Run Quick Scan" lander via ghostwallguard[.]com, captured 2026-06-24 and absent from prior public sets.




Exhibit 03. A fake "Windows Security Alert" scareware page (recipient details redacted).




Exhibit 10. A fake McAfee "Run Quick Scan" scareware lander via hackshielded[.]com (recipient details redacted).

6.4 Tech-support browser-lockers (two, both newly documented)

- The first, captured 2026-06-23, cloned a Microsoft Support and "Windows Defender Security Center" page and locked the browser: fullscreen, a hidden mouse cursor (Pointer Lock API), looping beeping audio with a synthesized "your PC is locked" voice, and a fake scrolling "hacker" terminal. The lure was the scam number +1 (866) 315-0945 ("Microsoft Agent / Windows Tech Support"). No file downloads; the goal is to panic the victim into calling.
- The second, captured 2026-06-24, was an Azure-Front-Door-hosted fake "Windows Defender SmartScreen" scam pushing the scam number +1-833-306-1211, observed during residential testing.




Exhibit 12. A fullscreen fake "Windows Defender" tech-support browser-locker; scam number +1 (866) 315-0945.

6.5 Firefox search-hijacker

A fake "Firefox add-on" search hijacker ("modifies your default search provider, powered by Bing") arrives via ff.newtabsearch[.]net (cid=10972). It is a distinct advertiser from the Windows payload, which demonstrates

live per-visit rotation.




Exhibit 09. A fake Firefox add-on search-hijacker via ff.newtabsearch[.]net.

6.6 iOS App Store scam

On iPhone and iOS the chain serves a JavaScript "Your iPhone has been hacked" alert, then a fake "Apple security" page, then App Store fleeceware, via fortiproapp[.]com.




Exhibit 07. An iOS "Your iPhone has been hacked" alert.




Exhibit 08. A fake "Apple security" page leading to an App Store scam.

6.7 Browser push-notification hijack

A fake "Click Allow / I'm not a robot" prompt drives push-notification subscriptions (matrixgrowthforge[.]com, pushtorm[.]net).




Exhibit 02. A "Did You Mean / I'm Human" PUP lure.

Three distinct malicious advertisers (MSIX malware, a Firefox add-on, and a fake McAfee lander) were captured off this one domain inside 48 hours, and the slot is continuously refilled. Per-advertiser takedowns therefore cannot end the campaign (section 7).

7. The monetization and parking layer

Per-advertiser takedowns did not end the activity, and the captures show why. The entry domain is parked with a large domain-monetization ("direct search") provider. Across every captured payload, two hosts stayed constant beneath the rotating advertisers: a device-fingerprinting cloaker (brandsmat[.]com, served as sd559908.js.brandsmat[.]com) and a broker or redirect (jubaaa[.]com). Registry RDAP, DNS, and ASN resolution place both constant hosts on the provider's own registrar and network. The same sd559908.js.brandsmat[.]com cloaker was observed on at least eight other unrelated parked domains that resolve to the same provider, confirmed via urlscan.io and by direct capture, where those domains funneled to the same rotation of search-hijacker, fake-McAfee, and tech-support-scam landers.

In observed behavior, removing one advertiser did not stop the chain: the next visit routed to a different advertiser while the same constant hosts persisted. Removing the entry domain from the parked-domain network ended it here; Carnival did this on 2026-08-26. A separate, more detailed technical analysis of this infrastructure is in preparation.

Current status, checked 2026-09-10: the Carnival entry point is resolved, but the payload network it fed into is still live and has rotated since the earlier public reporting. get.privacykeepersite[.]com and securi-guard-browser[.]com return live phishing verdicts, and cablegaurdian[.]online (the operator's own misspelling of "guardian"), guardianrole[.]online, euob.northwavepoint[.]com, tratobid[.]com, and kineticharbor[.]com are cloaking. Hosts named in the earlier reporting have gone quiet: results.streamio[.]site is clean and

cint[.]browsingit[.]online is unreachable. The continuity shows in the naming, where the earlier cint[.]browsingit[.]online and the current cint2.scrtrgd[.]online and cinga.ngapp[.]online share the same cint-prefixed subdomain grammar on different apex domains. Of 12 sampled redirect-hop domains, 11 resolved to dead parked shells: six-letter lowercase .com names, bulk-registered in the same month on a single parking IP.

8. Scope: confirmed, clean, and unverifiable

Testing included server-side resolution of the recipient's full 16-year email archive at a fixed gentle rate, and the resolver's detection was validated against the known-bad 2026-06-13 link.

- The 2026-06-13 booking-confirmation "Players Club" casino email is confirmed to route live to cclpromos.com.
- The broad marketing and casino mailstream is clean. Across 541 deduplicated casino tracker links (2018 to 2026) and 241 post-takeover emails, none routed to cclpromos.com, and all currently-live archived links resolve to legitimate www.carnival.com.
- Older booking emails from 2024 to 2025, whose tracker tokens have since expired, are unverifiable. Those tokens now resolve to redirect_error.html, so their original destinations are unrecoverable and cclpromos routing for them can be neither confirmed nor ruled out.

Exposure window: 2024-11-22, when the domain left the brand's control, to present, with the live malware cloaker confirmed no later than 2026-03-12. That is the earliest archived snapshot containing it; the prior 2025-07-23 snapshot was still benign.

9. Relationship to the April 2026 Carnival data breach

This malvertising is a separate incident from the data breach Carnival disclosed in April 2026. Both the domain takeover (2024-11-22) and the cloaker (confirmed by 2026-03-12) predate the April 2026 breach detection, and the operation is generic, mass-scale ad-monetization abuse, which is a different profile from a targeted intrusion. There is no evidence connecting the two, and they should not be conflated.

10. Timeline of observations

Date	Event
2018 to ~Aug 2022	cclpromos.com is the brand's official "Players Club" casino site (Wayback)
~Aug 2022	Site content ceases
2024-11-22	Domain re-registered by a third party (SnapNames); parked

Date	Event
2025-07-23	Last benign parking snapshot, no cloaker (Wayback)
2026-03-12	FingerprintJS cloaker confirmed present (earliest archived proof)
2026-04-22	iOS-branch domain fortiproapp[.]com registered
2026-05-22	Desktop payload domain swatchapp[.]online registered
2026-06-09	Scareware domain andtenands[.]com registered
2026-06-13	Live malware chain confirmed from the booking-confirmation email
2026-06-15	Full de-cloaked HAR capture; NetGuard.Msix hashed
2026-06-16	Rotated advertisers captured (search-hijacker ff.newtabsearch[.]net; fake-McAfee hackshielded[.]com), 3 distinct advertisers in 48h
2026-06-19	Forced SafeWatch.msix re-activated via spot.swatchapp[.]online; fake-AV via gum-promos[.]club
2026-06-20	.appinstaller backdoor loaders captured (PrivacyKeeper via get.privacykeepersite[.]com; SecuredWeb via get.gosecuredweb[.]com)
2026-06-23	Domain also shows a "Buy this domain" parking lander (~\$4,888); fullscreen tech-support browser-locker captured (scam +1 (866) 315-0945); forced download still fired on residential visits
2026-06-24	Anti-analysis and anti-debug layer confirmed; search-hijacker re-served; residential capture of other parked domains funneling to the same rotation; ghostwallguard[.]com fake-McAfee and Azure-hosted tech-support scam (+1-833-306-1211) captured
2026-08-05	Formal certified notice (return receipt) to Carnival's General Counsel, Global CISO, and registered agent
2026-08-13	Carnival Deputy General Counsel engages; matter escalated to Global Cybersecurity Services
2026-08-25	Full technical package delivered to Carnival SecOps (live samples withheld pending request)
2026-08-26	Carnival re-acquires cclpromos.com; nameservers moved off the third-party network to Carnival-controlled DNS
2026-08-27	Malicious routing independently verified dead across all vantages (RDAP plus DNS plus multi-vantage scan); disclosure resolved

The operator rotated fresh downstream payload domains throughout while the cclpromos.com cloaker persisted, until Carnival re-acquired the entry domain on 2026-08-26, which severed the entire chain.

11. Indicators of compromise (defanged)

Category	Indicators
Entry and cloaker	cclpromos.com, ww38.cclpromos.com; cloaker script host sd559908.js.brandsmat[.]com (ob. and obs. subdomains); brokers jubaaa[.]com, tratobid[.]com
Windows-malware delivery	ngapp[.]online (cinga., bacon., der., file.), swatchapp[.]online (spot., file.), get.privacykeepersite[.]com, get.gosecuredweb[.]com; scareware andtenands[.]com
Fake-AV and scareware	hackshielded[.]com, gum-promos[.]club, ghostwallguard[.]com
Tech-support scam	Azure Front Door host *.z02.azurefd[.]net; scam numbers +1 (866) 315-0945, +1-833-306-1211
Search-hijacker	ff.newtabsearch[.]net (cid=10972)
iOS	fortiproapp[.]com
Push	matrixgrowthforge[.]com, pushtorm[.]net
Feed, broker, and lander rotation	brandsmat[.]com, zieweh[.]com, jetsetarmor[.]com, kineticharbor[.]com, buzzfighter[.]com, northwavepoint[.]com, webfervor[.]com, yfdpco4[.]com, cdn-fileserver[.]com, realtimesearchresults[.]com, herme-zdf[.]com, joiquei[.]com, qateyu[.]com, miulyi[.]com, validclick[.]net, 7proof[.]com
Other parked domains serving the same cloaker	globalmotorsport[.]com, dohertysoccer[.]com, transformersarmada[.]com, beecool[.]com, twithority[.]com, mixhost-desu[.]com, cazwall[.]com
Rotated update and delivery hosts (live 2026-09-10)	get.privacykeepersite[.]com, securi-guard-browser[.]com, cablegaurdian[.]online, guardianrole[.]online, euob.northwavepoint[.]com, cint2.scrtrgd[.]online, show.ngrapp[.]online, spons.lkguard[.]online, cinnabon.pkeeper[.]net, lpic.prvbrws[.]com, extensionsnewtab[.]com, pkeepers[.]online, pbrowsingapp[.]online, safe-guard[.]online
Earlier-reporting hosts, now quiet	results.streamio[.]site (clean), cint[.]browsingit[.]online (unreachable)
Brand skins (fake-app names)	NetGuard, SafeWatch, LeakGuard (Trinity Cyber); PrivacyKeeper, SecuriGuard, QuickBrowse, MyConverterHub, SecuredWeb (this report)
Legitimate (do not block)	click.carnivalcruiselineemail.com, carnivalcruiselineemail.com, www.carnival.com (Carnival's own infrastructure)

Payload hashes (SHA-256):

File	Size	SHA-256
SafeWatch.msix	147,724,942 bytes	D66895D8DA6D5EB1D8658647C80F66DCE40236C06BB600F1C62A44A657F923B3 (PhantomJack)
PrivacyKeeper.msix	148,652,394 bytes	8A9006CFAEE227415EEEF0D645183CA423C80B9A8181E35A57BEC71468E72DAA

File	Size	SHA-256
LeakGuard.msix		E9CFDCB18BB4C54802B7214A226D983C0015BB37E41A39A4D293D4E674FA94C8
NetGuard.Msix	147,462,271 bytes	24ec63f3976d04f5e7a7f229ae76301bd1ca6099016d65a727d1c33459853847
BelezaAI.PrivacyKeeper (.appinstaller)		B8F0C82894032766293B7F6E6FEAF287A366EDD4886B8E665C41BC5103D512CE (publisher CN=31E36F3F-1F26-4209-AA73-9C18172AA2E6)
BonitoApp.SecuredWeb (.appinstaller)		8DE1357454488BD0702F0F22E912B412AF41D814DFFDC7C76F416D75D7DC1DC6 (publisher CN=11069677-9767-4ED4-9C5B-619B79B9FE6B)

Campaign artifacts: cloaker `/js/fingerprint/iife.min.js`; params `fp=`, `tr_uuid=`; ms-appinstaller abuse; `dkw=cc1promos.com`; advertiser IDs `cid=9961`, `cid=10972`, `pid=2496`, `ext_name=NetGuard`; `utm_campaign=bgcasino_s1_t4_e1`, `utm_source=bgmarketing`.

Repository status (filed 2026-09-10): the two MSIX samples are on MalwareBazaar, and all four hashes are on ThreatFox. The two .appinstaller manifests are not retrievable from MalwareBazaar, which accepted them and then removed them in processing because they are small XML files rather than binaries, so the manifest block in section 6.2 is the primary evidence for the PseudoJack behavior.

The durable signature is the kit's URL path grammar. The domains rotate; these paths do not:

```
/lps/security-check/
/impression?c=<campaign>&ext_name=<brand>&cid=<id>
/downloadproxy/<campaign>/<clickid>/?ext_name=<brand>&cid=<id>&tag=<cid>_<YYYY-MM-DD>&file=true
/event/download, /event/download_button_clicked, /event/download_xhr_completed
/event/pageload, /event/incognito_status
appData=<fixed 21-character header, identical across unrelated brand domains>
```

The sanitized HAR captures behind these indicators accompany this report in the evidence package.

12. Disclosure and resolution

I reported this the ordinary way first, to everyone who could act on a piece of it. Salesforce, as the email provider, got the source tenant, the malicious link, and the original message (cases #03750689 and #03750690). The registrars and hosts of the payload domains got abuse reports (NameSilo #33009609, Namecheap NC-XML-0090, GoDaddy, the parking provider, and the iOS-branch host), and so did Cloudflare (#434953d6bcfe188e). I filed with the browser blocklists (Google Safe Browsing, Microsoft SmartScreen, and APWG) and with the FBI's IC3 (#09a67a4866704f3893e65610fe5ff548) and the FTC (#202927036).

Carnival itself was harder to reach. The abuse and security inboxes I wrote to on 2026-06-15 went unanswered. A staff member did confirm the issue had been flagged internally, but nothing followed it. The phone tree dead-ended at a switchboard, and even the hotline stood up for Carnival's earlier breach led nowhere for this. The report was credible and complete, and it still could not find a person who owned the problem. The parking provider, meanwhile, pulled individual advertisers as I reported them, but the entry domain stayed live and the email link kept routing to it.

A certified letter finally worked. On 2026-08-05 I sent a formal notice with return receipt to Carnival's General Counsel, Global CISO, and registered agent. The Deputy General Counsel engaged on 2026-08-13 and moved it to Global Cybersecurity Services. I delivered the full technical package on 2026-08-25 (the de-cloaked chain, HARs, per-vantage captures, and IOCs), holding the live malware samples back until they asked. On 2026-08-26 Carnival re-acquired cclpromos.com, moved it onto their own DNS, and pointed it at a real Carnival page. I ran a final multi-vantage check on 2026-08-27 and confirmed the malicious routing was dead from every vantage. Carnival acknowledged the disclosure. The vector is resolved.

For context, on 2026-06-18 the FBI's IC3 published PSA I-061826-PSA, "Cyber Criminals Redirecting Users to Fraudulent Websites with Malicious Traffic Distribution Systems" (<https://www.ic3.gov/PSA/2026/PSA260618>), warning about the same class of TDS redirection. I make no claim that my report prompted it; a federal advisory takes far longer than a few days to prepare. I note it only as independent confirmation that this attack class was active and recognized in the same window.

Closing

This started with an email I almost trusted. I was booked on a Carnival cruise, the Players Club message had my real booking number in it, and every technical check said it was genuine. I clicked the link the way any passenger would, and a fake security page tried to hand me malware. I did not believe it at first, so I ran it again, then from a second machine and a different inbox that had received the same email. Same result.

That is when it became a project. I loaded the link from public scanners, from my own servers in different parts of the world, and from my phone, and watched the results split. Scanners and datacenter IPs got a clean, empty page. Real phones and home connections got malware. That divergence is the whole trick, and it is why the reputation engines everyone relies on waved the link through for months. I chased it by hand at first, hours of scripts across datacenter, residential, and mobile networks, then built whack.sh to do it automatically, because nothing I could find already did.

Analyzing the malware was the straightforward part. Getting a large company to receive a credible report was not. The attacker behaved predictably. The unmonitored abuse inboxes, the internal acknowledgement that went nowhere, the switchboard dead-ends, and the dead breach hotline did not. It should not take a certified letter to a General Counsel to get malware pulled out of a company's own mail.

To Carnival's credit, once the report reached the right people they moved quickly and fixed the root cause instead of the symptom. They bought the domain back rather than chasing payload URLs one at a time. From the first report to a verified fix was about two and a half months. That is how responsible disclosure is supposed to end, and it rarely does. I am glad this one did.