



Measuring Adoption, Implementation, and Impact of Available Privacy Solutions in DNS Resolvers

Jonathan Magnusson

Faculty of Health, Science and Technology

Computer Science

DOCTORAL THESIS | Karlstad University Studies | 2026:33

Measuring Adoption, Implementation, and Impact of Available Privacy Solutions in DNS Resolvers

Jonathan Magnusson

Measuring Adoption, Implementation, and Impact of Available Privacy Solutions in
DNS Resolvers

Jonathan Magnusson

DOCTORAL THESIS

Karlstad University Studies | 2026:33

urn:nbn:se:kau:diva-109643

ISSN 1403-8099

ISBN 978-91-7867-716-0 (print)

ISBN 978-91-7867-717-7 (pdf)

<https://doi.org/10.59217/mwtl6138>

© The author

Distribution:

Karlstad University

Faculty of Health, Science and Technology

Department of Mathematics and Computer Science

SE-651 88 Karlstad, Sweden

+46 54 700 10 00

Print: Universitetstryckeriet, Karlstad 2026

WWW.KAU.SE

Measuring Adoption, Implementation, and Impact of Available Privacy Solutions in DNS Resolvers

JONATHAN MAGNUSSON

Department of Mathematics and Computer Science

Abstract

The Domain Name System (DNS) is a fundamental component of the Internet, yet its resolution process exposes highly sensitive information about client activity. DNS resolvers, positioned between clients and authoritative name servers, typically observe both query contents and client identifiers, making them a focal point for privacy risks and a natural target for privacy-enhancing mechanisms. This thesis investigates privacy challenges involving DNS resolvers and analyzes how proposed technical solutions are deployed and implemented in practice.

The thesis adopts an empirical approach, using active and passive Internet measurements from multiple vantage points to study resolver behavior. The results show substantial growth in the adoption of DNS privacy mechanisms such as Query Name Minimization (QMIN), largely driven by defaults introduced by major public resolvers. However, adoption remains uneven across regions and resolver types, and persistently low for mechanisms enabling encrypted resolver discovery among non-public resolvers.

We identify significant deployability challenges: strict QMIN configurations cause resolution failures, increased latency, and higher traffic overhead due to widespread misbehavior by authoritative servers, while DNS resolution over the Tor anonymity network provides strong anonymity guarantees at the cost of reduced performance. The thesis further reveals systematic deviations from standards that enable fingerprinting of resolver software, as well as widespread misconfigurations in encrypted DNS transport signaling that reinforce resolver centralization. Overall, while privacy-enhancing DNS mechanisms are increasingly deployed, their real-world effectiveness is constrained by implementation gaps, performance trade-offs, and centralization pressures, highlighting the need for improved interoperability, validation tools, and greater attention to operational realities.

Keywords: Domain Name System, Resolver, Privacy, Internet Measurements

Mätning av införande, implementering och effekter av tillgängliga integritetslösningar i DNS-resolvrar

JONATHAN MAGNUSSON

Institutionen för matematik och datavetenskap

Sammanfattning

Domännamnssystemet (DNS) är en grundläggande komponent i internetns infrastruktur, men dess uppslagningsprocess exponerar omfattande känslig information om klienternas aktivitet. DNS-resolvrar, som är placerade mellan klienter och auktoritativa namnservrar, observerar vanligtvis både innehållet i förfrågningarna och klientidentifierare, vilket gör dem till en central punkt för integritetsrisker och en naturlig plats för integritetsskyddande mekanismer. Denna avhandling undersöker integritetsutmaningar relaterade till DNS-resolvrar och analyserar hur föreslagna tekniska lösningar används och implementeras i praktiken.

Avhandlingen använder ett empiriskt tillvägagångssätt, med aktiva och passiva internetmätningar från ett flertal olika utsiktspunkter för att studera beteendet av resolvrar. Resultaten visar en betydande ökning i användningen av integritetsmekanismer såsom Query Name Minimization (QMIN), i stor utsträckning driven av standardinställningar som införts av stora offentliga resolvrar. Användningen är dock fortfarande ojämn mellan olika regioner och resolvertyper, och fortsatt låg för mekanismer som möjliggör upptäckande av krypterade anslutningspunkter bland icke-offentliga resolvrar.

Vi identifierar betydande utmaningar när det gäller implementerbarhet: strikta QMIN-konfigurationer orsakar fel i uppslagningen, ökad latens och högre trafikvolym på grund av felaktigt beteende hos auktoritativa servrar. DNS-uppslagning över anonymitetsnätverket Tor ger starka anonymitetsgarantier på bekostnad av prestandaförluster. Avhandlingen visar på systematiska avvikelser från standarder som möjliggör unik identifiering av mjukvara samt utbredd felkonfiguration i mekanismer för att signalera krypterad DNS-transport som förstärker centralisering. Även om mekanismer för integritetsskydd i DNS används i allt större utsträckning, begränsas deras faktiska effektivitet av brister i implementeringen, avvägningar runt prestanda samt centralisering. Detta understryker följaktligen behovet av förbättrad interoperabilitet, valideringsverktyg och större hänsyn till operatörers praktiska förutsättningar.

Nyckelord: Domännamnssystemet, resolver, personlig integritet, internetmätningar

Acknowledgements

Finally holding the printed dissertation after five years of research and studies feels surreal. It is more than a collection of black text on white pages; it represents my growth as a researcher, the collaborations and friendships I have formed, and my first of many contributions to making the Internet a better place.

This research was funded by the Swedish Internet Foundation. I am grateful for their support and commitment to fostering a better, safer Internet, which closely aligns with the goals of this work. I am deeply grateful to my supervisors, Tobias Pulls, Anna Brunstrom, and Johan Stenstam, for their guidance and support throughout this journey. I also want to thank everyone at the Department of Mathematics and Computer Science, especially the PriSec research group, for providing an inspiring and supportive environment. To all current and former students, thank you for the friendships, encouragement, and many laughs along the way.

I am very thankful to my parents and my sister for their unwavering support. To my friends, thank you for reminding me to take breaks, whether through adventures or games. And last but not least, Alexis, thank you for always having my back. I look forward to sharing the next chapter of life with you.

Some of you might not read this book cover to cover. Some of you might only glance at the colorful graphs. Some of you might simply place it on a shelf. And that is perfectly fine, just take good care of it, and know that it carries a piece of my journey.

Karlstad University , May 13, 2026

Jonathan Magnusson

List of Acronyms

ADoE	Authoritative DNS-over-Encryption
ANS	Authoritative Name Server
AS	Autonomous System
ccTLD	country-code Top-Level Domain
CDN	Content Delivery Network
CT	Certificate Transparency
DDoS	Distributed Denial-of-Service
DDR	Discovery of Designated Resolvers
DGA	Domain Generation Algorithms
DNS	Domain Name System
DNSSEC	DNS Security Extensions
DoE	DNS-over-Encryption
DoH	DNS-over-HTTPS
DoH3	DNS-over-HTTP/3
DoHoT	DNS-over-HTTPS-over-Tor
DoQ	DNS-over-QUIC
DoS	Denial-of-Service
DoT	DNS-over-TLS
DoTor	DNS-over-Tor
ECS	EDNS Client Subnet
EDNS	Extended Mechanisms for DNS
ENT	Empty Non-Terminal
FQDN	Fully Qualified Domain Name
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
ISP	Internet Service Provider
IXP	Internet Exchange Point

MAX MAX_MINIMIZE_COUNT

MIN MINIMIZE_ONE_LAB

NS Name Server

ODNS Oblivious DNS

ODoH Oblivious DNS-over-HTTPS

OS Operating System

PIR Private Information Retrieval

QMIN Query Name Minimization

QNAME Query Name

QTYPE Query Type

RA Recursion Available

RCODE Response Code

RFC Request For Comments

RPZ Response Policy Zone

RR Resource Record

RTT Round-Trip Time

SLD Second-Level Domain

SUDN Special Use Domain Name

SVCB Service Binding

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**. Privacy and Security of DNS Resolvers used in the Nordics and Baltics. Secure IT Systems: 30th Nordic Conference (NordSec 2025). Springer.
- IV. Jonathan Magnusson**, Rick Fontein, and Roland van Rijswijk-Deij. Measuring the Operational Impact of Minimizing DNS Queries on Real-World Data. Network Operations and Management Symposium (NOMS 2026). IEEE/IFIP.
- V. Jonathan Magnusson**, Alfred Arouna, Tor Avrill, and Tobias Pulls. DoHoT or Not? Tuning Tor Circuits for Anonymous DNS Queries. (Under Submission).
- VI. Jonathan Magnusson**. Here be DDRagons! Navigating the Unexplored Map of Encrypted DNS Transport. (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 I designed the adoption measurement study of DNS privacy and security features implemented on resolvers using RIPE Atlas probes. I performed all of the measurements, analyzed the data and wrote all of the paper as a single author, with feedback from my supervisors.

Paper IV In discussion with Roland and Rick, we designed an extensive measurement on the operational impact of QNAME Minimization. Using OpenINTEL as the measurement platform, Rick set up all of the resolvers and collected the data. My primary role in the study was to analyze the data and writing the paper together with Roland.

Paper V Tobias and I collaboratively designed the study and investigated methods for tuning circuits in Tor. During a DNS hackathon I pitched the study and recruited a group of interested people for some initial measurement tools and results. Alfred and Tor from the hackathon wished to continue collaborating on the study, focusing on building measurement tools for DoTor and ODoH. I performed the measurements, analyzed the data, and then wrote the majority of the paper with the help of Alfred, Tor, and Tobias.

Paper VI I designed an initial complementary study around measuring the adoption of DDR based on a recent publication which did not include resolvers hidden from Internet-wide scans. With my previous familiarity with RIPE Atlas, I explored ways to circumvent the lack of SVCB support and get a conservative lower bound for DDR adoption from the edge. Over discussions with my supervisors I added results from a second measurement platform: CAIDA Ark to complement my RIPE Atlas findings, as well as measurements on encrypted DNS transport on the authoritative side. 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.

- **Jonathan Magnusson**. Survey and Analysis of DNS Filtering Components. Preprint arXiv:2401.03864, 2024.

I designed and performed the literature study and wrote the paper, all with input from my supervisors.

- Leonardo Martucci, **Jonathan Magnusson**, Tobias Vehkajärvi and Jonas Karlsson. The Cyber Range Lite: Lightweight Infrastructure for Training and Education. World Conference on Information Security Education (WISE 2025). Springer.

I helped designing the Cyber Range Lite based on my own familiarity with Jeopardy-style Capture the Flag events and experience setting up an earlier cyber range. During the development I have mainly been focusing on challenge QA as well as the implementation of CTFd plugins for course-relevant features and cyber range integration. I wrote the paper together with Leonardo with input from Jonas and Tobias.

Contents

List of Acronyms	x
List of Appended Papers	xi
List of Other Contributions	xiii
INTRODUCTORY SUMMARY	1
1 Introduction	3
2 Background	4
2.1 Domain Name System	5
2.2 DNS Threat Model	7
2.3 Confidentiality	8
2.4 Integrity and Authentication	10
2.5 Availability	10
2.6 Anonymization	11
2.7 Data Minimization	12
3 Research Questions	14
4 Methodology	15
4.1 Repeatability	16
4.2 Replicability	16
4.3 Reproducibility	17
5 Contributions	18
6 Summary of Appended Papers	23
7 Related Work	26
7.1 QNAME Minimization	26
7.2 Centralization of Resolvers	27
7.3 DNS Client Anonymization	28
7.4 Encrypted DNS Transport	29
8 Conclusions and Future Work	30
PAPER I:	
A Second Look at DNS QNAME Minimization	41
1 Introduction	42

2	Background & Related Work	43
2.1	The Domain Name System	43
2.2	Query Name Minimization	44
2.3	Related Work	45
3	Active Measurements	46
3.1	Resolver Adoption Over Time	47
3.1.1	Method	47
3.1.2	Results	48
3.2	Adoption by Open Resolvers	50
3.2.1	Method	50
3.2.2	Results: Over Time and Location	51
3.2.3	Results: Conflicting Resolvers	52
3.2.4	Results: Unexpected Google	54
4	Passive Measurements	55
4.1	Method	55
4.2	Results	57
4.2.1	Results: Impact of Non-existing Domain Names	57
4.2.2	Results: Qmin Adoption in Detail	58
4.2.3	Results: Qmin Imperfections	59
5	Controlled Experiments	60
5.1	Method	60
5.2	Results	61
6	Discussion	62
6.1	Analysis of the Results	63
6.2	Improvements of Measurements Methods	63
6.3	Qmin Depth Limitation	64
7	Conclusion	65
PAPER II:		
Fingerprinting DNS Resolvers using Query Pat-		
terns from QNAME Minimization		71
1	Introduction	71
2	Background	73
2.1	Domain Name System	73
2.2	QNAME Minimization	74
2.3	QNAME Minimization Signatures	75
3	Method	75

4	Measurements	76
4.1	Establishing Signatures	76
4.2	Client Side Queries	77
4.3	Server Side Queries	78
4.4	Query Patterns and Signatures	79
5	Discussion	81
5.1	Signatures	81
5.2	Comparison with Previous Study	82
5.3	Query Amplification	82
5.4	Limitations	83
6	Related Work	83
7	Conclusion	84
A	Ethics	85
B	Query Amplification	85
C	Common Signatures	86

**PAPER III:
Privacy and Security of DNS Resolvers used in the
Nordics and Baltics** **95**

1	Introduction	96
2	Background	97
3	Method	99
4	Measurements	100
4.1	Routing Results	101
4.2	Security and Privacy Feature Adoption	102
4.3	Comparison by Country	104
4.4	Feature Correlation	106
5	Discussion	107
5.1	Limitations	109
6	Related Work	109
7	Conclusion	110
A	Query Filtering	111

**PAPER IV:
Measuring the Operational Impact of Minimizing
DNS Queries on Real-World Data** **117**

1	Introduction	117
2	Background and Related Work	119
2.1	QNAME Minimization	119
2.2	Related Work	120
3	Approach	121
4	Results	122
4.1	Dataset Demographics	122
4.2	Share of Successful Responses	123
4.2.1	The Umbrella List	123
4.2.2	The .com Namespace	125
4.2.3	Domains from CT logs	125
4.3	Impact on round-trip time	127
4.3.1	The Umbrella List	127
4.3.2	The .com Namespace	127
4.3.3	Domains from CT Logs	127
4.4	Traffic Overhead	128
4.4.1	The Umbrella List	129
4.4.2	Domains from CT logs	129
5	Discussion	129
6	Conclusion	131

PAPER V:
DoHoT or Not? Tuning Tor Circuits for Anonymous DNS Queries **137**

1	Introduction	137
2	Background and Related Work	139
2.1	Encrypted DNS Transport	139
2.2	Tor	139
2.3	Anonymous DNS Queries	140
2.3.1	Oblivious schemes	140
2.3.2	Tor-based schemes	142
2.3.3	Non-proxy schemes	142
3	Threat Models for Anonymous DNS and Tor	143
4	Method	144

5	Measurement Results	147
5.1	Query Time by Method (torrc vs. carml+Stem)	147
5.2	Query Time by Circuit (Specific Middle)	148
5.3	Query Time by Circuit (Without Middle)	148
5.4	Query Time by Circuit (DoTor carml+Stem)	149
5.5	Query Time by Approach	150
6	Discussion	150
6.1	Comparison with Related Work	150
6.2	Evaluation Dimensions of Anonymous DNS Schemes . . .	152
6.2.1	Latency	152
6.2.2	Threat Model	152
6.2.3	Deployability	152
6.3	Limitations	154
7	Conclusion	155
A	Ethical Considerations	155

PAPER VI:
Here be DDRagons! Navigating the Unexplored
Map of Encrypted DNS Transport **161**

1	Introduction	161
2	Background and Related Work	163
2.1	Encrypted DNS Transport	163
2.2	Discovery of Designated Resolvers	164
2.3	Authoritative DoE	164
2.4	Related Work	164
3	Approach	165
3.1	DoE at Edge Vantage Points	165
3.2	Authoritative DoE	166
4	Results	167
4.1	DDR Adoption at the Edge	167
4.2	DDR Configuration at the Edge	168
4.3	Do53 Forwarding at the Edge	169
4.4	Authoritative DoE Adoption	170
4.4.1	Public DoE Resolvers	170
4.4.2	Authoritative Name Servers	171
5	Discussion	172
5.1	The Sorry State of DoE	173
5.2	Limitations	173
6	Conclusion	173

Introductory Summary



1 Introduction

Domain names such as `www.kau.se` are so deeply embedded in our everyday use of the Internet that it is easy to overlook the complex, large-scale distributed system responsible for translating them into Internet Protocol (IP) addresses: the Domain Name System (DNS). On the web, end users type a name into their browser’s address bar, which directly and transparently leads them to their desired website. In reality these domain names must be resolved to addresses by the DNS, and this resolution is highly revealing and often comparable to a browsing history. If the answer is not already cached, the resolution process exposes the client’s destination (e.g., website) to multiple servers across the Internet. Between the client stub and the name servers authoritative for the name, storing the DNS records, lies a crucial intermediary: the *resolver*.

The resolvers perform iterative queries to the name servers on behalf of clients and temporarily cache responses to reduce load on the name servers. Due to their central role and visibility into client activity, resolvers have become a focal point for efforts to improve privacy and security. This has led to new standards aimed at ensuring confidentiality [1, 2, 3], integrity [4], and availability [5] (the CIA triad), as well as authentication [4], data minimization [6], and client anonymity [7].

Several recent proposals aim to mitigate the privacy risks inherent in DNS resolution by reducing information exposure, protecting transport confidentiality, or obscuring the relationship between clients and their queries. One such mechanism is Query Name Minimization (QMIN), which limits information leakage toward authoritative name servers by making the resolver query only the minimally necessary suffix of a domain name during resolution [6]. By design, QMIN prevents name servers from learning the full queried domain, thereby reducing passive observation of user activity in the top of the DNS hierarchy.

A complementary line of work focuses on hiding the client’s identity from the resolver itself. This can be achieved by routing DNS queries through one or more intermediaries. Such intermediaries may be purpose-built DNS proxies [8, 9, 10], or general-purpose anonymity systems such as the Tor network [11], which routes traffic through multiple relays. While these approaches offer strong anonymity guarantees, they also introduce additional latency. In parallel, encrypted DNS transports protect queries against on-path observers, but raise the practical question of how clients learn which endpoints to use for encrypted transport. The recently standardized RFC called Discovery of Designated Resolvers (DDR) [12] addresses this problem by allowing resolvers to advertise supported encrypted endpoints via plaintext DNS records, enabling upgrades to encrypted resolution. Together, these mechanisms illustrate the broad design space of DNS privacy enhancements, each addressing different adversaries and leakage vectors, and each presenting distinct trade-offs in deployability, performance, and ecosystem impact.

The *objective* of this thesis is to investigate privacy challenges involving DNS resolvers, and to analyze how proposed solutions are being adopted and implemented in the wild. This includes examining deployment trends, identifying adoption obstacles, and assessing the impact of solutions on clients, operators, other features, and the broader DNS ecosystem. To pursue this objective, the thesis relies primarily on Internet measurements, an approach that offers valuable empirical insights but also introduces methodological constraints that must be carefully considered.

We measure the adoption of emerging features using longitudinal, one-off, and regional measurements. Our results reveal a strong overall growth in QMIN adoption, substantial variation across regions and between resolver types, and similar patterns for other data minimization features. We also find low adoption of DDR among non-public resolvers. We identify significant deployability challenges to some of the proposed features. QMIN in some configurations induces resolution failures, higher latency, and packet overhead. DNS resolution over the Tor anonymity network is commonly associated with considerable performance overhead. Although circuit configuration can be optimized to reduce latency, these gains are limited by the uneven geographic distribution of Tor relays. Our analysis further uncovers systematic deviations between resolver implementations and recommendations in standards, widespread failures caused by authoritative servers mishandling minimized queries, and misconfigured records in DDR that incorrectly outsource encrypted DNS transport. Finally, we demonstrate that deployments by popular public resolvers largely explain observed adoption gains, where public resolvers in general are quick to implement privacy and security features. Misused DDR deployments risk exacerbating centralization, but Tor can mitigate privacy risks of DNS centralization by distributing trust across a volunteer network. We also show that design and configuration choices can substantially reduce performance penalties.

The remainder of this thesis is organized as follows: Section 2 provides the background necessary to understand the context of the appended papers. The research objective and research questions are outlined in Section 3 followed by an overview of the research methods in Section 4. Section 5 lists our contributions, and the appended papers are summarized in Section 6. Section 7 positions these contributions within the context of related work in the area. Section 8 concludes this thesis, followed by the appended papers.

2 Background

This section introduces the fundamentals of DNS (Section 2.1), outlines the capabilities and objectives of adversaries with respect to DNS privacy (Section 2.2), and presents the privacy-enhancing solutions in the scope of this thesis, grouped by their respective goals (Section 2.3-2.7).

2.1 Domain Name System

For devices, such as laptops, smartphones, PCs, and servers, to communicate on the Internet, they use IP addresses [13] to identify each other. These addresses are easy for computers to process but difficult for humans to remember. Domain names are therefore mapped to the IP addresses, allowing users to associate a service with a name such as `www.kau.se`. In the early days of the Internet, this mapping was maintained manually in a single file on each host. However, as the number of devices connected to the network rapidly increased, a scalable naming system became necessary, leading to the creation of the Domain Name System (DNS) [14, 15].

Each client in the DNS ecosystem has a local program called a *stub resolver* (hereinafter referred to as “stub”) that assists applications on the operating system (OS) by resolving domain names to resources, such as IP addresses, against the name servers in DNS. A DNS query contains a Query Name (QNAME) that consists of a sequence of labels; for example, `www.kau.se` contains the labels `www`, `kau`, and `se` (and technically an empty label at the end). Along with the QNAME, a DNS query also specifies a Query Type (QTYPE), which indicates what kind of information is being requested. Common QTYPES include A (IPv4 address records), AAAA (IPv6 address records), NS (authoritative nameserver records), MX (mail exchange records), and TXT (arbitrary text records). Different record types allows DNS to support not only basic hostname-to-address lookups but also email routing, delegation, and various security and configuration mechanisms.

A zone is an administrative portion of DNS for which a specific operator is authoritative. A zone contains the resource records needed to resolve names within a defined part of the namespace, such as a domain and its subdomains. While a domain refers to a subtree of the DNS namespace, a zone refers to how that subtree is operationally managed: a single domain may be split across multiple zones if authority is delegated to other operators. Zones are connected through delegation, which allows a parent zone to hand off responsibility for a subdomain to a child zone. This is done using NS records in the parent zone that point to the child’s authoritative servers, often accompanied by glue records (IPv4/IPv6 address records) to bootstrap resolution. Within a zone, records are stored in a zone file or managed dynamically, and authoritative servers return final responses to queries for names in that zone. This structure enables the DNS to scale globally, distributing administrative control while allowing resolvers to follow a chain of authority.

The stub is preconfigured to use one or more *recursive name servers* (hereinafter referred to as “resolvers”) that iteratively query name servers in the DNS hierarchy until an answer is returned to the requesting client (see Figure 1). Resolvers typically serve many clients in a network and cache responses for a period of time to speed up subsequent queries. When resolving the A record for a domain name e.g., `www.kau.se`, the resolver first checks its cache. If the answer is not cached, it sends a query

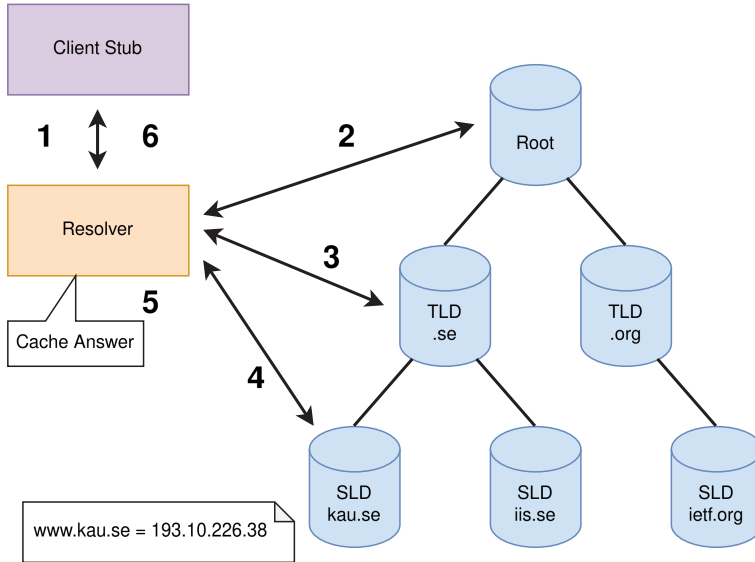


Figure 1: A client stub queries a resolver, which iteratively follows DNS referrals (Root \rightarrow TLD \rightarrow SLD) until the authoritative server returns the requested resource, which is then cached and sent to the client.

to the closest cached parent of the QNAME, and if nothing is cached it will send a query to the top of the DNS hierarchy: the *root* name servers. The root servers do not have direct mappings for `www.kau.se`, but they do store references to the Top-Level Domain (TLD) name servers responsible for domains ending in `.se`. Instead of returning an IP address, the root servers therefore return a referral (NS) to the `.se` TLD servers. The resolver then queries one of the TLD servers, which in turn responds with a referral to the Second-Level Domain (SLD) name servers authoritative for `kau.se`. Finally, the resolver queries the SLD server, which returns the resource record associated with the requested domain. In the case of a query for `www.kau.se. A`, the authoritative name server responds with the IPv4 address assigned to that QNAME. The resolver stores any intermediate records it encounters, along with the final resolved answer, in its cache before sending the final response back to the client. In DNS, Empty Non-Terminals (ENTs) are domain names that exist implicitly in the DNS tree structure but do not have any resource records of their own. They appear when a delegated or existing domain has child labels, yet the parent label itself contains no data. For example, an A record for `git.cs.kau.se` exists but `cs.kau.se` has no such records, so the latter is an ENT.

Resolvers can be offered by different providers and categorized by accessibility and network proximity to the client. Some resolvers are intentionally open to the public and are known as *public* resolvers. Popular

examples include Google [16], Cloudflare [17], Cisco [18], and Quad9 [19]. If a resolver is reachable due to misconfiguration, it is referred to as an unintentionally *open* resolver. A resolver accessible on a public IP address but only to a specific group of clients, such as those within an Internet Service Provider (ISP), company, organization, or school, is known as a *non-public* or *shared* resolver and is typically operated within the same Autonomous System (AS) as the client. Finally, resolvers operating on private IP addresses and serving only a small number of clients within the same household are described as *private* or *local* resolvers.

DNS, as a distributed hierarchical key-value store, was originally designed to be highly decentralized, with authority delegated across many independently operated zones. In practice, however, the modern DNS ecosystem has trended toward increasing centralization. Large hosting providers now co-locate vast numbers of authoritative zones on shared infrastructure [20], while a small set of popular public resolvers handle an ever-growing share of global query traffic [21]. This consolidation is reinforced by default resolver configurations embedded in web browsers, mobile OSes, and Internet-of-Things devices, which routinely direct users to the same handful of operators [22]. As a result, the operational reality of DNS is drifting away from its ideal decentralized architecture, raising concerns about resilience, diversity of authority, and the concentration of trust in a limited number of infrastructure providers.

2.2 DNS Threat Model

Before introducing resolver privacy mechanisms, we define a threat model that captures, for this thesis, the main adversaries and their capabilities in the DNS resolution process. We focus on adversarial capabilities relevant to the design and evaluation of privacy-enhancing DNS resolver solutions, namely those affecting (i) the confidentiality of DNS messages, (ii) the integrity and authenticity of DNS data, (iii) the availability and operational robustness of resolution, (iv) the unlinkability between a client identity and its queried domains, and (v) the minimization of exposed data across the DNS hierarchy. While confidentiality, unlinkability, and data minimization constitute the primary privacy objectives, integrity, authenticity, and availability are considered insofar as they influence the correctness, assumptions, or practical deployment of privacy mechanisms.

We classify adversaries by where they observe or influence the DNS protocol: (i) on-path between the client stub and the resolver, (ii) on-path between the resolver and authoritative name servers, and (iii) at the resolver itself. This reflects the distinction between threats mitigated by encrypted transport on the wire and threats arising from the resolver’s privileged position.

A *passive* on-path adversary can monitor and analyze the traffic, e.g., observing plaintext DNS queries and responses, or observe the timing, and message sizes of encrypted traffic. Such an adversary can, among

other things, infer browsing behavior from observed QNAMEs, correlate queries, and profile users over time [23, 24]. This adversary motivates *Confidentiality* mechanisms (Section 2.3) that encrypt DNS traffic between the stub and resolver and, where applicable, between resolvers and authoritative name servers. However, encryption does not necessarily prevent traffic analysis [25, 26, 27, 28], and the query by necessity is decrypted once it reaches the resolver. An *active* on-path adversary can modify, inject, replay, or drop messages. In the context of DNS, this includes response tampering, spoofing, downgrade attempts, and malicious redirection. These threats motivate *Integrity and Authentication* mechanisms (Section 2.4), to enable resolvers and clients to detect tampering and authenticate DNS data via a chain of trust.

The resolver is typically treated as trusted. However, from a privacy perspective, the resolver is a natural choke point: it can observe both client identifiers and query contents and thus profile users over time. Importantly, encrypting traffic between the stub and resolver does not protect against this adversary. This motivates *Anonymization* mechanisms (Section 2.6) beyond transport encryption. Some anonymization designs introduce intermediaries (e.g., proxies/relays) to separate client identity from query content. In these settings, an important threat is *collusion* among infrastructure entities, such as proxy-resolver collusion. A resolver adversary may also actively manipulate responses before relaying them to clients [29].

Despite the use of encryption and anonymization in DNS transport, query data may be unnecessarily revealed between DNS zones. We therefore consider threats related to excessive data propagation: sending full QNAMEs to higher levels of the hierarchy, disclosing client network prefixes, or exposing resolver behavior to clients via oversized or information-rich responses. These threats motivate *Data Minimization* mechanisms (Section 2.7), which reduce the amount that upstream servers learn about client queries and the amount that downstream clients learn about resolver state.

Availability is not a privacy goal per se, but DNS privacy mechanisms may interact with operational resilience. Encrypted transports can be blocked, downgraded, or misconfigured; minimization strategies can trigger resolution failures when authoritative servers behave incorrectly; and privacy-enhancing mechanisms can increase latency or load. We therefore consider availability threats and misconfiguration risks in Section 2.5, including operational mechanisms such as anycast and redundancy that affect the deployability and reliability of privacy-enhancing DNS solutions.

2.3 Confidentiality

While most traffic on the modern web with the Hypertext Transfer Protocol (HTTP) is now encrypted using Transport Layer Security (TLS) [30]

(commonly known as HTTPS), DNS traffic has, for most of its history, been transmitted in plaintext. This has enabled on-path adversaries to passively monitor queries both from stub to resolver and from resolver to authoritative name server. In recent years, several proposals have emerged to add TLS-based protection to DNS messages. Collectively, these protocols are referred to as DNS-over-Encryption (DoE), or *Encrypted DNS Transport*. On the path between stub and resolvers, the so-called “last mile”, deployment of DoE has accelerated in recent years, driven largely by modern web browsers that ship with support for encrypted DNS transport on the application layer.

DNS-over-TLS (DoT) [1], introduced in 2016, encapsulates DNS messages within TLS sessions carried over TCP. While effective in providing confidentiality and integrity, this approach introduces additional latency due to extra handshake overhead and relies on a dedicated port (TCP/853) rather than the traditional UDP/53 (often referred to as Do53). This dedicated port is making it easy for network operators to detect and block in order to downgrade connections back to plaintext.

DNS-over-HTTPS (DoH) [2], standardized in 2018, solves this problem by tunneling DNS messages over HTTPS. This allows encrypted DNS traffic to blend in with ordinary web traffic, significantly reducing the risk of blocking or interference. It also has a negative impact on tools and services using this traffic as a source for threat intelligence [25, 26, 27, 28]. However, like DoT, DoH inherits the performance costs of using TCP and TLS.

DNS-over-QUIC (DoQ) [3], standardized in 2022, introduces a more efficient approach by using the QUIC [31] transport protocol, which operates over UDP and integrates TLS. The combination of UDP, stream-based multiplexing and elimination of head-of-line blocking in QUIC, as well as 1-RTT or 0-RTT handshakes in TLS 1.3, offers substantial performance improvements compared to DoT and DoH [32, 33]. However, the use of 0-RTT also introduces exposure to replay attacks, as early data can be captured and retransmitted by an adversary before the handshake completes [34].

Despite rapid adoption, most encrypted DNS traffic still flows to a small number of large public resolvers because web browsers default to trusted resolvers [35]. To address this centralization, the Discovery of Designated Resolvers (DDR) standard [12] enables any resolver operator to advertise which DoE protocols they support, allowing clients to automatically discover encrypted endpoints offered by preconfigured Do53 resolvers.

Resolver-to-authoritative communication is operationally more complex because resolvers query many different name servers across the DNS hierarchy (see Figure 1). Two primary approaches for Authoritative DoE (ADoE) are under consideration. The first, *Unilateral Probing* [36], allows resolvers to probe common encrypted endpoints to determine whether an authoritative server supports encrypted transport. The second, *Trans-*

port Signaling [37], resembles DDR by allowing authoritative servers to explicitly signal their encrypted transport capabilities.

The most common deployment of DoH is typically using HTTP/2. DNS-over-HTTP/3 (DoH/3) preserves the DoH message format and HTTP exchange model but runs over HTTP/3 [38], i.e., over QUIC rather than TCP. Bielefeld *et al.* [39] report promising performance for DoH/3, but it has not yet seen widespread adoption. Since HTTP/3 transport aspects are outside the scope of Paper V and Paper VI, we use HTTP/2 when referring to DoH throughout this thesis. In Paper VI we investigate adoption and configuration of DDR on non-public resolvers. We also probe authoritative name servers for the adoption of DoT, DoH, and DoQ.

2.4 Integrity and Authentication

DNS Security Extensions (DNSSEC) [4] adds integrity and origin authentication to DNS through the use of public-key cryptography. Resource records are signed, and validating resolvers can verify these signatures against a chain of trust rooted, very fittingly, in the DNS root zone. This ensures that it is possible to verify if the received records have been manipulated in transit (integrity) and if they were published by the legitimate authoritative name server (authentication). The primary security objective of DNSSEC is therefore to protect against cache-poisoning attacks [40] and impersonation of authoritative name servers.

It is important to emphasize that DNSSEC does *not* provide confidentiality. DNSSEC and encrypted DNS transports such as DoT, DoH, and DoQ solve complementary, rather than overlapping, problems. While not a privacy mechanism, DNSSEC plays an essential role in enabling *authenticated* encrypted DNS transport. Mechanisms such as DDR and emerging proposals for discovering ADoE rely on DNSSEC-signed records to ensure that clients and resolvers can trust the advertised endpoints. Without DNSSEC, an adversary could spoof discovery records and mislead clients into using rogue encrypted DNS services, thereby undermining the privacy and security guarantees of DoE entirely.

In Paper III, we measure the adoption of DNSSEC validation of resolvers used in the Nordics and Baltics, in the context of other features.

2.5 Availability

Operational resilience mechanisms do not directly enhance privacy, but they play an important role in maintaining resolver performance, service continuity, and deployment flexibility. These factors can indirectly influence the adoption, reliability, and observable behavior of privacy-enhancing DNS mechanisms. A privacy feature that reduces availability or introduces instability is unlikely to be deployed at scale, and misconfigurations of security or privacy standards could cause

availability issues themselves. A key technique for improving availability in both recursive and authoritative infrastructures is *anycast* [5]. Anycast allows multiple geographically distributed servers to share the same IP address, enabling traffic to be routed to the topologically closest functioning instance. This approach increases resilience, reduces latency, and provides natural load-balancing and failover. For example, although the DNS root zone consists of only 13 named root servers, they are deployed across more than 2,000 global instances [41, 42], ensuring high availability even during network disruptions or traffic spikes. Redundancy is also a fundamental principle of authoritative name server operation: zones typically designate both primary and secondary servers to provide fault tolerance and distribute load. On the client side, stubs often include multiple preconfigured resolvers, ensuring queries succeed even if one resolver is unreachable. Similarly, broader adoption of IPv6 alongside IPv4 improves overall connectivity and enables clients to reach resolvers across multiple transport families, further enhancing robustness. These availability mechanisms interact with privacy and security standards in several ways. Highly distributed infrastructures, such as anycast-based resolvers, may accelerate the deployment of encrypted transports or automated discovery mechanisms by reducing latency penalties and increasing the geographic proximity of privacy-supporting endpoints. Conversely, misconfigurations in privacy features can negatively affect availability, leading to resolution failures.

If a client stub is configured with multiple resolvers and even one of them does not implement QMIN, the privacy benefits of QMIN are effectively negated, regardless of the behavior of the others. This mismatch leads to the *conflicting resolver* behavior identified in Paper I. In our regional study of resolvers (Paper III), we use IPv6 capability as an indicator of resolver availability in our feature adoption measurements. Furthermore, in Paper IV, we quantify the operational impact of QMIN and show that the strict-mode configuration results in increased resolution failures, primarily due to misbehaving authoritative name servers, thereby reducing availability for affected domains.

2.6 Anonymization

In the context of DoE, anonymity refers to a resolver which is oblivious to the identity of the client stub. Even if DNS traffic is encrypted between the stub and the resolver, thereby preventing on-path adversaries from monitoring queries, the resolver operator still observes both the source IP address of the client and the DNS query contents, such as the QNAME. This privileged position can be abused for monetary, personal, or political purposes, and also represents a high-value target for attackers seeking to exfiltrate sensitive data. Separating the client identity (source IP address) from the queried domain can help mitigate these privacy and security risks. Several proposals introduce proxies between the stub and resolvers

to ensure that no single entity learns both pieces of information [8, 9, 43]. These designs require that the query is encrypted end-to-end between the stub and the resolving entity; otherwise, the proxy simply becomes a new point of exposure.

Oblivious DNS (ODNS) [8] proposes using a dedicated TLD server for the domain `.odns`, which acts as an oblivious resolver. Clients encrypt the QNAME and a session key using the public key of the ODNS server before appending the `.odns` suffix. The encrypted query is sent to any resolver, which forwards the modified QNAME to the ODNS server. The ODNS resolver decrypts the inner query, performs resolution on the client’s behalf, encrypts the response with the session key, and returns it to the stub via the intermediate resolver. As a result, the ODNS resolver never learns the identity of the client, and the intermediate resolver never learns the true query contents.

Oblivious DoH (ODOH) [7] builds on similar principles by using HTTPS forwarders to relay encrypted DoH queries. In this approach, the DoH resolver only sees the source IP address of the forwarder, while the encrypted query is protected end-to-end between the client and the resolver. This prevents either party from correlating client identity with the QNAME unless they collude.

μ ODNS [10] proposes a multi-relay DNS anonymization framework that extends single-relay systems like ODOH by introducing a trusted next-hop relay, mutual sharing of dedicated relays, and randomized multi-hop paths to obscure a client’s identity even when some relays collude with the resolver. It preserves strong anonymity with minimal performance overhead by mixing traffic across users and preventing any single colluding node from reliably linking queries to their origin.

DNS-over-HTTPS-over-Tor (DoHoT) [43] tackles the collusion problem by tunneling DoH queries over the Tor anonymity network [11], a general-purpose anonymity system that routes traffic through three proxies, sampled from thousands of servers operated by volunteers. While this approach provides strong anonymity guarantees, it incurs significantly higher latency compared to other approaches due to Tor’s multi-hop design [9, 10]. As with ODNS and ODOH, the resolver is exposed only to the proxy’s source IP. Here that proxy is the third hop in the Tor circuit, the exit relay, rather than the client’s own IP address.

In the context of anonymous DNS resolution, we evaluate DoHoT under a realistic anonymization threat model and show that Tor circuit tuning can mitigate performance overheads, while still offering stronger resistance to collusion compared to single-proxy approaches such as ODOH (Paper V).

2.7 Data Minimization

The aim of data minimization is to reduce unnecessary exposure of information without impairing the functionality of a system. Each component

in the DNS hierarchy requires only a limited subset of information to perform its role: the root name servers do not *need* the full Fully Qualified Domain Name (FQDN) `www.kau.se` to provide a referral to the appropriate TLD servers; the authoritative server for `kau.se` does not *need* parts of the client’s source IP address to serve resources; and the client stub does not *need* visibility into which authoritative name servers were queried by the resolver. Minimizing data exposure, therefore, improves privacy while also supporting efficiency and simplicity.

Instead of sending the full QNAME to each server in the hierarchy, a resolver may send only the minimum number of labels required at each step, known as QNAME minimization (QMIN) [6]. This approach limits the leakage of subdomain information while still allowing successful resolution when authoritative servers behave correctly. In practice, QMIN is typically implemented with configurable fallback behaviors, referred to as *relaxed* and *strict* mode [44, 45]. If the resolver is configured to minimize queries in relaxed mode, it will send the full QNAME to the authoritative name server if problems occur. In strict mode, the resolver will terminate the resolution of the domain, prioritizing minimization over success.

Another mechanism relevant to data exposure is EDNS Client Subnet (ECS), which is defined as an EDNS(0) option (Extended Mechanisms for DNS [46]). With ECS, resolvers may include a client network prefix in queries to authoritative servers, enabling them to return geographically optimized responses, an approach commonly used by Content Delivery Networks (CDNs). Originally introduced to improve latency and service quality, subsequent studies have shown that ECS reduces user privacy by revealing coarse client location and increasing linkability across queries [47]. Consequently, disabling ECS is recommended to improve privacy.

A final aspect of data minimization concerns the responses the resolver returns to the client. When a client requests a specific resource record, any additional information, such as authority data or details about servers contacted during resolution, is unnecessary and provides no benefit to the stub. Returning only the requested resource record, commonly referred to as a *Minimal Response*, reduces the amount of data exposed to local clients and limits opportunities for information leakage. In addition, smaller responses reduce the potential for DNS amplification and reflection attacks and narrow the attack surface by decreasing the volume of data the stub must parse. Finally, avoiding the inclusion of unsolicited records prevents clients from inferring aspects of the resolver’s cache state, which could otherwise be exploited to learn about previous queries or the behavior of other clients sharing the resolver.

QMIN as a tool to limit information leakage is directly related to our studies tracking its increase in adoption (Paper I and III) and resolver fingerprintability due to deviations in its implementation (Paper II). Furthermore, the distinction between relaxed and strict modes in QMIN is critical to the thesis’s investigation of operational impact, where mea-

surements quantify how the strict mode can cause resolution failures and increase traffic overhead (Paper IV). ECS and Minimal Responses are explored through regional measurements in the Nordics and Baltics, revealing that while most ISPs minimize data, centralization through providers like Google often reintroduces privacy risks, such as ECS (Paper III).

3 Research Questions

In order to investigate the privacy challenges involving DNS resolvers, and analyze how proposed solutions are being implemented, we structure the scope of this thesis around three research questions. Rather than addressing privacy in isolation, the questions emphasize how feature designs, operational practices, and deployment realities intersect to shape user privacy, resolver behavior, and ecosystem-wide trends. These questions highlight where current approaches succeed, where they introduce new challenges, and how measurable Internet-scale observations can reveal gaps between intent and real-world outcomes.

RQ1: *How have the adoption rates of DNS privacy features in resolvers evolved over time globally and regionally, and what infrastructural or technical barriers limit their deployability?*

This question examines how privacy-enhancing DNS mechanisms are being adopted across diverse regions and operator types. Adoption trends often depend on factors such as resolver software defaults, operator resources, regulatory pressures, and the maturity of local infrastructure. By studying deployment at both global and regional scales, this thesis seeks to identify where adoption accelerates, where it stagnates, and which technical, operational, or resource constraints create persistent gaps. Understanding these barriers is essential for evaluating whether current privacy features are realistically deployable at Internet scale.

RQ2: *To what degree do real-world DNS resolver configurations adhere to privacy feature standards, and how do deviations in implementation affect the observable behavior of the DNS ecosystem?*

Even when privacy features are standardized, real-world implementations often diverge due to software variation, configuration choices, or operator-specific optimizations. This question focuses on how resolvers actually behave in the wild, and how frequently they comply with intended privacy protections. Examining deviations, both accidental and intentional, reveals how feature design translates into operational practice and what side effects or measurement limitations arise when implementations only partially or inconsistently follow standards.

RQ3: *What are the quantifiable performance penalties and levels of centralization associated with varying DNS resolver privacy features,*

and can these mechanisms be configured to balance the inherent trade-offs in privacy, performance, and centralization?

Privacy features often introduce additional cryptographic steps, more complex resolution logic, or increased query indirection, which may impose overhead on resolvers and authoritative name servers. Centralization plays a dual role in DNS privacy: large providers can drive rapid deployment of new standards, but concentrated control also amplifies risks to privacy, resilience, and trust. This question analyzes how different configurations affect latency, bandwidth, and resolver load. It also explores how privacy mechanisms influence traffic flows and whether they unintentionally reinforce reliance on a small number of public providers. Mechanisms that may counterbalance this trend are also included in this question. This includes mechanisms for decentralization and anonymization. By characterizing trade-offs, the goal is to determine whether practical configurations can deliver privacy guarantees without degrading user experience, operational efficiency, or risks associated with centralization.

4 Methodology

Several alternative research methods were considered but found unsuitable for the objective of this thesis. Evaluating feature behavior solely through reference or experimental implementations allows for in-depth analysis of software design choices, but it fails to capture how features behave once deployed across diverse operational environments. In practice, feature behavior is shaped by heterogeneous resolver implementations, authoritative server configurations, and operational decisions made by independent operators [48]. To capture these effects, this thesis prioritizes empirical observations of real-world deployments and their interactions. Similarly, formal methods such as symbolic verification or model checking are well-suited for proving correctness and security properties of mechanism designs under clearly defined assumptions. However, these approaches intentionally abstract away operational realities such as partial deployment, legacy systems, misconfigurations, and implementation deviations. This thesis focuses on the operational deployment, performance, and adoption of DNS privacy mechanisms to expose the practical challenges and failure modes observed in the wild.

Internet measurement studies face inherent methodological challenges. The Internet is a dynamic, distributed, and continuously evolving system, where feature behavior, infrastructure deployment, and operator practices can change over time. These characteristics directly affect the *repeatability* (Section 4.1), *replicability* (Section 4.2), and *reproducibility* (4.3) of results [49]. Repeatability is ensuring that the same researchers can run their own experiments and measurements again and yield consis-

tent results. Replicability enables other researchers to re-run experiments and measurements using available tools and artifacts. Reproducibility allows other researchers to achieve the same conclusions using different methods and setup. The research methods used throughout this thesis are designed with these principles in mind.

4.1 Repeatability

Testbeds provide a controlled environment for emulating DNS operations and observing resolver behavior consistently in isolation. They allow researchers to deploy and evaluate multiple software implementations from different vendors and versions while controlling all components of the resolution path: clients, resolvers, and authoritative name servers. Testbeds may be fully air-gapped or connected to the Internet, depending on the measurement goals. Running experiments locally offers ease of access, reduces operational overhead, and, when air-gapped, ensures that no additional load is placed on the live Internet infrastructure. Shared or standardized testbeds can also improve replicability. However, a key limitation is that testbeds may not accurately reflect real-world deployments [50]. Their representativeness depends on the experimental design, software configuration, and degree of similarity to production environments, which may limit the generalizability of the results. We employ testbeds to a limited extent to establish fingerprints of popular open source resolvers in Paper II.

Automation and scripts enable systematic and consistent data collection and analysis. Assuming that the scripts function as intended, this minimizes human errors. It also ensures that experiments can be repeated under the same configurations, even if the DNS ecosystem has changed over time. Such automation also complements written documentation, which should begin early in the research process and clearly describe the measurement setup, assumptions, methodology, and limitations. We use notated scripts throughout all studies for both measurement orchestration and data analysis.

4.2 Replicability

Ensuring replicability requires that datasets, analysis tools, and measurement scripts should be made publicly available whenever possible. This also supports peer verification and collaboration [51, 52]. Published documentation should be transparent, use commonly understood terminology within the community, and be accessible to other researchers to facilitate replicability. Open-source availability not only strengthens the study’s credibility but also enables the broader research community to build on, extend, or independently validate the findings. We make our tools publicly available where possible to enhance transparency, replicability, and

collaboration¹. Using common Internet measurement platforms also supports replicability, and we leverage RIPE Atlas [53], CAIDA Archipelago (Ark) [54], and OpenINTEL [55] in the thesis.

RIPE Atlas is a global, community-driven measurement platform operated by the RIPE NCC [56]. It consists of tens of thousands of lightweight probes hosted by volunteers across residential networks, data centers, universities, and ISPs. These probes can perform measurements such as DNS queries, traceroutes, and pings. We use RIPE Atlas in Paper I, II, III, and VI.

CAIDA Ark is a measurement infrastructure composed of strategically located “Ark monitors” in research networks, Internet Exchange Points (IXPs), and data-center environments. Ark is best known for large-scale topology and connectivity measurements, but it is also commonly used for DNS studies through targeted probing of resolvers and authoritative servers. CAIDA Ark has fewer probes and less geographic diversity compared to RIPE Atlas and limited availability for customized measurements, unless part of an approved research collaboration. We use CAIDA Ark in Paper VI to complement our measurements with RIPE Atlas.

OpenINTEL is a large-scale, longitudinal DNS measurement platform designed to systematically monitor the configuration and evolution of the global DNS ecosystem. Instead of relying on distributed vantage points, OpenINTEL performs highly structured daily measurements against tens of millions of domain names across major TLDs from a single vantage point. Each day, OpenINTEL queries a stable, predefined set of zones, collecting detailed information on authoritative name servers, DNSSEC deployment, mail security records, and other DNS configuration properties. This approach produces datasets that allow researchers to track operational trends, detect configuration changes, study deployment life-cycles of new DNS standards, and identify systemic issues on a global scale. In Paper IV we collaborate with people from OpenINTEL and use their system for our active measurements.

4.3 Reproducibility

Reproducibility requires that other researchers can confirm a study’s conclusions using their own infrastructure and measurement designs. Because Internet conditions evolve continuously, reproducible measurement studies must rely on methods that remain valid even when performed at different times, from different locations, and using different tools. In practice, this involves employing both *active* and *passive* measurements from a variety of *vantage points* to ensure that findings hold under real-world conditions.

Active measurements involve generating traffic from one or more vantage points. In this thesis, these vantage points are either client stubs

¹<https://github.com/Arcnilya>

or resolvers. Depending on the measurement design, the results may be observed at the client (as responses from the resolver), at the resolver (as queries from the client or responses from authoritative servers), or at authoritative name servers (as queries arriving from resolvers). The latter requires that the name server is authoritative for a name under the researchers’ control. Active measurements may offer strong advantages, including fine-grained control over the QNAME, the ability to design experiments that isolate specific behaviors, and clean separation of measurement traffic from unrelated queries. However, they must be conducted responsibly: queries should be rate-limited to avoid unnecessary load, and researchers should follow ethical guidelines such as announcing measurement goals and providing contact information to enable transparency and opt-out opportunities [57]. We use active measurements and collect data from the clients in all of our studies. In Paper I, II, III, and VI we also collect data from the authoritative side using a domain under our control. We also collect data at the resolver in Paper IV.

Passive measurements, in contrast, observe traffic that is not generated for the purpose of the study. This allows researchers to examine how clients and resolvers behave under real-world operational conditions. Passive measurements, however, comes with important limitations. Researchers lack control over the QNAME, the origin, and the timing of the observed queries. Moreover, passive measurement raises privacy concerns due to the potentially sensitive nature of DNS traffic, requiring careful handling, minimization, and, where appropriate, anonymization. We only perform passive measurements in Paper I, and the data collected by the operators was minimized and anonymized before being sent for analysis.

Vantage points are often biased toward specific geographic regions, networks, or resolver infrastructures. Increasing the diversity of vantage points can improve generalizability, but resource constraints may limit the number available. When relying on a single vantage point, it is essential to state this clearly, explain the rationale, and discuss how this may influence the interpretation of results. Measurements using RIPE Atlas and CAIDA Ark draw on multiple vantage points by design. In contrast, the Internet-wide measurements of open resolvers in Paper I and II relied on a single university vantage point. Paper IV issued queries via OpenINTel, which also constitutes a single vantage point. Finally, the latency measurements in Paper V relied on a single client located in Sweden.

5 Contributions

Figure 2 provides an overview of how the individual papers address the research questions, showing the relationship between each research question and the corresponding contributions summarized below.

C1: **Adoption.** *We look at the adoption of DNS privacy features using longitudinal, one-off, and regional measurements, showing that: (i)*

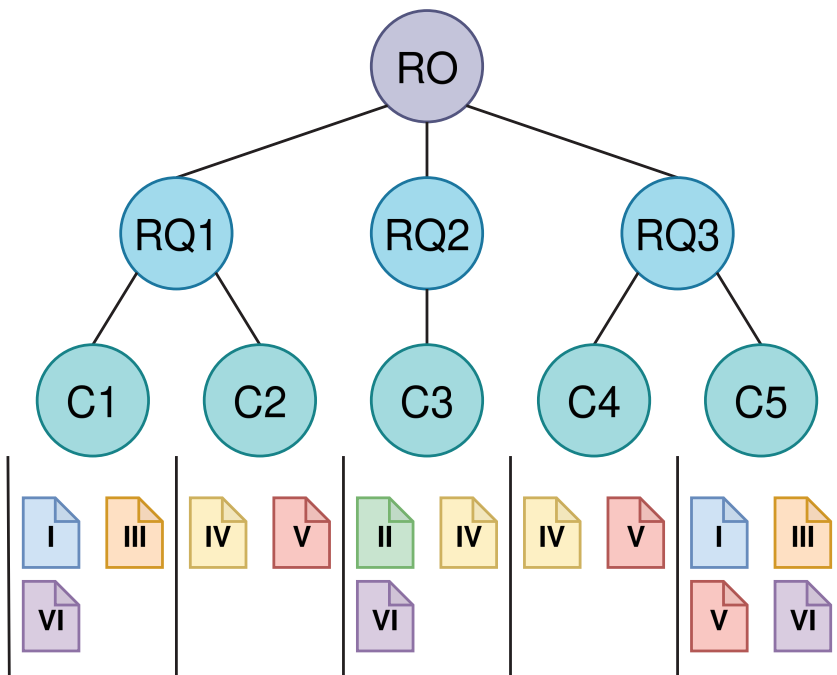


Figure 2: Overview of the Research Objective (RO), Research Questions (RQ), Contributions (C), and Appended Papers (I-VI). The papers contribute to answering the research questions, which structures the scope of the research objective.

the use of QMIN has grown significantly over time, largely due to public resolvers; (ii) feature uptake varies widely across regions and resolver types; and (iii) DDR adoption among non-public resolvers remains low, indicating a deployment gap despite standardization.

This contribution directly addresses RQ1 by empirically examining how DNS privacy features are adopted over time and across heterogeneous operational environments. In Paper **I**, we conduct a longitudinal analysis of QMIN adoption and observe a sharp increase from 10% in 2018 to 64% in 2022, demonstrating that deployment is not uniform but accelerates once widely used resolvers enable the feature by default. Complementing this global perspective, Paper **III** provides a regional view by studying the Nordic and Baltic countries, where the average adoption of QMIN, DNSSEC, ECS, Minimal Responses, and IPv6 exceeds 70% overall, yet differs markedly between countries and resolver types. Finally, Paper **VI** highlights persistent deployment barriers by focusing on DDR adoption among non-public resolvers, revealing that only 12%-23% are DDR-enabled when measured from the edge. Together, these stud-

ies show that while some DNS privacy mechanisms achieve broad adoption under favorable conditions, others remain constrained, underscoring the uneven and incomplete deployability of DNS privacy features at Internet scale.

- C2: **Barriers.** *We quantify deployability issues with (i) QMIN in strict-mode, causing resolution failures, latency increase, and packet overhead, and (ii) DoHoT suffering from latency increase due to the geographic concentration of Tor relays in the world.*

This contribution addresses RQ1 by identifying concrete technical and infrastructural barriers that hinder the deployability of DNS privacy mechanisms in real-world settings. In Paper IV, we evaluate the operational impact of QMIN in strict mode and find that, contrary to its intended privacy benefits, it introduces severe availability and performance issues. Resolutions of domains with long labels experience failure rates of up to 65%, average latency increases of up to 500 ms, and packet overhead exceeding 100%. While relaxed QMIN configurations substantially reduce these costs, they introduce fallback behavior that weakens privacy guarantees and leaves deployments vulnerable to downgrade attacks. These results demonstrate that strict QMIN is not universally deployable under current authoritative server behavior. Complementing these findings, Paper V examines the deployability of anonymized DNS resolution over Tor under a resolver-side adversary threat model. While routing DNS queries over Tor provides strong unlinkability guarantees without requiring new infrastructure, the default Tor configuration incurs substantial latency overhead. By tuning the Tor circuit construction average DoHoT latency can be reduced, while maintaining resistance to collusion. Nevertheless, performance remains influenced by the uneven global distribution of Tor relays, which limits the effectiveness of such optimizations for clients in under-represented regions and thus constrains the practical deployability at Internet scale.

- C3: **Deviations.** *We compare standardizations with their real-world deployment and show that: (i) resolver software consistently deviates from the recommended QMIN configuration in the RFC, enabling fingerprinting; (ii) strict-mode QMIN fails largely due to authoritative servers mishandling ENTs; and (iii) DDR-related SVCB records of non-public resolvers incorrectly point to resolvers maintained by other operators.*

This contribution addresses RQ2 by examining how deployed name server implementations diverge from standardized DNS privacy specifications, and how these deviations manifest in measurable ecosystem behavior. In Paper II, we demonstrate that resolver software consistently departs from the QMIN configuration recommended in RFC 9156 [6], producing distinct query patterns

that enable reliable fingerprinting of resolver software and versions. These findings show that even subtle deviations from standards can introduce new side channels observable through measurement. Paper [IV](#) further reveals how strict-mode QMIN interacts with non-compliant authoritative servers, uncovering that widespread mishandling of ENTs, contrary to RFC 8020 [\[58\]](#), leads to incorrect response codes and resolution failures. These failures are exacerbated in nested namespace structures where implicit assumptions about the QNAME being the FQDN no longer hold. Together, these behaviors illustrate how deviations across independently operated components undermine the intended semantics of standardized privacy mechanisms. Finally, Paper [VI](#) highlights similar misalignment between designed standard and deployment practice in DDR. Although the specification (RFC 9462 [\[12\]](#)) requires SVCB records to reference encrypted endpoints operated by the same resolver entity, our analysis shows that most non-public resolvers instead direct clients to public resolvers. This widespread misconfiguration alters trust relationships and resolver usage in ways not anticipated by the mechanism design. Collectively, these results demonstrate that standardized DNS privacy mechanisms are frequently implemented in ways that deviate from their specifications, and that these deviations have observable, and sometimes unintended, consequences.

C4: *Privacy-Performance Trade-offs.* *We investigate configurations of privacy features and their impact on performance. We show that: (i) relaxed-mode QMIN has less performance penalties compared to strict-mode; and (ii) DoHoT latency can be cut in half by shortening circuits and favoring nearby relays.*

This contribution answers RQ3 by quantifying the performance costs introduced by DNS privacy mechanisms and by demonstrating how configuration choices affect the trade-off between privacy and efficiency. In Paper [IV](#), we evaluate the operational impact of QMIN under real-world conditions and show that strict-mode deployment imposes considerably higher performance overhead than relaxed mode. Specifically, resolving domains with many labels increases average latency by up to 500ms, while the additional packet exchanges introduced by strict minimization increase packet overhead and susceptibility to loss on congested paths. These results show that stronger privacy semantics come at a tangible cost to resolution performance and availability. Paper [V](#) examines a complementary performance trade-off in the context of DNS resolution over the Tor anonymity network. Tor’s default latency overhead stems from routing traffic through multiple relays selected for anonymity rather than proximity, which incurs cryptographic processing at each hop and often results in long geographic paths.

By shortening circuits and biasing relay selection toward geographically nearby relays, we show that this overhead can be substantially reduced, cutting median DoHoT latency from 107 ms (p5–p95: 63–289 ms) to 37 ms (p5–p95: 20–144 ms). This demonstrates that performance penalties are not inherent to the approach itself, but are strongly influenced by configuration choices.

- C5: **Centralization.** *We analyze how popular resolvers shape the effectiveness of privacy features and how the associated privacy risks of centralization are mitigated. We find that (i) the major increase in QMIN deployment was driven by public resolvers enabling it by default; (ii) public resolvers almost universally adopt privacy and security features; (iii) DoHoT reduces centralization risks by distributing trust across a large volunteer relay network; and (iv) DDR deployments worsen centralization due to feature misuse.*

This contribution addresses RQ3 by examining how different DNS privacy mechanisms influence, and are influenced by, centralization. In Paper I, we show that the rapid increase in QMIN deployment beginning in late 2020 was largely driven by Google Public DNS enabling the feature by default. This finding demonstrates how a small number of dominant public resolvers can simultaneously accelerate the adoption of privacy features while reinforcing their central role in the DNS ecosystem. A similar pattern emerges in Paper III, where public resolvers almost universally adopt a broad set of privacy and security mechanisms, in contrast to smaller and non-public operators. While this uniform adoption improves baseline privacy and security for many users, it also highlights a structural imbalance in which advanced features are disproportionately deployed by a few large providers. Methodological limitations further underscore this point: Google’s implementation of QMIN was not directly observable in our measurements, illustrating how dominant actors can shape the ecosystem in ways that are harder to observe and evaluate empirically. In contrast, Paper V explores how anonymization techniques can counteract centralization effects. By routing DNS queries through Tor’s globally distributed volunteer relay network, DoHoT distributes trust across multiple independent entities and reduces the impact of proxy-resolver collusion compared to single-proxy designs. Although this approach introduces performance trade-offs, it demonstrates that decentralization and strong privacy guarantees need not rely on a small set of dominant providers. Finally, Paper VI shows that feature misuse can unintentionally exacerbate centralization. Most DDR-enabled non-public resolvers direct clients to Google Public DNS via their SVCB records. So instead of decentralizing encrypted DNS transport by allowing any resolver to advertise their own endpoint, DDR became the very thing that it swore to destroy.

6 Summary of Appended Papers

In this section, we summarize the appended papers and map their key results to the corresponding thesis contributions (C1-C5).

Paper I – A Second Look at DNS QNAME Minimization

This paper examines the global adoption and operational impact of QMIN in the DNS ecosystem from 2018 to 2022. Building upon earlier measurement studies, it combines extended active measurements using RIPE Atlas probes and open resolvers with refined passive measurements at root and TLD servers. The active measurements reveal a substantial rise in QMIN deployment, with QMIN-enabled RIPE Atlas resolvers increasing from 2.5k in 2018 to 14k in 2022, and open resolvers from 18k to 80k. Passive measurements also show positive trends, with QMIN use rising from 0.6% to 2.5% at one root server, and from 35.5% to 57.3% at the .nl TLD. These findings demonstrate that QMIN adoption has grown significantly and that observable deployment accelerates once large, widely used resolvers enable the feature (C1 and C5). Improved filtering of non-valid queries and incorporation of additional datasets allow a clearer understanding of the real-world impact of QMIN. The study also identifies “conflicting resolvers” that inconsistently apply QMIN due to parallel forwarding configurations. Finally, the paper discusses privacy-performance trade-offs and proposes using the public suffix list to guide minimization depth, offering a practical balance between privacy and resolver efficiency.

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

This study evaluates a proposed method for fingerprinting the software and version of DNS resolvers based on query patterns induced by QMIN. While Paper I establishes that QMIN adoption has increased significantly over time, this work shifts the focus from whether resolvers deploy QMIN to how closely their implementations follow the standard (C3). By analyzing minimization behavior from the perspective of authoritative name servers, we find strong correlations between specific query patterns and particular open-source resolver software versions. Notably, none of the observed resolvers fully implement the algorithm recommended by RFC 9156, revealing a significant gap between specification and practice. The paper also documents substantial query amplification, likely due to interactions between QMIN and resolver forwarding chains. These results highlight both fingerprinting risks and opportunities: resolvers can be identified based on their QMIN behavior, but operators can also use this knowledge to detect misconfigurations or security-relevant patterns in the DNS ecosystem.

Paper III – Privacy and Security of DNS Resolvers used in the Nordics and Baltics

This paper investigates the adoption of several DNS features related to privacy and security (QMIN, DNSSEC, ECS, Minimal Responses, and IPv6 support) across preconfigured resolvers on RIPE Atlas probes in countries located in the Nordics and Baltics (C1). The analysis also considers resolver categories based on proximity (private, shared, public) and explores correlations between features. This paper complements Paper I by narrowing the scope from global, longitudinal adoption trends to a detailed regional analysis. While Paper I demonstrates that QMIN deployment has increased significantly over time at Internet scale, this work shows that even within relatively homogeneous geographic regions, QMIN adoption varies markedly depending on resolver type and local operational practices. Together, the two studies illustrate that global growth in QMIN deployment masks substantial regional and structural differences in how and where the protocol is adopted. Overall adoption in the Nordic and Baltics exceeds 70% across features, though the average adoption by public resolvers are negatively influenced by Google Public DNS for certain capabilities. Ignoring Google, public resolvers adopt all privacy and security features, whereas smaller and non-public resolvers show more uneven deployment (C5). Country-level comparisons show that Sweden and Denmark rely more heavily on private resolvers, while Estonia and Norway primarily use shared resolvers located within the probe's AS. Among resolvers within the same AS, arguably the category most representative of the country, Norway shows the highest adoption of Minimal Responses and IPv6, Denmark achieves full DNSSEC adoption, and Estonia leads in QMIN deployment. Notably, no resolvers in this category support ECS, an encouraging finding for privacy. The study also finds strong correlations between QMIN and Minimal Responses, and between DNSSEC and IPv6.

Paper IV – Measuring the Operational Impact of Minimizing DNS Queries on Real-World Data

This study measures the operational impact of QMIN on DNS resolution success, latency, and packet overhead using over 158 million domain queries drawn from a popularity list, the full .com zone, and long domain names from certificate transparency logs. Four popular resolver implementations of QMIN and their fallback modes are evaluated. The results show that QMIN in relaxed mode (with fallback) achieves a near-baseline success rate, while strict mode (no fallback) leads to failure rates of up to 65% for long-domain queries. Latency increases with label depth, reaching average round-trip times of 570ms for long-tail domains. Packet overhead typically ranges from 2-17% for popular domains but can reach 102% for BIND in strict mode when querying long domains from certificate logs (C4). Based on these results, relaxed-mode QMIN is recom-

mended for general-purpose resolvers, as it balances privacy with reliability. The study also identifies misconfigurations on the authoritative side related to empty non-terminals and nested domains that create interoperability problems under QMIN. While Paper I demonstrates that QMIN adoption has increased significantly over time at Internet scale, this paper explains why such adoption often defaults to relaxed configurations: strict-mode QMIN is not reliably deployable in today’s ecosystem (C2) due to authoritative-side non-compliance (C3).

Paper V – DoHoT or Not? Tuning Tor Circuits for Anonymous DNS Queries

This paper evaluates DNS-over-HTTPS-over-Tor (DoHoT) as an alternative to ODNS and ODoH for anonymous DNS. By separating the client identity and the query contents, it is possible to reduce the privacy risks of using centralized resolvers (C5). While DoHoT has historically been dismissed due to the high latency introduced by Tor’s three-hop circuits, we argue that the threat model for anonymous DNS resolution enables redesigned Tor circuits that use fewer, more geographically proximate relays. This configuration is more resistant to collusion compared to single-relay proxy approaches like ODoH while significantly improving performance compared to default Tor circuits. Experimental evaluation shows that tuned DoHoT circuits reduce median query round-trip time by two thirds, from 107 ms (p5–p95: 63–289 ms) to 37 ms (p5–p95: 20–144 ms) (C4). Despite significant latency reductions, performance remains dependent on the uneven geographic distribution of Tor relays, meaning that users in underrepresented regions may still experience degraded performance (C2). The paper emphasizes that performance alone should not drive comparisons of anonymous DNS schemes; threat model assumptions, adversary capabilities, and deployability also have critical implications for practical end-user privacy.

Paper VI – Here be DDRagons! Navigating the Unexplored Map of Encrypted DNS Transport

This paper extends prior measurement work on DDR, used to advertise encrypted DNS endpoints (DoT, DoH, DoQ). Utilizing CAIDA Ark and RIPE Atlas probes as vantage points, we analyze DDR adoption on pre-configured resolvers. The results show that very few non-public resolvers implement DDR (C1), and those that do merely redirect users to large public resolvers (C3), reinforcing the centralization DDR was meant to reduce (C5). Because RIPE Atlas lacks support for SVCB queries (used by DDR), we use ANY queries as a best-effort approximation and carefully validate the resulting lower-bound estimates. On the authoritative side, we observe that only a single major public resolver operator encrypts its queries to our authoritative server. Probing root, TLD, and authoritative servers for popular domains reveals that support for encrypted DNS

transports remains extremely limited on the authoritative side.

7 Related Work

In this section we position the contributions in this thesis within the context of related work in four areas: QMIN, centralization, anonymization, and encrypted DNS transport.

7.1 QNAME Minimization

Wang [59] provides a performance-oriented evaluation of QMIN, showing that its label-by-label resolution significantly increases query volume, largely due to NXDOMAIN responses for intermediate labels. The study demonstrates that applying NXDOMAIN optimization can substantially reduce this overhead while mitigating the amplification risk introduced by QMIN’s iterative querying. Wang also highlights that broken ENTs remain common and can lead to misresolution when combined with QMIN, underscoring the need for proper ENT handling in real deployments.

De Vries *et al.* [60] conduct the first large-scale measurement of QMIN adoption, combining passive observations at root and TLD servers with active probing of resolvers. Their results show modest but growing adoption and reveal significant variation in resolver implementations.

Deccio *et al.* [61] present a formal analysis of QMIN. Their study evaluates interoperability, performance characteristics, and the effects of minimization on both resolver and authoritative server behavior. Using passive DNS traffic collected from a university campus network, they assess how different minimization strategies perform in real operational settings. They observe that QMIN adds some privacy gains but that it should be considered alongside its operational costs.

Despite these foundational insights, several gaps remain that necessitate the new studies presented in this thesis. First, most previous work on QMIN was conducted within the framework of RFC 7816 [62], which has since been superseded by RFC 9156 [6]. This newer standard introduced critical changes, such as opening up the use of QTYPEs other than NS for minimized queries, such as the less error-prone A/AAAA, and introducing tunable parameters for label prepending, which fundamentally altered resolver behavior. Second, earlier passive measurements at the root servers were heavily skewed by noise from Google’s Chrome browser, which generated millions of non-existing single-label domain names that previous studies did not fully filter out. Furthermore, while Wang and Deccio *et al.* explored operational costs, their evaluations often relied on purpose-built test domains or limited campus traffic. There is a clear need for large-scale measurements against millions of real-world domains to quantify how “long-tail” domains with ten or more labels are affected by resolution failures and latency spikes. Finally, since the landscape of DNS centralization shifted dramatically in 2020 when major providers like Google

Public DNS enabled minimization by default, new longitudinal and fingerprinting studies are necessary to characterize how actors deviate from standards to balance performance and privacy. These updated studies are essential to bridge the gap between theoretical standards and operational reality in an evolving DNS ecosystem. Paper **I** and Paper **II** explicitly validate and extend the results from De Vries *et al.* by investigating the adoption and variations in implementation of QMIN. Paper **III** investigates QMIN adoption among other resolver features across countries and between resolver types. In Paper **IV** we look at the performance impact of operating a QMIN-enabled resolver, and confirm the findings of Wang and Deccio *et al.* using a large set of real-world domains.

7.2 Centralization of Resolvers

Moura *et al.* [21] present early empirical evidence of resolver centralization using traffic from two major ccTLDs and B-Root. They find that a small group of cloud and content providers contributes over a third of all incoming queries, revealing clear consolidation. Their measurements also show heterogeneous resolver behavior across providers, i.e., IPv6 adoption, DNSSEC validation, and QMIN, indicating uneven uptake of modern security features. The work highlights both the reach of large operators and the systemic risks introduced by growing concentration.

Doan *et al.* [63] use RIPE Atlas probes to assess the growing shift toward public DNS services and its contribution to resolver centralization. They find substantial reliance on a few large operators, particularly Google Public DNS, and show that public resolvers can outperform local ISP resolvers in some regions while performing worse in others. The study highlights how uneven global deployment creates region-dependent incentives and reinforces concentration around major public providers, raising concerns about privacy and infrastructural dependence.

Ververis *et al.* [64] evaluate real-world deployments of the DDR standard and find that the small fraction of open resolvers advertising DDR overwhelmingly delegates to Google and Cloudflare, meaning DDR’s current use reinforces rather than reduces centralization. The study highlights how design intentions can be undermined by operational realities, resulting in limited privacy benefits and greater reliance on dominant providers.

The contributions of this thesis extend these foundational works by moving beyond passive monitoring and public-facing scans to investigate the operational realities at the Internet’s edge. While Moura *et al.* used passive measurements at the root and TLD levels to identify centralization as well as the 2020 surge in QMIN driven by Google Public DNS, Paper **I** utilizes longitudinal active measurements to reveal the technical nuances of this adoption, such as the partial depth limits and implementation choices used by centralized providers. Similarly, this thesis builds upon the work of Doan *et al.* by introducing a more granular catego-

rization of resolvers in Paper III. By mapping the forwarding paths of preconfigured resolvers of RIPE Atlas probes, categorized by network proximity, we find that even those operated by local ISPs frequently funnel traffic toward a few dominant public services. Finally, whereas Ververis *et al.* focused on publicly reachable resolvers to show how DDR reinforces centralization, Paper VI extends this analysis to non-public infrastructure. By probing resolvers invisible to traditional scans, we demonstrate that the trend of “encryption outsourcing” is systemic, with 89% of DDR-enabled non-public resolvers signaling a reliance on major public providers rather than hosting their own encrypted endpoints.

7.3 DNS Client Anonymization

While DoH encrypts transport, it still centralizes queries at major providers. One solution to the privacy risks associated with centralized DNS resolvers is to add a proxy between the stub and the resolver.

Schmitt *et al.* [8] propose ODNS to prevent any single entity from seeing both client source IP and the queried domain. ODNS splits knowledge across two parties: a resolver that sees only encrypted domains, and an ODNS resolver that sees plaintext domains but learns nothing about client identity. In the study they argue that Tor, even though it provides client anonymity, is not practical for DNS query anonymization primarily due to the introduction of substantial network latency.

Muffett presents DoHoT [43] as a practical, deployable model that significantly strengthens anonymity without requiring new DNS protocols. Routing DoH over Tor removes the link between the client identity and the query contents, by obscuring the client source IP from the resolver. After extensive testing of DoHoT in a household setting during the pandemic, Muffett argues that the approach maintained acceptable performance for everyday activities despite the increased latency.

Singanamalla *et al.* [9] extend the oblivious-resolution model to encrypted DNS with ODoH, introducing a proxy-and-target architecture that separates client identity from query contents. The proxy observes only the client’s IP address, while the target resolver can decrypt queries but sees only the proxy’s address. This prevents any single operator from reconstructing user histories, addressing privacy risks posed by consolidation among popular providers of encrypted DNS transport. In a comparison of ODoH to other approaches in terms of latency, DoHoT ends up being the slowest alternative.

Kurihara *et al.* [10] introduce μ ODNS, addressing a core weakness in relay-based anonymity systems such as ODNS, ODoH, and DNSCrypt: the possibility of relay-resolver collusion. μ ODNS applies a multi-relay, Tor-like design to make deanonymization significantly harder even if multiple infrastructure components cooperate. By generalizing the oblivious model and distributing trust across several relays, μ ODNS offers stronger anonymity guarantees while maintaining feasible performance

for real-world use. In the study, Kurihara *et al.* mentions using Tor as a simple and generic approach to client anonymity in DNS, but highlights the significant performance loss due to the multi-layered encryption, volume of traffic, and increased round-trip time caused by sparsely distributed nodes geographically.

Most of the comparisons between approaches in anonymizing DNS clients above have been focusing on latency [8, 9, 10], since proxy-based solutions inherently increase the distance between the client and the resolver. In Paper V we argue that factors such as threat model and deployability should also be carefully considered when comparing solutions. We also argue that it is possible to reduce latency of Tor circuits in DoHoT for this use case, while still having a stronger threat model than single-proxy approaches.

7.4 Encrypted DNS Transport

There have been many studies around encrypted DNS transport. In this section we will focus on those that look at the adoption and discovery of these encrypted endpoints in accordance with our research questions, instead of performance evaluation or analysis of encrypted traffic.

Deccio & Davis [65] provide an early Internet-wide measurement of DoT and DoH deployment across open resolvers and authoritative servers. Despite increasing interest in encrypted DNS, they find that real-world support in 2019 was extremely sparse: only a tiny fraction of open resolvers offered DoT, almost exclusively large public operators, and DoH support was even rarer. Authoritative deployment was virtually nonexistent. Their work emphasizes the slow and uneven adoption of encrypted transport protocols and highlights how early deployment was concentrated among a few major providers.

Li *et al.* [66] present the most comprehensive global study of encrypted DNS reachability. By assembling a vetted list of operational encrypted DNS endpoints and querying them from thousands of global vantage points, they reveal how factors such as network policy, regional censorship, and transport-layer behavior shape actual accessibility. Their measurements show wide geographic variability and underscore that protocol support alone does not guarantee reachability, making encrypted DNS performance heavily dependent on local network conditions and governance environments. At the time, DDR was still an Internet-Draft, but they identified it as a promising mechanism for enabling clients to discover encrypted resolvers. Scanning the IPv4 address space, they found roughly 317,000 resolvers advertising DDR, with more than 90% operated by just three major public providers.

Ververis *et al.* [64] investigate the operational state of the DDR standard, focusing on how open resolvers advertise and negotiate encrypted DNS transports. Their measurements reveal that widespread deployment is low and that most resolvers do not present valid DDR-compatible cer-

tificates, meaning clients cannot reliably verify or upgrade to encrypted transports. They also reveal severe resolver consolidation, finding that over 97% of DDR-enabled resolvers delegate to just four major providers, which raises significant concerns about DNS centralization and privacy risks.

In Paper VI we complement the studies from Li *et al.* and Ververis *et al.*, investigating non-public resolvers using two distributed Internet measurement platforms: CAIDA Ark and RIPE Atlas. Results showed that DDR adoption was low among the non-public resolvers. Paper VI also extends the work of Deccio & Davis by probing authoritative name servers for commonly used encrypted transport protocols, adding DoQ to the previously tested DoT and DoH.

8 Conclusions and Future Work

Throughout this thesis, we conducted an Internet-scale, measurement-driven investigation into the privacy challenges surrounding DNS resolvers. Our analysis moved from mechanism intent to real-world deployment practices, and finally to the broader implications for clients, operators, and the DNS ecosystem. Our results reveal a consistent theme: while privacy-oriented mechanisms are gaining traction, their deployment is uneven and often fragile in practice.

We observed a substantial longitudinal increase in QMIN adoption, though the rate varied significantly across regions and operator types. Much of this growth was driven by large public resolvers, whose defaults exert considerable influence on the ecosystem. Cloudflare and Google Public DNS that enable QMIN may promote broader adoption while simultaneously concentrating operational control. In the Nordics and Baltics in particular, several privacy and security features achieved high average adoption rates, though variation across countries and resolver types persists. In contrast, DDR adoption among non-public resolvers remains modest, highlighting a persistent gap between formal standardization and actual deployment at the network edge.

At the implementation level, our findings point to structural weaknesses. No evaluated open-source resolver applied QMIN with the recommended settings in the RFC, leaving them vulnerable to fingerprinting via query-pattern analysis. Our operational assessment further revealed that authoritative servers frequently mishandle ENTs by returning incorrect Response Codes (RCODEs), which causes breakage when QMIN is used in strict mode. This widespread misconfiguration makes standards-aligned privacy mechanisms brittle in real deployments. Similarly, DDR SVCB records were often misused, pointing to endpoints hosted by other operators, effectively reinforcing centralization.

Performance considerations introduce further trade-offs. Combined with parallel resolver forwarding configurations, QMIN can cause unnecessary query amplification, putting extra load on authoritative name

servers. Using QMIN in strict mode increases failure rates, latency, and packet overhead, posing challenges for operators seeking to adopt it safely. DoH over Tor can be optimized to roughly halve latency, but these improvements are unevenly distributed due to global relay imbalances. Thus, while privacy enhancements are technically feasible and increasingly deployed, operational fragility, performance costs, and centralization pressures collectively limit how effectively they translate into user-facing privacy. This work identifies where today’s DNS privacy measures succeed, where they fail, and the mechanisms behind those outcomes.

Looking ahead, several opportunities exist to strengthen privacy deployments. Expanding regional coverage and performing continuous longitudinal scans would help capture policy fluctuations and seasonal effects. Developing robust fallback policies for QMIN, together with explicit signaling, could mitigate downgrade risks without compromising reliability. Additional contributions include providing authoritative server tests for ENTs, publishing conformance suites, and “known-bad” behavior profiles. Proposing specification clarifications, enforcement mechanisms, or validation tools for DDR’s “same-entity” binding would also be beneficial. Complementing these technical efforts with operator-focused user studies could reveal additional barriers to practical adoption of privacy features.

In sum, Internet-scale measurements show that DNS privacy is advancing, yet current progress relies on fragile implementations and increasingly concentrated infrastructures. The path forward requires not only new standards but also attention to operation in the wild, tooling for validation, and improved interoperability.

References

- [1] 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. <http://www.rfc-editor.org/rfc/rfc7858.txt>.
- [2] P. Hoffman and P. McManus. DNS queries over HTTPS (DoH). RFC 8484, RFC Editor, October 2018. <http://www.rfc-editor.org/rfc/rfc8484.txt>.
- [3] C. Huitema, S. Dickinson, and A. Mankin. DNS over dedicated QUIC connections. RFC 9250, RFC Editor, May 2022. <http://www.rfc-editor.org/rfc/rfc9250.txt>.
- [4] P. Hoffman. DNS security extensions (DNSSEC). RFC 9364, RFC Editor, February 2023. <http://www.rfc-editor.org/rfc/rfc9364.txt>.
- [5] J. Abley and K. Lindqvist. Operation of anycast services. BCP 126, RFC Editor, December 2006. <http://www.rfc-editor.org/rfc/rfc4786.txt>.

- [6] S. Bortzmeyer, R. Dolmans, and P. Hoffman. DNS Query Name Minimisation to Improve Privacy. RFC 9156, RFC Editor, November 2021. <http://www.rfc-editor.org/rfc/rfc9156.txt>.
- [7] E. Kinnear, P. McManus, T. Pauly, T. Verma, and C.A. Wood. Oblivious DNS over HTTPS. RFC 9230, RFC Editor, June 2022. <http://www.rfc-editor.org/rfc/rfc9230.txt>.
- [8] Paul Schmitt, Anne Edmundson, and Nick Feamster. Oblivious DNS: Practical privacy for DNS queries. *arXiv preprint arXiv:1806.00276*, 2018.
- [9] Sudheesh Singanamalla, Suphanat Chunhapanya, Marek Vavruša, Tanya Verma, Peter Wu, Marwan Fayed, Kurtis Heimerl, Nick Sullivan, and Christopher Wood. Oblivious DNS over HTTPS (ODOH): A Practical Privacy Enhancement to DNS. *arXiv preprint arXiv:2011.10121*, 2020.
- [10] Jun Kurihara, Toshiaki Tanaka, and Takeshi Kubo. μ ODNS: A distributed approach to DNS anonymization with collusion resistance. *Computer Networks*, 237:110078, 2023.
- [11] The Tor Project. Tor Project. <https://www.torproject.org/>. Accessed: 2026-04-24.
- [12] T. Pauly, E. Kinnear, C. A. Wood, P. McManus, and T. Jensen. Discovery of designated resolvers. RFC 9462, RFC Editor, November 2023. <http://www.rfc-editor.org/rfc/rfc9462.txt>.
- [13] Jon Postel. Internet protocol. RFC 791, RFC Editor, September 1981. <http://www.rfc-editor.org/rfc/rfc791.txt>.
- [14] P. Mockapetris. Domain names - concepts and facilities. RFC 1034, RFC Editor, November 1987. <http://www.rfc-editor.org/rfc/rfc1034.txt>.
- [15] P. Mockapetris. Domain names - implementation and specification. RFC 1035, RFC Editor, November 1987. <http://www.rfc-editor.org/rfc/rfc1035.txt>.
- [16] Google. Public DNS | Google for Developers. <https://developers.google.com/speed/public-dns/>. Accessed: 2026-04-24.
- [17] Cloudflare. What is 1.1.1.1? | Cloudflare. <https://www.cloudflare.com/learning/dns/what-is-1.1.1.1/>. Accessed: 2026-04-24.
- [18] OpenDNS. Setup guide | OpenDNS. <https://www.opendns.com/setupguide/>. Accessed: 2026-04-24.
- [19] Quad9. Quad9 | a public and free DNS service for a better security and privacy. <https://www.quad9.net/>. Accessed: 2026-04-24.

- [20] Luciano Zembruzki, Arthur Selle Jacobs, Gustavo Spier Landtreter, Lisandro Zambenedetti Granville, and Giovane CM Moura. dnstracker: Measuring centralization of DNS infrastructure in the wild. In *International Conference on Advanced Information Networking and Applications*, pages 871–882. Springer, 2020.
- [21] Giovane CM Moura, Sebastian Castro, Wes Hardaker, Maarten Wullink, and Cristian Hesselman. Clouding up the internet: How centralized is DNS traffic becoming? In *Proceedings of the ACM Internet Measurement Conference*, pages 42–49, 2020.
- [22] Austin Hounsel, Paul Schmitt, Kevin Borgolte, and Nick Feamster. Designing for tussle in encrypted DNS. In *Proceedings of the 20th ACM Workshop on Hot Topics in Networks*, pages 1–8, 2021.
- [23] T. Wicinski. DNS Privacy Considerations. RFC 9076, RFC Editor, July 2021. <http://www.rfc-editor.org/rfc/rfc9076.txt>.
- [24] Selling access to resolver logs privacy and wiretap issues. <https://dn.org/selling-access-to-resolver-logs-privacy-and-wiretap-issues/>. Accessed: 2026-05-12.
- [25] Sandra Siby, Marc Juarez, Claudia Diaz, Narseo Vallina-Rodriguez, and Carmela Troncoso. Encrypted DNS -> privacy? a traffic analysis perspective. *arXiv preprint arXiv:1906.09682*, 2019.
- [26] Yaser M Banadaki and S Robert. Detecting malicious DNS over HTTPS traffic in domain name system using machine learning classifiers. *Journal of Computer Sciences and Applications*, 8(2):46–55, 2020.
- [27] Rikima Mitsuhashi, Yong Jin, Katsuyoshi Iida, Takahiro Shinagawa, and Yoshiaki Takai. Malicious DNS tunnel tool recognition using persistent DoH traffic analysis. *IEEE Transactions on Network and Service Management*, 20(2):2086–2095, 2022.
- [28] Marta Moure-Garrido, Celeste Campo, and Carlos Garcia-Rubio. Detecting malicious use of DoH tunnels using statistical traffic analysis. In *Proceedings of the 19th ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, & Ubiquitous Networks*, pages 25–32, 2022.
- [29] 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.
- [30] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, RFC Editor, August 2018. <http://www.rfc-editor.org/rfc/rfc8446.txt>.

- [31] J. Iyengar and M. Thomson. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000, RFC Editor, May 2021. <http://www.rfc-editor.org/rfc/rfc9000.txt>.
- [32] Mike Kosek, Trinh Viet Doan, Malte Granderath, and Vaibhav Bajpai. One to rule them all? A first look at DNS over QUIC. In *International Conference on Passive and Active Network Measurement*, pages 537–551. Springer, 2022.
- [33] Mike Kosek, Luca Schumann, Robin Marx, Trinh Viet Doan, and Vaibhav Bajpai. DNS privacy with speed? Evaluating DNS over QUIC and its impact on web performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 44–50, 2022.
- [34] Marc Fischlin and Felix Günther. Replay attacks on zero round-trip time: The case of the tls 1.3 handshake candidates. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 60–75. IEEE, 2017.
- [35] Alexandra Nisenoff, Ranya Sharma, and Nick Feamster. Understanding user awareness and behaviors concerning encrypted DNS settings. *arXiv preprint arXiv:2208.04991*, 2022.
- [36] D. K. Gillmor, J. Salazar, and P. Hoffman. Unilateral Opportunistic Deployment of Encrypted Recursive-to-Authoritative DNS. RFC 9539, RFC Editor, February 2024. <http://www.rfc-editor.org/rfc/rfc9539.txt>.
- [37] IETF Internet-Draft. Authoritative DNS Transport Signaling. <https://datatracker.ietf.org/doc/draft-johani-dnsop-transport-signaling/>. Accessed: 2026-04-24.
- [38] M. Bishop. Http/3. RFC 9114, RFC Editor, June 2022. <http://www.rfc-editor.org/rfc/rfc9114.txt>.
- [39] Philipp Bielefeld, Felix Hoffmann, Steffen Sassalla, vasilis ververis, and Vaibhav Bajpai. The Future of DNS Privacy: A Comparison of DNS over QUIC and DNS over HTTP/3. In *International Conference on Passive and Active Network Measurement*, pages 202–228. Springer, 2026.
- [40] D. Atkins and R. Austein. Threat Analysis of the Domain Name System (DNS). RFC 3833, RFC Editor, August 2004. <http://www.rfc-editor.org/rfc/rfc3833.txt>.
- [41] Internet Assigned Numbers Authority. Root Servers. <https://www.iana.org/domains/root/servers>. Accessed: 2026-04-24.
- [42] Root Server Technical Operations Association. root-servers. <https://root-servers.org/>. Accessed: 2026-04-24.

- [43] Alec Muffet. DoHoT: making practical use of DNS over HTTPS over Tor. <https://github.com/alecmuffett/dohot>. Accessed: 2026-04-24.
- [44] NLnet Labs. Unbound Documentation: qmin strict. <https://unbound.docs.nlnetlabs.nl/en/latest/manpages/unbound.conf.html#unbound-conf-qname-minimisation-strict>. Accessed: 2026-04-24.
- [45] ISC. BIND Documentation: options. <https://bind9.readthedocs.io/en/latest/reference.html#namedconf-statement-qname-minimization>. Accessed: 2026-04-24.
- [46] C. Contavalli, W. van der Gaast, D. Lawrence, and W. Kumari. Client subnet in DNS queries. RFC 7871, RFC Editor, May 2016. <http://www.rfc-editor.org/rfc/rfc7871.txt>.
- [47] 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.
- [48] Martijn de Vos, Georgy Ishmaev, Johan Pouwelse, and Stefanie Roos. A deployment-first methodology to mechanism design and refinement in distributed systems. In *2023 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, pages 472–477. IEEE, 2023.
- [49] 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.
- [50] David R Choffnes and Fabian E Bustamante. Pitfalls for testbed evaluations of internet systems. *ACM SIGCOMM Computer Communication Review*, 40(2):43–50, 2010.
- [51] Vern Paxson. Strategies for sound internet measurement. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*, pages 263–271, 2004.
- [52] Wilhelm Hasselbring, Leslie Carr, Simon Hettrick, Heather Packer, and Thanassis Tiropanis. Open source research software. *Computer*, 53(8):84–88, 2020.
- [53] RIPE NCC. RIPE atlas. <https://atlas.ripe.net/>, 2010. Accessed: 2026-04-24.
- [54] CAIDA. Ark - CAIDA. <https://www.caida.org/projects/ark/>. Accessed: 2026-04-24.

- [55] Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. A High-Performance, Scalable Infrastructure for Large-scale Active DNS Measurements. *IEEE Journal on Selected Areas in Communications*, 34(6):1877–1888, 2016.
- [56] RIPE NCC. Welcome to the RIPE NCC. <https://www.ripe.net/>, 1997. Accessed: 2026-04-24.
- [57] 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.
- [58] S. Bortzmeyer and S. Huque. Nxdomain: There really is nothing underneath. RFC 8020, RFC Editor, November 2016. <http://www.rfc-editor.org/rfc/rfc8020.txt>.
- [59] Zheng Wang. Understanding the Performance and Challenges of DNS Query Name Minimization. In *17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications / 12th IEEE International Conference On Big Data Science And Engineering, TrustCom/BigDataSE 2018, New York, NY, USA, August 1-3, 2018*, pages 1115–1120. IEEE, 2018.
- [60] 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 *International Conference on Passive and Active Network Measurement*, pages 147–160. Springer, 2019.
- [61] Casey Deccio, Robert Richardson, Nathaniel Bennett, and Nathan Craddock. Modeling DNS Queries and Caching to Evaluate the Merits of QNAME Minimization. In *2025 IEEE 33rd International Conference on Network Protocols (ICNP)*, pages 1–11. IEEE, 2025.
- [62] S. Bortzmeyer. Dns query name minimisation to improve privacy. RFC 7816, RFC Editor, March 2016. <http://www.rfc-editor.org/rfc/rfc7816.txt>.
- [63] Trinh Viet Doan, Justus Fries, and Vaibhav Bajpai. Evaluating public DNS services in the wake of increasing centralization of DNS. In *2021 IFIP Networking Conference (IFIP Networking)*, pages 1–9. IEEE, 2021.
- [64] Vasilis Ververis, Steffen Sassala, Felix Roth, and Vaibhav Bajpai. Path to Encrypted DNS with DDR: Adoption, Configuration Patterns, and Privacy Implications. *Proceedings on Privacy Enhancing Technologies*, 2025.

- [65] Casey Deccio and Jacob Davis. DNS privacy in practice and preparation. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, pages 138–143, 2019.
- [66] Ruixuan Li, Baojun Liu, Chaoyi Lu, Haixin Duan, and Jun Shao. A worldwide view on the reachability of encrypted DNS services. In *Proceedings of the ACM Web Conference 2024*, pages 1193–1202, 2024.



Measuring Adoption, Implementation, and Impact of Available Privacy Solutions in DNS Resolvers

The Domain Name System (DNS) is a core Internet infrastructure that mainly translates domain names into network addresses, yet its operation reveals sensitive information about client activity. This thesis investigates the privacy challenges involving DNS resolvers, which occupy a central position between clients and authoritative name servers, and examines how privacy-enhancing DNS mechanisms are deployed and function in real-world settings. Using large-scale active and passive Internet measurements from multiple vantage points, the thesis analyzes the adoption, performance, and interoperability of modern DNS privacy technologies. The results show increasing deployment of mechanisms largely driven by major public resolvers, alongside persistent challenges including uneven adoption, performance penalties, specification deviations, and growing centralization. Overall, the thesis highlights the tension between improved DNS privacy and operational constraints, underscoring the need for more robust, interoperable solutions.

ISBN 978-91-7867-716-0 (print)

ISBN 978-91-7867-717-7 (pdf)

ISSN 1403-8099

DOCTORAL THESIS | Karlstad University Studies | 2026:33
