Temporal Evaluation of Secure DNS Resolvers Against Malicious Domain Threat Feeds

Jaqueline Gonçalves de Souza[1], **Diego Nunes Molinos**[1], **Thiago Pirola Ribeiro**[1]

[1]Faculdade de Computação (FACOM) – Universidade Federal de Uberlândia (UFU)
Caixa Postal 593 – 38400-902 – Uberlândia – MG – Brazil

{jaqueline.goncalves, diego.molinos, tpribeiro}@ufu.br

Abstract. *Public recursive DNS resolvers with security filtering provide a transparent protection layer, but their blocking behavior depends on threat-intelligence sources and update policies. We temporally evaluated Quad9 Malware, CleanBrowsing Security, Cloudflare Security, and AdGuard DNS against CERT.pl and URLhaus corpora. Quad9 and CleanBrowsing showed high agreement with CERT.pl, whereas agreement with URLhaus was lower and varied over time. Because CERT.pl discloses that Quad9 uses its Warning List, the results indicate resolver-feed alignment rather than feed-independent superiority. The study provides an auditable protocol for corpus-specific analysis and discusses limitations related to feed provenance, sampling, and external observability.*

1. Introduction

In Brazil, Internet access has expanded rapidly among populations with heterogeneous levels of technical literacy [IBGE 2024]. This broader connectivity increases the importance of security mechanisms that do not depend on users identifying malicious infrastructure. Malicious domains support phishing, malware distribution, credential theft, and online fraud [Manasrah et al. 2022]. Because malicious campaigns frequently rotate domains to evade detection, users may reach malicious services before traditional security mechanisms react [Antonakakis et al. 2010].

In this context, public recursive Domain Name System (DNS) resolvers equipped with native security filtering emerge as a transparent and proactive protection strategy. By blocking requests to malicious domains directly at the DNS layer, these services can interrupt malicious communication at the earliest stage of the network connection, without requiring user interaction or additional software installation [Manasrah et al. 2022].

Recent studies have begun to examine Protective DNS deployment, response rewriting, and temporal blocking behavior [Liu et al. 2024, Fujii and Sato 2025]. However, comparative blocking rates remain difficult to interpret when benchmark feeds overlap with providers' threat-intelligence inputs. The effects of feed provenance and provider-feed alignment therefore require explicit treatment in corpus-based evaluations.

To address this measurement gap, we evaluate four public security-focused DNS resolvers against two threat-feed corpora. The contributions are: (a) a documented protocol for classifying heterogeneous DNS blocking responses; (b) a temporal characterization of resolver-feed agreement; and (c) an analysis of how feed provenance and sampling constrain comparative blocking rates. The study is a corpus-specific audit, not a universal ranking of resolver security.

2. Background and Related Work

This section summarizes the technical foundations of DNS-based security filtering, discusses the characteristics of threat intelligence feeds and domain obsolescence, and reviews prior studies on public DNS resolver evaluation.

2.1. DNS-based Security Filtering

The Domain Name System (DNS) is crucial for Internet communications and serves as an important point for network security. DNS-based filtering operates at the recursive resolver (RDNS) level, intercepting client queries before they connect to a destination server [Manasrah et al. 2022]. When a user or an infected device attempts to resolve a domain name, the filtering mechanism checks the request against threat intelligence, static blocklists (DNSBLs), or machine learning models. If the domain is deemed malicious, such as one associated with phishing or a botnet, the resolver either returns a Non-Existent Domain (NXDOMAIN) response or redirects traffic to a controlled server, a technique known as DNS sinkholing [Voloshina 2024].

2.2. Threat Intelligence Feeds

Threat intelligence feeds provide continuously updated repositories of malicious domains and URLs used by security systems for automated detection and blocking [Voloshina 2024]. Institutional feeds, such as CERT.pl [CERT Polska 2025], typically rely on formal validation processes and exhibit higher stability and lower false-positive rates. In contrast, collaborative platforms such as URLhaus [abuse.ch 2026] prioritize rapid, community-driven reporting of emerging malware distribution infrastructures, resulting in highly dynamic datasets characterized by short-lived malicious domains. The distinct operational characteristics of these repositories make them highly suitable for comprehensively evaluating both consistency and responsiveness of DNS security filtering services.

Feed provenance is particularly important in comparative protective-DNS measurements. CERT Polska states that its Warning List is used by private DNS providers, including Quad9 and NextDNS [CERT Polska 2025]. Consequently, testing Quad9 against the CERT.pl list may involve direct or indirect overlap between the benchmark corpus and the resolver's threat-intelligence inputs.

2.3. Evasion Techniques and Domain Obsolescence

Modern cybercriminal infrastructures rely on evasion techniques to bypass static security mechanisms and delay threat detection. Phishing and malware campaigns frequently rotate domains, IP addresses, and hosting providers, while Domain Generation Algorithms (DGAs) create large numbers of short-lived domains, rapidly making traditional blocklists obsolete [Antonakakis et al. 2010].

This dynamic behavior directly impacts DNS-based security filtering, whose effectiveness depends on fast threat intelligence synchronization and blocklist updates. Additionally, recursive DNS providers adopt heterogeneous blocking strategies, such as NXDOMAIN responses and sinkhole redirection, complicating direct comparisons between services [Voloshina 2024]. These characteristics make temporal and reproducible evaluations essential for understanding the real-world effectiveness of secure DNS resolvers.

2.4. Related Work

Prior evaluations of public DNS resolvers primarily addressed performance and availability. [Silva et al. 2023] systematized performance and resilience metrics, while [Silva et al. 2024] showed how caching, geolocation, and domain popularity affect latency. These studies provide relevant benchmarking practices but do not measure malicious-domain filtering.

More directly, [Liu et al. 2024] analyzed 28 Protective DNS providers and identified heterogeneous rewriting policies, including NXDOMAIN, special-use addresses, secure sinkholes, CNAME redirection, and empty responses. [Fujii and Sato 2025] compared ten services, measured blocking coverage and false blocking, and followed blocking and unblocking behavior over 32 days. Their results demonstrate that coverage and retention differ substantially among providers and over time.

[Liberato et al. 2025] provides an additional methodological comparator because it infers DNS blocking from categorized domain corpora, although its focus is parental control. Unlike these studies, our analysis explicitly contrasts institutional and community-driven corpora and examines how disclosed provider-feed overlap affects interpretation. Because CERT Polska states that Quad9 uses its Warning List [CERT Polska 2025], we treat blocking rates as corpus-specific agreement rather than intrinsic resolver quality.

3. Experimental Methodology

This section describes the methodological pipeline used to evaluate the cyber threat-blocking efficacy of security-focused DNS resolvers. The workflow encompasses provider selection, structuring a *dataset* of malicious domains, and executing parallel automated tests. Figure 1 shows the pipeline used in this work.

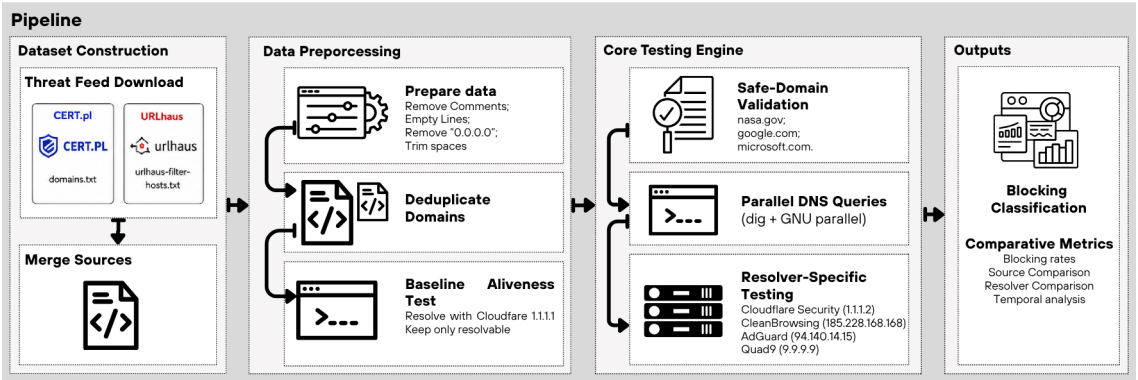

Figure 1. Pipeline.

3.1. Dataset Construction and Data Processing

The test dataset was constructed using raw data extracted from two reference repositories in the security community: (a) **CERT.pl [CERT Polska 2025]:** An institutional threat feed maintained by the *Computer Emergency Response Team* of Poland, characterized by formal validation and high stability of domains restricted to verified malicious activities, and (b) **URLhaus: [abuse.ch 2026]** A collaborative, open platform managed by abuse.ch, focused on the real-time, dynamic sharing of URLs associated with malware distribution.

A preprocessing stage was implemented to normalize and validate the collected threat feeds before executing the DNS benchmark. Initially, raw domain lists were cleaned by removing comments, empty lines, and sinkhole prefixes (e.g., `0.0.0.0`) using *sed* scripts. The resulting records were then merged and subjected to a textual deduplication process using *awk*. Domains without valid IP responses from Cloudflare's standard resolver (`1.1.1.1`) were excluded. This heuristic may discard domains because of transient or resolver-dependent failures and introduces selection bias; future replications should combine independent resolvers and authoritative lookups.

3.2. Selected DNS Providers and Inclusion Criteria

We purposively selected four free, registration-free, security-filtering resolvers with documented public endpoints. Other services, including DNS4EU, dns0.eu, NextDNS, and OpenDNS FamilyShield were excluded to keep the short paper benchmark tractable. The sample is therefore not exhaustive and does not represent the entire protective-DNS ecosystem. Table 1 summarizes the evaluated providers.

Table 1. Evaluated Secure DNS Resolvers

Resolver	IPv4	Description
Cloudflare Security	1.1.1.2	Filters known malicious domains while maintaining high network performance.
Quad9	9.9.9.9	Focused on enterprise-grade cybersecurity and threat intelligence supported by global security coalitions.
AdGuard DNS	94.140.14.14	Specializes in ad and tracker blocking with integrated malware security layers.
CleanBrowsing Security	185.228.168.9	Provides essential security filtering and malicious code blocking without parental control features.

3.3. Core Testing Engine

Experiments ran on Ubuntu 22.04 LTS using an automated Bash benchmark adapted from Kris Lowet's architecture [1]. *GNU Parallel* was used to issue concurrent queries for each domain-resolver pair. NXDOMAIN and identified sinkhole responses, including `0.0.0.0`, were classified as blocking, whereas successful resolutions were classified as non-blocking.

The original execution did not separately record query repetition, timeout, retry, and concurrency parameters. Consequently, the measurements do not estimate transient response variability or possible rate-limiting effects and should be interpreted as descriptive blocking proportions. The available source code and datasets are published at [2].

4. Experimental Results and Discussion

This section reports resolver-feed agreement across six snapshots collected between April 8 and May 16, 2026. The analysis compares observable blocking behavior across feeds and dates without inferring provider-internal update mechanisms.

[1] `https://gist.github.com/KrisLowet/675ba34e682c6d2afbc53fc317b41e85`
[2] `https://github.com/jaquelinegon/DNS-Resolver-Analysis`

4.1. Threat Feed Analysis and Temporal Evolution

The blocking rate measures agreement between each resolver's observable DNS behavior and the selected feed; it does not isolate detection capability from feed partnerships, ingestion policies, or update schedules. Quad9 and CleanBrowsing showed agreement near 99% with CERT.pl across the observed dates. For Quad9, this is consistent with CERT Polska's disclosure that the resolver uses its Warning List and is not independent evidence of superior detection.

Cloudflare Security and AdGuard showed approximately 11%–14% agreement with CERT.pl. These values indicate lower overlap with this corpus under our classification procedure, not necessarily inferior protection across other feeds, regions, or user profiles.

Table 2. Corpus Volumetrics: Active Threat Feeds and Raw Downloaded Totals

Collection Date	CERT.pl	URLhaus	Total Active Corpus	Total Downloaded (Raw)
Apr 08, 2026	58,120	7,066	65,186	157,119
Apr 10, 2026	59,041	7,504	66,545	157,543
Apr 14, 2026	59,995	8,004	67,999	158,573
Apr 17, 2026	60,101	6,867	66,968	157,165
Apr 20, 2026	58,022	6,807	64,829	152,988
May 16, 2026	56,726	5,636	62,362	149,405

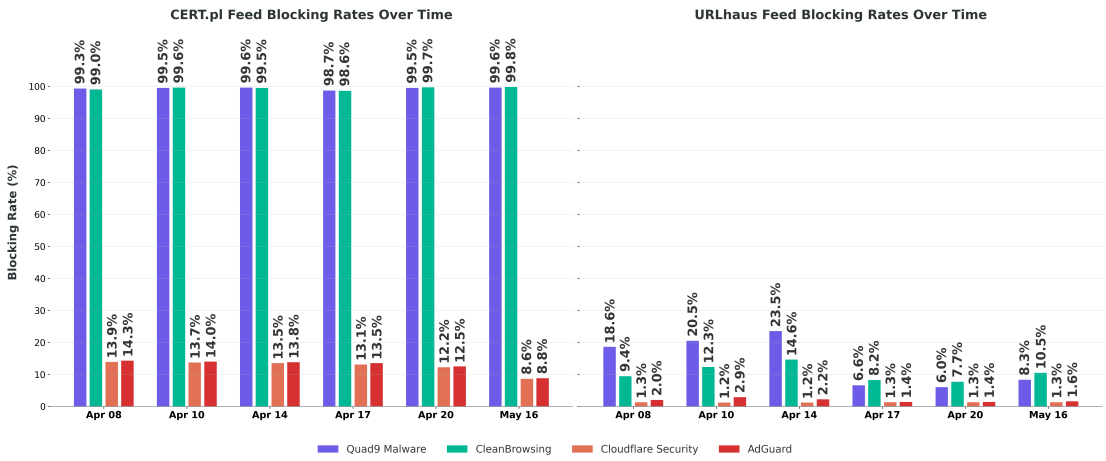

Figure 2. Detailed DNS Blocking Rate Comparison by Threat Feed and Resolver.

The *URLhaus* corpus produced a different pattern, as illustrated in Figure 2. In this analysis, even the resolvers that achieved the best performance on the *CERT.pl* baseline fell below 25%. Quad9 Malware reached its highest result on April 14 (23.51%), then declined to 6.57% on April 17, and further to 5.98% on April 20. A similar trend was observed for CleanBrowsing, which obtained 14.62% (April 14) and decreased to 7.67% on April 20. Even with the turn of the month (May), both achieved results below 10.6%.

URLhaus agreement was lower and varied across dates. Because the experiment did not track feed membership, authoritative liveness, hosting status, or provider-side

ingestion, the decline cannot be attributed specifically to takedown, record purging, conservative policies, or blocking latency.

The measurements are limited to four resolvers, two feeds, one vantage point, and six irregularly spaced snapshots, including a 26-day gap. A representative benign corpus was not evaluated; therefore, false-positive trade-offs remain unmeasured.

4.2. Operational Implications

The evaluation provides a feed-aware perspective for interpreting Protective DNS measurements. Separating institutional and community-driven corpora avoids reducing resolver behavior to a single aggregate score and reveals how observed coverage changes with threat-source characteristics. The normalization of NXDOMAIN and sinkhole responses also enables comparison among providers that implement different blocking mechanisms.

For practitioners, the results indicate that resolver selection should be aligned with the relevant threat model, documented intelligence sources, update policies, and acceptable false-positive trade-offs. Protective DNS can provide a transparent first layer of protection, particularly for users and small organizations that cannot maintain local security infrastructure. However, it is most effective when combined with browser, endpoint, and e-mail security controls.

For researchers and network administrators, the proposed workflow offers a compact auditing approach that can be applied periodically to versioned threat feeds. Comparing results across feeds and dates can help identify changes in observable coverage, support evidence-based resolver selection, and guide future evaluations involving additional providers, geographic vantage points, and benign-domain corpora.

5. Conclusion

This work measured the temporal agreement of four protective DNS resolvers with CERT.pl and URLhaus corpora. Quad9 and CleanBrowsing showed high agreement with CERT.pl, whereas Cloudflare Security and AdGuard showed lower overlap. Quad9's result must be interpreted in light of CERT Polska's disclosure that it uses the Warning List. URLhaus agreement was lower and variable, but the experiment cannot determine whether this resulted from feed changes, domain inactivity, or provider policies.

The findings characterize the evaluated feeds, dates, and execution conditions rather than universal resolver quality. By making feed provenance and heterogeneous blocking responses explicit, the study provides an auditable basis for recurring comparisons and evidence-based resolver selection. Future work will use regular sampling, multiple vantage points, additional independent feeds, repeated queries, and a benign corpus to estimate false positives and distinguish corpus aging from provider update behavior.

Generative AI Statement

The AI tools Google Gemini and Perplexity were utilized in this paper to assist with grammatical refinement, summarization, and text translation.

References

abuse.ch (2026). URLhaus – Sharing Malware URLs for Detection and Mitigation. https://urlhaus.abuse.ch/. Accessed: 2026-05-20.

Antonakakis, M., Perdisci, R., Dagon, D., Lee, W., and Feamster, N. (2010). Building a dynamic reputation system for {DNS}. In *19th USENIX Security Symposium (USENIX Security 10)*.

CERT Polska (2025). Warning list. https://cert.pl/en/warning-list/. Accessed: 2026-03-22.

Fujii, S. and Sato, T. (2025). 10 pdnss 10 colors: A measurement study of protective dns. *Journal of Information Processing*, 33:608–618.

IBGE (2024). Em 2023, 88,0% das pessoas com 10 anos ou mais utilizaram internet. https://agenciadenoticias.ibge.gov.br/agencia-noticias/2012-agencia-de-noticias/noticias/41026-em-2023-87-2-das-pessoas-com-10-anos-ou-mais-utilizaram-internet. Agência de Notícias do IBGE.

Liberato, M., Affinito, A., Meijerink, B., Jonker, M., Botta, A., and Sperotto, A. (2025). To block or not to block? evaluating parental controls across routers, dns services, and software. In *2025 9th Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–10.

Liu, M., Zhang, Y., Li, X., Lu, C., Liu, B., Duan, H., and Zheng, X. (2024). Understanding the implementation and security implications of protective dns services. In *NDSS*.

Manasrah, A. M., Khdour, T., and Freehat, R. (2022). Dga-based botnets detection using dns traffic mining. *Journal of King Saud University-Computer and Information Sciences*, 34(5):2045–2061.

Silva, M., Franco, M., Scheid, E., Zembruzki, L., and Granville, L. (2023). Desempenho de resolvedores de dns públicos: Uma análise do estado da arte. In *Anais da XX Escola Regional de Redes de Computadores*, pages 43–48, Porto Alegre, RS, Brasil. SBC.

Silva, M. A. d., Franco, M. F., Scheid, E. J., Zembruzki, L., and Granville, L. Z. (2024). Perfresolv: A geo-distributed approach for performance analysis of public dns resolvers based on domain popularity. In *International Conference on Advanced Information Networking and Applications (AINA 2024)*, volume 200 of *Lecture Notes on Data Engineering and Communications Technologies*, pages 35–47, Cham, Switzerland. Springer.

Voloshina, I. (2024). Domain name analysis with machine learning: Enhancing efficiency and reducing analyst strain. Master's degree in computer science - ict innovation cyber security, Università di Trento & University of Twente.