TECHNICAL WHITEPAPER · PROPOSED FRAMEWORK · VERSION 1.0

From Detection to Enforcement

A measurable architecture for DNS abuse response: evidence provenance, proportionate action, public accountability and rapid appeal — designed to be tested rather than merely asserted.

Published 2 August 2026 · PhishDestroy Research · Public consultation draft · Not peer reviewed

PD-WP-2026-01 · Version 1.0 · Published 2 August 2026
Permanent URL: https://phishdestroy.io/dns-abuse-enforcement-framework

ABSTRACT

INDEPENDENT PROPOSAL · NOT AN ICANN PROGRAMME

Direct conclusion. The gTLD contractual ecosystem contains a measurable gap between technically validated malicious use of a domain and effective mitigation. Published studies show that many phishing campaigns operate for hours or days, while contractual-compliance procedures can unfold over business-day timescales. The disclosed scenario model shows how delay increases expected campaign returns; the 60–85% range is the paper's central testable estimate of post-notice harm, not an established statistic for all global fraud. The $100–150 engineering estimate concerns collection and first-pass processing of a defined corpus of zone, Certificate Transparency and public-feed data, not the full cost of legal review or enforcement. The paper proposes signed evidence, measurable response targets, automatic escalation of validated cases, proportionate remedies, public accountability and rapid appeal. It does not propose giving reporters direct access to domain-suspension controls.

PRIMARY SOURCE
Directly supported by an official document or published study.

ILLUSTRATIVE MODEL
A reproducible calculation whose inputs are assumptions, not market averages.

POLICY PROPOSAL
A design offered for consultation, testing and revision.

01 Research question, scope and standard of proof

The question is narrower than "who failed?": **which observable delays between validated evidence and proportionate action are preventable, and which intervention reduces victim harm without creating unacceptable false positives or collateral damage?**

The scope is the contracted gTLD ecosystem: registrants, resellers, registrars, registries and ICANN's contractual compliance function. Hosting providers, CDNs, browsers, payment services, advertising platforms, national CERTs and law enforcement are included only where they affect the evidence or response chain. Country-code TLDs have different governance and are not assumed to be subject to ICANN's gTLD contracts.

Status of the central theses

- The paper analyzes ICANN as a private governance and contractual-compliance institution within the gTLD ecosystem. ICANN is not a government regulator or law-enforcement body, but it accredits registrars, contracts with registry operators and enforces those agreements; that limited remit does not remove institutional accountability.
- The 60–85% range is a central quantitative estimate in the paper's harm model. It is presented as a testable estimate derived from explicit assumptions and case observations—not as an already established statistic for all global fraud. Publishing an incident-level dataset that aligns detection, notification, mitigation and harm timestamps is the next evidentiary step.
- The $100–150 figure is a testable engineering budget estimate for continuous collection and initial triage over a defined set of accessible zone, Certificate Transparency and public-feed data on commodity infrastructure. It is not the cost of complete global visibility, legal review, appeals, high availability or enforcement. The paper therefore treats it as a benchmark to reproduce with a published workload, source inventory and cost log—not as an already proven universal price. That distinction narrows the claim without erasing the cost asymmetry being tested.
- The paper's three-week comparison concerns compliance latency, not a universal domain-takedown SLA. ICANN's Contractual Compliance FAQ says that most complaint types can pass through three successive five-business-day response periods before formal escalation, while many malicious campaigns operate for hours or days. Contracts separately require prompt and appropriate mitigation after actionable evidence. [2] [4]
- For a verified, single-purpose malicious registration, `serverHold` is among the most disruptive DNS-level remedies because it removes delegation from the zone. It is not universal: it can also disrupt mail, subdomains and legitimate services; cached answers may remain until TTL expiry; and compromised or shared services can require narrower mitigation. [4] [7]

> RESEARCH CONTRIBUTION
>
> The proposal is falsifiable. A pilot can compare treatment and control cohorts on time-to-triage, time-to-mitigation, victim-loss proxies, recurrence, false positives, reversal time and collateral impact.

02 What the primary sources establish

$165.1M

ICANN FY27 operations funding

Adopted FY27 budget. Contextual scale; not DNS-abuse spending. [1]

191,561

2025 IC3 phishing/spoofing complaints

Largest IC3 complaint category by count; $215.8M directly assigned losses. [10]

24 h

Median lifetime in one 2021 study

Dataset-specific result across 1,288 phishing domains; not a universal lifetime. [11]

No fixed SLA

Universal takedown deadline

The 2024 amendments use “prompt” and “appropriate” mitigation, not one global deadline. [2] [3]

PROPOSITION	STATUS	WHAT CAN BE STATED
Phishing is DNS Abuse	VERIFIED	Phishing is explicitly one of the five categories covered by the 2024 gTLD contractual amendments. [2]
Registrars must act	VERIFIED	Registrars must act promptly when they have actionable evidence that a sponsored name is being used for DNS Abuse. [2]
Registries must act	VERIFIED	Registry operators must act promptly where they reasonably determine, based on actionable evidence, that a registered name is being used for DNS Abuse; the appropriate measure remains context-dependent. [3]
`serverHold` removes delegation	VERIFIED	It is a server-set EPP status. Cached DNS answers may remain until expiry, and the action affects the whole registered domain. [7]
All abuse can be seen in CZDS or CT	FALSE	CZDS and Certificate Transparency are valuable sensors, not complete or real-time views of all domain use. [8] [9]

ICANN's first six-month enforcement report under the 2024 amendments recorded 192 investigations, more than 2,700 suspended domains, more than 350 disabled phishing pages and two Notices of Breach. Those figures do not prove that present controls are sufficient, but they rule out the absolute claim that the system takes no action. [5]

03 Why response latency remains a legitimate research target

Some phishing campaigns operate on a timescale of hours or days. Bijmans et al. reported an average observed lifetime of 45 hours and a median of 24 hours for 1,288 phishing domains in one 2021 study. [11] A 2026 study of newly registered phishing domains reported a highly skewed distribution: for 14,112 domains with measurable lifetimes, the mean was 8.6 days and the median one day. [12] Neither sample can be generalized to every threat class, but both show why a process measured only in business days may miss short-lived campaigns.

THE MEASURED TEMPORAL MISMATCH

These studies do not establish the share of financial harm that occurs after notice, but they do establish a temporal mismatch: a process whose early informal-compliance stages can consume up to fifteen business days cannot by itself serve as incident response for campaigns with median lifetimes near one day. The 21-day point below is therefore a compliance-latency scenario, not a label for a universal “ICANN SLA.”

FIGURE 1 · ILLUSTRATIVE CUMULATIVE CAMPAIGN OUTCOME

MODEL · LOG-LIKE TIME SPACING

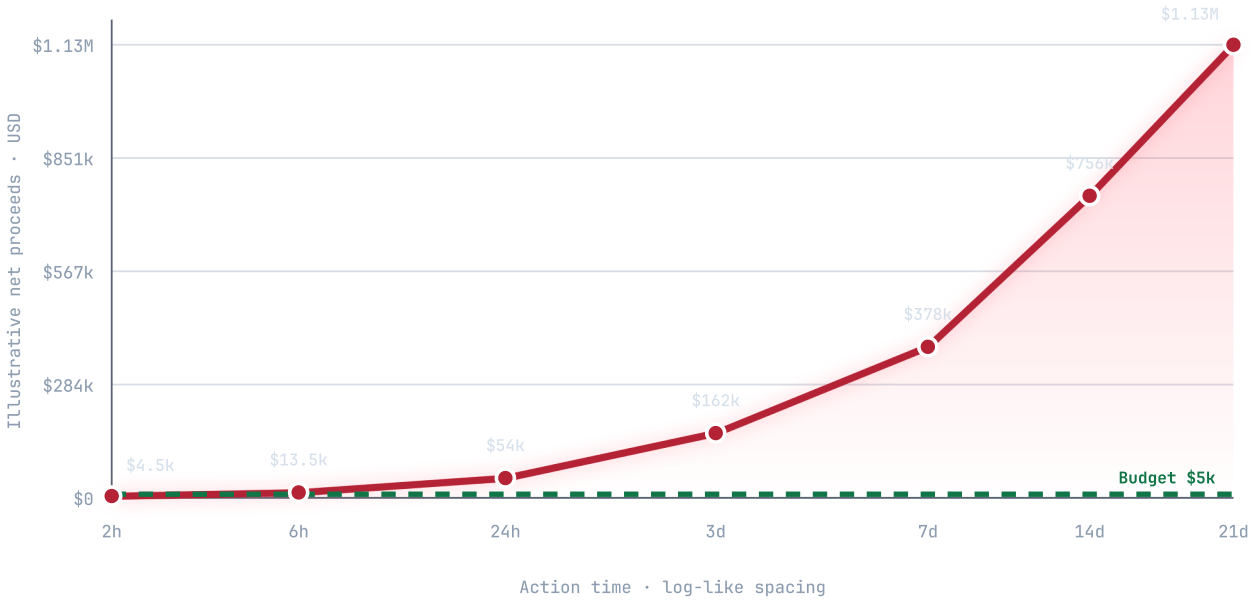

The curve visualizes the scenario in Section 4; it is not observed market data. Horizontal placement compresses long intervals for readability. Source: reproducible model with disclosed inputs.

CAUSAL CAUTION

A shorter domain lifetime does not automatically equal prevented financial loss. Attackers may migrate, use compromised sites, change delivery channels or cash out before detection. The causal effect must be measured, not inferred from takedown speed alone.

04 Campaign economics as a reproducible scenario

The economically relevant asset is often not the registration fee but the traffic, creative, infrastructure and trust accumulated around a domain. That intuition is plausible; its magnitude varies by campaign. The calculator below therefore exposes every input instead of presenting a hypothetical result as an observed average.

MODEL DEFINITION

```
visits(T)   = daily_traffic × T / 24
victims(T)  = visits(T) × conversion_rate
proceeds(T) = victims(T) × average_loss × (1 - cashout_attrition)
profit(T)   = proceeds(T) - campaign_budget
```

Scenario calculator

Change any assumption. Values are illustrative and are not estimates of the typical phishing market.

ILLUSTRATIVE MODEL

Campaign budget (USD)

5000

Daily visits

2000

Victim conversion (%)

1.5

Average loss per victim (USD)

2400

Cash-out attrition (%)

25

Action time (hours)

24

2,000

30

$54,000

$49,000

2.2 h

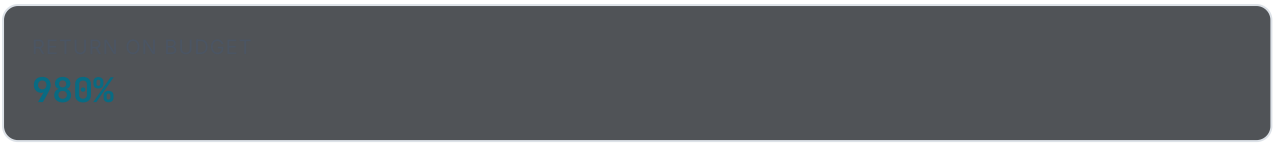

Expected victims may be fractional because the model expresses an expectation. It omits traffic decay, detection feedback, migration, repeat victims and uncertainty distributions.

▶ View model source table

05 Measuring preventable harm without inventing precision

“Loss after detection” can become a useful outcome only if detection, delivery and harm are defined consistently. A blacklist timestamp is not proof that the registrar received actionable evidence. A wallet transaction is not automatically attributable to one domain. A domain becoming unreachable does not identify which actor caused the change.

Minimum event schema

EVENT	REQUIRED TIMESTAMP	MINIMUM PROOF
First observation	`t_obs`	Sensor, request/response hash, capture method, clock source.
Technical validation	`t_valid`	Reproducible indicator, reviewer identity, threat category and confidence.
Notice delivered	`t_notice`	Authenticated delivery receipt and evidence bundle version.
Responsible actor acknowledges	`t_ack`	Case ID and destination role: host, registrar, registry, platform or authority.
Mitigation observed	`t_mitigate`	DNS, HTTP, payment or platform state, measured from multiple vantage points.
Appeal / restoration	`t_restore`	Decision basis, changed evidence and restoration verification.
Harm event	`t_harm`	De-identified victim report, attributable transaction or another documented proxy.

PROPOSED OUTCOME METRIC

$$\text{Post-notice harm share} = \frac{\text{attributable harm with } t_{harm} > t_{notice}}{\text{all attributable harm in the observation window}}$$

The numerator and denominator must use the same attribution rule. Results should report the cohort, confidence interval, right-censoring, missing sources, duplicate-victim handling and sensitivity to the observation window. The label “preventable” should be reserved for a causal design or, at minimum, a matched comparison — not every event after notice.

CENTRAL MODELED ESTIMATE · 60–85%

The paper retains 60–85% as its central estimate for the share of attributable harm occurring after detection and notice in the class of incidents being modeled. It is a scenario estimate derived from disclosed assumptions and case observations—not a measured statistic for all fraud worldwide. It should be used for sensitivity analysis until an incident-level dataset aligns observation, validated evidence, delivery, mitigation and harm timestamps. Publishing that dataset is the evidentiary test of the thesis, not a reason to delete the thesis.

06 The institutional map: capability is distributed

No single actor sees or controls the full incident. The registrar knows the registrant relationship; the registry controls registry-set domain status; the host and CDN control content delivery; browsers and security vendors control warnings; payment and wallet providers can interrupt transfers; authorities can compel preservation or seizure. ICANN writes and enforces contracts in the gTLD ecosystem but does not operate a universal EPP control plane.

01 · SIGNAL

Researchers & sensors

Observe content, infrastructure and victim indicators; preserve provenance.

02 · SERVICE

Host / CDN / platform

Can remove content or access rapidly, often without changing the domain.

03 · SPONSOR

Registrar

Reviews abuse, contacts the registrant and can set client-side status.

04 · ZONE

Registry operator

Controls registry-set EPP status and the TLD zone delegation.

05 · CONTRACT

ICANN Compliance

Investigates contracted-party compliance; does not adjudicate every fraud case.

06 · LAW

CERT / authorities

Coordinate response and use jurisdiction-specific legal powers.

Distributed capability does not mean dissolved responsibility. A sponsoring registrar does not shed its contractual duties by delegating abuse handling to a reseller; the registry controls server-set EPP status; ICANN enforces contractual compliance. The audit trail must record not only who acted, but who received actionable evidence, redirected the case, failed to respond or failed to perform an assigned duty. Legal authority remains with the actor empowered to apply a measure; accountability also covers documented non-performance.

Funding concentration and a testable conflict-of-incentives hypothesis

$165.1M

FY27 ICANN Operations funding

Adopted FY27 base-case budget. [1]

≈98.1%

registry and registrar fee streams

Approximately $162.0M of the $165.1M funding plan. [1]

70.7%

transaction-based fees alone

$116.8M tied to billable domain transactions. [1]

The funding structure is therefore highly concentrated in payments from the contracted parties whose agreements ICANN enforces, and a large share varies with billable registration transactions. That establishes a structural dependency and a legitimate conflict-of-incentives question. It does *not*, by itself, prove intentional under-enforcement or completed “regulatory capture.” The capture hypothesis should be tested against case-level response times, enforcement outcomes, sanctions, repeat abuse and the role of funded stakeholders in policy and implementation decisions.

07 The 2024 contractual baseline — and what it does not specify

---

The April 2024 global amendments to the Registrar Accreditation Agreement and Base Registry Agreement created express obligations to mitigate DNS Abuse. Registrars must act promptly when they have actionable evidence that a sponsored name is being used for DNS Abuse. Registry operators must act promptly where they reasonably determine, based on actionable evidence, that a registered

name is being used for DNS Abuse. In both cases the appropriate measure depends on context, severity and collateral harm. [2] [3]

Three timelines that are often conflated

TIMELINE	WHAT IT APPLIES TO	WHAT IT DOES NOT MEAN
24 hours	Review, under RAA §3.18.3, of well-founded reports of Illegal Activity submitted to the dedicated contact by law-enforcement, consumer-protection, quasi-governmental or similar authorities.	It is not a universal 24-hour suspension deadline.
"Promptly"	Taking appropriate mitigation after the applicable actionable-evidence threshold is met.	It is not a fixed number of hours and does not require the same remedy in every case.
21 days	Can appear as a cure period after a formal contractual Notice of Breach.	It is not the ordinary lifetime granted to every reported malicious domain.

This flexibility has benefits: a compromised university domain should not be handled like a newly registered single-purpose phishing domain. The accountability problem is that "prompt" and "appropriate" are difficult to compare without published timestamps, case categories and outcome codes. The proposal below adds measurement and review without pretending that one deadline fits every case.

08 The boundary problem: phishing is covered; some financial deception may not be

Phishing is explicitly within the contractual definition of DNS Abuse. The harder boundary concerns sites that deceive users under an original name rather than impersonating an identifiable third party: some fake investment platforms, fraudulent stores, recovery scams and wallet-signature schemes. Depending on facts, these may be treated as website-content abuse, consumer fraud or another legal category rather than contractual DNS Abuse.

> PROPOSED RESEARCH LABEL
>
> **Verified Financial Harm Abuse (VFHA)**: documented use of a domain to obtain funds, payment credentials, private keys or seed phrases through deception, regardless of brand impersonation. VFHA is not current ICANN terminology and does not itself create contractual authority.

A policy consultation could ask whether a narrowly evidenced category like VFHA belongs in future contracts, a cross-sector referral framework, or national law. Expansion should require a precise harm test, reliable evidence, proportional remedies, jurisdictional review and appeal. Vague "scam" labels are insufficient.

09

What the $100–150/month monitoring claim must prove—and what it cannot prove

A commodity server can download accessible snapshots of participating gTLD zone files and compute changes locally, ingest selected Certificate Transparency events and open feeds, normalize strings, compute hashes and prioritize candidates. That is a useful, testable engineering benchmark—not "complete monitoring of the Internet." In Q2 2026, Verisign reported 401.6 million registrations across all TLDs; CZDS covers participating gTLD zone files rather than all TLDs, and CT records publicly logged certificates or precertificates rather than every active or harmful domain. [8] [9] [17]

CAPABILITY	COMMODITY PROTOTYPE	PRODUCTION PUBLIC-INTEREST SERVICE
Zone/CT/feed ingest	Feasible for a defined source set	Redundant collectors, source contracts, gap monitoring
String/hash triage	Feasible	Benchmarking, drift detection, false-negative studies
Browser rendering	Small sampled subset	Isolated fleet, regional vantage points, malware containment
Legal determination	Not included	Qualified reviewers, jurisdiction and authority mapping
High availability / evidence retention	Usually absent	HSM-backed signatures, audit logs, backups and incident response
Appeals and restoration	Not included	24/7 operations, independent escalation and service targets

FIGURE 2 · ORDER-OF-MAGNITUDE CONTEXT — LOGARITHMIC WIDTH · NOT LIKE-FOR-LIKE COSTS

Logarithmic widths keep small values visible. The first three figures are disclosed engineering or programme assumptions that require a published benchmark and are not vendor quotations. ICANN Operations funding is a verified institutional-scale figure. These are deliberately non-equivalent budgets: the comparison tests whether basic collection and triage are an inherently prohibitive cost, not whether a complete enforcement programme costs $100 per month.

10 Proposed Verified Abuse Response Framework

The proposed framework is a reference architecture, not an automatic global kill switch. It standardizes the case envelope, decision record and timestamps while allowing the authorized operator to choose the least disruptive effective action.

01 · INGEST

Acquire & preserve

Collect reports and observations; hash raw artifacts; synchronize clocks.

02 · PROVENANCE

Normalize evidence

Record origin, method, handling history and source dependencies.

03 · VALIDATE

Corroborate

Reproduce the harmful behavior and distinguish independent evidence.

04 · ASSESS

Risk & collateral

Classify malicious registration, compromised service, shared hosting and criticality.

05 · RESPOND

Notify & mitigate

Assign the responsible actor, target time and least disruptive effective remedy.

06 · REVIEW

Audit & appeal

Publish outcome metadata, measure impact and support rapid reversal.

Evidence tiers

E1 **Unverified signal**

Single feed, lexical match or reputation claim. Suitable for observation, never sufficient alone for suspension.

E2 **Reproduced behavior**

Timestamped capture showing credential theft, malware delivery or a deceptive transaction path.

E3 **Independent corroboration**

At least two genuinely independent sources or one reproducible technical proof plus ownership/context checks.

E4 **Qualified confirmation**

Authority, affected service, brand owner or validated victim evidence, with privacy-safe chain of custody.

INDEPENDENCE RULE

Three feeds that copy the same upstream blacklist are one source, not three. The provenance graph must expose shared origin, synchronization and vendor re-publication.

11 A minimal evidence envelope and interoperable API

The API should create a verifiable case and start an auditable response timer; a reporter must not be able to issue a hold command directly. If an E3/E4 case concerns a single-purpose malicious registration and expires without a reasoned decision or effective mitigation, the system automatically escalates it to the registry operator or another actor with contractual or legal authority. A `serverHold` command can originate only from an authorized actor. Automatic escalation and registry authority to act on expiry are proposed policy changes, not claims about existing authority. Every state transition is signed and appended to the audit log.

POST /V1/CASES · EXAMPLE

```
{
  "indicator": {"type": "domain", "value": "example.invalid"},
  "alleged_category": "credential_phishing",
  "observed_at": "2026-08-02T06:14:22Z",
  "evidence": [
    {
      "type": "http_capture",
      "sha256": "a3f1c9…",
      "collection_method": "isolated_browser_v2",
      "source_id": "reporter:ed25519:7c2a…"
    }
  ],
  "provenance_graph": "ipfs-or-object-store:sha256:91bd…",
  "collateral_context": {
    "registered_domain_scope": true,
    "shared_service": false,
    "mail_observed": true,
    "suspected_compromise": false
  },
  "requested_action": "review",
  "reporter_signature": "ed25519:…"
}
```

A successful response returns a case ID, the receiving actor, evidence tier, completeness errors, target review time and a public transparency URL. It must not promise a specific remedy before a competent actor evaluates authority and collateral impact.

12 Proposed response targets, not universal takedown deadlines

Service-level objectives should separate *acknowledgement*, *triage*, *decision* and *effective mitigation*. This makes performance measurable without assuming that suspension is always the right outcome.

CASE PROFILE	ACKNOWLEDGE	TRIAGE TARGET	DECISION TARGET	TYPICAL RESPONSE OPTIONS
E3/E4 payment credential or seed-phrase theft on a single-purpose malicious registration	15 min	1 h	4 h	Content block; registrar/registry hold where authorized; browser and payment warnings.
E3 brand phishing or malware delivery	30 min	2 h	6 h	Host/CDN removal, domain mitigation, sinkhole or warning according to control point.
Compromised legitimate domain or shared SaaS tenant	1 h	4 h	12 h	Path/account isolation and owner recovery; avoid registered-domain suspension where possible.
Financial deception with contested classification	4 h	12 h	24 h	Preservation, platform/payment intervention, authority referral and reasoned decision.
E1 signal or incomplete report	Auto	24 h	None until validated	Monitor, enrich and request evidence.

These are pilot targets. The correct values should be derived from observed threat lifetimes, staffing, error cost and legal constraints, then published with attainment distributions rather than a single average.

13 Safety, due process, privacy and failure modes

Fast response without safeguards can amplify bad data into censorship, commercial sabotage or infrastructure outages. A legitimate architecture therefore treats false positives and collateral harm as first-class security failures.

FAILURE MODE	REQUIRED CONTROL	AUDITABLE MEASURE
Malicious report against a competitor	Reporter authentication, source reputation, independent reproduction and sanctions for abuse	Rejected-report rate by reporter; confirmed manipulation cases
Feed circularity masquerading as consensus	Provenance graph and upstream-source de-duplication	Independent-source count before/after lineage correction
Compromised legitimate domain suspended wholesale	Registration-intent classification and path-level mitigation preference	Compromised-domain share; affected legitimate services
Victim data or exploit details leaked publicly	Tiered disclosure, redaction, sealed evidence and retention limits	Privacy incidents; redaction review outcomes
Wrongful action persists	24/7 appeal intake, independent reviewer and authenticated restoration path	Median and 95th-percentile reversal time
Registry hold causes email/subdomain outage	Collateral inventory, remedy proportionality and post-action monitoring	Services disrupted per action; rollback rate

Minimum due-process guarantees

1. Reason code, evidence tier and responsible decision-maker are recorded for every restrictive action.
2. A registrant can obtain a privacy-safe statement of reasons and submit counter-evidence.
3. Urgent appeals are reviewed by a person who did not make the original decision.
4. Reversal propagates through the same signed channel as the original action.
5. Aggregate error and restoration statistics are public; sensitive evidence remains access-controlled.

14 From a “toxicity score” to an accountable Response Quality Index

A public registrar score can improve accountability, but naïve rankings are distorted by portfolio size, feed coverage, customer mix and whether domains were maliciously registered or later compromised. A score must never justify collective blocking of all customers of a registrar.

PROPOSED QUALITY VECTOR

```
RQI = report as a vector, not a single opaque rank:

  coverage-adjusted validated-abuse rate
  median / p90 acknowledge, triage and decision times
  mitigation effectiveness and recurrence
  false-positive and reversal rates
  median / p95 appeal-resolution time
  evidence completeness and transparency rate

Always publish N, observation window, confidence intervals,
feed coverage and malicious-registration / compromise split.
```

If a composite is required for governance, weights should be set before evaluation, sensitivity-tested and backtested against held-out cases. Results should be stratified by TLD, registrar size, threat class and evidence source. The output is an accountability signal — not an automated domain verdict.

15 A tiered Public Abuse Transparency Database

A transparency registry can eliminate disputes over whether a notice existed, when it was delivered and what action followed. Publishing every artifact, however, could expose victims, personal data, investigative methods and live exploit paths. The solution is tiered access.

P **Public record**

Case ID, indicator, broad category, key timestamps, evidence tier, actor roles, outcome code, appeal state and hashes of sealed artifacts.

C **Contracted-party view**

Reproducible technical evidence, contact channel, collateral context and preservation instructions.

R **Restricted authority view**

Victim data, financial attribution, active-investigation material and unredacted chain of custody.

Each update is append-only, timestamped and signed. Corrections do not erase history; they add a superseding record. Public search should obey retention rules, prevent bulk victim discovery and offer

redress for inaccurate personal data.

16 A six-month pilot with a pre-registered evaluation

A useful pilot should be small enough to govern and large enough to test the causal chain. Suggested scope: three volunteer registrars, two registry operators, one hosting/CDN partner, two independent research feeds and a qualified appeal panel. The study protocol and outcome definitions should be registered before cases are assigned.

MONTH 1

Protocol and legal mapping

Define categories, authorities, evidence tiers, privacy impact assessment and stop conditions.

MONTH 2

Instrument the existing process

Measure baseline timestamps and outcomes without changing response behavior.

MONTH 3

Run the evidence envelope

Introduce signed provenance, de-duplication and collateral classification.

MONTH 4

Test response targets

Randomize eligible cases or use a stepped-wedge rollout; monitor errors daily.

MONTH 5

Exercise appeals and failures

Conduct restoration drills, red-team malicious reports and audit access controls.

MONTH 6

Independent evaluation

Publish effect sizes, confidence intervals, missingness, adverse events, costs and replication materials.

Primary pilot outcomes

- Median and 90th-percentile time from validated evidence to effective mitigation.
- Difference in attributable post-notice harm or a pre-specified victim-exposure proxy.

- False-positive rate, wrongful-action severity and median restoration time.
- Recurrence at the same registrant, infrastructure cluster and campaign level.
- Direct operational cost per validated case and per effective mitigation.

> DECISION RULE
>
> Scale only if the pilot shows materially faster effective mitigation without exceeding pre-registered limits for false positives, privacy incidents, appeal delay or collateral service disruption.

17 Conclusions and testable recommendations

DNS abuse response is neither a purely technical filter nor a problem one institution can solve by itself. Existing 2024 contracts established meaningful mitigation duties, but they leave substantial discretion over evidence, timing and remedy. That discretion is necessary for proportionality; without comparable records, it also makes performance difficult to evaluate.

1. **Standardize the evidence envelope.** Adopt common provenance, timestamp and outcome fields across reporters, registrars, registries and infrastructure providers.
2. **Measure stages separately.** Publish acknowledgement, triage, decision, mitigation and appeal times by threat class and evidence tier.
3. **Distinguish malicious registration from compromise.** Prefer path/account remediation for compromised legitimate services and reserve domain-wide measures for cases where they are proportionate.
4. **Pilot response targets.** Test hour-scale objectives for high-confidence, single-purpose malicious registrations rather than declaring a universal deadline.
5. **Publish privacy-safe accountability data.** Use a tiered transparency registry with signed histories and restricted evidence.
6. **Evaluate causal impact.** Do not equate faster suspension with saved money until a pre-registered study measures victim-impact outcomes and displacement.

The constructive conclusion is not that the original problem disappears once its claims are qualified. It is that the central claims can now be tested: the 60–85% post-detection harm range is an explicit quantitative hypothesis; $100–150 is a testable engineering budget assumption for a defined sensor-and-triage layer; and hour-scale response targets can be piloted for high-confidence, single-purpose malicious registrations. Publish the workload, costs, timestamps and outcomes, compare them with a baseline, and measure displacement and error. Confirmation would justify contractual and policy change; rejection would identify which assumption failed.

Editorial and methods disclosure. PhishDestroy is an independent, non-commercial threat-intelligence project. This publication is a policy and engineering proposal, not legal advice, an ICANN programme or a peer-reviewed paper. Budget ranges labeled as models are planning assumptions. Primary-source claims are linked below. Corrections can be submitted through the site's contact and responsible-disclosure channels.

FIRST-PERSON POSITION · ACCOUNTABILITY MUST CONSTRAIN BOTH INACTION AND OVERREACH

A personal boundary on enforcement power

I am not asking for — and do not want — direct access to a domain-blocking tool. That power requires independent review, public audit logs, a rapid appeal and reversibility.

I believe there are many capable and responsible companies that could be trusted with parts of this work. Based on my experience and the record I have published, I do not place NameSilo in that group.

I would not advise legitimate US companies to entrust sensitive abuse or security functions to specific contractors whose conflicts, competence or conduct cannot withstand independent scrutiny. This is about documented actions, not nationality: Russian origin by itself proves nothing, just as an American corporate address proves nothing.

I want justice and meaningful accountability for registrars that earned revenue from registering and renewing domains documented as part of repeated scam campaigns. Where evidence points beyond negligence to possible knowing facilitation—whether direct operational involvement or acting as an informed intermediary whose function was to keep a harmful domain registered and online despite repeated warnings—that conduct should be independently investigated and, if substantiated, carry consequences. I do not present suspicion as established fact; but intermediary status must not confer immunity from scrutiny.

Review the underlying investigation ↗

A

References and source notes

[1] **ICANN, Adopted FY27 Budget (2026)** · https://www.icann.org/en/system/files/files/adopted-icann-budget-fy2027-published-2026-en.pdf

Source for $165.1M FY27 operations funding and the composition of funding streams.

[2] **ICANN, 2024 Global Amendment to the Registrar Accreditation Agreement** · https://www.icann.org/en/system/files/files/registrar-accreditation-agreement-global-amendment-05apr24-en.pdf

Primary contractual source for registrar DNS Abuse obligations.

[3] **ICANN, 2024 Global Amendment to the Base gTLD Registry Agreement** · https://itp.cdn.icann.org/en/files/registry-agreements/base-registry-agreement-global-amendment-05-04-2024-en.html

Primary contractual source for registry-operator mitigation duties.

[4] **ICANN Contractual Compliance, DNS Abuse Obligations Advisory (2024)** · https://www.icann.org/en/contracted-parties/advisories/documents/advisory-compliance-with-dns-abuse-obligations-in-the-registrar-accreditation-agreement-and-the-registry-agreement-05-02-2024-en
Operational interpretation, report requirements and examples of appropriate mitigation.

[5] **ICANN, Enforcement of DNS Abuse Mitigation Requirements (2024)** · https://www.icann.org/en/system/files/files/icann-enforcement-of-dns-abuse-mitigation-requirements-08nov24-en.pdf
Six-month implementation report, investigation and mitigation counts.

[6] **ICANN SSAC, SAC115: Report on an Interoperable Approach to Addressing Abuse Handling (2021)** · https://www.icann.org/en/system/files/files/sac-115-en.pdf
Background on abuse-report interoperability and evidence quality.

[7] **Hollenbeck, RFC 5731: EPP Domain Name Mapping (2009)** · https://www.rfc-editor.org/rfc/rfc5731.html
Technical definition of EPP domain status values, including server-set status.

[8] **Laurie et al., RFC 9162: Certificate Transparency Version 2.0 (2021)** · https://www.rfc-editor.org/rfc/rfc9162.html
Authoritative protocol specification and limits of CT as a sensor.

[9] **ICANN, Centralized Zone Data Service** · https://www.icann.org/en/contracted-parties/registry-operators/services/centralized-zone-data-service
Scope and access model for participating gTLD zone files.

[10] **FBI Internet Crime Complaint Center, 2025 IC3 Annual Report** · https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf
Complaint and reported-loss figures, including phishing/spoofing counts.

[11] **Bijmans et al., Catching Phishers By Their Bait (USENIX Security, 2021)** · https://www.usenix.org/system/files/sec21-bijmans.pdf
Dataset-specific measurements of phishing infrastructure and observed lifetimes.

[12] **Newly Registered Domains and Phishing Lifetimes, Journal of Cybersecurity (2026)** · https://doi.org/10.1093/cybsec/tyag020
Recent empirical evidence on newly registered phishing domains and lifetime distribution.

[13] **Inferential Analysis of Maliciously Registered Domains — INFERMAL (2024)** · https://www.icann.org/en/system/files/files/inferential-analysis-maliciously-registered-domains-08nov24-en.pdf
ICANN-funded analysis of malicious registrations and ecosystem factors; relevant limitations on causal claims.

[14] **Chainalysis, 2026 Crypto Scam Trends** · https://www.chainalysis.com/blog/crypto-scams-2026/
Vendor estimate of 2025 on-chain scam inflows; not treated as a domain-phishing total.

[15] **APWG, Phishing Activity Trends Report, Q1 2026** · https://docs.apwg.org/reports/apwg_trends_report_q1_2026.pdf
Observed phishing attacks/URLs; metrics are not equivalent to unique registered domains.

[16] **Cybercrime Information Center / Interisle, Phishing Landscape 2025** · https://www.cybercrimeinfocenter.org/phishing-activity-numbers-may-april-2025
Separates observed attacks, unique domains and maliciously registered domains.

[17] **Verisign, Domain Name Industry Brief, Q2 2026** · https://blog.verisign.com/domain-names/q2-2026-domain-name-industry-brief-quarterly-report/
Source for the reported 401.6M registrations across all TLDs.