Cybersecurity Mechanisms in DNS Resolvers

An Internet Measurement Perspective

Jonathan Magnusson

Faculty of Health, Science and Technology

Computer Science

LICENTIATE THESIS | Karlstad University Studies | 2025:1

Cybersecurity Mechanisms in DNS Resolvers

An Internet Measurement Perspective

Jonathan Magnusson

LICENTIATE THESIS | Karlstad University Studies | 2025:1

Cybersecurity Mechanisms in DNS Resolvers - An Internet Measurement Perspective

Jonathan Magnusson

LICENTIATE THESIS

Karlstad University Studies | 2025:1

urn:nbn:se:kau:diva-102373

ISSN 1403-8099

ISBN 978-91-7867-518-0 (print)

ISBN 978-91-7867-519-7 (pdf)

https://doi.org/10.59217/rbhs6890

Distribution:
Karlstad University
Faculty of Health, Science and Technology
Department of Mathematics and Computer Science
SE-651 88 Karlstad
+46 54 700 10 00

Print: Universitetstryckeriet, Karlstad 2025

WWW.KAU.SE

Cybersecurity Mechanisms in DNS Resolvers: An Internet Measurement Perspective

Jonathan Magnusson
Department of Mathematics and Computer Science

Abstract

Using the Internet today, both end-users and automated systems rely on the Domain Name System (DNS) to translate human-readable domain names to IP addresses for communication between machines. This system from 1985 has only in recent years seen Internet standards addressing security and privacy concerns. In the position as a machine-in-the-middle between the client and the distributed hierachical system of authoritative name servers, we find the DNS resolver. Due to its purpose of forwarding, looking up, and caching queries and responses, in addition to its location between the clients and the name servers, the DNS resolver becomes a critical point for implementing these security and privacy features. The widespread adoption of these features, their variation in implementation, and impact on both clients and other name servers remain as interesting topics in the research community. The goal of this thesis is to analyze servers in the wild and conduct a comprehensive investigation into the security and privacy mechanisms configured on DNS resolvers. Using an Internet measurement approach, we explore the trends in the adoption and implementation of these features by generating and observing our own queries to and from the resolvers. We also investigate how clients and the DNS ecosystem as a whole are impacted by resolver configurations. We use and improve methods for measuring adoption of various security and privacy related features. Based on these measurements we report the current level of adoption and adoption over time, investigate anomalies, and identify limitations with measurement approaches. We fingerprint the software and version of popular open-source DNS resolvers by classifying query patterns. Comparing the ingress and egress resolvers we analyze forwarding behaviors and their impact on the availability and effectiveness of security and privacy features. We also cross-analyze features in DNS resolvers to find correlations, which could help us understand obstacles and find solutions to feature adoption.

Keywords: Domain Name System, Resolver, Security, Privacy, Traffic Analysis, Internet Measurements

Cybersäkerhetsmekanismer i DNS-resolvrar: Ett perspektiv med internetmätningar

Jonathan Magnusson
Institutionen för matematik och datavetenskap

Sammanfattning

För att maskiner ska kunna kommunicera på Internet idag så bygger det på att domännamnssystemet (DNS) översätter domännamn till IP-adresser för både användare och automatiska system. Internetstandardiseringar som behandlar säkerhet och personlig integritet i detta system från 1985 har huvudsakligen dykt upp på senare år. Mellan klienter och den distribuerade hierakin av auktoritativa namnservrar finner vi DNS-resolvern. På grund av dess syfte att vidarebefodra och slå upp klienternas frågor samt cacha svar, och dess position som en låda-i-mitten blir den en kritisk punkt för säkerhet och personlig integritet. Hur antagna dessa mekanismer är, deras variation vid implementering samt påverkan på både klienter och andra namnservrar är fortfarande intressanta ämnen i forskningsvärlden. Målet med denna avhandling är att analysera DNS-resolvrar på internet för att genomföra en omfattande utvärdering av relaterade mekanismer runt säkerhet och personlig integritet. Vi utforskar trenderna i antagandet och och implementeringen av dessa funktioner, och analyserar hur de påverkar klienter och ekosystemet som helhet genom att observera trafik från DNS-frågor. Vi använder och förbättrar metoder för att mäta antagning av olika funktioner relaterade till säkerhet och personlig integritet. Baserat på dessa mätningar rapporterar vi den nuvarande nivån av antagning och antagning över tid. Vi undersöker även intressanta avvikelser i resultaten och identifierar begränsningar med mätmetoderna som används. Genom att klassifiera trafikmönster, lyckas vi identifiera versioner av programmvara från populära DNS-resolvrar med öppen källkod. När vi observerar resolvrar involverade i en uppslagning så analyserar vi hur de vidarebefodrar och hur detta påverkar tillgängligheten och effektiviteten av olika mekanismer. Vi undersöker även samband mellan olika mekanismer, vilket skulle kunna leda till en djupare förståelse om utmaningar och lösningar till högre antagande.

Nyckelord: Domännamnssystemet, Uppslagningstjänst, Säkerhet, Personlig Integritet, Trafikanalys, Internetmätningar

Acknowledgements

Time passes so quickly when you are doing things that you love, find interesting, or meaningful. So while one part of me is wondering where all the time went, the other is confident that I am doing exactly what I am supposed to do at this point in time. This thesis in your hands is a testament to the effort that I have put down so far in my academic endeavour, and I could not have done it alone.

The work in this thesis was funded by the Swedish Internet Foundation, which in addition has opened doors to many interesting people in the DNS community. I would like to thank my supervisors, Tobias Pulls, Anna Brunstrom, Johan Stenstam, and my former supervisor Ulrich Wisser, for their guidance, support, expertise, and contributions. Next I would like to thank the people at the Department of Mathematics and Computer Science, especially the PriSec research group, for a captivating working environment with plenty of opportunities for discussions and brainstorming. I would also like to thank all of the students that I have had the pleasure to meet and exchange ideas with.

On a personal note, I want to thank my parents and my sister for supporting me and celebrating my victories, however small. I want to thank my friends for helping me take breaks between intense workdays. Last but not least, thank you Alexis for having my back and encouraging me on this adventure.

Remember that time is passing whether you like it or not, so make sure to spend it on something you love or find facinating. At least, spend it on something that you find meaningful.

Karlstad University , December 9, 2024 Jonathan Magnusson

List of Acronyms

ANS Authoritative Name Server

AS Autonomous System

ccTLD country-code TLD

CDN Content Delivery Network

DDoS Distributed Denial-of-Service

DGA Domain Generation Algorithms

DNS Domain Name System

DNSSEC DNS Security Extensions

DoH DNS over HTTPS

DoQ DNS over QUIC

DoT DNS over TLS

ECS EDNS Client Subnet

EDNS Extended Mechanisms for DNS

IP Internet Protocol

ISP Internet Service Provider

MAX MAX_MINIMIZE_COUNT

MIN MINIMIZE_ONE_LAB

NS Name Server

ODNS Oblivious DNS

ODoH Oblivious DoH

QMIN Query Name Minimization

QNAME Query Name

QTYPE Query Type

RA Recursion Available

RCODE Response Code

RFC Request For Comments

RR Resource Record

SLD Second-Level Domain

TLD Top-Level Domain

TLS Transport Layer Security

TTL Time to Live

List of Appended Papers

I. **Jonathan Magnusson**, Moritz Müller, Anna Brunstrom, and Tobias Pulls. A Second Look at DNS QNAME Minimization. International Conference on Passive and Active Network Measurement (PAM 2023). Springer.

II. **Jonathan Magnusson**. Fingerprinting DNS Resolvers using Query Patterns from QNAME Minimization. Secure IT Systems: 29th Nordic Conference (NordSec 2024). Springer.

III. **Jonathan Magnusson**. SweDNS: Evaluating Privacy and Security of DNS Resolvers used in Sweden. (Under Submission).

Comments on my Participation

Paper I After exploring previous studies around security and privacy of DNS resolvers together with Anna and Tobias, we decided to build upon a study measuring the adoption of QNAME Minimization since we had similar datasets at our disposal thanks to the Swedish Internet Foundation. At one point, after discussions about the measurement details with a few of the previous authors, we decided to invite Moritz to collaborate with us. Moritz performed the passive measurements at the authoritative name servers and also wrote those sections in the paper. I did the active measurements, the controlled experiments and most of the writing.

Paper II I designed an extension study based on a one-off measurement in a study around the adoption of QNAME Minimization. The one-off measurement was investigating query patterns from resolvers implementing QNAME Minimization. I performed all of the measurements and wrote the entire paper, with feedback from my supervisors.

Paper III Together with my supervisors, we came up with the idea of expanding the adoption measurements of QNAME Minimization to include more security and privacy features while narrowing the focus to clients in Sweden. I designed the study with input from my supervisors and the Swedish Internet Foundation. I performed all of the measurements and wrote the entire paper, with feedback from my supervisors.

List of Other Contributions

Throughout my PhD studies, I have also contributed to the following:

- Matthias Beckerle, **Jonathan Magnusson**, and Tobias Pulls. Splitting Hairs and Network Traces: Improved Attacks Against Traffic Splitting as a Website Fingerprinting Defense. Proceedings of the 21st Workshop on Privacy in the Electronic Society (WPES 2022). ACM.

  I mainly contributed to the study with measurement results from my Master's Thesis about testing split-path defenses against website fingerprinting on encrypted traffic.

- Leonardo Martucci, **Jonathan Magnusson**, and Mahdi Akil. On-Campus Hands-On Ethical Hacking Course. International Symposium on Human Aspects of Information Security and Assurance (HAISA 2023). Springer.

  In addition to setting up the cyber range and participating in discussions, I mainly contributed to the paper by writing the discussion section about advantages and disadvantages with online and on-campus approaches, as well as proposed changes for next iteration of the course.

Contents

Introductory Summary

1 Introduction

When using the Internet today, we rely on the Domain Name System (DNS) to translate the human-readable domain names to Internet Protocol (IP) addresses, allowing computers to communicate. This hierarchical key/value store consisting of name servers is queried on our behalf by DNS resolvers. The main *research objective* of this thesis is to conduct a comprehensive investigation into the security and privacy of DNS resolvers. A secure and privacy-preserving resolver is a complex system that should provide legitimate users with correct information when requested and avoid leaking potentially sensitive information about the client to others. We aim to identify and analyze the trends in the adoption, variations in implementation, and impact of privacy and security features in DNS resolvers to understand how these measures affects the clients and the overall system.

The remainder of this thesis is organized as follows. Section 2 provides background to help the reader understand the context and preliminaries of the appended papers. We outline our research questions and overall objective in Section 3. Section 4 gives an overview of our research methods, and our contributions are described in Section 5. Section 6 summarizes the appended papers: Paper I, II, and III. Section 7 positions our contributions within the context of related work, and Section 8 concludes and briefly discusses future directions. This is followed by the appended papers.

2 Background

This section introduces DNS and its purpose, focusing on the resolver. Additionally, we outline relevant proposed and occasionally implemented security and privacy features associated with DNS resolvers.

2.1 Domain Name System

For the Internet to work, devices need to be assigned IP addresses [39], which are hard for humans to remember. This is solved by mapping these addresses to human-readable names for user-friendly navigation. Initially, ARPANET, the Internet's predecessor, maintained this mapping in a file on each computer. However, with the network's rapid growth, this approach became impractical. In 1985, the hierarchical DNS was introduced [29, 30], designed to map domain names to resource records. The DNS namespace is divided into distinct segments called zones, which specific organizations or administrators manage. DNS servers, also known as name servers, typically operate on UDP port 53 and can be either *authoritative*, storing records for specific zones, or *resolvers*, retrieving and caching data from authoritative servers on behalf of clients. The caching allows subsequent queries for the same resource to reuse the result for a short time, specified by the resouce's Time to Live (TTL). The most

common DNS queries are for `A` and `AAAA` records, corresponding to IPv4 and IPv6 addresses, respectively. Other typical Query Types (QTYPEs) include `MX` for mail servers, `TXT` for associating arbitrary text, and `NS` for identifying authoritative name servers for a zone. DNS operates using a hierarchical structure, with zones maintained locally and connected in a tree structure, with each zone served by multiple authoritative name servers, including root, Top-Level Domain (TLD), and Second-Level Domain (SLD) name servers.

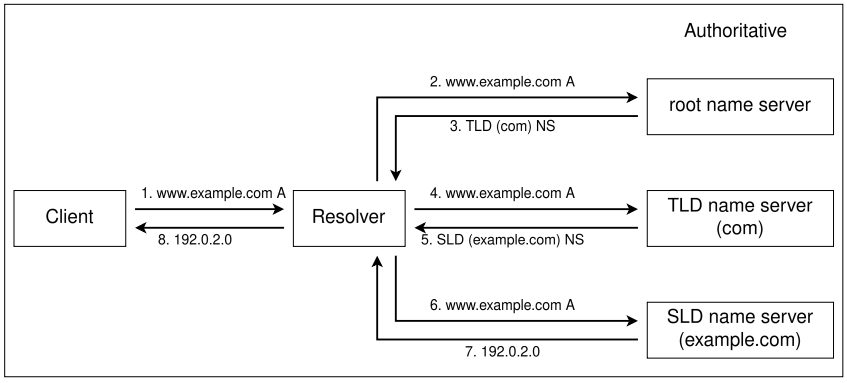

Figure 1: A client using a resolver to query for the IPv4 address (A record) of www.example.com. (Diagram from Paper II).

2.2 DNS Resolver

In Figure 1 we can see the eight steps from client query to resolver response. When querying for the IPv4 address of a domain name, a DNS client sends a query to a preconfigured DNS resolver, requesting the `A` record. The resolver then queries a root name server, which responds with a referral (`NS` record) to a TLD name server. The process continues through multiple referrals until the resolver obtains the correct `A` record from the authoritative server of the requested name. The resolver then forwards the response to the client. When receiving a query from a client, the preconfigured resolver may forward the query to another resolver, forming a resolver chain. This forwarding can improve the cache hit in terms of performance and benefit from the potential filtering of malicious domains. To increase resiliency, the preconfigured resolver may also forward the query to multiple resolvers in parallel.

Resolvers may specialize in serving specific client groups. Private resolvers handle internal queries on private networks, while resolvers operated by an Internet Service Provider (ISP) serve clients connected to their Autonomous System (AS). Public resolvers, such as Google Public DNS [14], Cloudflare DNS [5], OpenDNS [35], and Quad9 [41], offer global services. Private and AS resolvers may forward queries to other resolvers to leverage their caching as well as security and privacy features.

2.3 Resolver Security and Privacy

The CIA triad (i.e., Confidentiality, Integrity, and Availability) is a widely accepted model for information security policies. It defines how information is protected against unauthorized access (confidentiality) and unauthorized modification (integrity) while remaining accessible to authorized users (availability) [1]. Researchers and developers have proposed and implemented several features to address these security aspects in the context of DNS resolvers.

Several transport encryption methods ensure the confidentiality of DNS queries and responses, protecting them from interception or monitoring by unauthorized entities. An unauthorized entity in this threat model is a third-party evavesdropper, e.g., network operators, ISPs, or hackers intercepting traffic. If the resolver is untrustworthy, the security of the delivery path is irrelevant. DNS over TLS (DoT) [19] uses Transport Layer Security (TLS) to secure communication between the client and DNS resolver, though it is easily blocked due to its unique port (853/TCP). DNS over HTTPS (DoH) [18] encrypts DNS traffic using HTTPS, making DNS queries indistinguishable from regular web traffic, thus protecting the content and making it harder to selectively block or filter based on query. DNS over QUIC (DoQ) [20] uses the QUIC transport protocol to encrypt DNS communications, enhancing confidentiality while providing improved performance and reduced latency compared to the other two previous methods.

The integrity of DNS data is crucial to prevent unauthorized modifications and ensure that users are directed to the correct IP addresses. DNS Security Extensions (DNSSEC) [17] adds cryptographic signatures to DNS data, allowing clients and resolvers to verify the authenticity and integrity of the responses they receive. By validating these DNSSEC signatures, the resolver can ensure that the data has not been altered or tampered with during transit and come from the correct source. DNSSEC protects against attacks such as cache poisoning and fake name servers [2].

Maintaining the availability of DNS services is vital for ensuring continuous access to Internet resources. Anycasting [37] involves using multiple geographically distributed servers with the same IP address to handle DNS queries. This technique improves the availability and resilience of DNS services by reducing latency, balancing the load across multiple servers, and ensuring that DNS services remain accessible even if some servers experience failures. Another security related feature is the capability to resolve queries for resources located on authoritative name servers available on IPv6. If the resolver only has IPv4 capabilities, this could result in unavailability of certain resources.

Query Name Minimization (QMIN) [4] is a technique that reduces the amount of information sent in DNS queries to the authoritative name servers. This method limits the data exposure to the authoritative name servers by sending only the minimal necessary portion of the query to each name server in the resolution process.

A part of Extended Mechanisms for DNS (EDNS) called EDNS Client Subnet (ECS) [6] enables resolvers to share part of the client's IP address to the authoritative name server. While its intended use is to provide the client with a service geographically closer to them, this is a *negative* privacy feature. Turning off ECS is regarded as a privacy-friendly configuration for DNS resolvers [24].

Oblivious DNS (ODNS) [43] adds a layer of privacy by separating the DNS query from the client identity by introducing a proxy. The query is encrypted and sent to a proxy resolver, which forwards the encrypted query to another resolver which can decrypt the query. The first resolver will only know the identity of the client, and the second resolver will only know the content of the query. A combination of ODNS and DoH [23] has also been proposed. In addition, it is possible to route DoH through the Tor anonymity network [45] to anonymize the client [31].

A resolver can also be configured to send minimal responses to the client. This means that the client will only receive the resource records they requested and nothing else. Other information about the authoritative name servers is in most cases useless to the end-user and is usually intended for other DNS resolvers. Turning on minimal responses is a form of data minimization in the context of DNS queries.

DNS filtering helps protect users from malicious websites and prevents malware from communicating from an infected machine [12, 21, 46]. Denying access to known phishing websites prevents sensitive information such as usernames, passwords, and personal data from being exposed to unauthorized parties. Selectively blocking traffic from an infected client restricts the malware's capabilities and can prevent much harm. However, the exact same mechanism can be used for censorship [16]. By dropping queries, this tool becomes a threat to availability, and redirecting users to other servers becomes a threat to integrity.

3 Research Questions

The main research objective of this thesis is to *conduct a comprehensive investigation into the security and privacy of DNS resolvers.* We formulate the following research questions to investigate security and privacy aspects in DNS resolvers:

1. *What are the trends in the adoption and implementation of privacy and security features in DNS resolvers?*

   As with many systems from the early Internet, security and privacy were rarely expressed in the core design of DNS. Features improving on these aspects have, therefore, been added as an afterthought. By analyzing the trends in the adoption and implementation of these features we can (1) establish snapshots of the current state of DNS resolver privacy and security, and (2) explore variations in implementation across software and versions.

2. *How do privacy and security features in DNS resolvers impact the clients and the ecosystem as a whole?*

   The adoption, implementation and combination of these features on DNS resolvers could impact clients by affecting their data privacy and security both positively and negatively. For the ecosystem—specifically other resolvers and authoritative name servers—these changes could influence traffic patterns, trust dynamics, and interoperability, potentially guiding the evolution of standards across DNS services.

4 Research Methods

Our studies employ a set of methods to investigate the features and operational behaviors of DNS infrastructure. To effectively capture the diverse aspects of DNS technologies, we utilize active and passive measurements, testbed simulations, and we leverage the benefits of open-source tools and datasets. Each method is chosen to address specific research questions.

Active and Passive Measurements: Active measurements involve generating traffic to study how real-world systems behave. The advantages of using active measurements include real-time data collection, direct control over test conditions, and measurement customizability. It is essential to be aware and explicit about the measurement vantage point and the tool or system used to perform the measurements. The limitations of using active measurements include biased vantage points, affecting the system under evaluation, and due to the nature of the ever-changing Internet, low repeatability, replicability, and reproducability [3]. Passive measurements entails monitoring real traffic in a system. The advantage of this approach is that it is possible to analyse the system under normal conditions without impacting its behavior. However, this approach may result in limited insights depending on the traffic or vantage point, and may raise privacy concerns about the collected data.

In the context of DNS, active measurements allow us to examine query behaviors from both the client-side and from an authoritative name server under our control. Conversely, passive measurements involve analyzing naturally occurring traffic to observe DNS behavior under normal operational conditions. When performing active measurements it is important to spread out the traffic over time to not put too much load on the networks and transparently communicate the nature of the measurement with contact information and ways to opt out [10].

When analyzing traffic patterns and resolver behavior in DNS measurements, understanding the effects of caching is crucial. Caching can obscure the actual behavior of resolvers, as repeated queries for the same domain are answered quickly from the cache. Generating unique queries that are unlikely to be cached is essential to assess resolver behavior and

traffic patterns accurately [8, 36]. This approach ensures that each query requires the resolver to perform a complete DNS resolution, allowing for a detailed analysis of the resolver's interactions with authoritative servers.

RIPE Atlas is a globally distributed network of probes and anchors designed to measure Internet connectivity and real-time performance [32]. These probes can perform network measurements, including ping, traceroute, DNS queries, and TLS measurements. Due to their diverse nature, we can query DNS resolvers which are usually unreachable from our own vantage points, such as private resolvers and resolvers operated by ISPs. When exploring what types of DNS queries we could send from these probes, we were unable to perform adoption measurements for encrypted DNS transport capabilities.

Testbeds: While active and passive measurements provide valuable insights, testbeds allow for emulating DNS operations in a controlled environment. One of the biggest advantages with testbeds is reproducability since they are isolated from the unpredictable nature of the Internet [34, 11]. However, testbeds come with certain limitations. They might not fully capture the complexity and variability of real-world conditions, leading to results that might not generalize well to the actual Internet environment. The controlled nature of testbeds can sometimes oversimplify the intricate interactions and diverse traffic patterns observed in live networks. Despite these limitations, testbeds are invaluable for initial hypothesis testing and controlled experimentation, providing foundational insights that can be further validated through live measurements.

Automation, Documentation, and Open-source: To ensure the reproducibility and reliability of our findings, we meticulously document each step of our research and automate processes wherever possible. This automation minimizes human error and maximizes efficiency, leading to more consistent and systematic data collection and analysis. We further promote transparency and reproducibility by making datasets and tools publicly available where possible [38, 15]. This practice allows peer verification and fosters global collaboration, enhancing innovation and the development of DNS research.

Application of Methods in Our Studies: Initially, we conducted a literature survey on DNS in the context of security and privacy to identify trends, gaps, and a good starting point based on our available capabilities. The results from this survey led us to design a second look at the adoption of QMIN at DNS resolvers (Paper I). In Paper I, we combine active and passive measurements based on methodologies from prior studies, leveraging publicly available tools for reproducibility. In Paper II, we utilize testbed simulations to establish DNS query patterns based on methodologies from prior studies before applying active measurements. In Paper III we performed various active measurements on DNS resolvers to

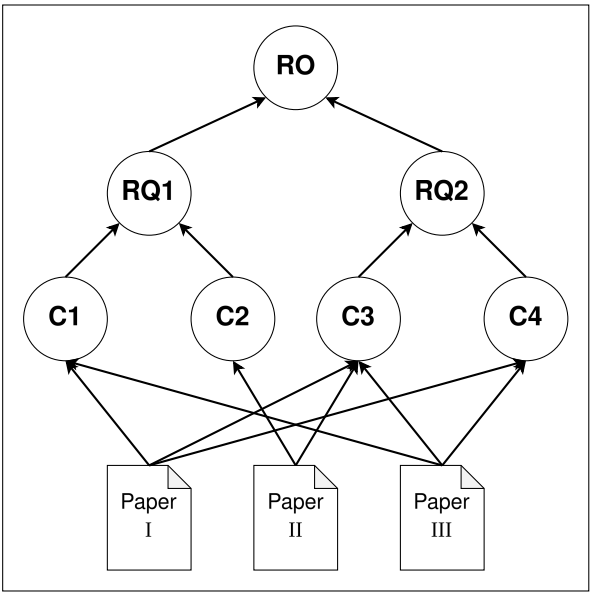

Figure 2: Overview of the research objective, research questions, contributions, and appended papers.

assert whether they adopted security and privacy features. In all papers we automated some of the measurements alongside detailed documentation to improve repeatability. We also made the testbed in Paper II publicly available on GitHub[1] for replicability.

5 Contributions

The key contributions of this thesis are summarized below. A depiction of their connection to our research questions and the associated papers is illustrated in Figure 2.

1. *Show the current level of adoption, or adoption over time, of security and privacy features in DNS resolvers.*

   In Paper I, we used active measurements from RIPE Atlas probes and open resolvers, revealing a significant increase in QMIN adoption, from 2.5k RIPE Atlas probes in 2018 to 14k in 2022 and from 18k open resolvers in 2018 to 80k in 2022. The rapid increase in 2020 could be attributed to Google Public DNS, a popular provider of DNS resolution, enabling the feature. Additionally, we extended passive measurements at root and TLD name servers over a longer period, filtering out non-valid queries and incorporating additional

[1] https://github.com/Arcnilya/resolver-lab

data sources. We reported that QMIN adoption increased at one root server from 0.6% (2018) to 2.5% (2022) and at the .nl TLD from 35.5% (2019) to 57.3% (2022).

In Paper III, we used active measurements from RIPE Atlas probes located in Sweden to measure adoption of five security and privacy related features at their preconfigured resolvers. We reported that 86% of the resolvers supported IPv6, 85% validated DNSSEC signatures, 85% minimized queries using QMIN, 86% had turned off ECS, and 93% of the resolvers returned minimal responses. Similar measurements in this context have not been done before. In a study by Nosyk *et al.* [33], less than 18% of open resolvers, and only 37.4% of closed non-forwarding resolvers on IPv4 validated DNSSEC signatures. Comparing QMIN adoption in this study with Paper I (64% of resolvers used by RIPE Atlas probes) and the study before it by De Vries *et al.* [8] (10%), we see a relatively high adoption of 85%.

2. *Analyze variations and changes in the implementation of security and privacy features in DNS resolvers across software and versions.*

   In Paper II, we established query patterns from four popular open-source resolvers implementing QMIN (Bind [22], Unbound [27], Knot Resolver [7], and PowerDNS Recursor [40]). We also compared patterns from each resolver across different versions. Every provider had a distinct signature that could be identified in our active measurements. We could also establish that two resolvers in our study (Bind and PowerDNS Recursor) had changed their implementation of QMIN, resulting in new unique patterns which allowed us to identify old versions in the wild.

3. *Investigate how the client choice of DNS resolver, and resolver forwarding behavior, affects the availability and effectiveness of security and privacy features.*

   In Paper I, we identified a subset of resolvers, which inconsistently adopted QMIN. We discovered this by sending multiple queries to each open resolver and put correctly responding resolvers in three categories: QMIN (10.8%) not-QMIN (72.8%) and conflicting (16.4%). This result showed that the previous adoption measurements of QMIN could be missing crucial complexities that reflect the true behavior of resolvers.

   In Paper II, we looked at query patterns from resolvers implementing QMIN to fingerprint software and in some cases versions. When looking at the query patterns, we noticed that some resolvers forwarded to multiple other resolvers in parallel. We highlighed instances where QMIN was undermined because a resolver forwarded the query to both a minimizing and a non-minimizing resolver, effectively negating the privacy feature. This combination of QMIN and parallel

forwarding also contributed to instances of query amplification, which could put stress on authoritative name servers.

In Paper III, we investigated the adoption of various security and privacy features at resolvers used by RIPE Atlas probes in Sweden. We also categorized the resolvers by their network proximity to the probe, and analyzed their forwarding behaviors. We found that the availability of features varies depending on these categories and forwarding behaviors. The IPv6 capabilities were only at 44% when using a private resolver without forwarding. This implies that a substantial amount of home routers are unable to connect via IPv6. However, we observed that these private resolvers always provided minimal responses back to the client. We found that IPv6 and DNSSEC were almost always adopted when using a public resolver (over 98%) , but there was less adoption of privacy-focused features (around 70%). Looking closer at the three most popular providers of public DNS, we saw that Google Public DNS did not seem to use QMIN at all (a limitation of our measurement method when observing queries at the Second-Level Domain). Google Public DNS was the only resolver sharing part of the client's IP using ECS, and did not provide any minimal responses.

4. *Cross-analyze security and privacy features in DNS resolvers by looking at correlations in adoption.*

   In Paper I, we looked at QMIN in the context of other Internet standards, including IPv6 and DNSSEC support. The passive analysis at .nl ccTLD name servers shows that resolvers with IPv6 capabilities and signaling support of DNSSEC by setting the DO-bit are likelier to implement QMIN. In contrast, those not following these practices have a lower adoption rate of QMIN.

   In Paper III, we looked closer at the adoption correlation between the five security and privacy features in the scope of the study. After normalizing the dataset to only include unique resolvers to prevent popular resolvers to skew the analysis, we found positive correlations between the privacy and data minimization related features (QMIN, not using ECS, and minimal responses). We also observed positive correlations between IPv6 capabilities and DNSSEC, and between DNSSEC and QMIN. The correlations between IPv6 capabilities and privacy related features were negligible, indicating that the former is adopted in isolation from privacy considerations.

6 Summary of Appended Papers

In this section we summarize each of the appended papers.

Paper I – A Second Look at DNS QNAME Minimization

In this paper we examine the adoption and impact of QMIN in DNS. We build upon the work of De Vries et al. [8] by extending both active and passive measurement methodologies to provide a comprehensive view of QMIN adoption from 2018 to 2022. Through active measurements using RIPE Atlas probes and open resolvers, we observe a significant increase in QMIN adoption. Specifically, we find that the number of QMIN-enabled resolvers has grown from 2.5k RIPE Atlas probes in 2018 to 14k in 2022, and from 18k open resolvers to 80k during the same period. This highlights a marked improvement in the implementation of QMIN across the Internet. Our passive measurements at the root and TLD name servers also show a positive trend, with QMIN adoption increasing from 0.6% to 2.5% at one root server and from 35.5% to 57.3% at the .nl TLD. These measurements are enhanced by filtering out non-valid queries and incorporating additional data sources, providing a clearer picture of the impact of QMIN. Additionally, we identify "conflicting resolvers" that inconsistently apply QMIN, revealing variability in the forwarding configurations. We discuss the trade-off between privacy and performance inherent in QMIN and propose using a public suffix list to set the depth limit for minimizing labels. This approach aims to balance privacy needs with performance considerations. Overall, our paper highlights the significant progress in QMIN adoption, improvements in resolver performance, and the challenges in achieving consistent QMIN implementation across different DNS resolvers.

Paper II – Fingerprinting DNS Resolvers using Query Patterns from QNAME Minimization

This study aims to evaluate if patterns in queries from DNS resolvers—implementing QMIN as a privacy enhancing feature—can reveal their characteristics such as their software and versions. We examined the query patterns of minimizing resolvers at the authoritative name server side, and our findings indicate that distinct patterns correlate with specific open-source resolver software versions. Notably, none of the resolvers fully follow the recommended QMIN algorithm outlined in RFC 9156, suggesting a discrepancy between recommendations and real-world implementations. We also identified high rates of query amplification, possibly caused in part by the combination of minimization and forwarding configurations. Our research contributes to understanding the current state of the DNS ecosystem, highlighting the potential for fingerprinting to enhance Internet security by identifying and addressing resolver-related risks.

Paper III – SweDNS: Evaluating Privacy and Security of DNS Resolvers used in Sweden

In this study, we examine DNS resolvers used by clients in Sweden, conducting active measurements to assess the adoption of security and privacy standards. We categorize the resolvers based on their network proximity to the client, allowing for more in-depth analysis. We utilize the RIPE Atlas network of volunteer-run probes for our measurements in October 2024 and reveal that 86% of the identified resolvers supported IPv6, 85% were validating DNSSEC, 85% implemented QMIN, 86% avoided using ECS, and 93% returned minimal responses to the client. A deeper analysis of private, within-AS, and public (outside-AS) resolvers shows varying levels of feature adoption across these categories. We also identified strong correlations between privacy-focused features and links between DNSSEC and both QMIN and IPv6 support.

7 Related Work

In this section we enumerate work related to this study to highlight our contributions to the research area.

7.1 Security and Privacy Feature Adoption

De Vries *et al.* [8] conducted a study on the adoption of QMIN from April 2017 to October 2018. They used active and passive measurements to show that QMIN adoption was slow but steady, improving query privacy at root and TLD name servers. They developed a method to measure QMIN adoption accurately by introducing a wildcard label to ensure unique queries in active measurements. In our studies we build upon their work by extending the scope to more resolvers, and look at QMIN adoption over time. In 2021, Yajima *et al.* [47] did a large-scale analysis of security mechanisms at authoritative name servers. They put DNS resolvers and their related mechanisms outside of their scope. Similar to their study, we also investigate features in DNS in a cross-sectional manner in one of our studies, but we focus on both privacy and security related features, and limit our scope to resolvers. Previous approaches to identify DNSSEC-validating resolvers either used paid services, had access to privileged data, or limited coverage until Nosyk *et al.* [33] proposed a novel technique utilizing active measurements. Setting up an authoritative name server with misconfigured domains, they classify resolvers by analyzing the Response Codes (RCODEs) from their active measurements. A less extensive version of this technique is used in one of our studies to measure adoption of DNSSEC validation. Saluja *et al.* [42] used RIPE Atlas to investigate the causes behind the lack of IPv6 capabilities in DNS. They identified vantage point *islands* and *peninsulas* where IPv6 were not available. These findings could explain the low adoption of IPv6 capabilities for private resolvers in one of our studies.

7.2 DNS Software Fingerprinting

DNS software fingerprinting mainly involves two methods. The first is the *fpdns* tool [25], which, though innovative initially, is now outdated and less effective for modern DNS software. It operates from the client side, identifying software and versions based on responses to crafted queries. The newer *dnssoftver* tool [48] serves as an alternative. The second method uses `version.bind` queries, originally specific to Bind but adopted by other servers under similar labels like `version.server`. This method provides server software information, although many servers obfuscate it for security. Studies have applied these methods to analyze DNS software.

Sisson [44] combined them to fingerprint a significant subset of the IPv4 address space. After the "no match" category, they found Bind9 as the most common software but noted the potential for false results. Using fpdns, Glynn [13] identified software on `.ie` domain servers, highlighting older, less secure Bind versions. Kührer *et al.* [26] employed `version.bind` and `version.server` queries to classify public resolvers over a year, revealing a landscape dominated by various Bind versions, Unbound, Dnsmasq [9], PowerDNS recursor, and Microsoft DNS [28]. When we fingerprinted resolver software we found that Bind remains popular even today. By scanning the entire IPv4 address space, Park *et al.* [36] uncovered over 3 million resolvers. They tracked behavior by running an authoritative server and monitoring query responses. Many of these resolvers were found to deliver incorrect or malicious responses. De Vries *et al.* [8] introduced a new approach to fingerprint DNS resolvers specifically. They analyzed query patterns, or *signatures*, from minimizing resolvers used by RIPE Atlas probes. These signatures varied across resolver software. In their 2018 study, they sent unique queries to RIPE Atlas probes to establish QMIN signatures and mapped these to Bind 9.13.3, Knot 3.0.0, and Unbound 1.8.0. Their analysis showed that real-world implementations often differed from RFC 7816, aiming to reduce errors and improve performance. We built our fingerprinting study based on the methodology by De Vries *et al.*, improved the analysis of the QMIN patterns, and extended their measurements.

8 Conclusions and Future Work

The research objective of this thesis has been to conduct a comprehensive investigation into the security and privacy of DNS resolvers. Using active and passive measurements, isolated testbeds, and specialized authoritative name servers, we have analyzed the adoption and implementation of security and privacy features of DNS resolvers on the Internet. By analyzing query patterns, query responses, and forwarding strategies, we have established current adoption levels of security and privacy features, found measurement limitations, identified privacy-compromising config-

urations, discovered variations in feature implementation, and observed adoption correlations between features. The nature of the studies in this thesis has been mainly explorative, measuring and analyzing how resolvers behave on the Internet today. Moving forward, we aim to use our findings to help the ecosystem with adoption and addressing misconfigurations compromising security and privacy in DNS.

References

[1] Information technology - Security techniques - Information security management systems - Overview and vocabulary. ISO/IEC 27000:2018(E), International Organization for Standardization (ISO), Geneva, CH, 2018. https://www.iso.org/standard/73906.html.

[2] D. Atkins and R. Austein. Threat analysis of the domain name system (dns). RFC 3833, RFC Editor, August 2004.

[3] Vaibhav Bajpai, Anna Brunstrom, Anja Feldmann, Wolfgang Kellerer, Aiko Pras, Henning Schulzrinne, Georgios Smaragdakis, Matthias Wählisch, and Klaus Wehrle. The dagstuhl beginners guide to reproducibility for experimental networking research. *ACM SIGCOMM Computer Communication Review*, 49(1):24–30, 2019.

[4] S. Bortzmeyer, R. Dolmans, and P. Hoffman. DNS query name minimisation to improve privacy. RFC 9156, RFC Editor, November 2021.

[5] Cloudflare. What is 1.1.1.1? | cloudflare. https://www.cloudflare.com/learning/dns/what-is-1.1.1.1/. Accessed: 2024-05-29.

[6] C. Contavalli, W. van der Gaast, D. Lawrence, and W. Kumari. Client subnet in DNS queries. RFC 7871, RFC Editor, May 2016.

[7] CZ-NIC. Knot resolver. https://www.knot-resolver.cz/. Accessed: 2024-11-15.

[8] Wouter B de Vries, Quirin Scheitle, Moritz Müller, Willem Toorop, Ralph Dolmans, and Roland van Rijswijk-Deij. A first look at QNAME minimization in the domain name system. In *Passive and Active Measurement (PAM): 20th International Conference*, pages 147–160. Springer, 2019.

[9] Dnsmasq. Dnsmasq - network services for small networks. https://dnsmasq.org/doc.html. Accessed: 2024-11-25.

[10] Zakir Durumeric, Eric Wustrow, and J Alex Halderman. ZMap: Fast internet-wide scanning and its security applications. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 605–620, 2013.

[11] Sarah Edwards, Xuan Liu, and Niky Riga. Creating repeatable computer science and networking experiments on shared, public testbeds. *ACM SIGOPS Operating Systems Review*, 49(1):90–99, 2015.

[12] Martin Fejrskov, Jens Myrup Pedersen, and Emmanouil Vasilomanolakis. Using netflow to measure the impact of deploying DNS-based blacklists. In *Security and Privacy in Communication Networks: 17th EAI International Conference, SecureComm 2021, Virtual Event, September 6–9, 2021, Proceedings, Part I 17*, pages 476–496. Springer, 2021.

[13] William J Glynn. Measuring DNS vulnerabilities and DNSSEC challenges from an irish perspective. *Proceedings of SATIN*, 2011.

[14] Google. Public DNS | google for developers. https://developers.google.com/speed/public-dns/. Accessed: 2024-05-29.

[15] Wilhelm Hasselbring, Leslie Carr, Simon Hettrick, Heather Packer, and Thanassis Tiropanis. Open source research software. *Computer*, 53(8):84–88, 2020.

[16] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. How great is the great firewall? measuring china's DNS censorship. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3381–3398, 2021.

[17] P. Hoffman. DNS security extensions (DNSSEC). RFC 9364, RFC Editor, February 2023.

[18] P. Hoffman and P. McManus. DNS queries over HTTPS (DoH). RFC 8484, RFC Editor, October 2018.

[19] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. Hoffman. Specification for DNS over transport layer security (TLS). RFC 7858, RFC Editor, May 2016.

[20] C. Huitema, S. Dickinson, and A. Mankin. DNS over dedicated QUIC connections. RFC 9250, RFC Editor, May 2022.

[21] Hikaru Ichise, Yong Jin, and Katsuyoshi Iida. Policy-based detection and blocking system for abnormal direct outbound DNS queries using RPZ. In *International Workshop on Computer Science and Engineering*. WCSE, 2022.

[22] ISC. Bind 9. https://www.isc.org/bind/. Accessed: 2024-11-15.

[23] E. Kinnear, P. McManus, T. Pauly, T. Verma, and C.A. Wood. Oblivious DNS over HTTPS. RFC 9230, RFC Editor, June 2022.

[24] Panagiotis Kintis, Yacin Nadji, David Dagon, Michael Farrell, and Manos Antonakakis. Understanding the privacy implications of ecs. In *Detection of Intrusions and Malware, and Vulnerability Assessment: 13th International Conference, DIMVA 2016, San Sebastián, Spain, July 7-8, 2016, Proceedings 13*, pages 343–353. Springer, 2016.

[25] kirei. Github: kirei/fpdns. https://github.com/kirei/fpdns. Accessed: 2024-05-29.

[26] Marc Kührer, Thomas Hupperich, Jonas Bushart, Christian Rossow, and Thorsten Holz. Going wild: Large-scale classification of open DNS resolvers. In *Proceedings of the 2015 Internet Measurement Conference*, pages 355–368, 2015.

[27] NLnet Labs. Unbound. https://nlnetlabs.nl/projects/unbound/about/. Accessed: 2024-11-15.

[28] Microsoft. Domain name system (DNS) | microsoft learn. https://learn.microsoft.com/en-us/windows-server/networking/dns/dns-top. Accessed: 2024-11-25.

[29] P. Mockapetris. Domain names - concepts and facilities. RFC 1034, RFC Editor, November 1987. http://www.rfc-editor.org/rfc/rfc1034.txt.

[30] P. Mockapetris. Domain names - implementation and specification. RFC 1035, RFC Editor, November 1987. http://www.rfc-editor.org/rfc/rfc1035.txt.

[31] Alec Muffet. DoHoT: making practical use of DNS over HTTPS over Tor. https://github.com/alecmuffett/dohot. Accessed: 2024-10-29.

[32] RIPE NCC. RIPE atlas. https://atlas.ripe.net/, 2010. Accessed: 2024-06-13.

[33] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. Guardians of DNS integrity: A remote method for identifying DNSSEC validators across the internet. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1470–1479. IEEE, 2023.

[34] Lucas Nussbaum. Testbeds in computer science. In *Reproducible Research Webinars*, 2017.

[35] OpenDNS. Setup guide | OpenDNS. https://www.opendns.com/setupguide/. Accessed: 2024-05-29.

[36] Jeman Park, Aminollah Khormali, Manar Mohaisen, and Aziz Mohaisen. Where are you taking me? behavioral analysis of open DNS resolvers. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 493–504. IEEE, 2019.

[37] Craig Partridge, Trevor Mendez, and Walter Milliken. Host anycasting service. RFC 1546, RFC Editor, November 1993. http://www.rfc-editor.org/rfc/rfc1546.txt.

[38] Vern Paxson. Strategies for sound internet measurement. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*, pages 263–271, 2004.

[39] Jon Postel. Internet protocol. RFC 791, RFC Editor, September 1981. http://www.rfc-editor.org/rfc/rfc791.txt.

[40] PowerDNS. PowerDNS recursor. https://www.powerdns.com/powerdns-recursor. Accessed: 2024-11-15.

[41] Quad9. Quad9 | a public and free DNS service for a better security and privacy. https://www.quad9.net/. Accessed: 2024-05-29.

[42] Tarang Saluja, John Heidemann, and Yuri Pradkin. Differences in monitoring the dns root over ipv4 and ipv6. In *2022 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*, pages 194–203. IEEE, 2022.

[43] Paul Schmitt, Anne Edmundson, and Nick Feamster. Oblivious DNS: Practical privacy for DNS queries. *arXiv preprint arXiv:1806.00276*, 2018.

[44] Geoffrey Sisson. DNS survey. http://www.open-spf.org/surveys/201010/dns_survey_2010.pdf, October 2010. Accessed: 2024-12-03.

[45] The Tor Project. Tor Project. https://www.torproject.org/. Accessed: 2024-11-15.

[46] Norman Wilde, Lauren Jones, Robert Lopez, and Travis Vaughn. A DNS RPZ firewall and current american DNS practice. In *Information Science and Applications 2018: ICISA 2018*, pages 259–265. Springer, 2019.

[47] Masanori Yajima, Daiki Chiba, Yoshiro Yoneya, and Tatsuya Mori. Measuring adoption of dns security mechanisms with cross-sectional approach. In *2021 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6. IEEE, 2021.

[48] yevheniya nosyk. Github: yevheniya-nosyk/dnssoftver. https://github.com/yevheniya-nosyk/dnssoftver. Accessed: 2024-05-29.

Cybersecurity Mechanisms in DNS Resolvers

Using the Internet today, both end-users and automated systems rely on the Domain Name System (DNS) to translate human-readable domain names to IP addresses for communication between machines. This system from 1985 has only in recent years seen Internet standards addressing security and privacy concerns. In the position as a machine-in-the-middle between the client and the distributed hierachical system of authoritative name servers, we find the DNS resolver. Due to its purpose of forwarding, looking up, and caching queries and responses, in addition to its location between the clients and the name servers, the DNS resolver becomes a critical point for implementing these security and privacy features. The widespread adoption of these features, their variation in implementation, and impact on both clients and other name servers remain as interesting topics in the research community. The goal of this thesis is to analyze servers in the wild and conduct a comprehensive investigation into the security and privacy mechanisms configured on DNS resolvers. Using an Internet measurement approach, we explore the trends in the adoption and implementation of these features by generating and observing our own queries to and from the resolvers. We also investigate how clients and the DNS ecosystem as a whole are impacted by resolver configurations. We use and improve methods for measuring adoption of various security and privacy related features. Based on these measurements we report the current level of adoption and adoption over time, investigate anomalies, and identify limitations with measurement approaches. We fingerprint the software and version of popular open-source DNS resolvers by classifying query patterns. Comparing the ingress and egress resolvers we analyze forwarding behaviors and their impact on the availability and effectiveness of security and privacy features. We also cross-analyze features in DNS resolvers to find correlations, which could help us understand obstacles and find solutions to feature adoption.

ISBN 978-91-7867-518-0 (print)

ISBN 978-91-7867-519-7 (pdf)

ISSN 1403-8099

LICENTIATE THESIS | Karlstad University Studies | 2025:1