

I See DNS People: DNS Resolver Security, Through Operator Perspectives and Practices

Wisdom Obinna
Georgetown University, US
wko7@georgetown.edu

Florian Alt
LMU Munich, Germany
University of the Bundeswehr Munich, Germany
florian.alt@ifi.lmu.de

Katharina Barlage
LMU Munich, Germany
katharina.barlage@ifi.lmu.de

Harel Israel Berger
Ariel University, Israel
Harelb@ariel.ac.il

Abstract

We study how DNS resolver operators and security practitioners make DNS security decisions. Prior research has measured DNS security adoption rates and misconfiguration, but offers limited insight into why operators make the choices they do. To address this gap, we conducted semi-structured interviews with 24 DNS resolver operators and security practitioners from 14 countries, covering service providers, enterprises, and government organizations. Our analysis identifies three interconnected barriers impeding wider adoption. First, there is a gap between how complex DNSSEC is perceived and how complex it is in practice: operators frequently describe DNSSEC as intimidating prior to deployment, yet majority who deployed it reported the process simpler than expected. Two factors help explain why this gap persists. DNS suffers from systematic deprioritization: because DNS “just works,” organizations tend to pay less attention to DNS security until failure. Secondly, expertise gaps, combined with limited resources and coordination costs, amplify the prior two factors. Grounded in practitioners’ own recommendations, we suggest actionable interventions to reduce the complexity of DNSSEC, such as clearer documentation and visible success stories, while also addressing its deprioritization through compliance mandates or security-as-a-mission framing.

1 Introduction

The Domain Name System translates every web address into a routable IP address, yet repeated incidents expose how vulnerable this critical infrastructure remains. In October 2016,

a DDoS attack on DNS provider Dyn took down Twitter, Netflix, and dozens of other services for several hours^{1,2}. Most recently, in October 2025, a race condition in AWS’s internal DNS management system triggered a 14-hour outage, disrupting services across multiple countries³. Each incident exposes a different dimension of DNS fragility, from infrastructure resilience to response authenticity and the cascading risks that arise when DNS is treated as invisible infrastructure unworthy of sustained investment. Taken together, they reveal a deeper, systemic vulnerability at the core of the modern internet.

Technologies to mitigate some of these threats have been standardized for 30 years. Yet, adoption remains limited. DNSSEC provides response authentication. DoH and DoT protect query confidentiality. Rate limiting and DNS firewalls mitigate abuse, while monitoring improves operational visibility. Yet despite their availability, deployment lags. In March 2026, NIST released its first major update to secure DNS deployment guidance since 2013⁴. Studies have documented this gap extensively [2, 7, 25]. They show what operators deploy, but not why those decisions are made. The organizational constraints, perceived risks, and human decision-making processes sustaining this gap remain largely unexplored.

We investigate the organizational and human factors shaping DNS security decisions in practice. Specifically, we ask:

- **RQ1:** What factors shape how DNS operators prioritize security relative to other operational concerns?
- **RQ2:** How do operators perceive and experience specific DNS security mechanisms?

¹Cyber attacks briefly knock out top sites: <https://www.bbc.com/news/technology-37728015>, last accessed: August 26, 2026

²How to not break the Internet: <https://www.kaspersky.com/blog/attack-on-dyn-explained/13325/>, last accessed: August 26, 2026

³Amazon reveals cause of AWS outage: <https://www.theguardian.com/technology/2025/oct/24/amazon-reveals-cause-of-aws-outage>, last accessed: August 26, 2026

⁴Secure Domain Name System (DNS) Deployment Guide: <https://csrc.nist.gov/pubs/sp/800/81/r3/final>, last accessed: August 26, 2026

- **RQ3:** What conditions enable DNS security adoption? What interventions do operators deem most effective?

To answer these questions, we conducted semi-structured interviews with 24 DNS professionals from 23 organizations, including service providers, enterprises, and government bodies across 14 countries. Through thematic analysis following Braun and Clarke [4], we examine how DNS security decisions are shaped by organizational context, perceived complexity, and resource constraints in ways that deployment measurements alone cannot reveal.

Our analysis identifies three interconnected phenomena that impede DNS security adoption. First, DNS security is systematically deprioritized despite near-universal acknowledgment of DNS’s critical importance. 18 of 24 participants described organizational environments in which DNS “just works” and therefore remains invisible until it fails. This reliability creates a self-reinforcing cycle of underinvestment. Availability consistently outranked security in operational priorities, and participants evaluated security measures primarily in terms of their potential to disrupt service.

Second, we identify a perception-reality gap surrounding DNS security mechanisms, most clearly visible in discussions of DNSSEC. Participants frequently described DNSSEC as intimidating prior to deployment. Yet those who had implemented it often characterized the process as more manageable than expected. Perceived complexity, shaped by dense documentation, unfamiliar cryptographic terminology, and peer narratives of failure, appears to function as a barrier partly independent of actual technical difficulty.

Third, constraints in expertise, staffing, resources, and cross-organizational coordination create barriers that no single intervention can resolve. Eight participants reported delegating DNS security substantially to commercial vendors, often without verifying the protections actually provided. This reflected a pattern of assumed rather than assessed security.

These findings situate established technology adoption constructs within a distinctive security domain. Rogers low-observability barrier [28] and the facilitating-conditions construct in UTAUT [32] both help explain our results, but DNS security presents a structural challenge. Unlike productivity tools whose benefits are immediately visible, DNS security creates value by preventing incidents that may never occur, and unlike other security investments, that value cannot be captured unilaterally: DNSSEC requires coordinated action across zone operators, registrars, and resolvers, with no enforcement mechanism analogous to browser-enforced HTTPS to compel adoption. Security investments therefore remain structurally disadvantaged relative to operational improvements with clear short-term returns.

We outline implications for three audiences. For tool developers and standards bodies, practice-oriented documentation that separates “how to set it up” from “how it works,” combined with improved visibility into security posture, may nar-

row the perception-reality gap more effectively than protocol simplification alone. For DNS operators, managed security services and incremental deployment pathways can reduce perception barriers, while structured vendor evaluation practices can mitigate the risks of unexamined security delegation. For policymakers, compliance mandates can break the cycle of neglect that stalls voluntary adoption — participants in regulated sectors described them as effective forcing functions.

Contributions. This paper makes three contributions: (1) We provide the first qualitative investigation of why recursive resolver operators make particular DNS security decisions. Prior measurement studies document what is deployed; we reveal the organizational and perceptual drivers quantitative methods cannot capture. (2) We identify a perception-reality gap around DNSSEC, in which operators overestimate deployment difficulty relative to post-implementation experience. We identify how this gap interacts with low prioritization and reliance on vendors to sustain underdeployment. (3) We derive practitioner-grounded design implications: documentation redesign and improved observability for tool developers, incremental deployment and vendor evaluation practices for operators, and forcing functions for policymakers.

2 Background and Related Work

2.1 Current State of DNS Security

DNS security encompasses mechanisms protecting the Domain Name System from attacks, including cache poisoning, spoofing, and man-in-the-middle interception, with DNSSEC serving as the primary cryptographic protocol for ensuring data authenticity and integrity through digital signatures anchored in a hierarchical chain of trust [1]. Measurement studies consistently reveal limited adoption despite protocol maturity. Chung et al. [6] conducted a 21-month longitudinal study probing over 1.5 million domains and 60,000 resolvers, finding that DNSSEC signing was present in 0.6% of .com domains and 1% of .org domains, indicating minimal deployment despite years of standardization. Among resolvers, only 12% performed validation, and critically, 31% of signed zones contained validation errors, including expired signatures and missing records. The study further revealed that 39% of domains used insufficiently strong key-signing keys.

Recent measurements show modest improvement. Misell [22] analyzed DNSSEC bootstrapping signals using RFC 9615 mechanisms, finding adoption increased to 5.5%, though many of the 303,000 identified secure islands remain hindered by DS record upload barriers, a challenge partially alleviated by automated bootstrapping deployed by providers such as Cloudflare. Complementing these findings, Ashiq et al. [2] analyzed 1.1 million DNSViz diagnostic snapshots from 2020–2024 and found that NSEC3 misconfigurations, delegation failures, and expired signatures account for over

70% of validation failures, with 18% of affected domains remaining broken for extended periods.

Resolver-side validation presents an additional gap in the security chain. Although most open resolvers are DNSSEC-enabled, fewer than 18% of IPv4 resolvers (and 38% of IPv6 resolvers) actually perform cryptographic validation of received responses [25]. This disparity means that properly signed domains often go unvalidated, undermining the security benefits DNSSEC provides. The validation gap persists because operators must balance security against performance overhead, legacy system compatibility, and the risk of blocking legitimate traffic due to upstream misconfigurations.

Beyond DNSSEC, encrypted DNS protocols such as DNS over HTTPS (DoH) and DNS over TLS (DoT) address confidentiality risks on the path between clients and recursive resolvers. Deployment has been driven primarily by browser vendors: for example, Mozilla enabled DoH by default for U.S. Firefox users in February 2020⁵, with similar support subsequently integrated into other major platforms [16]. While encrypted DNS improves user privacy, it introduces operational challenges for enterprise networks. Prior work and industry guidance highlight that DoH and DoT can bypass traditional DNS-based security controls, reducing network visibility and complicating policy enforcement, traffic monitoring, and threat detection⁶ [21]. Hence, enterprise deployments typically favor tightly managed, organization-controlled encrypted resolvers rather than unrestricted use of third-party DoH services, reflecting an ongoing tension between privacy guarantees and operational security requirements.

DNSSEC's inherent complexity also poses security risks even in mature implementations. Heftrig et al. [15] demonstrated that cryptographic agility, while critical for enabling migration to stronger algorithms, introduces severe vulnerabilities. The KeyTrap attack, disclosed in 2023 (CVE-2023-50387), revealed algorithmic complexity exploits that can cause denial-of-service conditions through specially crafted responses. Such findings underscore that DNS security mechanisms can themselves become attack vectors when their complexity is not carefully managed.

2.2 Security Technology Adoption

Explaining why organizations adopt or resist security technologies requires theoretical frameworks accounting for both individual decision-making and organizational contexts. The Technology Acceptance Model (TAM) posits that adoption is

primarily determined by perceived usefulness and perceived ease of use [10]. This applies to cybersecurity as well, where tools perceived as cumbersome or offering unclear benefits face resistance [5, 14]. Still, TAM's individual-level focus limits explanatory power for organizational security decisions involving stakeholder and infrastructure dependencies.

To address these limitations, the Unified Theory of Acceptance and Use of Technology (UTAUT) synthesizes prior models into a framework incorporating performance expectancy, effort expectancy, social influence, and facilitating conditions [32]. Modified UTAUT models applied to information security find that social influence shapes adoption intention, while facilitating conditions shape actual use [19, 20]. Complementing individual-level theories, the Technology-Organization-Environment (TOE) framework examines organizational-level adoption through technological, organizational, and environmental contexts [30]. However, Wallace et al. [33] found the traditional TOE framework insufficient for the cybersecurity context, proposing an extended framework with additional dimensions, including cybercatalysts and practice standards, to better capture the range of factors influencing cybersecurity adoption decisions.

Diffusion of Innovation (DOI) theory offers an additional lens by identifying attributes such as relative advantage, compatibility, and complexity as determinants of adoption rates [28]. DOI and combined models have explained the stages and outcomes of security innovation adoption in organizational settings [13]. Notably, empirical DNS security measurements support DOI's complexity predictions: Yajima et al. [35] found adoption rates across eight DNS security mechanisms inversely correlate with configuration difficulty, with simple mechanisms like SPF achieving far greater deployment than complex ones like DNSSEC.

2.3 Technical Complexity and Usability

DNSSEC introduces operational complexity that extends beyond initial deployment to continuous maintenance. Zone operators must generate and manage cryptographic key material, perform key rollovers, maintain valid signatures, and coordinate DS record publication with parent zones to preserve a complete chain of trust [1, 23]. These tasks impose operational burdens, particularly when distributed across stakeholders.

This multi-party coordination frequently fails in practice. 40% of domain owners who initiated DNSSEC deployment using third-party DNS operators did not complete the registrar-side steps required to establish a valid chain of trust, highlighting registrar-operator coordination as a persistent source of deployment friction [7]. Misconfigurations remain prevalent even among domains that attempt full DNSSEC deployment. Van Adrichem et al. [31] analyzed DNSSEC-enabled domains across six top-level domains and found that over 4% exhibited configuration errors, with nearly 75% of misconfigured domains becoming unreachable from DNSSEC-validating

⁵Firefox continues push to bring DNS over HTTPS by default for US users: <https://blog.mozilla.org/en/firefox/firefox-continues-push-to-bring-dns-over-https-by-default-for-us-users/>, last accessed: August 26, 2026

⁶NSA Recommends How Enterprises Can Securely Adopt Encrypted DNS: <https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/2471956/nsa-recommends-how-enterprises-can-securely-adopt-encrypted-dns/>, last accessed: August 26, 2026

resolvers. The dominant failure mode—missing DS records at the parent zone—accounted for 68% of observed failures, underscoring the fragility of inter-zone coordination.

Large-scale measurements indicate that these issues persist and disproportionately affect domains managed by a limited number of hosting and DNS providers. Nosyk et al. [24] show that DNSSEC validation failures commonly arise from missing or mismatched DNSKEY and DS records, expired signatures, and unsupported algorithms, with errors often clustered at the provider level. Critically, such failures may remain invisible to operators: non-validating resolvers continue to resolve affected domains correctly, masking deployment errors unless active DNSSEC monitoring is performed.

While DNSSEC exemplifies the coordination and complexity barriers to adopting DNS security, these challenges extend across multiple security mechanisms. Table 1 summarizes deployment rates and key adoption barriers for eight DNS security mechanisms.

3 Methodology

We conducted semi-structured interviews with 24 DNS resolver operators and IT professionals across enterprise IT departments, Internet Service Providers, public DNS service providers, and academic/research networks. Participants were recruited through professional networking platforms: (1) personal contacts; (2) technical mailing lists (e.g., NANOG, DNS-OARC); (3) DNS-OARC Mattermost channel; and (4) private LinkedIn messages. The following subsections detail our recruitment process, data collection, and analysis.

3.1 Participant Recruitment and Screening

We employed a mixed-methods approach, recruiting from July–December 2025 to secure a diverse pool of experts for in-depth interviews on perceptions and views of DNS threats.

Recruitment Recruitment began in July 2025 using personal contacts from both academia and industry. This initial approach yielded 2 participants: their interviews served as a pilot study to validate and refine our interview protocol, ensuring its representativeness and alignment with our research methodology. Recognizing the need for a broader reach, we implemented two additional strategies:

1. **Community Outreach:** We posted recruitment announcements on the DNS-OARC Mattermost channel⁷ and its associated mailing list. The initial posts included a link to a Google Form detailing the methodology and gathering demographic information. Three respondents replied saying they would not answer this kind of inquiry

⁷DNS-OARC Mattermost channel: <https://chat.dns-oarc.net/>, last accessed: August 26, 2026

or open external links. Thus, we subsequently adapted our approach to rely on direct professional networking.

2. **Targeted LinkedIn Sourcing:** Between August and December 2025, we shifted to highly targeted recruiting on LinkedIn for IT professionals and DNS experts globally.

- *Search Criteria:* We initially searched for “DNS” and refined this by looking for job titles and expertise keywords such as “DNS Engineering” and “DDI” (DNS, DHCP, IPAM). We did not set a minimum years-of-experience requirement, to capture views of both fresh and experienced individuals across large and small organizations. We combined these terms with country names (e.g., India, United States, UK) to ensure geographic diversity. We also included relevant profiles surfaced through LinkedIn’s continuously updated “People You May Know” suggestions, which helped broaden our sample beyond keyword-based search results.
- *Expert Vetting:* To ensure relevance, we meticulously reviewed profiles, looking for skills like “DNS,” “DNS management,” and “DNS server” to confirm potential participants possessed the practical experience necessary to discuss DNS security.
- *Connection Invitations:* This vetting process resulted in identifying approximately 1000 suitable candidates. A connection invite was sent to those individuals from one of the research team members’ established profiles, which featured top-tier publications on DNS and security.

3. **Industry Outreach:** In parallel, we contacted several DNS resolution service providers and DNS-related companies through their public “Contact Us” channels, briefly describing the study and inviting operational staff to participate. This approach yielded one participant (P2). This person was then contacted through email.

Screening and Interview Engagement For LinkedIn users, once a connection was established, a detailed recruitment message (provided in the Appendix) was sent, inviting the contact to participate in a 45–60-minute interview on perceptions of DNS threats. A \$20 Amazon gift card was offered as gratitude for their time, though most participants expressed low interest in this incentive.

- **Initial Contact:** We sent this first message to 700 individuals who accepted the connection request.
- **Follow-up:** Following a low initial response rate, a second follow-up message was sent after several days.

The screening involved self-selection and formal steps:

Table 1: DNS Security Mechanisms: Deployment and Barriers

Mechanism	Deployment Status	Key Adoption Barriers
DNSSEC	5.5% domains, 12% validate [6, 22]	Coordination failures (40%) [7]; Configuration complexity (4% errors) [31]; Key management overhead
DoH/DoT	Browser-driven; Limited enterprise [16, 21]	Bypasses controls [21] ⁶ ; Reduces visibility; Privacy-ops tension
QNAME Minimization	Not measured	Unclear tradeoffs; Performance concerns [11]
Port Randomization	Not measured	Legacy compatibility issues [18]
0x20 Encoding	Not measured	Limited effectiveness [9]
Redundancy	17% (G2000) [8]	Infrastructure costs; Multi-provider coordination
Registry Lock	<25% (G2000) [8]	Low urgency perception; Admin overhead
RRL	Not measured	Tuning complexity; Traffic blocking risk

Note: “Not measured” = no peer-reviewed quantitative data, justifying qualitative investigation.

- **Self-Screening:** A number of respondents screened themselves out, primarily due to a lack of qualification or being bound by a Non-Disclosure Agreement (NDA), preventing them from discussing relevant experiences.
- **Formal Screening:** Ultimately, 100 people formally agreed to be interviewed. These individuals were sent a link to a demographics form to gather essential background data. The form concluded with a Calendly link for convenient interview scheduling.
- **Interview Conversion:** From 100 people who agreed, approximately 50 people proceeded to set up a scheduled interview slot. Of those who scheduled, 24 people successfully attended and completed the full interview. The high number of no-shows was attributed to a heavy workload, as participants reassured us in follow-up communications (both via email and LinkedIn), where we offered a polite opportunity to reschedule.

Final Sample: The final sample of 24 people skews experienced and mid to late-career: most participants fall between their mid-20s and early-60s, with many reporting 15–30 years of DNS exposure. Representation spans Europe, Asia, North America, and Africa, with the largest representation from India and Germany. Occupations cluster around security architecture, engineering, and DNS operations, indicating a technically specialized sample. Organizational contexts varied, but enterprise environments dominated, with some respondents coming from ISPs, government, or private operators. Participants’ details are summarized in Table 2 (Appendix C).

3.2 Procedure

All interviews were conducted remotely via Zoom. Each interview was conducted by a member of the research team. Written informed consent was obtained from all participants. With participants’ consent, interviews were audio-recorded and transcribed. Participation was voluntary, and participants were informed that they could decline to answer any question or withdraw from the study at any time without penalty. Transcripts were reviewed for accuracy and anonymized prior

to analysis, and the data were securely stored on password-protected, access-controlled institutional storage accessible only to the research team.

Participants responded to semi-structured questions focusing on their operational context, DNS infrastructure responsibilities, perceived DNS security threats, implemented and considered security measures, organizational and usability factors influencing adoption, and future challenges in DNS security. Questions were designed to elicit both descriptive accounts of current practices and reflective explanations of decision-making processes. Example questions included: “How do operational priorities such as availability, performance, or security influence DNS security decisions in your organization?” and “Can you describe factors that have made it easier or harder to implement DNS security measures such as DNSSEC or encrypted DNS?” The complete interview protocol is provided in Appendix B.

3.3 Analysis

Braun and Clarke [4] identified six steps for effectively analyzing qualitative data through thematic analysis: (1) familiarization, (2) coding the data, (3) generating initial themes, (4) reviewing the themes, (5) naming and defining the themes, and (6) writing the report. We leveraged this framework to code and classify the interview transcripts thematically using MAXQDA, a qualitative data analysis application that can aid in making sense of vast amounts of textual information [27]. The techniques involved in thematic analysis aided the detection of patterns or themes in the qualitative data from the interviewees [26]. We chose thematic analysis for our research because it is a viable method for identifying patterns and themes that can transform a significant volume of transcribed interview text into a highly organized summary of essential outcomes [12].

We combined deductive and inductive coding. Two researchers independently coded an overlapping subset of interviews using an initial codebook developed deductively from the interview guide and interview notes. Both researchers identified emergent codes inductively. The researchers then met to compare codes, resolve disagreements through dis-

cussion, and consolidate a shared codebook. This process produced a refined codebook, which was then applied in a second coding round. We monitored the emergence of new codes throughout the coding process, and after approximately 13 interviews, no new theme-level codes were generated. The final codes were grouped into higher-level themes based on our research questions, and findings are presented within these themes. The complete codebook, including code definitions and example segments, is available as supplement⁸.

3.4 Study Limitations

Our study has limitations that stem both from the researchers' perspectives and experiences, and from the study design.

Qualitative research is shaped by researchers' backgrounds, assumptions, and prior experiences. Our research team has academic backgrounds in network security and usable security, but we do not have direct, day-to-day responsibility for operating production DNS resolver infrastructure. This positioning offered advantages—we approached interviews without strong prior commitments to particular operational practices and could probe for explanations and trade-offs—but it also introduces the possibility that we missed tacit operational nuances that experienced operators might treat as implicit. To mitigate this risk, we kept reflexive memos throughout data collection and analysis, documented evolving assumptions, and revisited emerging themes when participant accounts challenged our expectations.

Our 24 interviews enable depth rather than population-level representativeness. The sample is enterprise-heavy, hence findings may not transfer directly to other smaller operators. A full breakdown by organization type and size is provided in Table 2. Participants were recruited through professional networks and technical communities, which may overrepresent operators who are more engaged, connected, or more willing to discuss security. Hence, we treat statements about prevalence (e.g., “most participants”) as descriptive of our sample, not as estimates of ecosystem-wide rates.

Interviews capture perceptions, recollections, and explanations of practice; they may be affected by recall limitations or social desirability pressures, particularly around incident experience, security investments, and organizational decision processes. Accordingly, our claims are framed around reported practices and perceived barriers rather than verified deployments or objectively measured outcomes.

We focused on recursive resolver operators and closely related security practitioners because they face distinctive availability/performance obligations and upstream dependency risks. This scope fits our research questions on resolver-side security decision-making, but it limits transferability to other

DNS ecosystem roles (e.g., authoritative operators, registrars/registries) that face different constraints and workflows.

4 Results

4.1 Participant Overview

The 24 participants had substantial experience in DNS operations: 11 had 10–20 years, 7 had over 20 years, and 6 had under 10 years. Organizationally, 6 worked for service providers (ISPs, DNS vendors, managed service providers), 14 for enterprises, and 4 as independent consultants or in private practice; one participant also held a government role. Participants spanned 14 countries across Europe, Asia, Africa, and the Americas, with the largest representations from India (6) and Germany (5). Role types included security engineers and architects (9), DNS operators and engineers (8), and other technical leadership positions (7). A full description of the participants can be found in Appendix C.

4.2 RQ1: What Factors Shape DNS Security Adoption and Prioritization

4.2.1 Threat Perception and Awareness

Participants described DNS as a critical component of organizational infrastructure, with 18 of 24 characterizing DNS as business-critical. Fourteen participants reported that the DNS threat landscape has intensified over time, citing expanding attack surfaces and increasingly sophisticated adversaries:

“The attack landscape has substantially increased. We also know that anything which is connected to the internet has a chance of exploitation.” (P20)

The most commonly cited threats were phishing and malicious domains (14 participants), DDoS attacks (14 participants), cache poisoning (13 participants), and DNS tunneling/data exfiltration (11 participants). Participants also identified amplification attacks (9 participants) and man-in-the-middle attacks (8 participants) as significant concerns. Ten participants specifically highlighted AI-generated threats as an emerging concern:

“With the introduction of artificial intelligence and the introduction of cloud virtualization and all of that, we expect the number of malicious domains and problematic sites to increase as well” (P17)

Disconnect between threat awareness and incident experience. More than half of the participants (13 of 24) reported never experiencing a significant DNS security incident

⁸Access to codebook, code definitions and example segments: https://osf.io/ausg5/overview?view_only=f2fb3e9cb69f449d93394352d0d029d2, last accessed: August 26, 2026

directly. This suggests two interpretations: successful prevention (potentially through vendor-managed security) or under-detection due to monitoring gaps. Both interpretations are consistent with the organizational dynamics we now describe.

4.2.2 DNS Security Deprioritization

The most pervasive finding for RQ1 was security deprioritization, reported by 18 of 24 participants. Participants described DNS as only receiving attention only during outages:

“Even people from other teams may joke that ‘what kind of work does a DNS or DHCP team do?’ They never understand the criticality until it goes down.” (P20)

This perception creates a self-reinforcing cycle: because DNS “just works,” leadership assumes limited need for investment. P24 described how only a crisis could break this cycle: *“The trigger that forced us to change our solution and address the DNS issue was the security issue, precisely. DNS was a service that was not very appreciated at the time.”*

Availability and operational flexibility consistently out-ranked security in participants’ priorities. Ten identified up-time as primary concern, with security framed as a supporting function. P22 articulated the tension: *“Sometimes, some customer wants to secure everything at the time. But if we do very tight security, your service will not be flexible.”* Availability and performance failures trigger immediate response, while security failures may go undetected. This hierarchy has concrete consequences for deployment decisions: the perceived risk of implementing security measures (potential service disruption) [3] often outweighs the perceived risk of not implementing them (potential future attack), cf. Section 4.3.

Resource Constraints and Organizational Silos Resource constraints were common: ten participants cited staffing limitations, seven mentioned budget constraints, and six reported poor collaboration between teams—particularly between network operations and security. Small teams managing critical infrastructure were a common pattern:

“We are actually two people working in our company, managing the whole DNS or the whole global DNS for our company. Very short staff, but we use automation scripts to manage.” (P19)

P9 summarized the coordination challenge: *“The biggest enablers are clear ownership, cross-team cooperation, and leadership support. The main obstacles are silos between network and security teams.”* DNSSEC implementation exemplifies this problem, requiring cooperation between DNS operators, security teams, and external registrars.

These constraints compound one another: without dedicated personnel, expertise cannot develop; without budget,

automation tools cannot be acquired; without cross-team coordination, multi-party mechanisms like DNSSEC cannot be deployed. The result is that security improvements compete with day-to-day operational firefighting for attention.

Security Measure Deployment Patterns Adoption varied across security mechanisms. The most widely deployed were monitoring and logging (18 of 24 participants) and access control lists (14), which are basic perimeter controls requiring minimal configuration. More complex mechanisms saw lower adoption: encrypted DNS protocols such as DoH and DoT (12), rate limiting (9), DNS firewalls (7), and QNAME minimization (5). 13 participants deployed some form of DNSSEC: 10 configured validation on their recursive resolvers and 5 signed their authoritative zones, with some participants doing both.

This adoption gradient from simple/reactive to complex/proactive mirrors prior quantitative findings [35]. Our data illuminate *why* it persists: monitoring and logging align with availability-focused priorities (“if something breaks, we need to see it”), while mechanisms like DNSSEC require upfront investment with benefits that remain invisible unless an attack occurs. Some participants viewed DNSSEC not only as complex but as a potential liability. P6 described how DNSSEC’s larger response sizes create amplification risks:

“The responses are 107 times bigger than normal DNS responses. When you ask the DNS server in a special way, you can shut down the DNS server with the certificates [DNSSEC signatures].” (P6)

Vendor delegation as security strategy. Eight participants relied substantially on commercial security products for DNS protection. In most cases, this reliance was based on trust in the vendor rather than independent confirmation of the protections provided. Three participants noted that the specifics of their vendors’ protections were not fully visible to them, and participants generally did not describe independently testing vendor security claims. As P20 explained:

“For external resolution, I’ve been using third-party tools. And those third-party tools have their own things in place for this kind of attacks.” (P20)

P20’s inability to describe vendor-provided protections and using the vague phrase “have their own things in place” suggests assumed rather than verified security [34]. P21 described the passive nature of this: *“Personally, I just rely on vendors a lot because vendors always want to tell you about all the new stuff. You always get the newsletters and things from them.”*

4.3 RQ2: How Do Operators Perceive and Experience DNS Security Mechanisms?

4.3.1 Expertise Gaps and Perceived Complexity

Lack of expertise emerged as one of the most significant barriers, cited by sixteen of 24 participants. This gap manifested both as insufficient internal expertise and difficulty in finding qualified personnel. P2 captured the paradox: *“It’s an understanding gap. I am shocked constantly with how little is known about the DNS by the security industry in general. Even organizations that say ‘oh yes, we do DNS’ — I ask the most basic questions and they don’t know. The expertise is extremely thin in truly understanding DNS security.”*

The expertise gap reflects DNS’s contradictory status: perceived as simple (“it just resolves names”), yet operationally complex. P4 identified a structural cause: *“it’s hard to make a career in DNS.”* When DNS expertise is neither valued nor developed as a career path, fewer people develop deep DNS skills, reinforcing this deprioritization (cf. Section 4.2.2).

Fourteen participants expressed concerns about the complexity of security mechanisms, and an equal number cited performance concerns. The two are interrelated: complex mechanisms are perceived as more likely to introduce latency or new failure modes. Six participants explicitly described fear of causing outages as a barrier to implementing security changes. P17 captured this: *“Usually the limitation comes from expertise or technical competence. A lot of people have fears that if you touch something that is working, you might break it and you will not be able to fix it again.”*

4.3.2 DNSSEC: Perceptual and Operational Barriers

Participants described DNSSEC challenges in two categories: perceptual barriers that shape expectations before deployment, and operational barriers that arise during and after.

Perceptual barriers. Eight participants described DNSSEC as complex or intimidating before attempting implementation. P4 recalled that early DNSSEC documentation *“ran to seventy pages before a practitioner could begin”*. P21 noted this problem persists in current documentation: *“It looks a little bit intimidating when you start trying to read all the documentation. There should be more documentation about ‘here’s just how you set it up,’ because setting it up is not complicated.”*

Operational barriers. Two recurring challenges emerged: key management and DS record coordination.

Key management failures are the most common source of post-deployment DNSSEC outages. P7 explained this: the zone signing key rolls monthly, while the key signing key turns over annually, which is a longer cycle often missed by operators: *“Most of the time they forget... When that time comes to expire, your zones don’t get answers.”* P6 noted the

downstream consequences: *“You’ll see a lot of failures with DNSSEC. It’s human error because of wrong management, wrong handling, wrong implementation, wrong key rollover.”*

A second challenge involves DS record coordination, where DNSSEC’s multi-party architecture requires operators to coordinate with registrars or parent zones. P1 named it directly: *“One of the biggest challenges is publication of the DS record. You have to publish the DS record from your zone when you sign it... The DS record you have to publish in the parent zone.”* This multi-party coordination requirement means DNSSEC deployment depends on external actors whose responsiveness and capabilities vary. P17 was precise: *“The other aspect is the signing. If your parent zone has not been signed, then obviously, you, the domain owner, cannot sign yours and have a full chain of trust up to the root. Those are two different things.”* The dependency multiplies across every registrar relationship maintained by the operator. P9 illustrated one such case: *“We had situations where if you wanted to update something for this specific zone, the operator requested a piece of paper with a signature from the CEO. As a DNS administrator, you have to climb the stairs all the way to the group CEO of this massive 100,000-employee organization to get a signature. The signature has to meet a timeframe...say it takes you two months to get the signature, and then you provide it to them. If the delay is too massive, they won’t accept it because the document is already outdated. Stuff like that is actually the main challenge nowadays with regards to DNSSEC, but not the actual technical implementation.”*

4.3.3 The Perception–Reality Gap

Eight of the thirteen participants who had deployed DNSSEC described the initial setup as simpler than expected. Six of these eight had initially characterized it as intimidating. P21’s experience illustrates this contrast:

“Then I went to actually set it up, and I was like ‘you know what? It is really easy to set up.’ I’m not sure I actually needed to know exactly how it worked in order to set it up. Configuring it is pretty straightforward.” (P21)

P4 similarly emphasized how modern tooling lowered implementation barriers: *“If you have open-source software and all the commercial software these days has DNSSEC support. If you want to have a service provider, you can go to anyone, be it GoDaddy or Cloudflare; they will support DNSSEC.”*

Participants’ accounts suggest three primary drivers of this perception–reality gap. First, *documentation design*: participants who described DNSSEC as intimidating frequently cited the volume and theoretical depth of available documentation as a source of anxiety, even when the practical steps proved straightforward (P21, P13, P7, P19). P21’s distinction between needing to “know exactly how it worked” and being able to “set it up” suggests current documentation con-

flates conceptual understanding with operational setup. Second, *cryptographic unfamiliarity*: several participants without prior experience described DNSSEC's terminology (key-signing keys, zone-signing keys, DS records, chain of trust) as more intimidating than the underlying mechanisms warranted (P5, P22). Third, *peer accounts of failure*: participants cited community discussions and conference presentations that emphasized DNSSEC failures and misconfigurations (P6, P4), which shaped expectations of difficulty before experience could provide calibration.

The perception–reality gap should not be interpreted as evidence that DNSSEC is objectively simple. Even those who reported easier-than-expected deployment identified burdens. Key rollover management was cited by eight participants as challenging. Six described registrar coordination around DS records as cumbersome. Eight emphasized the overhead of monitoring signature expiry and validation failures.

The gap appears concentrated in initial deployment, which modern tooling and managed services simplified, rather than in ongoing maintenance, which retains genuine operational complexity. This distinction matters. Telling operators that “DNSSEC is easy” would be misleading. A more accurate framing is that initial setup is often more manageable than anticipated, particularly when supported by managed services, while long-term maintenance still requires sustained attention.

Although DNSSEC was the clearest case, four participants described similar dynamics with commercial vendor platforms: systems that appeared complex during evaluation became easier once deployed. This suggests that the perception–reality gap may extend beyond DNSSEC, though our evidence is strongest in that context.

4.3.4 Knowledge and Resource Gaps

Participants relied on diverse sources for DNS security guidance. Industry conferences and training were most frequently cited (14 participants), followed by vendor documentation (11) and independent reading (7). Nobody mentioned academic research or security conferences; the venues referenced were practitioner-oriented (DNS-OARC, NANOG).

Despite the availability of information, participants described meaningful gaps. Eleven reported that existing tools lacked necessary features. Five characterized documentation as poor or overwhelming, and six identified missing or unclear best-practices guidance. P13 captured the tension between technical possibility and organizational capacity:

“Technically, the documentation is there. Every modern DNS server platform on the planet supports DNS security options now. The challenge is getting the time, the money, and the personnel necessary for implementation.” (P13)

These constraints extend beyond knowledge access. The absence of independent comparative evaluations, such as head-

to-head assessments of major DNS security products, leaves operators dependent on vendor claims. P10 encountered this gap after losing DNS connectivity in Azure and filing a support ticket, only to learn: *“Azure Private DNS, we don’t do query logs. How in the hell can you not do query logs? It’s one of the most common things you use in the DNS world for diagnosing situations and issues.”*

4.4 RQ3: Enablers and Interventions

4.4.1 Organizational Enablers

Despite pervasive barriers, participants identified factors that facilitated security adoption. Adherence to best-practice standards and strong leadership support were each cited by 16 participants, followed by proactive security mindsets (13) and compliance mandates (8).

Leadership support. P20 described how organizational culture shapes security investment: *“It is more on the support from the management and the proactive gesture or the culture in an organization which inculcates this thinking where you actually don’t wait for the house to get burned, but you rather have some preventive measures in place before you see the smoke.”* P24 emphasized how leadership awareness transforms DNS’s organizational status: *“Above all, awareness has been created in management and decision-makers that DNS is an extremely critical service and that it really becomes a security hole if proper attention is not paid to it.”*

Compliance as forcing function. While cited by fewer participants (8 of 24), compliance requirements served as powerful adoption drivers. P22 described compliance-driven implementation: *“In Singapore, we have [what] they call the Singapore government agency... they make the policy and document for the DNS. We just follow them.”* Compliance mandates effectively externalize the cost-benefit calculation: when non-compliance carries penalties, the invisibility–deprioritization cycle is interrupted by regulatory pressure.

4.4.2 Conditions for Sustained Prioritization

While the described barriers were pervasive, a small number of participants described organizational contexts where DNS security received sustained attention and investment. Examining these illuminates conditions that enable prioritization.

Mission alignment. P2’s organization (a public DNS provider) positioned security as a core value proposition rather than overhead: *“Security is one of the big reasons that people use our service because they get free cybersecurity. That’s hard to say no to.”* In this context, the invisibility–deprioritization cycle does not apply: DNS security is the product, making its benefits visible by definition. However,

this pathway is structurally limited to organizations whose business model centers on DNS services.

Experiential learning. P24 described how a prior incident transformed organizational attitudes: *“A trigger to be able to change a solution or address the DNS issue was the security issue, precisely.”* The organization subsequently invested in more robust solutions. However, incident-driven prioritization is inherently reactive. P24 noted that post-incident urgency can fade once the immediate crisis passes, returning to baseline deprioritization unless structural changes—such as dedicated roles or recurring budgets—are established.

Regulatory compulsion. P1 and P18 operated in regulated sectors (financial services and healthcare) where compliance frameworks mandated specific DNS security controls. P18 elaborated: *“If it’s a large enough client which actually needs to go by DORA or NIS2 or public frameworks, regulatory frameworks, it’s easier to make them understand through fear of penalties from the state the importance of investing in security measures.”* Compliance was the most reliable pathway to sustained investment in our sample, though participants also cautioned that compliance-driven security risks devolved into checkbox compliance rather than genuine risk reduction.

Implications of exceptional cases. The rarity of sustained DNS security prioritization in our sample underscores why voluntary adoption alone may be insufficient. Of the three pathways identified (mission alignment, experiential learning, and regulatory compulsion), only regulatory compulsion operates independently of organizational circumstance, suggesting its potential as a scalable intervention (cf. also Section 5).

4.4.3 Anticipated Needs

Looking ahead, 19 of 24 participants identified improved tooling, automation, or AI-assisted threat detection as a critical need, making it the most frequently requested. P3 described the need for standards that would allow automated configuration exchange between services, reducing manual integration: *“We need to create a standard so that app authors can publish the details so that services like ours can ingest [those] details to automatically make that work.”*

Eight participants expressed interest in predictive intelligence. The desire for automation reflects the compounding constraints documented above: with limited staff and expertise, participants see automation as the primary means of maintaining security at scale without increasing headcount.

5 Discussion

Our results show that DNS security adoption is shaped more by organizational constraints, perceptual barriers, and resource misalignments than by technical limitations. We syn-

thesize these findings within existing frameworks to draw implications for operators, tool developers, and policymakers.

5.1 RQ1: DNS Security Remains Deprioritized

The most consistent finding across our interviews is systematic deprioritization of DNS security, despite near-universal acknowledgment of DNS’s criticality. Eighteen of 24 participants described this pattern. Its roots are structural rather than attitudinal: operators understand DNS security importance but work in organizational contexts that reward attention to visible, immediate problems over invisible, preventive ones.

This finding aligns with Rogers’ claim that low observability slows technology adoption [28]. In DNS security, the observability problem is compounded. DNS itself is largely invisible during normal operation, and the benefits of securing it are even harder to see. DNS fades into the background until failure, and DNS security is noticed mainly when it fails to prevent an incident. This differs from security domains where the protected asset is highly visible. Web applications, for example, sit directly in front of employees and customers, so investments in web security are easier to justify even when their benefits are similarly preventive.

The TOE framework [30] captures many of the factors participants described, including leadership support, resource constraints, and compliance pressure. What it misses is their dynamic interaction. Participants described a feedback loop in which deprioritization reduces monitoring and visibility, limited visibility dampens perceived risk, and muted risk perceptions justify continued underinvestment. The cycle does not correct itself. It typically breaks only when something external intervenes, such as an incident, a regulatory requirement, or a leadership mandate.

Protection Motivation Theory [29] predicts that increased threat perception paired with confidence in mitigation should increase protective action. Our data complicate this prediction. Participants were aware of DNS threats. Eighteen described DNS as business-critical, and fourteen reported rising threats. Yet, threat awareness did not reliably translate into sustained investment. Three factors help explain the gap.

First, threats often remained abstract. Thirteen participants reported never having experienced a major DNS security incident, reducing the urgency PMT assumes drives action. Second, efficacy was frequently delegated. Twelve participants relied on vendors, which can satisfy the sense that “something is being done” without prompting internal investment. Third, response costs were asymmetric. Participants emphasized risks difficult to price, including service disruption, operational complexity, and the cost of security work displacing reliability tasks. When outages are immediately punished, and security failures are rarely detected, prioritizing availability becomes a rational organizational choice.

Vendor reliance is not a problem in itself. The risk is delegation without verification. Many operators assumed they

were protected without confirming what protections vendors actually enabled or maintained. The result is a gap between perceived and actual security posture. This mirrors a broader pattern in cloud security, where organizations often misread shared-responsibility models and overestimate provider-managed protections [17].

5.2 RQ2: Perception–Reality Gap

The perception–reality gap around DNSSEC complexity is our most actionable finding. Eight participants described DNSSEC as intimidating before deployment. Yet eight of those who had implemented it, six of whom had initially found it intimidating, reported the process was simpler than expected. Among those who found initial deployment manageable, seven operated in enterprise or vendor environments with commercial DNS platforms that abstract key management. By contrast, participants who described the process as burdensome were more likely to manage BIND or open-source infrastructure with manual key rollover. This suggests that the perception–reality gap is not uniformly distributed: tooling context substantially mediates it. The difficulty is mostly in the initial setup, not in ongoing maintenance. Participants attributed the intimidation to documentation design, unfamiliar cryptographic language, and peer narratives that emphasize failure over routine success.

Prior usability research shows that users can struggle to operate security tools correctly even when motivated [34]. Our findings suggest challenges faced by operators may be anticipatory rather than execution-based. What our data adds is specificity about *what* drives the misperception. Participants did not simply say “DNSSEC is hard.” They pointed to documentation that makes DNS setup appear to require deeper theoretical knowledge than deployment actually demands, to cryptographic terminology that sounds more intimidating than the operations it describes, and to community narratives that emphasize failure over success. Each of these drivers is addressable through design rather than protocol changes.

Prior studies documented DNSSEC’s technical challenges, including key management burden, DS record coordination failures [7], and configuration errors affecting a substantial fraction of signed zones [6]. Our findings do not contradict these findings. Rather, they suggest these technical challenges are often amplified by perceived complexity, shaped before operators gain direct experience. The implication is that reducing perceived complexity through task-oriented documentation, clearer success criteria, and better onboarding may yield adoption gains comparable to further protocol simplification.

The reported expertise gap reinforces the same cycle of underinvestment. DNS is often treated as simple infrastructure, “it just resolves names,” which reduces incentives to build or retain specialized skills. Without that expertise, operators struggle to assess whether mechanisms like DNSSEC are truly complex or merely unfamiliar. As P4 noted, it is “hard to

make a career in DNS.” When expertise is undervalued, it remains scarce. Documentation and tooling improvements help, but without organizational recognition that DNS operations demand specialized knowledge, adoption will be limited.

The availability–security tradeoff further amplifies perception barriers. Fourteen participants cited performance concerns, and a similar number raised complexity concerns. In practice, the two are intertwined. Complexity becomes a proxy for outage risk. For DNSSEC, the perceived irreversibility of deployment heightens caution, as rollback is assumed to be costly and fragile. Managed services have substantially reduced actual rollback risk, but those improvements have not yet reshaped operator perceptions. The result is another perception–reality gap, this time at the organizational level.

5.3 RQ3: Scalability and Voluntary Adoption

Our interviews point to three pathways to sustained attention to DNS security. First, mission alignment, which could address the invisibility problem by making DNS security the product itself rather than an overhead. When security is what customers pay for, its benefits become visible by definition, eliminating the deprioritization cycle. But this pathway has inherent limits. Only DNS service providers can position security as a core value proposition, and they represent a small fraction of organizations operating DNS infrastructure. For the vast majority of deployments that use DNS as supporting infrastructure, security will remain invisible regardless of how well the mechanisms are designed or documented.

Secondly, experiential learning, where incidents trigger sustained investment, appears promising but proves unreliable in practice. It depends on the harm that occurs, which is ethically problematic. Our data suggest the urgency it creates often fades. Post-incident prioritization decays without institutional reinforcement. Unless structural changes lock gains in place, such as dedicated roles, recurring budgets, or leadership mandates, organizations tend to revert to baseline deprioritization after the crisis. Moreover, 13 of our 24 participants had never experienced a significant DNS incident, suggesting that many organizations may never encounter this trigger.

Thirdly, regulatory compulsion emerges as the only pathway that scales independently of organizational circumstance. It works whether or not leadership understands DNS, whether staff have expertise, or whether an incident has occurred. This aligns with work on cybersecurity adoption: when security investments compete with operational priorities and benefits remain invisible, external mandates often help overcome organizational inertia [33]. Participants nevertheless cautioned that compliance-driven security can devolve into checkbox exercises, with organizations implementing minimum controls without understanding their purpose. Regulation alone is therefore insufficient. Effective approaches must be paired with mechanisms that reduce the barriers we identified, including clearer documentation to address perception–reality

gaps, tooling that makes security posture visible, ecosystem coordination to resolve multi-party dependencies, and automation capabilities. Without these complementary interventions, mandates may achieve nominal compliance without meaningful security improvement. The adoption history of other security protocols offers a useful reference point for what such a combination of enforcement and support can achieve.

Transport Layer Security (TLS) provides a useful contrast. Its adoption was accelerated by browser enforcement via warnings for non-HTTPS sites and by zero-cost automation via Let's Encrypt and the ACME protocol. Critically, a single operator can deploy TLS unilaterally. DNSSEC has none of these advantages: It offers no comparable enforcement pressure, and its chain of trust requires coordination among the domain operator, registrar, and parent zone, meaning no single party can complete deployment alone. This structural difference helps explain why our participants experienced DNSSEC as harder to adopt than its technical complexity suggests, and why compliance mandates, which are the closest DNS analog to browser-driven TLS enforcement, emerged in our data as a likely adoption driver. The transferable lesson is that the levers that moved TLS, namely automation that lowers deployment costs and enforcement that reduces the voluntary-adoption bottleneck, map directly onto DNS: automated DS-record bootstrapping and regulatory frameworks that classify DNS as critical infrastructure are signs of both.

5.4 Implications

5.4.1 Implications for Practice.

For Technology Developers and Standards Bodies. The perception–reality gap points to design opportunities. Documentation should separate task-oriented quick-start guidance from conceptual reference material and clearly state what operators do *not* need to understand to begin. Tools should also make security posture visible by surfacing blocked threats, validated responses, and prevented attacks, transforming security from abstract insurance into an observable service. Reducing deployment risk is equally important. Staging environments, gradual rollout, and reliable rollback mechanisms would lower the perceived stakes that deter experimentation.

5.4.2 For DNS Operators.

Monitoring and logging provide a natural foundation. Building visibility first allows operators to identify specific threats that justify further investment, shifting the case for security from abstract risk reduction to concrete incident prevention. The vendor delegation pattern we observed, which assumed rather than verified security, is a particularly tractable risk. Operators should audit what protections vendors actually provide rather than assuming comprehensive coverage. Addressing the expertise gap requires longer-term investment. Organizations need to treat DNS expertise as infrastructure, supporting

staff participation in practitioner communities and investing in dedicated training.

5.4.3 For Policymakers.

Regulatory compulsion emerged as the most scalable pathway to adoption, but mandates alone risk producing checkbox compliance. Effective frameworks should pair technical requirements with capacity-building support, including training resources, reference architectures, and subsidized tooling for smaller operators. Many barriers also reflect ecosystem fragmentation rather than individual organizational failures, including DS record coordination, registrar responsiveness, and multi-provider management. Policy interventions that standardize automation interfaces, create liability frameworks that incentivize cooperation, or establish shared infrastructure for key management would address constraints no operator can resolve alone.

6 Conclusion

This study examined organizational and usability barriers to DNS security adoption through interviews with DNS professionals. We find that security investments compete with availability, perceived complexity often exceeds actual difficulty, and organizational factors outweigh technical ones.

Three key findings emerge. First, DNS security is systematically deprioritized and often receives attention only during outages, creating structural barriers to proactive investment. Second, a perception–reality gap exists around DNSSEC complexity: while many view it as intimidating, those who implement it often find it manageable. Third, expertise gaps and resource constraints compound adoption barriers and cannot be addressed through technical improvements alone. Our findings highlight how invisible security benefits and limited organizational support weaken the link between threat awareness and protective behavior. For practitioners, incremental adoption and managed services may reduce entry barriers. For tool developers and standards bodies, clearer, practice-oriented documentation can narrow the perception gap. For policymakers, compliance mandates may accelerate adoption.

DNS remains foundational to the modern Internet, yet security adoption continues to lag. Addressing this gap requires coordinated attention to organizational culture, usability, expertise development, and ecosystem support.

Ethics Statement

The study was approved by the ethics committee of our University. Given that participants discussed organizational security practices, we took particular care to protect confidentiality. All identifying information was removed from transcripts before analysis. Participants are identified only by alphanumeric codes (P1–P24) in all reporting. We also avoided quoting

statements that, while not directly identifying, could enable inference of participant identity, given the relatively small DNS professional community.

Acknowledgments

We thank the North American Network Operators Group (NANOG) and DNS Operations, Analysis, and Research Center (DNS-OARC) communities for their assistance with participant recruitment. Special appreciation to the DNS resolver operators who generously contributed their time and insights to this research. This research has been funded by the Israeli Ministry of Innovation, Science, and Technology, the European Union – NextGeneration EU, and the dtec.bw – Center for Digitization and Technology Research of the Bundeswehr as part of the project *Voice of Wisdom*.

References

- [1] Roy Arends, Rob Austein, Matt Larson, Dan Massey, and Scott Rose. RFC 4033: DNS Security Introduction and Requirements. Technical report, RFC Editor, USA, 2005.
- [2] Md. Ishtiaq Ashiq, Olivier Hureau, Casey Deccio, and Taejoong Chung. Decoding DNSSEC Errors at Scale: An Automated DNSSEC Error Resolution Framework using Insights from DNSViz Logs. In *Proceedings of the 2025 ACM Internet Measurement Conference, IMC '25*, pages 242–257, New York, NY, USA, 2025. Association for Computing Machinery.
- [3] Adam Beaument, M. Angela Sasse, and Mike Wonham. The Compliance Budget: Managing Security Behaviour in Organisations. In *Proceedings of the 2008 New Security Paradigms Workshop, NSPW '08*, pages 47–58, New York, NY, USA, 2008. Association for Computing Machinery.
- [4] Virginia Braun and Victoria Clarke. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [5] Burcu Bulgurcu, Hasan Cavusoglu, and Izak Benbasat. Information Security Policy Compliance: An Empirical Study of Rationality-Based Beliefs and Information Security Awareness. *MIS Quarterly*, 34(3):523–548, 2010.
- [6] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1307–1322, Vancouver, BC, August 2017. USENIX Association.
- [7] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. Understanding the Role of Registrars in DNSSEC Deployment. In *Proceedings of the 2017 Internet Measurement Conference, IMC '17*, pages 369–383, New York, NY, USA, 2017. Association for Computing Machinery.
- [8] CSC Digital Brand Services. 2024 domain security report: How protected are global 2000 companies from domain attacks and digital threats? Technical report, CSC Digital Brand Services, 2024. Accessed: 2025-11-03.
- [9] David Dagon, Manos Antonakakis, Paul Vixie, Tatuya Jinmei, and Wenke Lee. Increased DNS Forgery Resistance through 0x20-bit Encoding: Security via Leet Queries. In *Proceedings of the 15th ACM Conference on Computer and Communications Security, CCS '08*, pages 211–222, New York, NY, USA, 2008. Association for Computing Machinery.
- [10] Fred D. Davis. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3):319–340, 1989.
- [11] Casey Deccio and Jacob Davis. DNS Privacy in Practice and Preparation. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT '19*, pages 138–143, New York, NY, USA, 2019. Association for Computing Machinery.
- [12] Christen Erlingsson and Petra Brysiewicz. A hands-on guide to doing content analysis. *African Journal of Emergency Medicine*, 7(3):93–99, 2017.
- [13] Mumtaz Abdul Hameed and Nalin Asanka Gamagedara Arachchilage. A Conceptual Model for the Organizational Adoption of Information System Security Innovations. In Ramesh C. Joshi and Brij B. Gupta, editors, *Security, Privacy, and Forensics Issues in Big Data*, pages 317–339. IGI Global Scientific Publishing, 2020.
- [14] Tahereh Hasani, Norman O'Reilly, Ali Dehghantanha, Davar Rezanian, and Nadège Levallet. Evaluating the adoption of cybersecurity and its influence on organizational performance. *SN Business & Economics*, 3(97), 2023.
- [15] Elias Heftrig, Haya Schulmann, Niklas Vogel, and Michael Waidner. The Harder You Try, The Harder You Fail: The KeyTrap Denial-of-Service Algorithmic Complexity Attacks on DNSSEC. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security, CCS '24*, pages 497–510, New York, NY, USA, 2024. Association for Computing Machinery.

- [16] 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*, HotNets '21, pages 1–8, New York, NY, USA, 2021. Association for Computing Machinery.
- [17] Wayne A. Jansen and Tim Grance. Guidelines on security and privacy in public cloud computing. NIST Special Publication 800-144, National Institute of Standards and Technology (NIST), Gaithersburg, MD, 2011.
- [18] Amit Klein. BIND 9 DNS cache poisoning. *Report, Trusteer, Ltd*, 3, 2007.
- [19] William W. Lidster. *Factors that Influence Selection of Frameworks for Information Security Program Management: A Correlational Study*. PhD thesis, ProQuest Dissertations and Theses, 2018. Order No. 13424640.
- [20] William W. Lidster and Shawon S. M. Rahman. Identifying Influences to Information Security Framework Adoption: Applying a Modified UTAUT. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 2605–2609, Atlanta, GA, USA, December 2020. IEEE Computer Society.
- [21] Minzhao Lyu, Hassan Habibi Gharakheili, and Vijay Sivaraman. A Survey on DNS Encryption: Current Development, Malware Misuse, and Inference Techniques. *ACM Computing Surveys*, 55(8), December 2022.
- [22] Q Misell, Florian Steurer, Johannes Zirngibl, Anja Feldmann, and Tobias Fiebig. Measuring the deployment of DNSSEC Bootstrapping Using Authenticated Signals. In *Proceedings of the 2025 ACM Internet Measurement Conference*, IMC '25, pages 1002–1009, New York, NY, USA, 2025. Association for Computing Machinery.
- [23] Moritz Müller, Jelte Jansen, Marco Davids, and Willem Toorop. DNSSEC Deployment Metrics Research. Technical report, Internet Corporation for Assigned Names and Numbers (ICANN), August 2022.
- [24] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. Extended DNS Errors: Unlocking the Full Potential of DNS Troubleshooting. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, IMC '23, pages 213–221, New York, NY, USA, 2023. Association for Computing Machinery.
- [25] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. Guardians of DNS integrity: A remote method for identifying DNSSEC validators across the internet. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1470–1479. IEEE, 2023.
- [26] Lorelli S. Nowell, Jill M. Norris, Deborah E. White, and Nancy J. Moules. Thematic analysis: Striving to meet the trustworthiness criteria. *International Journal of Qualitative Methods*, 16(1), 2017.
- [27] Stefan Rädiker and Udo Kuckartz. *Focused Analysis of Qualitative Interviews with MAXQDA: Step by step*, volume 125. MAXQDA Press, Berlin, Germany, 2020.
- [28] Everett M. Rogers. *Diffusion of Innovations*. Free Press, 5th edition, 2003.
- [29] Ronald W. Rogers. A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology*, 91(1):93–114, 1975.
- [30] Louis G. Tornatzky, Mitchell Fleischer, and Alok K. Chakrabarti. *The Processes of Technological Innovation*. Lexington Books, 1990.
- [31] Niels L. M. van Adrichem, Norbert Blenn, Antonio Reyes Lúa, Xin Wang, Muhammad Wasif, Ficky Fatturrahman, and Fernando A. Kuipers. A measurement study of DNSSEC misconfigurations. *Security Informatics*, 4(1):8, 2015.
- [32] Viswanath Venkatesh, Michael G. Morris, Gordon B. Davis, and Fred D. Davis. User acceptance of information technology: Toward a unified view. *Management Information Systems Quarterly*, 27(3):425–478, September 2003.
- [33] Steven Wallace, Karen Y. Green, Catherine Johnson, Joseph Cooper, and Collin Gilstrap. An Extended TOE Framework for Cybersecurity Adoption Decisions. *Communications of the Association for Information Systems*, 47, 2020.
- [34] Alma Whitten and J. D. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *8th USENIX Security Symposium*, pages 169–184. USENIX Association, 1999.
- [35] Masanori Yajima, Daiki Chiba, Yoshiro Yoneya, and Tatsuya Mori. Measuring Adoption of DNS Security Mechanisms with Cross-Sectional Approach. In *2021 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6. IEEE, 2021.

A Open Science

We acknowledge SOUPS' open science policy, which encourages authors to share research artifacts by default. The artifacts associated with this interview study include our interview protocol and qualitative analysis materials such as the codebook. To enable critical review and support replication, we make available our full semi-structured interview protocol, our codebook, and analysis documents via an anonymous GitHub repository.

We do not release the original raw transcripts because they may contain personally identifying information and are covered by our human-subjects research and informed consent commitments (e.g., IRB requirements). The shared transcripts have been carefully anonymized by removing or generalizing direct and indirect identifiers (e.g., names, organizations, locations, and other potentially identifying details) to reduce re-identification risk while preserving the substance necessary to interpret and validate our findings.

B Interview Protocol

B.1 Introduction

Thank you for agreeing to participate in our study! We're researching how IT professionals experience DNS security challenges, threats, and implementation practices. We are especially interested in how you manage DNS infrastructure, respond to security incidents, and how your security practices have evolved over time.

Your responses are confidential and will only be used for academic research. You may choose not to answer any question, and you may stop the interview at any point, for any reason. There are no right or wrong answers—we're just interested in your own views and experiences.

With your permission, I'd like to record this interview using Zoom or an audio recorder. The recording will be transcribed and anonymized. Do you have any questions before we begin? May I start the recording now?

B.2 Professional Experience

1. Could you tell me about your role and experience with DNS resolver operations?

- How long have you been working with DNS infrastructure?
- What are your primary responsibilities regarding DNS management?
- How large is the DNS infrastructure you manage, and what types of users/clients does it serve?

If not clear yet:

- How many queries per day/hour do you typically handle?
- What's the geographic distribution of your users?
- Do you serve internal users, external customers, or both?

3. What kinds of DNS-related tasks or challenges do you deal with regularly?

Note: the point of this question is to collect more specifics compared to the opening question. If necessary, ask them:

- Can you give me some examples of typical daily/weekly tasks?
- Are there aspects of DNS management you prefer to avoid or find particularly challenging? Why?

Note: if they start talking about general IT resource constraints, politely acknowledge then move on. Focus on DNS-specific challenges.

B.3 Threat Perception and Incident Experience

5. What DNS-based threats do you consider most critical to your systems? Why? If they answer in vague terms ("DDoS attacks"), push for specifics.
6. Can you rank these threats in order of concern for your specific environment?
7. Has your perception of DNS threats changed in recent years? If so, how?
 - What events or experiences influenced this change?
8. Have you ever experienced a significant DNS-related security incident? Can you tell me about an example?

If they haven't experienced a major incident:

- Can you imagine scenarios that would be particularly concerning for your infrastructure? What would worry you most?
- ...then, skip to Q11.

9. I'm going to refer to these as "DNS security incidents." How often have these incidents happened to you?
10. Were these incidents similar to the one you just described? What patterns have you noticed in the types of attacks or issues you face? (*phrase depending on response to 7*)
11. When you do encounter DNS security issues, how do you typically respond?
 - What's your incident response process?
 - Who gets involved in the response?

B.4 Technical Implementation and Organizational Context

12. What specific measures have you implemented against cache poisoning attacks? *If they don't explain:*
 - How effective have these been?
 - What challenges did you face during implementation?
13. How do you protect your resolvers from being used in DNS amplification attacks?
 - Are there specific configurations or tools you rely on?
14. How would you mitigate the exploitation of your DNS resolver using DNS tunneling?
 - Have you actually encountered DNS tunneling attempts?
15. How do you protect your resolvers from DNS DDoS attacks? (e.g., amplification, NXDomain, reflection, CacheFlush)
16. Have you implemented DNSSEC or QNAME minimization? What has been your experience?

Important: if they didn't address implementation challenges, ask about them:

 - So, thinking back to technologies like DNSSEC, what were the main barriers to implementation?
17. What organizational factors help you implement DNS security best practices?
 - Suppose you were advising a colleague at another organization about DNS security implementation. What organizational support would you tell them they need?
18. What organizational factors hinder implementing DNS security best practices?
 - Does it depend on the specific security measure? (If so, how / which ones?)
19. How do you balance security measures with performance considerations?
 - Can you give me a specific example where you had to make this trade-off?
20. What type of support do you receive for DNS security initiatives, and from whom?
 - *If not covered by the previous question:*

- Who in your organization makes decisions about DNS security investments?
- How do you justify DNS security spending to management?

21. In which areas of DNS security do you feel you need more information?
22. What would help you stay current with DNS security developments?

B.5 Future Outlook and Wrap-Up

23. How do you see DNS security evolving in the next few years?
 - What new threats are you preparing for?
 - What emerging technologies are you considering?
24. What tools or resources would help you better address DNS security challenges?
 - What's missing from current solutions?
25. What challenges do you face when implementing privacy-enhancing technologies in DNS resolvers?
26. Do you think different organizations face similar DNS security challenges, or does it vary significantly by industry/size?
 - How do you think your DNS security approach compares to industry standards?
27. If you had to advise someone just starting to secure their DNS infrastructure, what would be your top three recommendations?
 - What mistakes should they avoid?
28. How do you test the effectiveness of your DNS security measures?
 - What metrics do you use to evaluate DNS security?
29. As researchers, we might want to assess DNS security implementations across different organizations. How would you suggest we evaluate whether an organization has good DNS security practices?

Ask follow-up questions as necessary to dig into this. For example:

 - What indicators would you look for?
 - What questions would reveal the most about their security posture?
30. Is there anything else about your experiences with DNS security that you'd like to share?
 - Any lessons learned that might help other professionals?

B.6 Closing Statement

(Paraphrase, but don't read aloud—that sounds very impersonal and can come off as not genuine)

This is the end of the study, so I am going to stop recording now.

Thank you very much for your time and for sharing your technical insights. Your experiences are incredibly helpful for understanding how IT professionals approach DNS security and the real-world challenges they face. If you have any follow-up questions or thoughts, please email us.

Do you have any questions for me about the study, or anything else?

End of the Interview.

C Participants Characteristics table

In the table below, we describe our participants. Participants were recruited through three channels: LinkedIn sourcing ([†], n=20), direct industry outreach (^{*}, n=1), and community outreach via professional networks and conferences ([‡], n=3). Our sample represents DNS operators from diverse geographic regions across 14 countries, with DNS experience ranging from 1–2 years to 36 years (median: 15 years). Most participants hold technical degrees in engineering, computer science, or related fields, though educational backgrounds varied from high school with professional certifications to graduate degrees. Professional roles include security engineers, network engineers, DNS operators, and technical leadership positions (CTOs, architects, team leads). Organizations represented span Internet Service Providers (ISPs), enterprises, government agencies, DNS vendors, and private sector companies. The sample is predominantly male (22 male, 2 female), reflecting known gender imbalances in network operations and DNS technical communities. Ages range from 26 to 63 years (median: 43 years).

Table 2: Demographics and Professional Background of DNS Interviewees

ID	Age	Gender	Education	Degree in	Occupation	Country	DNS Exp.	Org. type	Org. Size
P1 [†]	46	male	Bachelors	Information Systems	Security Lead	UK	22 years	ISP	Large
P2 [*]	63	male	Masters	–	Chief Technology Officer	Switzerland	10 years	ISP, Enterprise	Small
P3 [†]	56	male	College	–	Head of Security Operations	UK / Germany	36 years	Enterprise, Private	Small
P4 [‡]	46	male	–	–	Security Engineer (Auth & Recursive)	Germany	25 years	Enterprise	Large
P5 [†]	58	male	College	–	Principal Security Architect	Germany	8 years	Enterprise	Small
P6 [†]	52	male	Training	–	CEO & Principal Security Architect	Germany	26 years	Government, Enterprise	Large
P7 [†]	44	male	College	Mechanical Eng.	Security Engineer	India	13 years	Vendor, Enterprise	Medium-Large
P8 [†]	26	male	College	–	DNS Engineer	India	5 years	Enterprise	Medium-Large
P9 [‡]	40	male	College	Net/Soft Eng. and IT Security	Eng. Training	Germany	16 years	Vendor, Enterprise	Medium
P10 [†]	50	male	College	Electronics and Eng.	Head of DNS Arch./Eng.	India	30 years	Enterprise	Large
P11 [†]	37	male	Masters	Cloud Eng.	Solution Engineer & CTO	Cyprus	15 years	Enterprise	Medium
P12 [†]	48	male	Bachelors	Engineering	Security Engineer	Argentina	27 years	Enterprise	Large
P13 [†]	43	male	Bachelors	Engineering	Security Engineer	USA	20 years	Enterprise, Private	Large
P14 [†]	31	male	Bachelors	Engineering	DDA Engineer	India	7 years	Enterprise, Private	Large
P15 [†]	44	male	Bachelors	Computer Networking	Security Engineer	Jamaica	28 years	Enterprise, Private	Small
P16 [†]	52	male	Bachelors	–	DNS Operator	Spain	20 years	Private	Small
P17 [†]	37	male	Bachelors	Computer Eng.	DNS Operator	Ghana	8 years	Private	Large
P18 [‡]	37	male	Bachelors	Network Security	Network Engineer	Romania	15 years	Vendor	Medium
P19 [†]	36	male	Bachelors	Electronics and Eng.	DNS Operator	India	5–6 years	Enterprise, Private	Medium
P20 [†]	35	male	Masters	Business Administration	DNS Operator	India	12 years	Enterprise, Private	Large
P21 [†]	46	female	High School	Certification	Staff Network Eng.	USA	20 years	Private	Large
P22 [†]	36	male	Bachelors	Physics / IT	Senior Engineer	Singapore	10 years	Private	Small-Medium
P23 [†]	42	male	Masters	Math	Technician	Armenia	1–2 years	Private	Small
P24 [†]	41	female	Bachelors	Eng.	DNS Operator	Mexico	10 years	Enterprise	Large