How Fast Do DNS Filters Block Fresh Phishing? A Seven-Day, Three-Vantage Comparison of Quad9, Cloudflare 1.1.1.2, and Public Blocklists

Abhijeet Apsunde[0009−0004−6618−7450]

Independent researcher

Abstract. Advice on DNS-level malware blocking often leans on older head-to-head tests that favour Quad9 over Cloudflare. We re-measure the question on current phishing. Every hour a collector on a cloud virtual machine outside our home network pulls newly listed phishing domains from two feeds (OpenPhish and the CERT Polska warning list), and asks Quad9 (9.9.9.9), Cloudflare's malware filter (1.1.1.2) and an unfiltered control (1.1.1.1) whether each domain still resolves, at 0, 1, 6, 24, 72 and 168 hours after first sighting. Four public blocklists used on a home resolver are matched offline against the same domains, and a fixed sample of 3,000 popular domains measures false positives. Domains hosted on shared platforms are analysed separately because DNS can block them only by blocking the platform name. In Atlanta (1,723 own-domain phishing domains first seen 31 October to 3 November, each followed for a week) Cloudflare blocks most at first sighting (17% against 6%), Quad9 overtakes it by six hours (89% against 29%), and the two are within seven points at 24 hours (97% and 91%). Blocklists add little until a day has passed. On shared-hosting phishing Cloudflare blocks 90% and Quad9 23% at 24 hours. Comparing Atlanta with Frankfurt and Singapore shows Quad9's early blocking is location-dependent (18%, 71% and 49% at one hour), while Cloudflare's is not. By 72 hours Quad9's blocks have fallen to 83 to 87% on the same domains while Cloudflare's rise to 92 to 94%. For a home resolver limited to one upstream, neither service dominates: Quad9 is better in the first day and Cloudflare thereafter and on shared hosting. At 168 hours Quad9 has levelled off (82 to 86%) while Cloudflare holds (92 to 94%), and the blocklists, which lag at first, reach 82 to 85%.

Keywords: DNS filtering · phishing · protective DNS · Quad9 · Cloudflare · blocklists · measurement

1 Introduction

A household or small network that wants DNS-level protection against malware and phishing must pick an upstream resolver and, optionally, local blocklists. Published comparisons are scarce, small, or vendor-supplied (Section 2), and phishing infrastructure and the feeds behind each service change quickly. The

practical question we answer is narrow: *for a resolver that can forward to exactly one upstream and also load local blocklists, which combination blocks the most fresh phishing, how quickly, and at what false-positive cost?*

We do not rank the services in general. Our contributions are:

1. a reproducible, hourly-sampled design that measures *block latency* (how block rates change from first sighting to one week), not only a single end-of-test block rate;
2. a cohort rule that counts a domain only while an unfiltered control resolves it, so dead domains cannot inflate any arm;
3. a split between phishing on a domain of its own and phishing on shared hosting, which DNS cannot treat alike;
4. one-week measurements for Quad9, Cloudflare 1.1.1.2 and four blocklists, with the threats to validity stated beside them.

2 Related Work

Independent measurements of protective DNS resolvers are limited. Quad9 publishes a post on third-party tests of its service [1]; it reports over 97% of listed hosts blocked in one test and itself notes that one of its threat-intelligence providers ingests the list that test used, which raised the result. A 2019 blog test of eleven public resolvers on 20 live domains (14 malware, 6 phishing; 31 March 2019) found that none blocked all of them: Quad9 blocked 55% (50% of the phishing domains) and Cloudflare's 1.1.1.1 none; the author tested only each provider's primary address and did not measure false positives [2]. A September 2026 benchmark by MUO, reported by ISPreview [3], tested Cloudflare, Quad9, NextDNS and AdGuard against 100 sites (50 malware, 50 phishing) drawn from a list of 41,000. Malware blocking was strong (94 to 100%) and phishing was the weak category: Cloudflare blocked 74% of phishing sites and Quad9 36%. The report does not state which Cloudflare service was used. Top-site rankings are a common source of benign samples; we use Tranco [4], which was designed for research use and resistance to manipulation.

Our design differs in sampling the same domains repeatedly over a week from a fixed vantage point, which exposes how block rates evolve.

3 Method

3.1 Vantage point

Ubuntu virtual machines in the Vultr Atlanta region (from 27 October) and, from 31 October, Frankfurt and Singapore, all outside the authors' home network, so results do not depend on a local resolver, cache or ISP. Main results use Atlanta. PeeringDB lists Vultr (AS20473) in a single Atlanta facility, QTS Atlanta-Metro [5]. The services we query are anycast, so we observe the nearest point of presence, not each service worldwide.

3.2 Arms

Three resolvers are queried over plain DNS with `dig`: Quad9 `9.9.9.9` [6], Cloudflare `1.1.1.2` with malware blocking [7], and Cloudflare `1.1.1.1` as the unfiltered *control*. A block is recognised by each service's own signature: Quad9 answers NXDOMAIN; Cloudflare answers `0.0.0.0` or an empty answer. Four public blocklists are matched offline: Hagezi PRO++ and TIF [8], AdGuard DNS [9], and Steven Black's unified hosts [10]. Lists are re-downloaded daily.

3.3 Samples and cohorts

Each hour the collector adds up to 400 newly listed domains per feed: OpenPhish [11] and the CERT Polska warning list [13]. The first pass over a bulk feed only seeds the set of seen domains. Each domain is re-queried at first sighting and again at 1, 6, 24, 72 and 168 hours. A domain counts at a checkpoint only if the *control* resolves it at that checkpoint. We call a domain *own-domain* when its registered domain belongs to the attacker, and *shared-hosting* when it is a subdomain of a platform that serves many unrelated sites (for example `pages.dev` or `github.io`); DNS can block the latter only by listing the individual host name or the whole platform.

Because later checkpoints cover only domains old enough to have been re-queried, columns that report "all live domains" have different populations. To compare like with like we also report a fixed cohort: domains live at first sighting with results at every checkpoint through 24 hours.

3.4 False positives

A fixed random sample of 3,000 domains from the Tranco top 100,000 [4] is queried daily on every arm. The sample is the same every day, so repeated runs do not add information about false positives.

4 Results

Data frozen 9 November 2026 (19:42 PST). Matched cohort: domains first seen from 31 October 13:44 UTC to 3 November 08:00 UTC at each vantage. Atlanta: 2,200 domains (1,723 own-domain, 477 shared-hosting; 834 from OpenPhish, 1,366 from CERT Polska). Frankfurt 2,217 and Singapore 2,498. The 45 newest own-domain samples at each vantage had not reached 168 hours at the freeze and are not in that column. Tables are generated by `analysis/final_tables.py` and `analysis/extra_tables.py`.

4.1 Own-domain phishing

Figure 1 follows one fixed cohort ($n = 1{,}305$). Cloudflare leads at first sighting; the two are level at one hour; Quad9 is far ahead by six hours (90% against 27%); and Cloudflare closes to six points at 24 hours (97% against 91%). Blocklists catch up only after a day. Table 1 adds the 72- and 168-hour columns: Quad9 has fallen back by 72 hours and then levels off, while Cloudflare holds (Section 6).

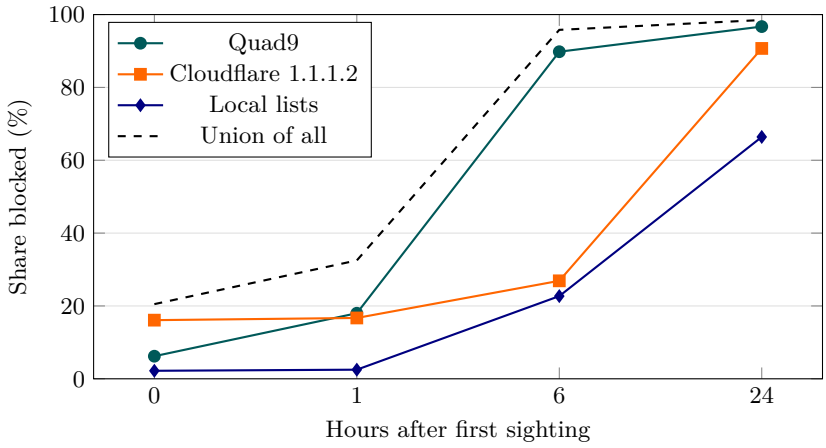


Fig. 1. Share of own-domain phishing blocked, by checkpoint, for a fixed cohort of 1,305 Atlanta domains (matched window) live on the control at every checkpoint from 0 to 24 hours. Checkpoints are spaced evenly (categorical axis).

4.2 Shared-hosting phishing

Cloudflare blocks 79.5% at first sighting rising to 89.5% at 24 hours (Atlanta, $n = 476$) and 89.9% at 72 hours ($n = 476$); Quad9 blocks 11.7% rising to 22.5% at 24 hours; local lists block under 2% through 24 hours (Table 2). Frankfurt and Singapore agree within a few points for the two services. The 72-hour list cell differs between vantages (23%, 6% and 20%). An earlier pull traced the difference to `pages.dev` domains matched by a platform-wide list entry in Atlanta and not in Frankfurt; we have not re-run that check on this pull. We do not read that cell as a blocking speed. At 168 hours ($n = 473$) Cloudflare blocks 90.9% and Quad9 18.2%, and the list cell is 3.2%, so the 72-hour spike did not persist.

4.3 False positives

Of 2,597 popular domains from Atlanta (Table 3), the lowest-error arm is Quad9 (0.19%), then Cloudflare 1.1.1.2 (0.39%) and Hagezi TIF (0.46%). Frankfurt and Singapore give 0.12 and 0.23% for Quad9 and 0.35 and 0.39% for Cloudflare. PRO++ and AdGuard also list advertising and trackers, so their rates are not like-for-like malware-filter error.

4.4 Vantage comparison

On 31 October we started the same collector on virtual machines in Frankfurt and Singapore. Table 4 compares the three using one cohort: domains first seen in the same window at every vantage. Cloudflare is steady across locations; Quad9 is not. At one hour Quad9 blocks 18% of own-domain phishing from Atlanta,

	Hours after first sighting					
Arm	0	1	6	24	72	168
Quad9	6.0	**18.3**	**89.1**	**96.6**	86.9	85.8
Cloudflare 1.1.1.2	**16.7**	17.7	28.9	90.7	**93.7**	**93.7**
Local lists (all four)	2.3	2.7	23.0	66.4	79.1	84.8
Union of all arms	21.0	33.8	96.0	98.5	97.1	97.2
n (live on control)	1,536	1,503	1,434	1,308	1,114	1,039

Table 1. Own-domain phishing blocked (%), Atlanta, for the matched cohort: domains first seen 31 October 13:44 to 3 November 08:00 UTC (1,723 own-domain domains), counted at a checkpoint when the control resolves them there, so n differs by column. The 72-hour column holds the 1,114 domains old enough to have reached it; the 168-hour column holds the 1,039 domains that were live on the control then (the newest 45 domains had not reached 168 hours at the freeze). Data frozen 9 November 2026.

	Hours after first sighting					
Arm	0	1	6	24	72	168
Quad9	11.7	12.6	17.4	22.5	18.7	18.2
Cloudflare 1.1.1.2	**79.5**	**80.3**	**85.5**	**89.5**	**89.9**	**90.9**
Local lists (all four)	0.0	0.0	0.0	1.5	22.9*	3.2
n	477	477	477	476	476	473

Table 2. Shared-hosting phishing blocked (%), Atlanta, matched cohort. *A platform-wide list entry (see text); not a blocking speed.

49% from Singapore and 71% from Frankfurt, and the gap closes by six hours (88 to 89%). Quad9 versus Cloudflare on the same domains differs significantly at one hour in Frankfurt and Singapore (exact McNemar $p < 10^{-49}$) but not in Atlanta ($p = 0.71$). We read this as a property of the nearest Quad9 site or its caches and not of the phishing, but have not tested why.

4.5 Control sensitivity

A domain counts at a checkpoint only if the control resolves it, so the control's own errors could shift every column. On 1 November we added two further controls on each vantage: Google Public DNS (`8.8.8.8`, no malware filter) and an Unbound resolver walking down from the root servers with no upstream and a 5-second cache. Across the checkpoints measured with all three (4,758 in Atlanta, 3,647 in Frankfurt, 3,581 in Singapore), 1.1.1.1 and Google disagreed on 9, 12 and 0 of them, and all three agreed on 98.8%, 98.7% and 100%. The local resolver is the one that departs: 50 and 44 times in Atlanta and Frankfurt, nearly always because it timed out on a domain the others resolved (for example a domain served by Alibaba's nameservers, which do not answer our Vultr addresses at

Arm	Blocked	Rate (%)
Quad9	5	0.19
Cloudflare 1.1.1.2	10	0.39
Hagezi TIF	12	0.46
Steven Black	25	0.96
AdGuard DNS	53	2.04
Hagezi PRO++	105	4.04

Table 3. Popular domains wrongly blocked, of 2,597 sampled from Tranco (Atlanta, 10 November UTC).

		Hours after first sighting					
Arm	Vantage	0	1	6	24	72	168
Quad9	Atlanta	6.0	18.3	89.1	96.6	86.9	85.8
	Frankfurt	15.0	70.5	88.7	96.4	86.2	85.3
	Singapore	12.1	48.7	87.8	93.1	83.0	81.6
Cloudflare 1.1.1.2	Atlanta	16.7	17.7	28.9	90.7	93.7	93.7
	Frankfurt	17.7	18.6	29.6	90.2	93.0	93.2
	Singapore	20.2	21.2	31.8	89.5	92.2	92.1

Table 4. Own-domain phishing blocked (%) by vantage, matched cohort (the same first-seen window at each vantage; 1,723, 1,738 and 1,806 domains). The 72-hour column has n = 1,114, 1,131 and 1,188; the 168-hour column has n = 1,039, 1,055 and 1,111.

all). We therefore keep 1.1.1.1 as the definition of "live" and treat the other two as sensitivity checks. The check only covers domains first seen after 1 November.

5 Threats to Validity

Feed circularity. Quad9 and phishing feeds may draw on overlapping intelligence, which would flatter Quad9 on these samples. CERT Polska is a second feed, but its independence from Quad9's sources is not established. **Few vantages.** Three virtual machines in three metros; Section 4.4 shows Quad9 differs by location, so results elsewhere may differ again. **List refresh.** Local lists rise from 15% to 67% by 24 hours, which may reflect the daily list update and not blocking speed; we have not verified list versions per checkpoint. **Cohort drift.** Columns in Table 1 count different numbers of domains, because a domain enters a column only if the control resolves it; only Figure 1 follows one fixed set of domains (an earlier cohort, so it differs from the table). **Small tail.** The 168-hour column rests on about 1,000 domains per vantage and leaves out the newest 45 of each cohort. **Control definition.** A domain that resolves on 1.1.1.1 may still be dead for users (parked or sinkholed); we do not fetch pages. **Resolution only.** We measure whether DNS blocks, not whether users would have been harmed.

One week only. Each domain is followed for seven days; blocks that arrive or lapse later are not seen. **Early correction.** An earlier analysis suggested Quad9's blocks decay by 24 hours (34%, 81%, 47%); that rested on 64 domains and did not hold at 226 (20%, 72%, 66%). A decay does appear later, between 24 and 72 hours, and holds on 1,114 to 1,188 domains per vantage (Section 6) and levels off by 168 hours.

6 Discussion

Speed depends on where you ask. In the first hour Cloudflare's filter blocks more than Quad9 from Atlanta (about 17% against 6% at first sighting), but the order flips by six hours (about 90% against 30%), and from Frankfurt Quad9 is already at 70% at one hour. Cloudflare's own rates vary by a few points across the three locations. A reader comparing the two services from a single city can therefore reach opposite conclusions about early blocking. By 24 hours the services are within about seven points of each other on own-domain phishing.

Platform-hosted phishing is a different problem. On shared hosting Cloudflare blocks 80 to 91% and Quad9 11 to 27% across the checkpoints. Table 2 cannot say whether Cloudflare is blocking the phishing pages or the whole platform name; our design cannot separate the two, and a household that uses those platforms legitimately would pay for the blocks. We report the false-positive rate (Table 3) but not the collateral cost on platforms.

Local lists lag; they matter mainly after a day. Four public lists blocked under 6% at first sighting and 60 to 70% at 24 hours and 82 to 85% at 168 hours, by when they are level with Quad9 and about 10 points behind Cloudflare. This may mostly reflect the lists' own daily refresh (Section 5), so we do not read it as a blocking speed. Their false-positive rates are higher (up to 4% for Hagezi PRO++, which also lists advertising and tracking domains).

Quad9's blocks fall away by 72 hours, then level off. Quad9 blocks 96.6, 96.4 and 93.1% at 24 hours and 86.9, 86.2 and 83.0% at 72 hours (Atlanta, Frankfurt, Singapore); Cloudflare rises from about 90% to 92 to 94%. This is not a change in which domains are old enough to be measured: on the same 1,113, 1,127 and 1,186 domains measured at both checkpoints, Quad9 stops blocking 111, 115 and 117 of them and starts blocking 7, 7 and 7, while Cloudflare stops blocking none and starts blocking 34, 33 and 34. We cannot say whether a Quad9 block expires or the domain is taken down and no longer matches. Seen from one place only, the 72-hour picture would favour Cloudflare, so the choice between the services depends on how long the user needs protection. From 72 to 168 hours it levels off: Quad9 blocks 85.8, 85.3 and 81.6% at 168 hours, and on the same 1,034, 1,048 and 1,106 domains it stops blocking 8, 8 and 9 and starts blocking 3, 4 and 4. Cloudflare holds at 93.7, 93.2 and 92.1%. Its lead over Quad9 at 168 hours is 8 to 11 points and significant at every vantage (exact McNemar $p < 10^{-14}$).

Reachability is a separate axis, and it is ISP-specific. Our vantages are cloud virtual machines with good transit, so they measure what each service

blocks and not whether a home line can reach it. At the author's home, Quad9 over DNS-over-TLS failed about four times a day on one fibre provider, whose upstream carrier confirmed it does not peer with Quad9, so traffic left the region (27 ms against 3 ms to Cloudflare). On the household's cable line the same service was reached as fast as any other. Because a resolver can forward to only one upstream, each failure sent queries to root recursion, which bypasses the filter altogether. A resolver with good blocking but a poor path to it can therefore protect less than a weaker one that is always reachable. Network quality varies from one ISP to the next, so it cannot be settled from our three vantages.

Control choice does not drive the result. Google, 1.1.1.1 and a from-the-root resolver agree with each other on about 99% of checkpoints (Section 4.5).

7 Future Work

This study is a pilot. A follow-up should (i) draw a feed-independent sample, for example newly issued certificates carrying brand keywords labelled by our own crawler, so that neither a resolver nor a feed picked the domains; (ii) measure harm directly, by fetching each candidate (headers and first 64 KB, no scripts) and reporting exposure coverage, the share of live-malicious hours during which a resolver was blocking; (iii) sample more densely between one and six hours, where Quad9's catch-up happens, and report time to first block with survival curves; (iv) add malware domains (URLhaus, ThreatFox) as their own class; (v) place benign canary pages on the same platforms to test whether Cloudflare blocks hosts or whole platforms; (vi) re-query a subset over DNS-over-TLS and record Extended DNS Errors, to separate a block from a real NXDOMAIN; (vii) replay the distinct names a household resolver saw in 30 days to measure false positives that matter to real users, publishing counts only; and (viii) pre-register hypotheses and decision rules before collecting, over four weeks with sampling stopped at day 21 so every domain reaches 168 hours.

8 Recommendation

The deployment constraint is one upstream plus local lists: forwarding to two upstreams selects by round-trip time, so which one answers a given query is effectively random. The decision is therefore Quad9 plus local lists against Cloudflare 1.1.1.2 plus local lists, and it turns on what the household wants protection from.

- **Quad9** blocks most own-domain phishing within six hours (88 to 89% at every vantage, against 29 to 32% for Cloudflare) and is ahead at 24 hours (93 to 97% against about 90%). Its behaviour depends on where the resolver is: at one hour it blocks 18% from Atlanta and 71% from Frankfurt. By 72 hours it has dropped to 83 to 87%, holding at 82 to 86% at 168 hours, and it blocks only 11 to 27% of phishing hosted on shared platforms.
- **Cloudflare 1.1.1.2** is slow on own-domain phishing in the first six hours but stable everywhere, holds 90 to 94% from 24 hours on, and blocks 80 to

91% of shared-hosting phishing. Its false-positive rate is slightly higher (0.35 to 0.39% against 0.12 to 0.23%), though both are low.
– **Local lists** add little in the first day (66 to 69% at 24 hours) and cost far more in false positives (Hagezi PRO++ 3.7 to 4.0%), so they supplement either service and do not replace one. By a week they reach 82 to 85%, level with Quad9.

For a household that mostly meets phishing through links it opens within hours of delivery, Quad9 is the better single upstream. For one that wants steady protection across days, and for traffic that reaches shared platforms, Cloudflare 1.1.1.2 is. We lean to Cloudflare 1.1.1.2 plus local lists as the default, because its coverage does not fall away and does not depend on location, and the 168-hour data show no Quad9 recovery (82 to 86%, against 92 to 94% for Cloudflare).

Before adopting either, test the path from your own line. Over a few days, query the candidate (over the transport you will use, for example DNS-over-TLS) every minute or so, and record failures and round-trip time next to a second candidate; also record how often your resolver falls back when the upstream fails. A candidate that fails more than a few times a day, or that is much slower than the alternative, should lose regardless of its blocking rate. A one-file script that does this, `pathtest.py`, is in the code archive [12].

9 Conclusion

Measured from three metros against about 2,200 newly listed phishing domains at each, public protective DNS does not have a single speed. Quad9 catches up quickly, but how quickly depends on where it is asked, and its blocks fade by three days and then hold. Cloudflare's malware filter is slower to start and steadier afterwards, and is far stronger on phishing hosted on shared platforms. Local lists lag both. A comparison from one vantage and one checkpoint would have reached a different verdict from each of these, which is the case for reporting block rates over time and across locations.

10 Data and Code Availability

The collector, report and analysis code (Python and SQLite) are archived on Zenodo [12] under the MIT licence, and the aggregate counts behind every table cell, with per-domain outcomes for the CERT Polska-sourced domains, under CC BY 4.0. We do not redistribute the raw lists of phishing domains. The CERT Polska warning list may be used by anyone for any purpose, commercial included, with no attribution required [13]; we still release only outcomes keyed by a salted hash of each domain, which is enough to reproduce the CERT Polska tables and not to recover the domains. OpenPhish's community feed permits academic, independent research but forbids making information from the feed available to third parties or creating derivative works without written permission [14]. We

therefore publish OpenPhish-sourced results as aggregate counts only, not per domain, unless OpenPhish permits more.

The code is developed at `https://github.com/abhiapsunde/dns-malware-protection-study`; the Zenodo archive is the citable, fixed copy. This paper is archived as a preprint at `https://doi.org/10.5281/zenodo.24098122` (this DOI always resolves to the latest version).

Ethics and Disclosure

The study sends ordinary DNS queries for publicly listed phishing domains and never fetches the sites. The authors have no commercial relationship with Quad9, Cloudflare, Vultr, or any list maintainer. This work was self-funded.

References

1. Quad9. "DNS blocking effectiveness: recent independent tests." `https://quad9.net/dns-blocking-effectiveness-recent-independent-tests/`. Accessed 9 November 2026.
2. "How good are public DNS providers in blocking malware & phishing?" Skadlig kod (blog); tests of 31 March 2019. `https://www.skadligkod.se/?p=3324`. Accessed 9 November 2026.
3. ISPreview. "New study benchmarks malware blocking by free third-party DNS providers," September 2026. `https://www.ispreview.co.uk/index.php/2026/09/new-study-benchmarks-malware-blocking-by-free-third-party-dns-providers.html`. Accessed 9 November 2026.
4. V. Le Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen. "Tranco: A research-oriented top sites ranking hardened against manipulation." In *Proc. NDSS*, 2019. `https://arxiv.org/abs/1806.01156`; list at `https://tranco-list.eu`.
5. PeeringDB. `https://www.peeringdb.com`. Queried 28 October 2026 (AS20473 facility records).
6. Quad9. `https://quad9.net`.
7. Cloudflare. "1.1.1.1 resolver documentation (malware-blocking 1.1.1.2)." `https://developers.cloudflare.com/1.1.1.1/`.
8. Hagezi DNS blocklists. `https://github.com/hagezi/dns-blocklists`.
9. AdGuard DNS filter. `https://adguard.com`.
10. S. Black. Unified hosts file. `https://github.com/StevenBlack/hosts`.
11. OpenPhish phishing feed. `https://openphish.com`.
12. A. Apsunde. Code and data for "How fast do DNS filters block fresh phishing?" Zenodo, 2026. `https://doi.org/10.5281/zenodo.24098038`.
13. CERT Polska. Domain warning list and data usage policy. `https://cert.pl/en/warning-list/`. Accessed 7 November 2026.
14. OpenPhish. Terms of use. `https://openphish.com/terms.html`. Accessed 7 November 2026.