Malanta:

MALANTA RESEARCH PAPER

The Anatomy of a Russian-Speaking Cybercrime Cluster

Malanta Research Team and Assaf Morag, Cyber Threat Researcher

24 SEP 2026

We did not create this cluster. It was already there, one of 80,000 clusters the platform already holds. We opened one of them.

This paper is the proof of what was inside. It is not the way a cluster is made. The walk below is how we explored that one cluster, so you can see what sits behind AI-Powered Threat Intelligence: infrastructure already grouped, including resources that are soon to be malicious and do not yet have an indicator of compromise.

Inside it, one domain had never served a page. The corpus around that domain was 1,585 names, and most of them were not attacking anyone on the day we counted them.

IOC VS IOPA

Indicators of compromise (IOCs) describe resources that are already malicious. Indicators of Pre-Attack (IoPAs) describe the resources while they are still being built, held, or wired. Pre-Attack Intelligence is how those IoPAs become lead time: time to watch the setup, and time to prevent access before the same resource becomes an IOC. Waiting for "malicious" is waiting until the lead time is gone. What follows is one cluster, already built, opened and read. The other 79,999 are not built by writing a paper like this one.

The full domain list for the August 2026 snapshot is a domain-only IoPA CSV on Malanta's public GitHub. Domains only. No scores, no sources, no other attributes.[3] The indicator set for the hosts, addresses, hashes, and kit atoms we name below is available by request.

PUBLISHED WITH THIS PAPER

All-audience cut, same day: **Most of this infrastructure was not attacking anyone yet**

Indicators on Malanta's public GitHub: **Russian-Speaking-Cybercrime-Cluster-2facheck**

Week after that: **Two malware domains. Twenty-one hosts behind them.**

And: **Eighteen drainers. One engine. A buyer who is not the builder.**

1,585
domains in the corpus

28%
live and serving

27 of 30
new domains with no external feed label

80,000
clusters the platform already holds

We started with a question that should have been small

A brand-impersonation site was distributing the AMOS infostealer. Pixel-perfect. The question we were asked was simple: had the operator made more copies?

We did not start on the impersonation page itself. We started on `2facheck[.]ltd`, a domain strongly connected to it. Registered on 27 Jul 2026. First observed two days later. It has never served content.

An IOC workflow would have stopped there, or never started. Nothing to sandbox. Nothing to hash. Nothing for a blocklist to bite on. We treated the empty domain as infrastructure, not as a verdict, and we asked what else shared its identity.

The empty domain was already inside a cluster

We asked the platform about `2facheck[.]ltd`. We did not assemble what came back. The cluster was already formed, and this seed was already in it. The platform returned 1,025 domains with a strong metadata connections to that seed and, through that, to the same malicious entity. Grouping had already happened. We read it. We inferred maliciousness from the seed, which had already been detected and marked malicious, and from what we found when we kept going through a corpus that was already there.

A shared registrar, a shared CDN, or a shared host is not what held this cluster together. Widely shared internet infrastructure is how legitimate companies run too. We are not going to describe the method that forms a cluster. That method is the platform. This paper is the proof of one result: what a single cluster contained when we opened it.

The cluster already held a second population. 560 domains, tied to each other by behavioral connections that are different from the metadata ties of the first 1,025. One domain, `mertamask[.]info`, is the bridge between the registration population and the tooling population. It sits in both. The nuxt batch inventory holds 561 names. Subtract that one bridge and the behavioral set is 560. Together: 1,585.[1] Those two populations are contents we found by opening the cluster. They are not a procedure for building the next one.

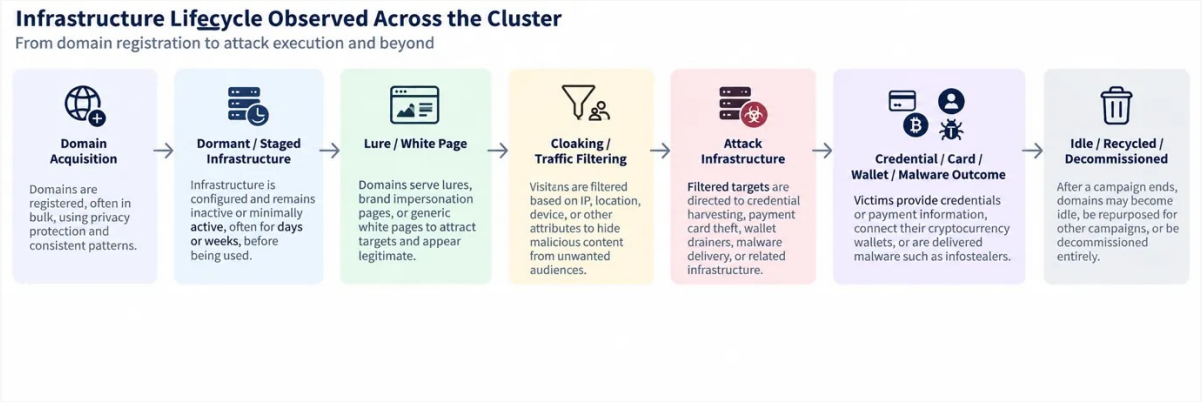

Figure 1. Platform cluster view for the seed domain

That 1,585 is the corpus we explored. It is one nexus, not 1,585 separate cases, and not a cluster we stitched together for the page. Phishing pages, stealer lures, wallet drainers, ticket scams, and the tools that mint the next site were already hanging off the same relationships. The individual campaigns are interesting. The infrastructure behind them is the finding. One of 80,000. The rest of this paper is what was behind this one.

The cluster itself was first built on 15 Jun 2026. Hundreds of members joined that week. The seed we started from was registered on 27 Jul 2026, first observed two days later, and carried a Malanta IoPA classification on 5 Aug 2026 while it was still content-free. That calendar gap is Resource Development lead time: the nexus existed weeks before this empty domain, and weeks before Initial Access or Delivery on the live set.

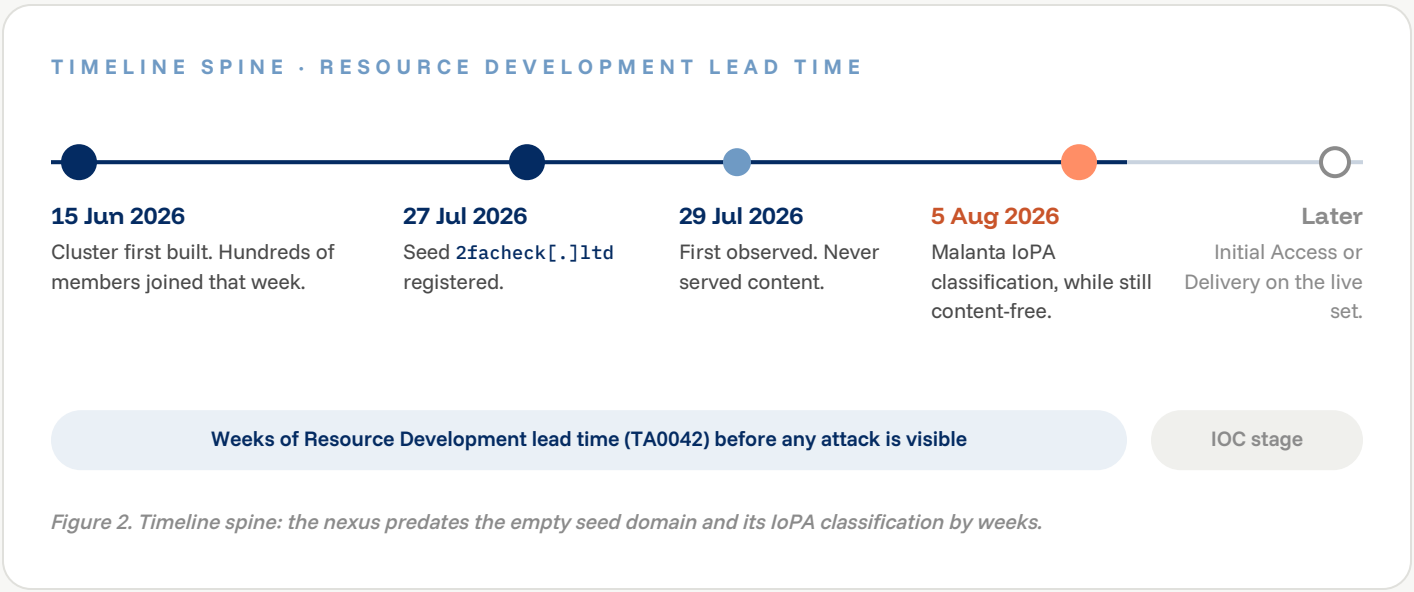

Figure 2. Timeline spine: the nexus predates the empty seed domain and its IoPA classification by weeks.

Later, the live IoPA feed can add domains and drop domains as classification moves. Those changes are churn after this snapshot. They do not rewrite the 1,585, and they do not rewrite the 30-domain expansion below.[4]

The names were templates, and the templates were the tell

Look at 1,585 domains as a list and they look like junk. Look at them as a corpus and the junk repeats. We found seven generation strategies running side by side: brand imitation, transactional language, names that mimic cloud backends, pronounceable nonsense, and strings that look random and are not. Without the cluster, each family looks unrelated to the next. Inside the cluster, they are product lines of one operation.

How we found them: we compared the names to each other, not to a reputation score. The score, when it exists, is an IOC conclusion. The template is visible earlier.

Pseudo-Italian

One of the largest groups almost sounds Italian. The consonant patterns recur (sc, ch, gli, str), and so do the fragments (ri-, dis-, stra-, -ato, -ini, -one, -i). `rabbacarti[.]com`, `rammostrin[.]com`, `riaccplaci[.]com`, `rispavivac[.]pro`, `scazzrispu[.]pro`, `copiaciuff[.]rest`, `semipacros[.]rest`, `scomblazza[.]com`. A single one of these is noise. A set of them, on cheap TLDs, inside this nexus, is a generator.

German banking and orders

These names are aimed at a person who expects a bank or a parcel. konto, kunden, zugang, hilfe, direkt. `de-ihresoforthilfe[.]info`, `de-infokonto[.]com`, `de-kontaktdienst[.]com`, `de-kundenberater[.]info`, `de-paketplus[.]com`, `eu-bank-kunden[.]info`, `eu-login53-bank[.]digital`, `eu-secure-konto[.]live`. A second cut of the same vocabulary: `go-bank-digital45[.]info`, `go-check-zugang32[.]live`, `kunden-hilfe[.]digital`, `kundenservice23-center[.]live`, `mein-check-direkt[.]digital`, `mein-intern-hilfe[.]digital`. A third cut is just orders: `bestell-193367[.]info`, `bestell-193551[.]info`, `bestellung-104229[.]info`, `bestellung-107492[.]info`. Same corpus. Different lure. The TLD is part of the mold: German account lures keep landing on `.info`, `.digital`, and `.live`.

Crypto

One family swaps `2x` and `x2` as if they were the same brick: `arb2x[.]com`, `arbx2[.]com`, `hedge2x[.]com`, `hedgex2[.]com`, `fin2x[.]com`, `mint2x[.]com`, `2xmint[.]com`. Doubling. Profit. Another family does not promise a multiple. It promises that the money is already yours: eligible, rewards, allocation. `eligible-bonkcoin[.]life`, `eligible-jup[.]live`, `eligible-jup[.]rocks`, `jupiter-rewards[.]app`, `allocation-flare[.]com`, `allocations-arcium[.]com`, `reallocation-flare[.]network`, `reallocation-sparkdex[.]com`, `proposal-flare[.]com`, `proposal-hachiko[.]com`. The drain is written into the name before anyone connects a wallet.

Fake backends

These do not impersonate a consumer brand. They impersonate a hostname a developer would ignore. api, gateway, prod, alpha, edge, portal, auth, secure. `com-alpha-edge-service-prod[.]pro`, `com-alpha-online-portal-gateway[.]pro`, `finance-secure-api-edge-dex[.]pro`. An IOC list full of "login" and "secure" will catch some of them after they phish. The pattern is visible while they still look like infrastructure.

Brands, including AI brands

`microsoft-live[.]ltd`, `citionline-us[.]com`, `citisupport-us[.]com`, `help-trustwallet[.]email`, `binance-careers[.]com`, `blackrock-gift[.]net`. And the newer shelf: `claude-pormo[.]com`, `claude-promax[.]com`, `claude-subscribe[.]com`, `gpt-promo[.]com`, `gpt-proplus[.]com`. Hold the Claude and GPT names. We come back to them when we can show they were staged, not born malicious.

Pronounceable inventions, and "random" that is not random

`celtrionavex-pardemolixu[.]live` and `zorquandelmis-praxelotivun[.]live` share a shape: two invented words, a hyphen, `.live`. Fragments repeat (vex, lix, vori, morti, pra, meli, vandi, xant, qu). This is not a classic high-entropy domain generation algorithm. It is a style. At the other end, `c155w5qw5rwwqrqpoq[.]com` and `hakdsiwqs281ks[.]com` look like keyboard garbage. One subset leans on q, w, c, r, 1, 2, 4, and 5. Another mixes asd, dsk, jsk, kak, ksa with numbers. We met this second family again inside the ClickFix stagers. The linguistic tell and the malware chain are the same corpus.

Tickets

`alcatraz-cruise[.]pro`, `angkortickets[.]today`, `carnivalmagic-tickets[.]pro`, `mahanakhon-skywalk[.]pro`, `tickets-alhambra[.]today`, `tivoli-dk[.]pro`, `visit-taronga[.]today`, `efteling-nl[.]pro`, `efteling-tickets[.]pro`. `.pro` or `.today`. A different victim, the same factory.

Across the 1,585, `.com` is still the largest single TLD, about 41%. `.pro`, `.info`, `.live`, `.rest`, and `.digital` together are about 42%. Invented Romance-looking words prefer `.rest`, `.news`, `.digital`, `.pro`. Synthetic two-word names pile up on `.live`. The suffix is part of the generator.

We could not tell from names alone whether this was an internal factory or a service sold onward. The templates are compatible with both. Behavior is what assigned roles. That was the next door.

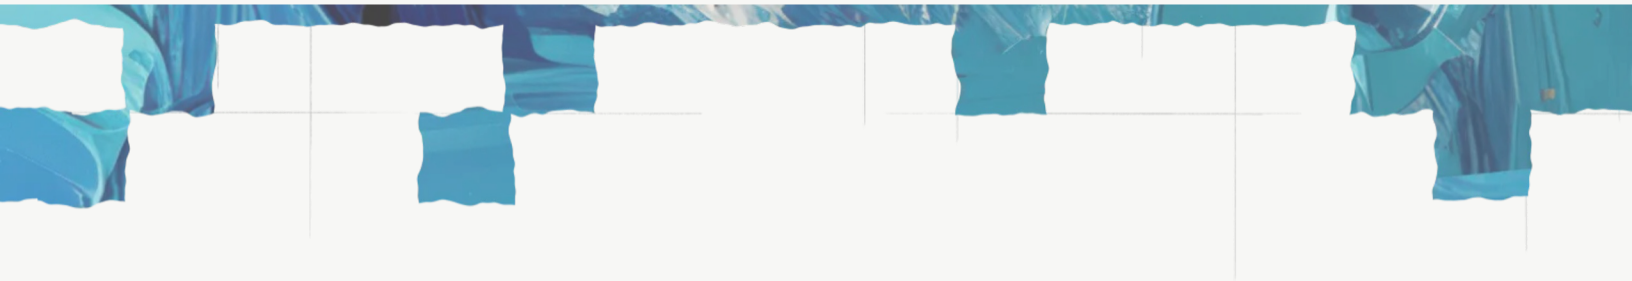

Only 28% of the corpus was attacking anyone

We classified every domain by what it was doing at the time we looked, not by what a feed had already decided.

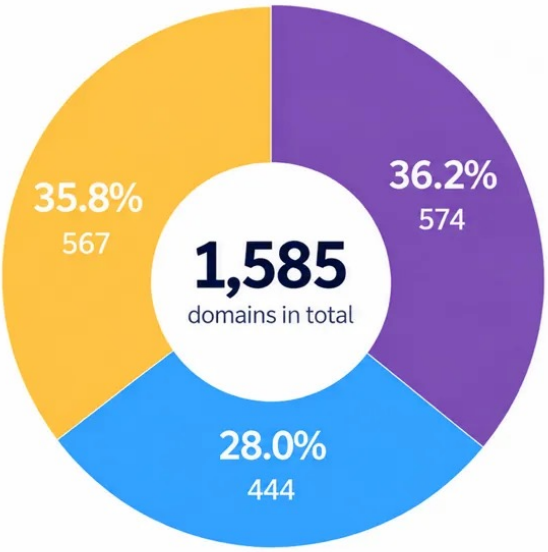

Category	Domains	Share
Decommissioned Infrastructure is no longer operational, including withdrawn, suppressed, expired, or sinkholed domains.	574	36.2%
Idle Infrastructure remains present but is not currently serving active malicious content.	567	35.8%
Live and serving Currently serving active malicious content or controls.	444	28.0%

Figure 3. Live, idle, and decommissioned distribution of the corpus

State	Domains	Share
Live and serving	444	28.0%
Idle	567	35.8%
Decommissioned	574	36.2%

Table 1. Corpus state at classification time

Seventy-two percent was idle or decommissioned. If a program only ingests IOCs, it is instrumented for the smaller slice. The lead time sits in the other 72%, and in the movement between the three states. A domain here can be unused, go live, go quiet, and later be reused. "Not malicious today" is a state. It is not a biography.

STATE IS NOT BIOGRAPHY

The lead time sits in the other 72%, and in the movement between the three states. A domain here can be unused, go live, go quiet, and later be reused.

"Not malicious today" is a state. It is not a biography.

Live, 444. Phishing, card theft, drainers, malware, scams, criminal services, or cloaking that shows the bad content only to the visitor it wants. Credential and payment-card phishing is 340 of the 444 (76.6%). Once a domain is switched on, it is in a race. We saw lifetimes from hours and days out to weeks and longer. We do not have one lifetime for the set.

Decommissioned, 574. DNS withdrawn, registration expired or suspended, infrastructure gone after it had served content, or a CDN standing in the way. At least 160 of these had been active before: 101 later suspended or expired, 59 where DNS was withdrawn after malicious activity. A CDN interstitial can block the visitor and leave the origin alive. Expiry is not a takedown. We counted 180 domains suspended by registrars and 65 simply allowed to expire. Set that against May 2026, when this operation registered 249 domains in a single month, and attrition looks like a cost of doing business, not a broken factory. On 31 of 104 domains sitting behind CDN blocks, Cloudflare showed "Suspected Phishing" or "Suspected Misleading Website." Useful corroboration. Not a dismantling. We found scattered public classifications of single domains. We did not find structured public reporting that had already tied them into this nexus.

Idle, 567. Still in the corpus. Not currently serving confirmed operator-controlled malicious content. Idle is a pile of different conditions, and we refused to flatten it into "the next attack."

What we actually saw	Domains	Share of 567
Reverse proxy armed, origin switched off	180	31.75%
Locked out (CDN block or challenge)	140	24.69%
Served, not attributable to this operator	105	18.52%
Already attacked, infrastructure still intact	55	9.70%
Placeholder, default, or panel page	45	7.94%
Answers, nothing published	22	3.88%
DNS withdrawn since our sweep	9	1.59%
Prior-owner history only	7	1.23%
Certificate does not match host	4	0.71%

Table 2. Idle (567): what we actually saw

WATCH, DO NOT OVER-CLAIM

The 180 reverse proxies are the sharpest "soon to be" population we have, and also the easiest to over-read. The front end is in place. The origin does not answer. That can be one switch from live. It can also be us misreading a dead upstream. We know this slice has already been used to serve malicious content. We do not know which of the 180 will be switched on, or when. An IoPA here is a reason to watch and to prevent access if the state changes. It is not a reason to tell a board that 180 attacks are scheduled.

The 55 are the history an IOC feed throws away. They resolve. They serve nothing. They have already run an attack. The content left. The infrastructure stayed.

`caixadirecta[.]life` is the clean example of how we found that. It still resolves on `46.151.182[.]154`, the same origin that served a 1.05 MB clone of Caixa Geral de Depósitos, on the same nginx/1.18.0 (Ubuntu). Urlscan shows the campaign live in early 2026. One February URL, `https[:]//caixadirecta[.]life/REQX7PJOTFXE5DIYV67KIK7GM8/pin`, was a CaixaDirecta phishing page. The server was still up when we looked. The phish was not. An IOC for that URL goes stale. The host, still in this cluster, is the thing Pre-Attack Intelligence keeps on the watchlist.

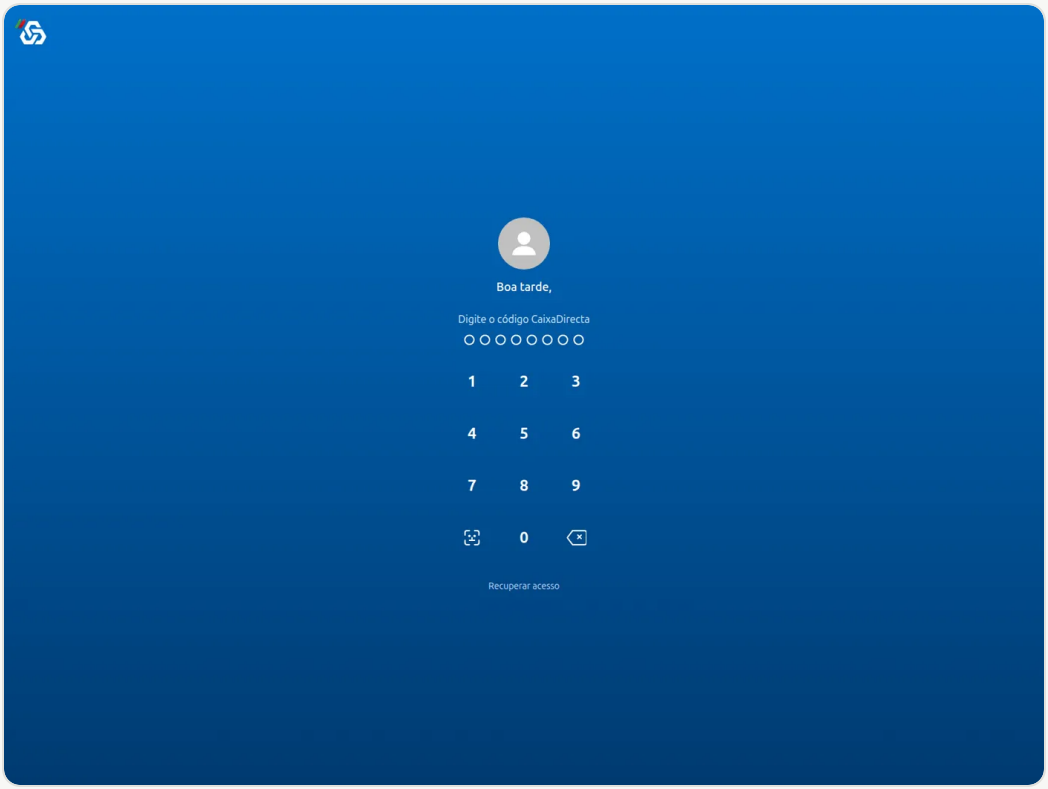

Figure 4. CaixaDirecta phishing capture on caixadirecta.life

`seg-social[.]life` is why we did not paint the whole idle set with that brush. It shows an nginx default page. We have no record that it ever displayed malicious content. It stays idle. Same corpus, different claim. Relationship is not the same thing as "this hostname is malicious right now."

Where we could see a serving pattern, the content usually sat on the apex or on a subdomain (702 of 779, 90.1%). Half the corpus, 806 domains, we never observed serving at all.

Pattern	Domains	Share of 1,585	Share of 779
Apex domain	492	31.0%	63.2%
Subdomain only	210	13.2%	27.0%
Path on apex	45	2.8%	5.8%
Subdomain and path	21	1.3%	2.7%
Subdomain root	11	0.7%	1.4%
Never observed serving	806	50.9%	
Total	1,585	100%	

Table 3. Where content sat when serving was observed

The same campaign, three clocks

The Claude lures show the corpus moving through states in front of us.

Around May 2026, a brand impersonation against Anthropic's Claude was live. By the time we wrote this up, `claude-pormo[.]com` and `claude-promax[.]com` were decommissioned. Urlscan still has what `claude-promax[.]com` displayed while it was alive.

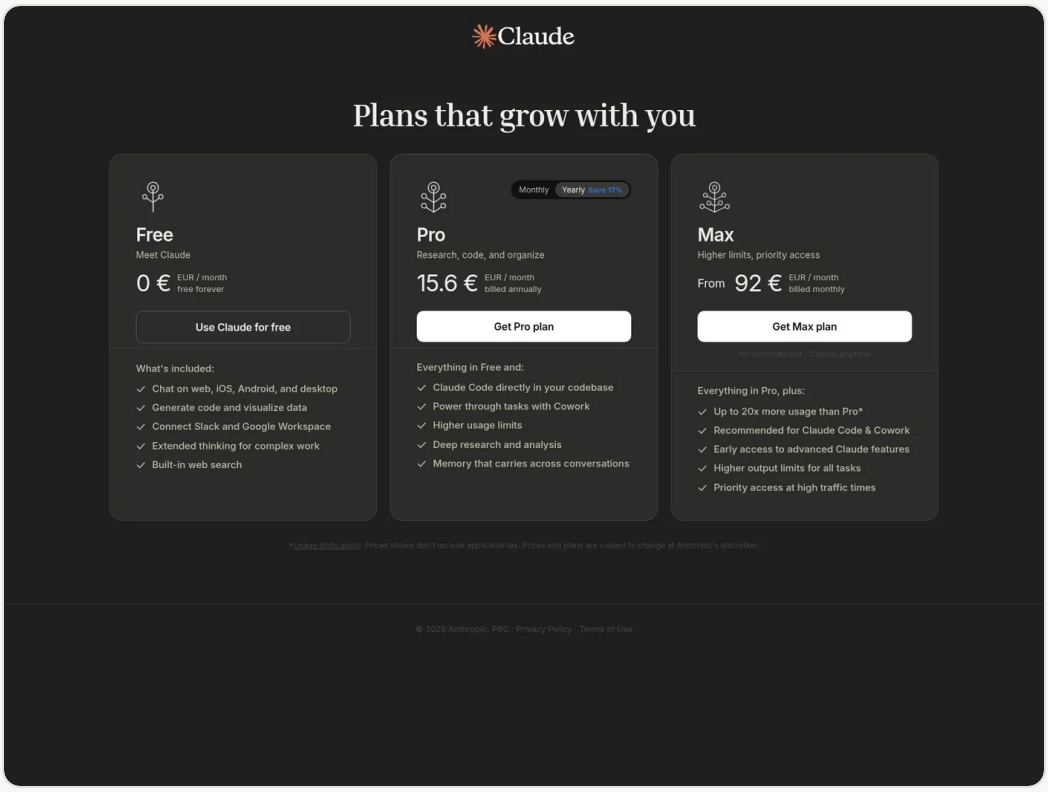

Figure 5. claude-promax.com while the Anthropic impersonation was live

`claude-subscribe[.]com` was still answering on the same IP and returning 404. That 404 might be our misclassification, or it might be idle. Either way it is the same campaign, still in the cluster, no longer a tidy IOC.

Figure 6. claude-subscribe.com returning 404 on the same campaign IP

`gpt-promo[.]com` and `gpt-proplus[.]com` had never served content. VirusTotal and the other engines we checked had never indexed them as malicious. They were in the corpus anyway, because of how they connect, not because a feed had convicted them. That is an IoPA in the strict sense: soon-to-be infrastructure, visible through the nexus, with no IOC yet. We infer that the operators register several domains for the same sector in advance and light them at different times. We do not have an activation date for those two. The inference stops where the evidence stops.

Three quarters of what was live was one business

Of the 444 domains actually serving, 340 were credential and payment-card phishing. The other quarter is not a second copy of the same campaign. It is a portfolio. The counts below are the populations we identified. They overlap in places, and they are not a partition we forced to add up to 444.

We went through them one business at a time. Each time the question was the same. What did we find, how did we find it, and how does it attach to the 1,585 rather than sit as a one-off?

Two malware domains, then twenty-one hosts behind them

Only two domains in the entire corpus were classified as malware delivery. Both are ClickFix: a fake check tells the victim to paste a command. If we had stopped at those two IOCs, we would have reported a small malware incident.

We followed the chain. Cloaking, telemetry, command delivery, staging, execution: 21 additional hosts. The visible lure is separated from the command and the payload on purpose. The loader filters security crawlers and AI crawlers, including GPTBot, ChatGPT, ClaudeBot, Anthropic, and CCBot. A scanner that is on the block list receives a benign page and files a benign verdict. That is how an IOC gets delayed, and how an IoPA (the cloaker, the stager, the payload host, already wired) stays invisible to anyone waiting for the lure to look malicious.

`openwork[.]life` is the macOS face, a developer job board. `nova-trade[.]world` is the Windows face, a fake trading app. `worldherald[.]life`, which shows up again in the two-week expansion, is a fake newspaper on the same cloaker, `corrykro[.]icu`. We walk that chain in full the day after this paper. The point for the corpus is smaller and sharper: the malware "campaign" is two doors on a much larger hallway, and the hallway is part of the same nexus as the German banks and the ticket shops.

Eighteen drainers, one engine, a buyer who is not the builder

We classified 18 domains as cryptocurrency wallet drainers. They look like eighteen attacks. Under the lures, the implementations repeat, bought and reskinned.

The way in was a mistake. `listedokx[.]net` left an Apache directory listing at the web root. We could read the loader, the wallet adapter, six compiled bundles, two injectors, and a `vote/` directory of per-token campaigns. Timestamps run from January through June 2026. From line 183, `loader.js` is obfuscated. Lines 1 through 182 are plaintext because the buyer has to edit them: their wallet, the minimum victim value, the Telegram alerts, the lure behavior. The comments call `targetWallet` the buyer's "split wallet."

```
// configuration - we do not recomend touching most of these.
// if they have [!] symbol, pls do not touch them.
var DEFAULT_CONFIG = {
  hideChanges: true,   // [!] bypass estimated changes. will show simulation
                       //     error when enabled. always keep true.
  spoofPrompt: true,   // [!] always keep true.
  customAirdropMint: "677jK6xujhRFtBcD7Szn32tt2MSv4k6qZ4v3nGMgpump",
  customAirdropAmount: 25,
  targetWallet: "Aa5tvMinJreJHGpxZPV4R7YGCeRyxT6hKTt1XYdEdpFe", // your split wallet.
  userChatId: "-5060889444",
  minValue: 0.1,
  className: "interact-button",
}
```

Figure 7. Buyer plaintext configuration in loader.js on listedokx.net

That is the supply chain, sitting inside our cluster. The person who built the engine and the person aiming it are not the same person. Taking down one lure domain raises the operator's cost by the price of the next domain. It does not remove the kit. We publish the mechanics two days after this paper, including EasyRug, which this operation uses and which we excluded from the factory counts because it predates the factory.

Six of the eighteen pretend to be security tools. Analyzr says "Decode Every Token." DevWatch says "Read every EVM token before you ape." The pitch is aimed at the person who learned to inspect a contract before buying. The tool asks them to connect a wallet. `analyzr[.]live` is 1.2 MB. `povtrader[.]world` is 6.9 MB. These are applications, not thin phish pages. A feed of malicious domains will eventually list some of them, after victims have connected wallets. The IoPA is earlier: the naming family (eligible, allocation, 2x), the shared kit atoms, and the fact that the domains already sit in this corpus.

Twenty-five fake ticket windows

We found 25 domains selling entrance tickets that are not tickets. Alcatraz City Cruises, Hersheypark, Tivoli Gardens, Niagara Parks, Taronga Zoo, Angkor, Efteling, plus a set of fake e-shops. The names we already pulled out of the linguistic pass (`alcatraz-cruise[.]pro`, `tickets-alhambra[.]today`, and the rest) are this business.

On `.pro` we found a straightforward phishing kit.

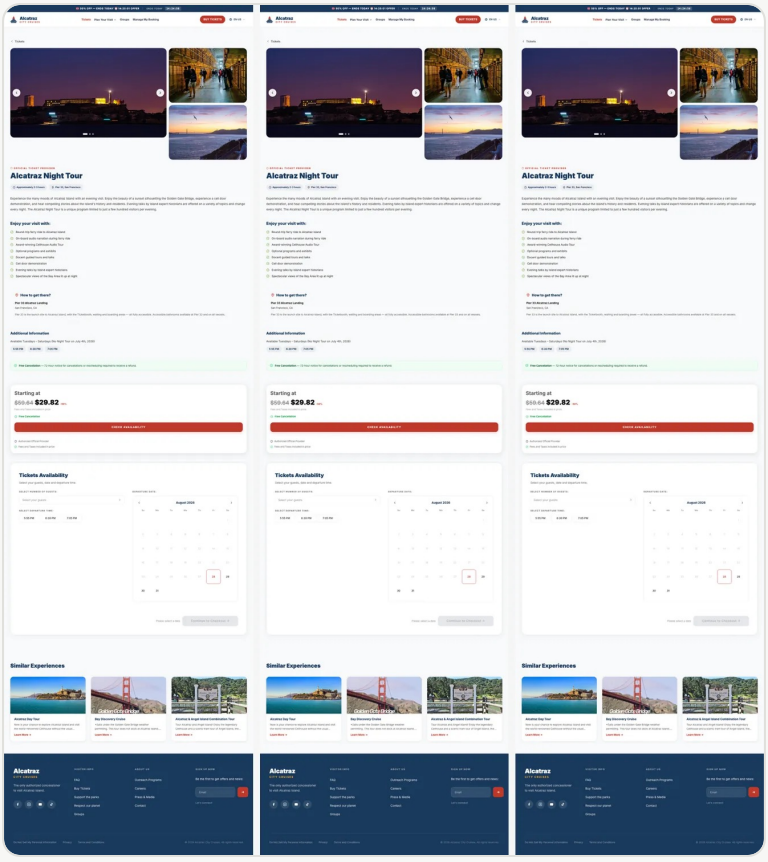

Figure 8. Alcatraz .pro phishing kit (combined view of the .pro kit pages)

On `.today` we captured the Andalusia lure at `tickets-alhambra[.]today`. The homepage is a bland visitor guide. The payment flow hides under `/tickets/<slug>`, with checkout, payment, and booking confirmation. Card number, expiry, CVC, cardholder, billing. The gateway path is built from a template the kit calls jespay, falling back to `jespay-default`. The request uses `keepalive: true`, so the card still submits if the victim closes the tab. We saw "Payment Declined" and "Try Again" screens. We could not confirm that a retry sends a second card.

Figure 9. Andalusia visitor-guide lure on tickets-alhambra.today

How it attaches: same TLDs, same corpus, different victim. A defender waiting for an IOC tagged "bank phish" will not be watching a zoo.

Forty-seven domains that are not for victims

We found 47 domains that look like the operator's own infrastructure. Proxy consoles, hosting and mail panels, a live self-hosted remote-support server whose client-installer path is a ready-made way to deliver an agent, and a hardened internal chat server. These will not look like phishing in an IOC feed, because they are not phishing. They are how the rest of the corpus gets run. Preventing access to a lure while leaving the panel up is how this operation survives a takedown.

Fourteen of the 47 front Atlascore (`atlascore-panel[.]net`), a Russian-language Scam-as-a-Service panel. Atlascore publishes a public GitBook of its affiliate playbook. An unauthenticated page in this infrastructure contained the word мамонт (mammoth). In the material we recovered, that word is used as Russian-language slang for the scammer, next to documentation for exchanges, a Polymarket clone, casino, and betting. Group-IB's Classiscam work and Brian Krebs's writing on affiliate economics already cover Atlascore as an operation. We are not re-investigating that reporting here. We are showing that 14 doors to it sit inside this cluster.

The machines that mint the next site

Two domains explain why so many decoys in this corpus look like real businesses.

`affspect[.]pro` is in the cluster and is a byte-identical mirror of Cloakflare (`cloakflare[.]io`). The site sells reverse-proxy cloaking. Cryptocurrency only. No identity checks. About eight dollars a month. It wraps itself in Cloudflare's look: a modified orange-cloud logo, a Cloud-to-Cloak pun, Cloudflare's orange. We have two readings and we did not pick one. Either these operators run the service as a side business, or the mirror is a campaign against Cloakflare's own customers. Both stay open.

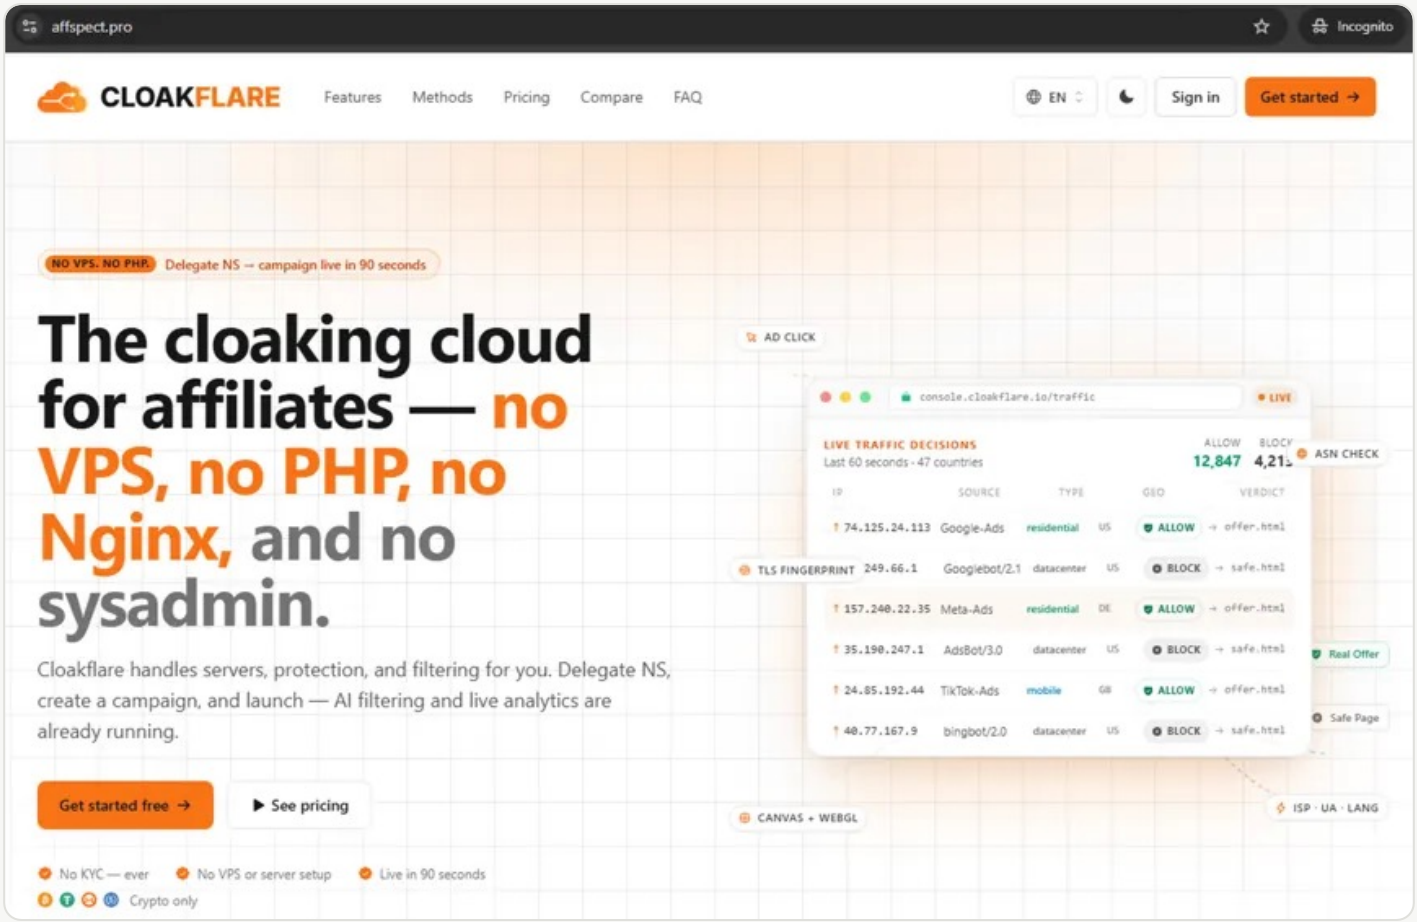

Figure 10. Cloakflare cloaking console (mirrored by affspect[.]pro in this cluster): live traffic decisions, bot/datacenter blocks, real offer vs safe page.

`whitepageflow[.]com` hosts Mainframe, a Russian-language service that describes itself as a website generator for traffic-arbitrage teams (генератор сайтов для арбитражных команд). Customers upload domains and topics, pick a language, watch generation, retry failures. Quota is per domain. The feature that matters for us is the 1:1 mode: it scans an example site for structure, styling, and legal pages, then generates a new site from those traits. Privacy, Terms, and Cookie pages are a product feature. Delivery is a ZIP, deployed wherever the customer wants.

That is how a domain in this cluster can look legitimate before it is malicious. The white page is the setup. The phish, the card form, or the drain is a later state of the same asset, or of the next asset minted the same way. An IOC appears when the malicious state is visible. An IoPA can appear when the generator, the cloaker, and the empty domain are already connected.

Three identities, and a funnel we could not finish

`fzo[.]info` clones the Russian state veterans' fund Defenders of the Fatherland: branding, navigation, branches, vacancies, documents. `kdha[.]world` wears the title "Qatar General Electricity & water Corporation." Cloudflare later put a "Suspected Phishing" warning on it.

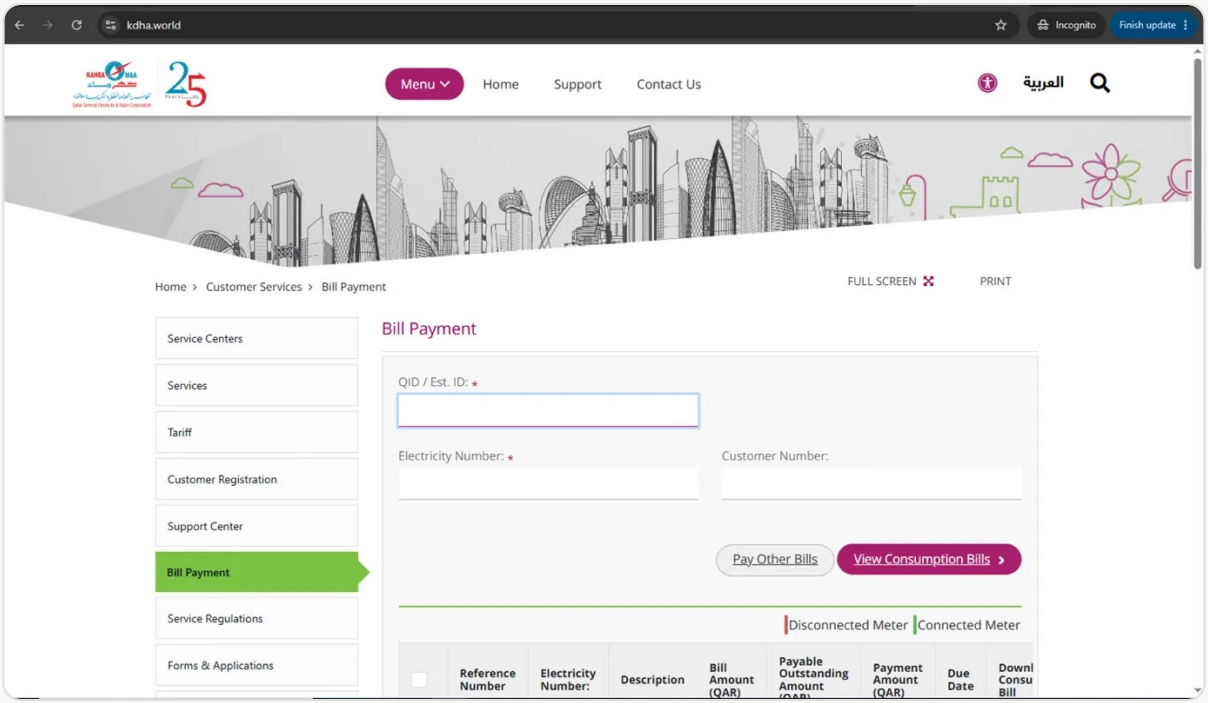

Figure 11. kdha.world utility impersonation page

`claritynow[.]life` uses a Russian television personality as the face of a six-step quiz: relationship breakdown, domestic problems, estranged or addicted children, widowhood, health, distrust of doctors, the evil eye, hexes, generational curses.

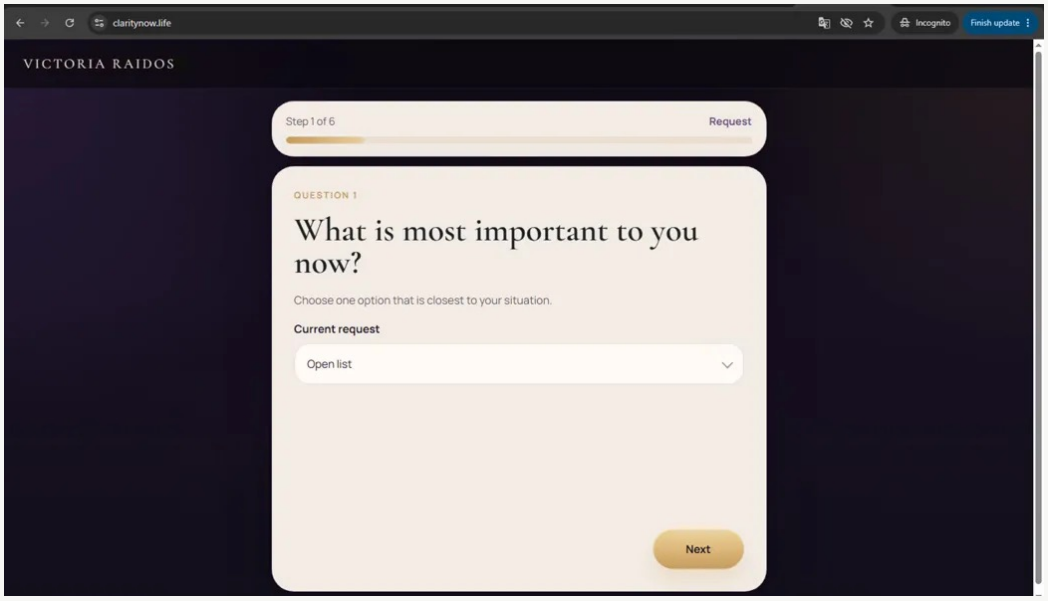

Figure 12. claritynow.life qualification funnel

We did not find a payment form on these pages. We think they are a qualification step, selecting people for a later approach. That is our speculation. We are not going to promote it to a confirmed fraud path just because the domains sit in the corpus.

The factory's pulse, then the two weeks we watched it grow

Registration, for the 1,025-domain metadata set, was not a spike. It was a production line. The monthly counts sum to 1,025.

Month	Domains	Mean per day	Busiest day	Days at zero
Nov 2025	37	3.70	12	1
Dec 2025	28	0.90	4	13
Jan 2026	31	1.00	6	20
Feb 2026	83	2.96	14	1
Mar 2026	91	2.94	11	8
Apr 2026	191	6.37	22	5
May 2026	249	8.03	30	2
Jun 2026	216	7.20	26	2
Jul 2026	61	1.97	10	10
Aug 2026	38	2.24	12	5

Table 4. Registration monthly counts (1,025-domain metadata set)

April through June is the surge. May peaks at 249. July and August slow down. They do not stop.

We then froze a baseline on 25 Aug 2026 and watched for two weeks. The cluster added 30 domains. None of the 30 were in the original 1,585.[2]

What we counted	Two-week snapshot	Rate
Newly observed infrastructure	19 domains	about 1.1 per day
Overall growth of the cluster	30 additions	about 1.8 per day

Table 5. Growth after 25 Aug 2026 baseline

New names arrived at roughly two a day, between one and three, mean of two. At the recheck, nothing had been removed.

Here is the lead time, as a count we can actually stand behind. Of the 30, 27 carried only our pre-attack classification. Three already had an external threat-feed label. Ninety percent of the new infrastructure was visible to us, through relationship to this corpus, before a conventional feed would have handed a defender an IOC.[2]

That is not a claim that we blocked those 27. We did not record a prevented incident, and we did not compute a Mean Time to Preempt. We recorded the order: the domain enters the cluster, the external malicious label is not there yet, the setup may already be real.

A redirect that ends on the real website

`profiledetails[.]info` and `753484474[.]info` are in that set of 30. We followed them and found a chain through `gumtree.mydetails-a49b90b5f15ce514473[.]top`, registered on 2 Sep 2026, and then out to the genuine Gumtree site. DNS worked. Certificates worked. The redirect worked. The last hop was legitimate.

SOON TO BE MALICIOUS

An automated control that only follows redirects will grade this as fine. A defender waiting for an IOC will see nothing to collect. We saw a delivery path, already built, already tied to the cluster, not yet delivering a phish. That is what "soon to be malicious" looks like when you are early enough. The IoPA is the chain and its connection to the nexus. Preventing access means being able to act on that chain before the last hop changes to a fake login. It does not mean the last hop was evil on the day we tested it. It was not.

`worldherald[.]life` is the other shape of the same fortnight. It did not sit idle. It arrived in the cluster and was already wired to `corrykro[.]icu/t.js`, the ClickFix cloaker. Sometimes the IoPA and the attack show up together. Lead time is not guaranteed on every name. On this one, the value was the link back to a cloaker we already had, which pulls a new lure into a chain we had already mapped.

`deklaracija-vmi[.]digital` and `mdeklaracia[.]digital` point at Lithuania's State Tax Inspectorate the way the German names point at banks. We call them a future target. We did not find a live phishing page on them at the time of writing. They are staged names inside the expansion set. `rootproxy[.]life` is a proxy on a campaign-registered domain. We probed it once. Unconfirmed. Monitor, do not block.

Russian is in the room. A name is not.

We call this a Russian-speaking cybercrime ecosystem. We mean that precisely. We do not mean a named group, a nationality, a city, or a person.

The strongest operator-side evidence is on `firstopengames[.]com`. We recovered 173 comments in Russian, written in an operator-to-self register, across four independent codebases: dated build notes, conversion results, telemetry files, API paths, offer URLs. We searched several internal build tags and got no public hits. These read as development leftovers, not as comments copied out of a kit anyone can download.

We also found Russian-language admin interfaces on internal services, including a VPN admin environment and the Mainframe generator.

Then we had to separate what the operators wrote from what they bought. Atlascore speaks the slang of the Russian-language scam scene. Other kits we recovered contain Russian vendor-to-customer setup notes. Mainframe and EasyRug are services this corpus consumes. A Russian comment inside a purchased kit tells us about the vendor. During the work we pulled some artifacts back out of the "operator" pile once we recognized them as vendor documentation. That correction matters. If we had treated every Cyrillic string as attribution, we would have named the wrong person.

Where we could identify victim-facing language, German was 58%, Russian 23%, then Japanese, Polish, Portuguese, Dutch, and others. The corpus impersonates German and Portuguese banks, agencies, utilities, crypto services, logistics companies, attractions, and consumer brands. Activity aimed at Europe sits next to a 23% Russian victim-facing share. We went into this expecting the familiar claim that Russian criminal actors, and state-sponsored actors, do not target Russians. The 23% does not sit quietly with that claim. It left us puzzled. We are not going to resolve a puzzle by inventing a sponsor.

What we will say: Russian is in the operator's own comments and in the working environment, and Russian-speaking vendors are part of the supply chain. Ecosystem. Not a flag.

The phase an IOC feed does not cover

Most of what we could see first sits in MITRE ATT&CK Resource Development, TA0042. The attack does not start when the phishing page or the malware becomes visible. Domain acquisition, certificates, servers, accounts, bought tools, staging, and cloaking leave signals while the resource is still soon-to-be. The techniques below are the ones we actually observed. They are not a checklist for every domain, and they are not evidence that one person runs every component.

The short version, tied to lead time:

- **Domains (T1583.001).** 1,025, then 1,585, then 30 more in two weeks. Peak month, 249. The idle and decommissioned piles mean they acquire domains beyond whatever is live today.
- **Servers and web services (T1583.004, T1583.006).** The lure is not the server. ClickFix taught us that. CDNs separate the name the victim sees from the origin.
- **Accounts (T1585.002).** We identified 182 registrant personas, recurring name patterns, on a public disposable-email service. These are not the operator's personal inboxes. They are how the factory registers the next name.
- **Tools (T1588.001, T1588.002).** ClickFix, drainers, cloaking, Mainframe, commercial bot detection, Scam-as-a-Service. In several places the artifact itself separates builder from buyer.
- **Certificates (T1588.004, T1608.003).** Issued as infrastructure is prepared. In one case, six certificates for unrelated impersonated brands landed within about 90 seconds. That burst is an IoPA. It does not require the pages to be malicious yet.
- **Staging (T1608).** The 567 idle domains. Prepared, or retained, between campaigns.

After Resource Development, the live set does what live criminal infrastructure does. These are the IOC stage. We are not arguing they are unimportant. We are arguing they are late.

Initial Access	Phishing: Spearphishing Link (`T1566.002`). Distribution we observed includes email, SMS, and advertising. Some kits expose malicious content only on a specific path. The root page can stay innocuous.
Execution	User Execution: Malicious Copy and Paste (`T1204.004`). Command and Scripting Interpreter: PowerShell (`T1059.001`) and Unix Shell (`T1059.004`). The macOS branch selects an architecture-specific payload and continues outside the original Terminal session.
Defense Evasion	Impersonation (`T1656`). Obfuscated Files or Information (`T1027`): obfuscated JavaScript, a virtual-machine-like implementation, encrypted payloads, dynamic API resolution, oversized binaries with padding. Subvert Trust Controls (`T1553`): the macOS ClickFix chain removes the quarantine attribute and applies ad-hoc signing. Masquerading (`T1036.005`): persistence under `~/Library/Caches/com.apple.securityd/com.apple.periodic`. Indicator Removal: Clear Command History (`T1070.003`). Virtualization/Sandbox Evasion: System Checks (`T1497.001`), including the crawler block list on the ClickFix loader.
Credential Access	Input Capture: Web Portal Capture (`T1056.003`): credentials, card number, expiry, CVC, cardholder, billing. Multi-Factor Authentication Interception (`T1111`) on some card flows, not all.
Command and Control	Dynamic Resolution (`T1568`), including DNS-over-HTTPS TXT lookups such as `apiconfig.nova-trade[.]world`. Web Service (`T1102`): Telegram for notifications, support, and drainer alerts. Telegram alone is not a cluster identifier.
Impact	Financial Theft (`T1657`): card data, cryptocurrency theft, and other fraud aimed at direct financial gain.

What we do with the lead time

Preventing access to a soon-to-be resource is not the same as blocking every neighbor of a bad domain. We learned that the hard way inside this corpus: shared services are full of tenants who are not this operation. The move is narrower.

Watch domains that connect to known malicious infrastructure even when they return 404, a default page, a redirect, or nothing. In this corpus, 55 idle domains had already attacked, and 27 of 30 new domains had no external label when they joined.

Alert when the state changes. DNS, certificate, HTTP status, a new form, a new redirect, a new upstream. A dormant name returning HTTP 200 with a login is a stronger signal than yesterday's reputation score.

Hunt the atoms that survive a rename: analytics IDs, ad conversion IDs, the phishing-kit hash that repeats across the 340-domain carding fleet, cloaker site IDs, certificate bursts, redirect chains. Those are how a new domain gets pulled into the nexus before it earns an IOC.

Test from more than one vantage point. This cluster hides from security crawlers and AI bots on purpose. One benign fetch is not a benign domain.

Use relationships as enrichment until behavior or a technical atom promotes the asset. Then, and only then, feed the confirmed domains, URLs, redirectors, payload hosts, and hashes into DNS filtering, the secure web gateway, EDR, the SIEM, and the threat-intelligence platform. Keep the dormant and the suspected under a different label, so they can be watched without being treated as confirmed malicious.

That is Pre-Attack Intelligence in practice. IoPAs extend the lead time. IOCs still matter once the resource is actually malicious. The mistake is using only the second list and calling the gap "detection." Act in the IoPA window: empty, idle, or staged names already in the nexus are how security teams can prevent access before Initial Access and Delivery.

THE BOTTOM LINE

We started with `2facheck[.]ltd`, a domain that never served content, because it was connected to an AMOS impersonation site. Metadata gave us 1,025 domains. The nuxt batch gave us 560 more. The corpus was 1,585. On the day we classified it, 28% was live, and 76.6% of that live set was stealing credentials or cards. The rest of the live set was a portfolio: ClickFix hallways, rented drainers, fake tickets, operator panels, a cloaker, and a machine that mints white pages.

The other 72% was idle or decommissioned, and those states move. Fifty-five idle domains had already been used. In the two weeks after 25 Aug 2026, 30 domains joined, and 27 of them were not in an external threat feed when they arrived. One pair already had DNS, certificates, and a redirect, and the redirect still ended on the real Gumtree.

The nexus was first built on 15 Jun 2026, weeks before the empty seed existed. The infrastructure comes before the attack. That gap is the lead time. IoPAs are how we see it. Waiting for the IOC is how we give it away.

None of that is a recipe. We opened one cluster the platform already held, out of 80,000, and we wrote down what was inside. That is the proof of what sits behind AI-Powered Threat Intelligence. The clustering itself happened before we started.

Endnotes

[1] August 2026 membership at the time of writing: 1,025 domains from the metadata inventory, plus 561 domains from the nuxt / behavioral batch, of which `mertamask[.]info` is the single overlap. Union: 1,585.

[2] Two-week expansion against the 25 Aug 2026 baseline: 30 domains added. At the recorded recheck, 27 carried only a Malanta IoPA label and 3 carried an external threat-feed IOC label.

[3] Domain-only IoPA CSV for the 1,585-domain snapshot, for Malanta's public GitHub.

[4] Domains added to or removed from the IoPA feed after this window are churn. They are not part of the 1,585 and they are not part of the 30.