THREAT INTELLIGENCE REPORT

The Real Enemy Isn't the Scammer. It's the Registrar Cashing His Checks.

Ref. **PD-TI-2026-36358** Published **10 July 2026** Source **phishdestroy.io/the-real-enemy-is-the-registrar**

Editorial — Op-ed

The scammer is a cheap, replaceable node in an economy of desperation. The registrar taking his crypto with a smile — and the ICANN that won't make him answer for it — is the disease.

PhishDestroy — Independent anti-phishing research operation — 11 min read

Let's be brutally honest about something the anti-fraud industry doesn't like to say out loud: most scammers are not criminal masterminds. They are not sophisticated. Most of them are barely competent.

Stop chasing the foot soldiers

We have tracked threat actors for years. We watch them mutate from one grift to the next: stealer operator today, drainer developer tomorrow, carding shop next quarter. We watch them migrate their fake identities and infrastructure — Switzerland this month, Turkey the next. They don't hide well. They leave trails a first-year analyst could follow. Most of them ended up in this "business" for one simple reason: they had no other exit and not enough brains for anything legitimate.

So why does the problem keep growing? Because the population of bad actors is enormous, and hunting them one by one is mathematically pointless.

Look at our own Telegram intelligence platform. Nearly 130 thousand malicious actors in our dataset alone — and that is one dataset.

4,678
tracked Telegram sources (4,394 alive)
12.4M+
messages collected
129,718
unique malicious actors profiled
10×
replacements spawned per takedown

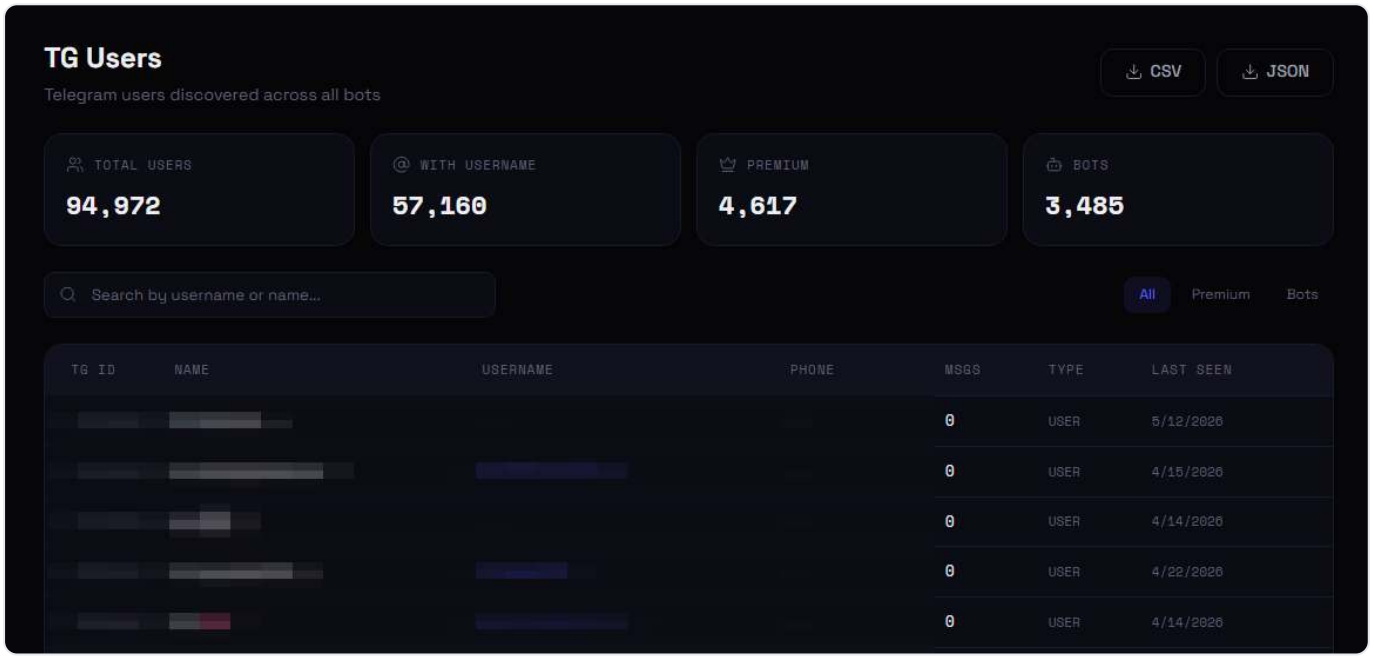

*A single sweep of monitored Telegram bots: **94,972** discovered accounts across **3,485** bots. Individual identifiers are **redacted** — surfacing in a monitored bot's user list is a lead, not a verdict .*

Add the thousands more identified by past investigations from other research teams, and the picture is clear: you cannot arrest your way out of a population this size. Take down one actor, one group — through months of legal work, cross-border jurisdiction nightmares, and case amounts that rarely interest law enforcement — and ten replacements spawn the same week. The individual scammer is a cheap, replaceable node.

The math only changes when you change the economics. As long as scam stays cheap — cheap domains, no-questions-asked registrars, crypto payments with zero verification — there will be a lot of it. The only strategy that scales is making scam expensive. And that leads directly to the people who currently make it cheap:

The registrars.

> **The numbers say it's a crisis. The industry pretends it's weather.**

This is not our opinion. This is the documented state of the internet:

FBI IC3 · 2024

$16B lost, +33% YoY

The FBI's IC3 received **859,532 complaints** in 2024 with reported losses exceeding **$16 billion** — a 33% increase from 2023 — and the top cybercrime by complaint volume was phishing/spoofing. Nearly 150,000 complaints involved digital assets, amounting to **$9.3 billion** in losses — a 66% increase from the previous year.

Interisle · 2025

Phishing up 180% since 2021

Interisle Consulting's annual study , analyzing nearly four million phishing reports between May 2024 and April 2025, found reported phishing reaching **nearly two million attacks** — an increase of over 180% since 2021. The same report highlights the ease with which criminals exploit **permissive industry policies and business practices** to acquire domain names.

Read that last line again. The leading independent research in the field says exactly what we say: the enabler is permissive industry practice. Not scammer genius. Permissiveness. And concentration proves it: Interisle found that in one campaign, a sample of 37,000 Unpaid Toll Scam domain names showed **65% were registered through a single Chinese registrar** .

The scam economy is not distributed randomly across the internet. It clusters around specific registrars — because scammers shop for silence.

Registrars are not neutral. They are paid participants.

There is a comfortable myth that registrars are "neutral infrastructure" — passive utilities with no stake in what happens on the domains they sell.

This is a lie. A registrar is a for-profit party that directly earns revenue from every phishing domain it declines to suspend. Every ignored abuse report is protected revenue. Every "we can't judge content" reply is a business decision dressed as a legal position.

Scammers don't choose a registrar for the pricing page or the design. They choose the registrar that won't ask questions, won't suspend domains, will take crypto without verification, and will sit silently on abuse reports. That is a market. That is a product. And the industry's own data shows who sells it. Per Interisle's findings , the top five gTLD registrars by raw phishing volume were **NameSilo, GoDaddy, GMO (Onamae), PublicDomainRegistry, and NameCheap** ; normalized by domains under management, the most abused were **NiceNIC, URL Solutions, Aceville, WebNic, and OwnRegistrar** — with NiceNIC seeing 45% of their gTLD portfolio reported for phishing.

Forty-five percent of a portfolio reported for phishing. At what point does "neutral infrastructure" become "criminal supply chain"?

Who regulates the registrars? Technically ICANN. Practically nobody.

Above the registrars sits ICANN. Above ICANN sits nobody.

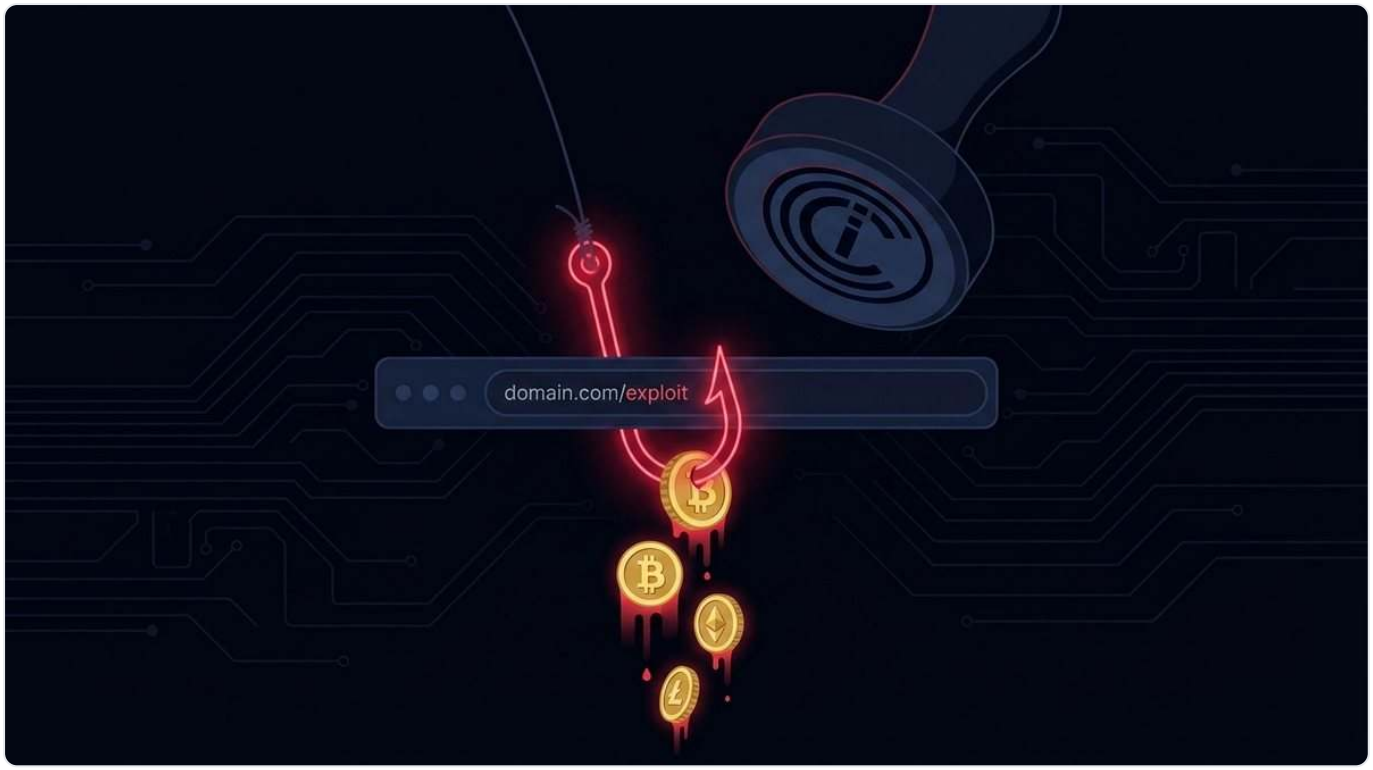

Accreditation as a rubber stamp: RAA §3.18 obliges registrars to act on abuse — but a compliance regime where compliance is optional isn't regulation.

ICANN's 2013 Registrar Accreditation Agreement , **Section 3.18** , contractually obligates every accredited registrar to maintain an abuse contact and take reasonable and prompt steps to investigate and respond appropriately to abuse reports. On paper — an enforceable duty.

In practice? From what we see on the front lines daily, we estimate a single abuse-friendly registrar violates 3.18 at minimum **~1,000 times per year** — ignored reports, phantom "investigations," phishing domains left alive for weeks while victims bleed. Scale that across the industry over five years and the true count of 3.18-class failures plausibly reaches a million ignored or mishandled reports. Nobody knows the real number, because nobody is required to count. That is the point.

And ICANN's enforcement toolkit? Effectively one nuclear button: terminate the accreditation. No fines. No graduated penalties. No victim restitution. And that button gets pressed almost exclusively against registrars that are already dead — shells that stopped paying fees. Why? Because ICANN is funded by fees from the very registrars it "regulates." Cutting off a large, profitable, high-abuse registrar means cutting off its own revenue. It is structurally incapable of wanting to enforce.

ICANN took the regulator's chair nominally — precisely to preserve its independence and keep governments out. The result: a rule no registrar fears. Yes, there are precedents — compliance notices and warnings (WebNIC's case among them). Look at the consequences. A sternly worded letter versus a business model generating millions. **A compliance regime where compliance is optional isn't regulation. It's theater.**

Case study: NameSilo — "fastest growing," fastest ignoring

NameSilo is a publicly traded company. It files reports for the stock exchange showing millions in revenue and markets itself as one of the fastest-growing registrars in the world. It also sits, per

independent research cited above, at the top of the raw phishing-volume rankings. Coincidence? Here is what we witnessed personally:

Documented

Twenty reports, then "no prior reports"

We reported a phishing site to NameSilo. Not once — **at least 20 documented reports** on that single domain, with evidence, screenshots, scans, and analysis. Nothing happened until we escalated the case publicly on Twitter/X. NameSilo's public response: that they had "received no prior reports." A publicly traded registrar, facing documented evidence of twenty ignored abuse reports, publicly claimed the reports didn't exist. That is either a broken abuse pipeline they refuse to fix, or a public lie. Neither is acceptable — and both are profitable.

How many requests for US users' protection did NameSilo directly ignore to avoid losing revenue? We'll answer from experience: not just many — we believe the majority. And there is zero accountability for it.

It gets better. NameSilo has publicly described helping a client — a client they claimed to have received no complaints about — get VirusTotal detections removed. Follow the logic:

- **Asked to suspend a phishing domain:** "We're not qualified to determine what's malicious."
- **Helping a paying client:** suddenly qualified to overrule detection verdicts from Fortune 500 security vendors.

You cannot claim technical incompetence as a shield and technical authority as a service. Pick one. This isn't an accident and it isn't an isolated case. It's a pattern: love the money, do nothing, answer to no one.

What an ICANN accreditation is actually worth: Trustname, right now

The conflict of interest is structural and simple: **a regulator should never draw its revenue directly from the parties it regulates.** ICANN does — and we understand its roughly 400 staff are, no doubt, occupied with matters more pressing than the safety of the users whose wallets these domains are built to drain.

Here is what that accreditation is worth in practice, live today. Consider Trustname (Fewmoretaps OÜ) — a registrar we documented in detail . On paper it is an Estonian company; operationally it is run out of **Belarus** by a single owner-operator — its 100% shareholder, CEO and only employee. Declared 2024 revenue: **€120** . It began selling domains in earnest only this year. And its ICANN accreditation (IANA #4318) is **still active even with the company in liquidation** . Of the ~7,641 domains it manages, **86% of the live ones are confirmed malicious** .

That is who ICANN handed a badge of trust. And this is not our characterization of how it handles abuse — it is Trustname's **own published abuse policy** :

> "According to ICANN guidelines, in most cases, to take action against a domain the registrar must be provided with a valid court order or the registrant's consent."

> “Reports received solely through automated threat-analysis systems are treated as investigative leads, not conclusive evidence… Trustname may also request confirmation that the alleged abuse remains active at the time of review.”
>
> “Abuse reports sent from free or anonymous email services (e.g., @gmail.com, @proton.me) are subject to additional verification.”
>
> “We do not determine the legality of a site’s content and will act only upon the direction of law enforcement or in clear cases of violation of our terms.”

Read it as a system and the design is unmistakable: a valid **court order** to touch a domain; real detections downgraded to mere “leads”; a demand that you re-confirm the scam is *still live* at review time — after they have stalled long enough; extra “verification” for anyone reporting from Gmail or Proton; and a blanket refusal to judge content, politely redirected to “contact your hosting provider.” Wrap that in registrar-owned privacy shields that take crypto and discard physical mail, and you do not have an abuse process. You have **bulletproof-by-paperwork** — and it is written on the letterhead of an ICANN-accredited registrar.

This indicts ICANN directly. A one-person, €120-revenue, Belarus-run entity — in liquidation — holds a live accreditation and openly publishes a policy engineered never to suspend anything. The strongest tool ICANN has in response is a letter; the response we have seen to such letters is a Terms update, not a suspension. **A regulator whose ultimate sanction is a letter — answered with a Terms edit — is not a regulator.** Meanwhile, every day, real people lose money on `.com` and other domains that pass through operators like this while the paperwork loops. Three court orders to take down a live phishing page is not due process; it is a farce, and it is a farce ICANN underwrites by accrediting the operators who demand it.

Trustname is not a freak accident; it is what happens when accreditation is cheap and unaccountable. Several EU jurisdictions — Estonia among them (home to a long list of such shells, Keitaro included), and the UK with its ultra-cheap corporate formation — sell corporate and service access at a price that strips out any real vetting. The result is an environment of built-in distrust, even for legitimate businesses, because nobody is checking and nobody is accountable. We have hit variants of the same court-order wall with certain country-code operators too; those we will document separately.

// The regulator that isn’t

What a functioning regulator has

- Graduated fines that scale with the harm done
- Penalties short of the death penalty
- Restitution routed to the victims
- Independent audits — not self-attestation
- Funding independent of the regulated
- Power to stop a live harm in hours

What ICANN has

- One sanction: a strongly-worded letter

- Terminate-only — used mostly on already-dead shells
- Zero fines. Zero restitution.
- Honor-system compliance and self-attestation
- Funded by fees from the registrars it “regulates”
- Abuse left alive for weeks behind “court-order” demands

€120 Trustname’s declared 2024 revenue
86% of its live domains, malicious
$0 fines ICANN can levy
1 enforcement tool — a letter
ICANN didn’t fail to regulate the domain economy.
It was never built to.

The fix is obvious: fines and liability

We are not asking for censorship powers. We are asking for what every other industry already has: financial consequences for contractual negligence.

Graduated fines

- A registrar ignores three substantiated abuse reports — evidence, scans, analysis attached — on the same malicious domain? Automatic penalty.

Liability to victims

- Fines payable to victims, or to the state where the crime occurred.
- If a US user loses funds to a phishing domain the registrar was warned about and did nothing — the registrar shares liability for failing its explicit duty to mitigate harm.

Public transparency

- Mandatory public abuse-report transparency — reports received, response times, actions taken — published quarterly and auditable.
- NameSilo could never claim "no prior reports" against a public log.

Independent audits

- Independent audits of abuse handling as a condition of accreditation renewal — not self-attestation.

The registrars will cry that they "can't be expected to determine what is phishing." Counter-question: you're capable of taking the money and signing the accreditation obligations — but incapable of honoring them? If a registrar genuinely cannot evaluate an abuse report with evidence attached, it has no business holding an accreditation that contractually requires exactly that.

What should replace it: an open, real-time abuse ledger

We are not proposing a second ICANN — we do not need one. Two of the internet’s load-bearing systems already show the shape of the answer: **TLS/SSL certificate transparency** and **WHOIS** — open, queryable, public records. Abuse handling should work the same way: **a clear, public registry of every abuse report a registrar receives and exactly what it did in response** — timestamped, auditable, and impossible to deny after the fact.

With that ledger in place, most of today's theatre collapses:

- **No offshore court order to stop a live threat.** A phishing page draining wallets is a here-and-now harm; blocking it should not require a ruling from a court on some island in an offshore zone. The evidence is in the report, and the public record shows whether the registrar acted.
- **Triage can be automated.** Registrars insist they are drowning in "attacks" and junk reports — which is exactly the kind of filtering an **AI layer** does well: separate the noise, surface the substantiated cases, and log every decision. A simple, honest proxy in front of the abuse pipeline would already outperform the current human non-process. You do not need an ICANN #2 for that.
- **The excuse expires.** A registrar is paid on every domain it sells; handling abuse is the other half of that transaction, not a courtesy. NameSilo's "we're not qualified" routine is a business decision, not a technical limit — and a company that genuinely cannot tell a phishing kit from legitimate traffic should be running its other businesses, not underwriting scam and phishing operations by default.
- **No more "we never received a report."** WebNic and NiceNic answered us with auto-replies demanding a screenshot that was already attached to the report; NameSilo told the public there had "never been a single complaint." None of those moves survive contact with a public, timestamped ledger.

And the ledger makes the one consequence that matters *enforceable* . When a victim loses funds to a domain **days after** a substantiated report was already on file — and the public record shows the registrar sat on it — that is no longer an unfortunate grey area. It is a documented failure with a timestamp on it, and the registrar should carry financial liability for that specific incident.

Put a price on the wallet a registrar let drain after we warned it, and every registrar on earth discovers competence overnight.
That is the whole mechanism. The moment ignoring a report costs more than acting on one, the "we're not qualified" routine ends, abuse desks suddenly find budget and staff, and the scammers who shop for silence discover there is none left to buy — and quietly abandon the domains they were counting on.

Transparency is the fix. You cannot quietly ignore a report that everyone can see you received.

Bottom line

$16 billion in reported losses in a single year. Phishing up 180% since 2021. 130,000 malicious actors in one monitoring dataset alone. And the entire pipeline still starts the same way: a cheap domain, sold by a registrar that will never answer for its silence.

The scammer is the symptom — a replaceable, mostly dim-witted node in an economy of desperation. The registrar taking his crypto with a smile, ignoring twenty abuse reports, and then lying about it publicly — that is the disease. And ICANN's decorative "regulation" is the immune system that decided not to show up.

Scam will remain massive exactly as long as it remains cheap. And it will remain cheap exactly as long as registrars stay silent — and pay nothing for their silence.

We will keep naming names. We will keep publishing evidence. Registrars are not neutral. Impunity is the business model. It's time to make it expensive.

Sources & references

1. FBI Internet Crime Complaint Center (IC3) — 2024 Internet Crime Report (859,532 complaints; $16.6B reported losses; phishing/spoofing #1 by complaint volume).
2. Interisle Consulting Group — Phishing Landscape 2025 (≈2M phishing attacks; +180% since 2021; registrar/TLD concentration; 37,000-domain toll-scam sample).
3. ICANN — 2013 Registrar Accreditation Agreement, §3.18 (abuse point of contact; duty to take reasonable and prompt steps to investigate and respond to abuse).
4. ICANN Contractual Compliance — Notice of Breach to Web Commerce Communications (WebNic), 29 Jul 2025 .

Figures are drawn from the primary sources above; interpretation and commentary are the authors'.

Name the enabler. Price the silence.

— **The PhishDestroy Team** — An independent anti-phishing research operation