

Post-Quantum DNSSEC: A Systematic Review of Transport Constraints, Security Vulnerabilities, and Migration Pathways

Maksim Iavich ^{1,2,*} , Audrius Lopata ³  and Nino Bagalishvili ⁴

¹ Department of Computer Science, Caucasus University, Tbilisi 0102, Georgia

² Department of Cybersecurity and Cryptology, Al-Farabi Kazakh National University, Almaty 050040, Kazakhstan

³ Faculty of Informatics, Kaunas University of Technology, 44249 Kaunas, Lithuania; audrius.lopata@ktu.lt

⁴ Scientific Cyber Security Association, Tbilisi 0102, Georgia; nbagalishvili@scsa.ge

* Correspondence: miavich@cu.edu.ge

Abstract

The Domain Name System Security Extensions (DNSSEC) provide the cryptographic foundation for DNS data integrity and origin authentication. Cryptographically relevant quantum computers (CRQCs) threaten this foundation, since RSA and ECDSA, the signature schemes underlying DNSSEC, are susceptible to polynomial time attacks via Shor's algorithm. NIST-standardized post-quantum cryptography (PQC) signatures structurally exceed the 1232-byte UDP payload limit of DNS, and no existing proposal simultaneously resolves all transport, security, and operational constraints. Following PRISMA 2020 guidelines, this paper systematically reviews 27 peer-reviewed works published between 2020 and 2025, providing the first unified analytical framework for post-quantum DNSSEC research across five dimensions: transport constraints, cryptographic agility, denial-of-service resilience, operational deployment, and standardization readiness. Three findings emerge. First, every NIST-standardized PQC scheme is structurally incompatible with UDP transport in non-minimal DNSSEC responses. Second, analytical modeling predicts that the DNSKEY-RRSIG validation product scales quadratically in the number of coexisting signature algorithms, and that the interaction of KeyTrap with PQC validation overhead compounds denial-of-service severity super-linearly. This prediction is formalized as a testable hypothesis pending experimental characterization, while the algorithm agility mechanism enabling PQC deployment simultaneously enables downgrade attacks. Third, the architecture of signature-based DNSSEC lacks retroactive protection against key compromise, which is the signature-domain equivalent of forward secrecy. No security proof exists for hybrid dual-signature constructions owing to IND-CMP composition problems. We identify five previously unaddressed research gaps and present a prioritized three-phase roadmap through 2029. Post-quantum DNSSEC constitutes a protocol-redesign problem, not an algorithm-substitution problem.



Academic Editor: Raymond Lee

Received: 19 June 2026

Revised: 11 July 2026

Accepted: 13 July 2026

Published: 15 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: DNSSEC; DNSSEC trust anchor; forward secrecy; ML-DSA; Mosca's theorem; post-quantum cryptography; SLH-DSA; systematic review; UDP fragmentation

MSC: 94A60; 68P25

1. Introduction

The advent of quantum computing, marked by Shor's algorithm [1], endangers the secure internet protocols that rely on public key cryptography. The Domain Name System

(DNS), defined in RFC 1034 and RFC 1035 in 1987 [2,3], serves as a fundamental component of the global internet infrastructure, converting human-readable domain names into machine-readable network addresses and supporting electronic mail, anti-spam systems, the Session Initiation Protocol, and enterprise directory services. DNS functions as a globally distributed hierarchical database, handling around one trillion queries daily among recursive resolvers, authoritative name servers, and stub resolvers, hence ranking among the most extensively utilized protocols in the Internet Protocol suite. Notwithstanding its essential role, DNS was developed without any authentication mechanisms, making it inherently susceptible to cache poisoning, spoofing, and man-in-the-middle attacks [4,5].

The Kaminsky attack of 2008 illustrated the worldwide ramifications of this vulnerability, as an attacker may seize arbitrary domain names by abusing the absence of response authentication in the foundational protocol [5]. Liu et al. [6] demonstrate that more cyberattacks involve the resolution of malicious domain names, confirming that the protocol is a prime attack vector and its integrity is crucial for the security of virtually all application-layer services. To address these vulnerabilities, a public key cryptographic layer was added to the DNS Security Extensions (DNSSEC) that provides data origin authentication and data integrity protection via a hierarchical chain of trust rooted at the DNS root zone [4]. The DNSSEC paradigm was first introduced in 1997 and was substantially revised in RFC 4033–4035 in 2005. It requires that each zone sign its resource records with a Zone Signing Key (ZSK) whose authenticity is guaranteed by a Key Signing Key (KSK) held by the parent zone, thereby creating a verifiable delegation chain from the root to any leaf record. [7]. Presently, DNSSEC employs the RSA and ECDSA signature algorithms, whose security relies on the discrete logarithm problem, solvable in polynomial time by Shor’s quantum method, as well as on the computational complexity of integer factorization. An opponent could acquire private signing keys from publicly available DNSKEY records by utilizing a Cryptographically Relevant Quantum Computer (CRQC). This would allow forging any DNSSEC signature, breaking DNSSEC signature verification, making cache poisoning attacks indistinguishable from legitimate answers, and invalidating the protocol’s main security guarantee across the entire global DNS hierarchy. Mallick et al. [8] also consider the “harvest now, decrypt later” (HNDL) threat where passive adversaries collect signed responses containing DNSKEY and RRSIG records as well as authorized DNS traffic for later cryptanalysis when CRQC becomes available. This means cryptography exposure is an imminent risk rather than a future one.

The theory of Mosca [9] states that if $t_{\text{protect}} + t_{\text{migrate}}$ is greater than the quantum threshold, cryptographic compromise occurs before the end of migration, thus quantifying the urgency of the migration process. Section 2.4 shows that this condition is met: Nal and Niemiec [9] estimate t_{protect} at approximately 20 years for critical DNS infrastructure, t_{migrate} at approximately 7 years based on the empirical history of DNSSEC algorithm transitions [10,11], and t_{quantum} at approximately 15 years according to expert consensus, placing the DNS ecosystem within the risk zone of the theorem.

The inaugural post-quantum cryptography (PQC) schemes approved for production deployment were ML-KEM, ML-DSA, and SLH-DSA, standardized in NIST’s 2024 release of FIPS 203, 204, and 205. The structural incongruity arising from the incorporation of these algorithms into DNSSEC is unparalleled in prior DNSSEC algorithm transitions. Prior to the addition of zone-specific resource records, NIST-standardized post-quantum cryptography (PQC) signature schemes produce outputs that considerably exceed the 1232-byte UDP payload limit of DNS. ML-DSA-44 signatures occupy 2420 bytes, SLH-DSA-128 signatures extend to 7856 bytes, and the associated public keys contribute further overhead to DNSKEY records, elevating complete DNSSEC responses significantly beyond the fragmentation threshold. [1]. The size discrepancy stems from the mathematical prin-

ciples underlying lattice-based and hash-based signature methods, necessitating larger algebraic structures to withstand quantum assaults. This issue is structural rather than a mere parameter-tuning problem and cannot be resolved solely through compression or encoding optimizations. Consequently, any deployment approach must either eliminate the traditional UDP transport layer, incorporate fragmentation along with the associated denial-of-service and reliability issues, or restructure the DNSSEC protocol itself, a constraint that is further analyzed in Section 2.4.

Other studies have considered some aspects of this issue, but no unified strategy exists. Müller et al. [12] provide an important empirical characterization of size incompatibility by looking at transport restrictions and finding the 1232-byte UDP limit to be the main deployment barrier. However, they do not consider the implications of denial-of-service attacks, or the lack of forward secrecy in signature-based architectures. Heftrig et al. [13–15] show that the signaling infrastructure supporting cryptographic changes in the context of DNSSEC algorithm agility mechanisms also allows adversary reversion to less secure algorithms, but do not account for how this attack surface grows with the simultaneous deployment of conventional and post-quantum algorithms. Rawat and Jhanwar [16] propose SL-DNSSEC, a signatureless design which reduces the Bandwidth Amplification Factor. Nonetheless, neither its security attributes within the entire DNSSEC chain-of-trust framework nor its ramifications for non-repudiation are expressly assessed. Collectively, these works tackle distinct subproblems; however, no previous research offers a cohesive analytical framework that concurrently delineates transport constraints, denial-of-service threats, cryptographic agility weaknesses, and forward secrecy deficiencies as an integrated entity, nor elucidates their interrelations. This paper addresses the existing gap.

This paper offers a systematic review in accordance with PRISMA 2020 principles of 27 peer-reviewed articles published from 2020 to 2025. The key innovation has three previously unstudied contributions: (i) a single analytic model that captures the post-quantum DNSSEC transition across five interrelated dimensions simultaneously, revealing interaction effects missed in domain-specific studies; (ii) a novel quantitative analysis of transport overhead and fragmentation that integrates published parameter data into a standardized cross-scheme comparison (Section 5.1); and (iii) a systematic cross-proposal evaluation using a nine-criterion framework and weighted scoring system that enables structured comparative analysis (Sections 5.9 and 5.10). The outstanding contributions are as follows.

A comprehensive analytical framework is created, characterizing the migration from classical to post-quantum DNSSEC along five axes, namely transport limitations, cryptographic agility and downgrade attacks, robustness to denial-of-service attacks, operational deployment, and standardization readiness, revealing the interdependencies between these axes that prior work has considered in isolation.

A systematic review reveals that none of the alternatives meets all the operational criteria at once. The six mitigation strategies evaluated in this work, namely TCP fallback, EDNS(0) extension, request-based fragmentation, QNAME-based fragmentation, SL-DNSSEC and MTL Mode, address some limitations but also present various residual attack surfaces. None of them achieves full compliance across the five analysis dimensions.

We formally identify five previously unexamined research deficiencies, including the structural lack of forward secrecy in signature-based DNSSEC (understood here as retroactive protection of authentication against future key compromise, the HNDL threat, rather than TLS-style ephemeral session-key secrecy), the projected amplification of denial-of-service severity arising from the interaction of KeyTrap (CVE-2023-50387) with post-quantum validation overhead, an analytically derived projection, formalized as testable hypothesis H2 in Section 6.6, whose super-linear form follows from protocol arithmetic

and whose magnitude is illustrated by the single-testbed $2.8\times$ per-operation estimate for ML-DSA-44 [16]) and the lack of a security proof for hybrid dual-signature frameworks due to IND-CMP composition difficulties between incompatible security parameter domains.

A research agenda is established with a focus that provides practical guidance for protocol designers, implementers and standardization bodies. The success of the project is based on the quantitative study of the size of the signature, the latency of the resolution and the Bandwidth Amplification Factor, deployed in three phases till 2029. The remainder of this document is organized as follows: Section 2 provides background information on DNSSEC architecture, operational processes, deployment status, and the quantum threat. The PRISMA-guided approach is outlined in Section 3. Section 4 analyzes recognized issues within conventional and quantum threat models, Section 5 reviews post-quantum suggestions, Section 6 identifies unresolved challenges, and Section 7 presents findings along with a focused research agenda.

2. Background

2.1. DNSSEC Security Model and Cryptographic Architecture

By generating digital signatures over resource record sets (RRsets), which are collections of all resource records in a zone that share the same owner's name, class, and type, and publishing those signatures as RRSIG records with the signed data, DNSSEC expands DNS [7]. Because of this design, DNS data is self-authenticating without depending on the security characteristics of intermediary network equipment, and a validating resolver may confirm the integrity and origin of any signed response regardless of the transport channel over which it was received.

According to RFC 4034 [7], the protocol uses two functionally different key types for each zone. All data RRsets within the zone are signed by the Zone Signing Key (ZSK, DNSKEY Flags = 256); only the DNSKEY RRset is signed by the Key Signing Key (KSK, DNSKEY Flags = 257, Secure Entry Point bit set), which authenticates the ZSK and binds it to the zone identity. Müller et al. established a specific operational function for this two-tier key design [17]. In contrast to the KSK, which rotates infrequently because any KSK change necessitates a corresponding Delegation Signer (DS) record update in the parent zone—an operation that crosses administrative boundaries and cannot be carried out unilaterally—the ZSK can rotate frequently to limit the window of potential compromise because it signs individual resource records and is thus exposed more broadly.

RFC 4034 [7] specifies the inter-zone delegation chain and signature process as follows. Each RRset is signed by the zone operator using the ZSK private key, creating an RRSIG record that includes the cryptographic signature, the covered record type, the algorithm identification, the original TTL, and the timestamps of the signature validity period. The ZSK public key is bound to the zone via an RRSIG (DNSKEY) record that is created when the KSK private key independently signs the DNSKEY RRset. The DS record, which the parent zone publishes as a cryptographic hash of the child zone's KSK public key, establishes the inter-zone delegation link; RFC 4034 [7] specifies SHA-256 (DNSSEC algorithm Algorithm 8) or SHA-384 (Algorithm number 14) for this digest computation. In the validation process, a recursive resolver obtains the RRset and its matching RRSIG record, verifies the signature using the ZSK public key, and then verifies the validity of the KSK by comparing its hash with the DS record that the parent zone has published [7]. The root KSK, which is pre-configured in validating resolvers as the trust anchor, is where this verification process ends after moving up the DNS hierarchy from leaf zones through TLD operators to the DNS root zone. The chain's only unwavering source of confidence is the trust anchor. The integrity of the entire global DNSSEC hierarchy would be compromised by an adversary with the ability to forge root zone DNSKEY records. As a result, the security

of root KSK storage and the integrity of the Root KSK rollover procedure become key infrastructure concerns.

A complete successful chain-of-trust validation sets the Authenticated Data (AD) flag to 1 in the resolver response; any verification failure at any level of the hierarchy produces a SERVFAIL response, withholding the data from the client entirely.

The scope of the security guarantees that DNSSEC provides precise characterization, as misidentification of this scope is a source of deployment errors in practice. DNSSEC provides data integrity and source authentication for DNS resource records exclusively; it does not encrypt DNS communications, does not authenticate the query itself, and does not protect the channel between the recursive resolver and the stub resolver. As Rose et al. [18] specify in NIST SP 800-81r3, DNS over TLS (DoT, RFC 7858), DNS over HTTPS (DoH, RFC 8484), and DNS over QUIC (DoQ, RFC 9250) provide channel confidentiality and transport-layer authentication for the resolver-to-client path but provide no authentication of the DNS data itself. These two security properties, data authentication and channel encryption, are orthogonal by design and must be deployed in combination to provide comprehensive protection. DNSSEC without encrypted transport leaves the resolver-to-client channel open to observation and injection, while encrypted transport without DNSSEC provides no assurance that the DNS data retrieved by the resolver is authentic. The interaction between these complementary mechanisms, and their joint behavior under post-quantum cryptographic schemes, is analyzed in Section 6.5.2. The operational procedures through which DNSSEC's security model is exercised at runtime zone signing, chain-of-trust traversal, and authenticated denial of existence are described in Section 2.2.

2.2. Operational Mechanism Zone Signing, Validation, and Denial of Existence

The zone signing and validation procedure specified in RFC 4035 [7] is initiated by the zone operator, who generates ZSK and KSK pairs and publishes both corresponding public keys as DNSKEY records in the zone. As stated in Section 2.1, the private KSK only signs the DNSKEY RRset; the private ZSK is used to sign every RRset in the zone, generating one RRSIG record per RRset. The operator creates the inter-zone delegation link that links the zone to the global chain of trust by computing a cryptographic hash of the KSK public key and submitting it to the parent zone for publishing as a DS record. The security of this delegation mechanism stems from the fact that the DS record is signed by the parent zone's own ZSK and not by the child: publishing the KSK within the child zone only would be self-attestation, whereas the parent-signed DS record binds the child's key identity to a key that the resolver has already authenticated one level higher in the hierarchy [19,20]. Note that the DS record is not the key itself, but a fixed length cryptographic digest of the DNSKEY. This keeps the parent zone small but is still fully verifiable because the resolver will fetch the child's DNSKEY on its own and recompute the digest to compare. [19]. The trust anchor that is pre-configured in validating resolvers is where this submission procedure ends after moving up the DNS hierarchy to the root zone. In modern deployments, this submission process can be automated using the CDS/CDNSKEY mechanism, whereby the parent polls the child zone for signaled key changes instead of relying on manual submission through registrar interfaces [21].

In order to retrieve RRsets and their matching RRSIG entries at every level of the hierarchy, a recursive resolver iteratively queries the root, TLD, and authoritative name-servers for the searched domain. Repeating upward until the trust anchor is reached, the resolver checks the RRSIG record against the zone's ZSK public key at each stage and validates the validity of that ZSK using the DS record released by the parent zone [7]. The validation model described in RFC 4035 is operationally realized in production deployments rather than just in controlled environments, as demonstrated by Nosyk et al.'s [22]

extensive empirical confirmation that DNSSEC-validating resolvers consistently carry out this chain-of-trust traversal across the global internet. A successful chain-of-trust traversal, as established in Section 2.1, sets the AD flag to 1 in the resolver response; any verification failure at any level produces SERVFAIL.

In order to prevent an adversary from suppressing a valid NXDOMAIN response and substituting a forged positive answer, DNS must also be able to authenticate the non-existence of a queried name. This is a structural gap left by the validation model previously described. DNSSEC uses authenticated denial-of-existence records to address this. A resolver can confirm that a query name falls within the relevant interval and does not exist as a signed record by using an NSEC record, which lists the next secure name in canonical lexicographic order within the zone [23]. The cryptographic authentication of non-existence stems from two constraints of the DNSSEC signing model. First, the private ZSK is deliberately kept away from the query path: signatures are generated offline at zone-signing time, so a nameserver cannot produce a freshly signed statement in response to an arbitrary query. Second, the set of nonexistent names is unbounded, so no finite pre-signing process can cover them individually. NSEC resolves both constraints by converting the infinite negative space into a finite set of signable assertions about intervals: the signer orders all existing owner names into the canonical sequence defined in RFC 4034 [20] and generates, for each consecutive pair, an NSEC record asserting that no name exists between them, each accompanied by its own RRSIG produced with the zone's ZSK [20,24]. The chain is circular: the NSEC record of the lexicographically last name points back to the zone apex, so every conceivable query name falls within exactly one signed interval [24]. Upon receiving a negative response, the resolver verifies the RRSIG over the returned NSEC record against the zone's DNSKEY, itself authenticated through the chain of trust described above, and confirms that the queried name lies lexicographically between the NSEC owner name and its Next Domain Name field; for a complete NXDOMAIN proof, RFC 4035 additionally requires an NSEC record demonstrating that no wildcard could have synthesized the answer [21]. A validated signature over the covering interval constitutes proof that the queried name cannot exist in the signed zone, since its presence would have partitioned the name space differently and produced a different signed NSEC chain [21,24]. The security argument is precisely that of the signature scheme: an adversary may replay legitimately signed NSEC records but cannot mint an interval of its own choosing without the private ZSK, and suppressing the NSEC records outright yields a validation failure (SERVFAIL) rather than an acceptable forged negative [21]. The NSEC type bitmap extends the same guarantee to a second class of negative answers, proving the non-existence of a specific record type at a name that does exist and thereby covering NODATA as well as NXDOMAIN responses [20]. While this mechanism provides cryptographically authenticated non-existence proof, according to Bau and Mitchell [25], it also permits zone walking, whereby an adversary can list every item in a signed zone by sending iterative queries that move through the NSEC chain from any starting point, gradually revealing all owner names, record types, and record counts inside the zone.

In order to address this zone enumeration vulnerability, NSEC3, as defined in RFC 5155 [26], was introduced. It replaces plaintext owner names in denial-of-existence records with iterated SHA-1 hash values that are calculated over the owner name concatenated with a zone-specific salt; the NSEC3PARAM record publishes both the iteration count and the salt value. This design was intended to make zone walking computationally infeasible by requiring an adversary to invert the hash function rather than simply traversing a lexicographically ordered chain. Bau and Mitchell [25], however, demonstrate that the public accessibility of the salt which is a necessary property of the NSEC3 design, since validating resolvers must be able to replicate the hash computation, means that

NSEC3 remains vulnerable to offline dictionary attacks and an adversary who retrieves the salt can precompute hash values for candidate domain names and recover original owner names within days on modern hardware [25]. The NSEC3 opt-out mechanism introduces an independent vulnerability, one that originates in an economic trade-off rather than a cryptographic weakness. In large delegation-centric zones such as TLDs, where the overwhelming majority of child delegations are unsigned, generating and signing an NSEC3 record for every unsigned delegation would inflate the zone by millions of records and proportionally increase signing time, zone size, and memory footprint; RFC 5155 therefore permits an NSEC3 record to carry the opt-out flag, which changes the semantics of its interval from “no name exists between these two hashes” to the strictly weaker assertion “no signed name exists between these two hashes” [26]. Unsigned delegations may thus legitimately appear inside an opt-out span without being represented in the NSEC3 chain at all. The security consequence follows directly from this weakened semantics. When a resolver receives a referral for a name whose hash falls between two consecutive NSEC3 hash values in an opt-out span, the covering NSEC3 record validates correctly that its RRSIG is genuine yet by design it makes no assertion about the queried name, so the resolver cannot distinguish a legitimate unsigned delegation from one fabricated by an adversary. Under the RFC 4035 validation-state model the response is classified as insecure rather than bogus: it is accepted and the resolution proceeds without DNSSEC protection, rather than being rejected with SERVFAIL [25,26]. An on-path or off-path adversary in the A_OFF position can exploit exactly this gap by injecting a forged referral redirecting the victim toward attacker-controlled nameservers for any name whose hash lands in an opt-out interval. Bau and Mitchell [25] demonstrate the attack in a controlled setting using a formal model of the protocol, and its defining characteristic is that it requires no cryptanalytic effort whatsoever: no hash inversion, no key recovery, and no signature forgery are involved, because every cryptographic artifact the resolver checks remains genuine. The vulnerability is thus a property of the opt-out design itself rather than of any implementation, and it persists for as long as a zone elects opt-out, trading enumeration-resistant coverage of unsigned space for operational scalability. In Sections 4.3 and 6.7, respectively, the security implications of these techniques and their compounded interaction with post-quantum RRSIG overhead and the SHA-1 weakening effect of Grover’s algorithm are examined.

2.3. Current Deployment Status

Adoption of DNSSEC is still far below its potential for security. Only 1% of .com, .net, and .org names are correctly signed, according to Chung et al. [27], who attribute this to inadequate registrar support and the operational difficulty of key lifetime management. Lian et al. [28] discovered that DNSSEC-signed domains can result in collateral resolution failures in valid domains due to fragmentation and validation issues after measuring the practical impact across more than 500,000 clients. Reference [29] demonstrates that while ECDSA reduces response sizes relative to RSA, at the same time it introduces substantially higher validation overhead at resolvers, since ECC signature verification is considerably more expensive than RSA verification, shifting computational load toward the network edge. DNSSEC increases response-to-query size ratios, which increases the potential for DDoS attacks [30]. ZSK and KSK rotation, RRSIG validity period management, and multi-signer zone cooperation are identified [17] as major administrative burden sources. A minimum disruptive migration path is essential for any successful plan because of the combination of these current deployment constraints with PQC performance requirements, as examined in Sections 4 and 5.

2.4. The Quantum Threat

The signature algorithms deployed in DNSSEC RSA and ECDSA derive their security from mathematical problems that are computationally intractable for classical hardware but efficiently solvable by quantum algorithms. RSA security is grounded in the presumed intractability of integer factorization; ECDSA security rests on the elliptic curve discrete logarithm problem. The General Number Field Sieve (GNFS), the best-known classical algorithm for integer factorization, runs in sub-exponential time with heuristic complexity $L_n[1/3, c] = \exp\left((c + o(1))(\ln n)^{1/3}(\ln \ln n)^{2/3}\right)$ where n is the modulus and $c = (64/9)^{1/3} \approx 1.923$ [31]. The sub-exponential character arises from the algorithm's structure: GNFS reduces factorization to the collection of multiplicatively smooth relations over a factor base followed by sparse linear algebra, and the $L[1/3]$ exponent expresses the optimal balance between the rarity of smooth values and the cost of sieving for them [31]. This complexity class lies strictly between polynomial and exponential: the running time grows faster than any polynomial in the bit length of n , yet the exponent grows only with the cube root of $\ln n$ rather than linearly. The practical consequence is that security grows sublinearly in key size NIST rates RSA-1024 at approximately 80 bits, RSA-2048 at approximately 112 bits, and RSA-3072 at approximately 128 bits of classical security [32], which is precisely what forced the successive migrations from 512-bit to 1024-bit to 2048-bit moduli, and why RSA keys must grow disproportionately to purchase each additional increment of security margin. Empirically, the largest publicly factored RSA modulus to date, RSA-250 (829 bits), required roughly 2700 CPU core-years of GNFS computation [33]; the sub-exponential formula places the work factor for a 2048-bit modulus many orders of magnitude beyond that record, far outside any realistic computational capability. The elliptic curve discrete logarithm problem underlying ECDSA occupies a qualitatively different position: elliptic curve groups expose none of the algebraic structures and no analog of smoothness over a factor base that index-calculus methods such as GNFS exploit, and no sub-exponential classical algorithm is known for the problem [34]. The best available attacks are generic: Pollard's rho method runs in fully exponential time $O(\sqrt{q})$, amounting to roughly 2^{128} group operations for the P-256 curve, and within the generic-group model no generic algorithm can perform asymptotically better [34]. This asymmetry is why a 256-bit curve delivers approximately 128 bits of classical security, which is a stronger guarantee than RSA-2048 provides, at one-eighth the key size [32,35]. Both problems are therefore computationally infeasible on classical hardware at deployed key lengths within any practical timeframe. Shor's algorithm solves both problems in polynomial time on a quantum computer [1], representing a superpolynomial speedup that renders RSA-2048, ECDSA, DSA, and Diffie–Hellman key exchange cryptographically broken on a sufficiently powerful Cryptographically Relevant Quantum Computer (CRQC). The algorithm recovers a non-trivial factor of the modulus by identifying the period of a modular exponentiation function via the Quantum Fourier Transform and, with high probability, extracts the private key from the publicly distributed public key [1].

For symmetric cryptography, Grover's algorithm provides a quadratic speedup for unstructured search, effectively halving the security strength of any symmetric key length [36]. In practice, this is mitigated by doubling key lengths: AES-256 retains adequate security in the post-quantum context while AES-128 does not [36]. Hash functions used in DNSSEC including SHA-256 in DS records and SHA-1 in NSEC3 are similarly weakened, with SHA-1 reduced to an effective security of approximately 80 bits, a consequence examined in the context of NSEC3 in Section 6.7. Quantum-enabled recovery of a ZSK private key from its publicly distributed DNSKEY record would permit an adversary to forge valid RRSIG records over arbitrary DNS data, constructing a chain of trust indistinguishable from a legitimate response to any validating resolver. Nal and Niemiec [9] characterize the

harvest-now–decrypt-later (HNDL) attack vector in the DNSSEC context because DNSSEC provides no forward secrecy: a single ZSK signs all RRSIG records throughout its validity window and an adversary archiving DNS traffic today gains retrospective access to all signatures produced during the classical-to-PQC migration period upon CRQC availability, enabling forgery of the entire archived corpus without any active network capability at the time of forgery.

HNDL and hybrid signatures: A frequently proposed transitional measure is the hybrid dual-signature scheme, combining a classical algorithm such as RSA-2048 with a PQC algorithm such as ML-DSA-44, to provide security against both classical and quantum adversaries simultaneously during the migration period. The mechanism operates through dual signing. The zone publishes DNSKEY records for both algorithms, the classical and the post-quantum key each carrying its own algorithm identifier, and every RRset is accompanied by two parallel RRSIG records, one produced with each private key; RFC 4035 in fact requires that a zone be signed with at least one key of every algorithm present in its apex DNSKEY RRset, so the dual-signature configuration is structurally native to the protocol rather than an extension of it [37]. Under a conjunctive (AND) validation policy, in which a resolver accepts a response only if both signatures verify, a successful forgery requires breaking both algorithms simultaneously: a classical adversary can break neither component, while a quantum adversary can forge the RSA signature via Shor’s algorithm but not the ML-DSA signature, whose underlying Module-LWE problem admits no known efficient quantum algorithm [38]. This is not merely a heuristic layering argument. For signature combiners of this form, the composition is provably unforgeable whenever at least one component remains unforgeable, and the guarantee has been established for precisely the adversary that matters during migration: a two-stage attacker that is classical during the signing epoch and acquires quantum capability only afterwards [38]. The composed scheme therefore remains secure throughout a transition period in which confidence in newly standardized PQC schemes is still maturing while classical schemes await quantum compromise, at the operational cost of carrying both public keys in the DNSKEY RRset and a full post-quantum RRSIG alongside the classical one in every response, so that the hybrid configuration inherits the transport-size regime of its larger component (Sections 4 and 6.1). Crucially, however, the combiner guarantee is conditional on two properties that DNSSEC as deployed does not provide. The first is conjunctive validation itself: RFC 6840 instructs validators to accept a response whenever any single supported algorithm yields a valid path [39], so the deployed policy is disjunctive, and under a disjunctive policy the security of the pair collapses to that of the weaker algorithm. The second is non-separability: the combiner analyses presuppose that the component signatures cannot be presented in isolation, whereas the two RRSIGs are independent records over the same RRset, and an on-path adversary can strip either one without disturbing the other the mix-and-match manipulation analyzed in Section 6.2.3 [38]. However, this approach does not resolve the HNDL threat when the classical component is targeted. If an adversary archives a DNSSEC response signed with both algorithms and subsequently recovers the RSA private key via Shor’s algorithm [1], the classical RRSIG becomes forgeable regardless of the presence of the PQC RRSIG. A counterfeit classical RRSIG may suffice to achieve validation under present resolver behavior, wherein any syntactically valid RRSIG can fulfill chain-of-trust authentication without necessitating an assault on the PQC component. Consequently, perfect quantum resistance requires the whole elimination of classical signature elements; hybrid methodologies reduce but do not entirely eliminate HNDL danger [9].

Achieving forward secrecy within the signature-based DNSSEC scheme necessitates a major architectural change. In signature-based DNSSEC, the absence of forward secrecy is

not a configuration restriction but an inherent architectural characteristic of the protocol. Forward secrecy in session-based protocols such as TLS 1.3 is achieved using ephemeral Diffie–Hellman key exchange, which ensures that the compromise of a long-term private key does not expose previous sessions because the encryption keys for each session are generated separately and subsequently discarded. DNSSEC does not use session keys; instead, a single ZSK generates all RRSIG records during the course of its rotation period, which are then published in plaintext on the worldwide internet for the duration of their validity. Because each RRSIG generated during any previous rotation period remains permanently forgeable once the accompanying ZSK private key is retrieved by a CRQC, even vigorous ZSK rotation minimizes the exposure window but cannot completely eradicate it. Therefore, the signature-based DNSSEC model cannot achieve forward secrecy without a fundamental redesign of its authentication architecture. For example, a KEM-based construction like SL-DNSSEC [16] can approximate forward secrecy at the session level by establishing a session-specific shared secret for each DNS query-response exchange. There is no formal security demonstration for the SL-DNSSEC design within the complete DNSSEC chain-of-trust model, and no proposal in the examined literature formally characterizes this architectural constraint as an independent research target [16]. Mosca’s theorem provides a mathematical measure of the urgency of migration [9].

$$t_{\text{protect}} + t_{\text{migrate}} > t_{\text{quantum}} \implies \text{compromise before migration completes}, \quad (1)$$

2.5. Formal Threat Model

This section delineates the adversarial assumptions that underpin the security analysis articulated in Sections 4–6. Six adversary classes are delineated, each characterized by capability, network position, and temporal relation to cryptographically pertinent quantum processing. The six adversary classes are not an unstructured list but the cross-product of two organizing axes: a computational axis, which fixes the cryptographic assumptions an adversary is permitted to defeat, and a temporal axis, which fixes the adversary’s position relative to the arrival of a Cryptographically Relevant Quantum Computer (CRQC). Separating these axes is what allows a single framework to govern the analysis of both the currently deployed protocol (Sections 4.3–4.6) and its post-quantum successor (Sections 6.1–6.3), because the same network capabilities recur under both computational regimes while the cryptographic ground shifts beneath them.

On the computational axis, the classical adversary is bounded by the sub-exponential algorithms characterized in Section 2.4. It cannot factor an RSA-2048 modulus, for which the General Number Field Sieve places the work factor many orders of magnitude beyond the largest published factorization, nor solve the elliptic-curve discrete logarithm problem for ECDSA P-256, for which only fully exponential generic attacks are known. Its entire repertoire is therefore protocol- and implementation-directed rather than cryptanalytic: cache poisoning and spoofing against the unauthenticated base protocol [4], side-channel-assisted defeat of port randomization of the kind demonstrated by the 2020 SADDNS attack [5], man-in-the-middle interception on unencrypted paths, and abuse of DNSSEC’s own algorithm agility signaling to induce resolvers to bypass validation [13–15]. Crucially, none of these capabilities is neutralized by migrating to a quantum-resistant signature, because none of them attacks the signature primitive in the first place; they attack the surrounding protocol. The quantum adversary occupies the complementary position. Endowed with a CRQC able to run Shor’s algorithm [1] against deployed key sizes, it recovers a Zone Signing Key private key directly from the publicly published DNSKEY record in polynomial time and can thereafter mint RRSIG records that are, to any validating resolver, indistinguishable from authentic signatures collapsing the chain of trust not at

one zone but potentially across the whole hierarchy. The two adversaries are thus defined by disjoint capabilities rather than by a difference in degree: a countermeasure effective against one is generically silent against the other. This is the structural reason the review treats transport hardening, downgrade resistance, denial-of-service resilience, and forward secrecy as separate dimensions rather than as facets of a single “PQC upgrade.”

On the temporal axis, each adversary is located relative to the quantum horizon, estimated by the expert consensus reported in Section 2.4 at approximately 10–15 years [9]. Two adversaries are defined precisely by their relation to this horizon. The pure quantum adversary is assumed active only after it, whereas the harvest-now–decrypt-later adversary straddles it: classical during the signing epoch, when it can do no more than passively archive signed responses, and quantum afterwards, when it retroactively recovers the archived keys and forges the entire captured corpus without any network presence at the time of forgery [8,9]. This two-stage attacker of classical now, quantum later is not an informal device; it is exactly the adversary for which the hybrid dual-signature combiner guarantee has been formally proved [38], and it is the adversary against which Mosca’s inequality is evaluated, since the theorem’s condition quantifies whether the archived material remains sensitive for longer than the time until the CRQC arrives [9]. Because signature-based DNSSEC provides no forward secrecy, the exposure created by this adversary is retroactive by construction: every RRSIG published in plaintext during the migration window is a standing liability the moment the corresponding key becomes recoverable, which is why the analysis frames cryptographic exposure as already accruing rather than as a purely future event [8,9]. The remaining classes of downgrade, on-path, off-path, and resolver-manipulation adversaries are all instantiations of the classical computational tier distinguished by network position and by the specific protocol weakness they exercise, and each is carried forward unchanged into the post-quantum setting, where several of them (notably the off-path fragment-tag-modification variant [15]) become more, not less, potent because PQC response sizes make fragmentation unavoidable.

Traditional antagonist (A_C): A_C functions on classical hardware and is constrained by sub-exponential algorithms. A_C is incapable of resolving the integer factorization problem for RSA-2048 or the elliptic curve discrete logarithm problem for ECDSA P-256 during the protocol’s duration. A_C is capable of executing cache poisoning, DNS hijacking, man-in-the-middle interception in the absence of encryption, and exploiting implementation flaws, including those associated with the attack. A_C is regarded as the principal adversary for the analysis of the existing DNSSEC deployment (Sections 4.3–4.6).

A_Q is a quantum opponent. A_Q has a Cryptographically Relevant Quantum Computer (CRQC) that can execute Shor’s algorithm [1] on the RSA-2048 and ECDSA P-256 key sizes used in functional DNSSEC. In polynomial time, A_Q can extract a Zone Signing Key (ZSK) private key from its publicly accessible DNSKEY record, making it easier to create RRSIG records over any DNS data that are identical to genuine signatures. According to expert consensus described in Section 2.4 [9], A_Q is assumed to be accessible at the quantum horizon $t_{quantum}$, which is estimated to be 10–15 years.

Harvest-now–decrypt-later adversary (A_{HNDL}): A_{HNDL} is a hybrid adversary that has postponed quantum capabilities and is presently classical. While maintaining the conventional ZSK public key and previously generated signatures, A_{HNDL} archives DNS traffic, including RRSIG, DNSKEY, and DS records. Without requiring any active network capacity during the forgery process, A_{HNDL} uses A_Q capabilities retroactively to retrieve archived private keys and build the complete stored corpus of signatures when CRQC is available [8]. The lack of forward secrecy in DNSSEC (Section 6.2.4) indicates that once DNS traffic is traceable, A_{HNDL} ’s threat will be immediate, making cryptographic exposure an urgent problem rather than a possible one.

Attacker downgrade (A_{DG}): A_{DG} is a traditional attacker skilled at altering or injecting DNS answers as they are being transmitted. A_{DG} allows conforming resolver implementations to get around signature verification by substituting unknown algorithm identities in RRSIG entries via the DNSSEC algorithm agility mechanism (Section 6.2.2). According to Heftrig et al. [15], the off-path form of A_{DG} uses the structural fragmentation that PQC response sizes naturally cause to change the method tag byte in a fragmented RRSIG record without requiring a privileged network position (Section 6.1.1). A_{DG} does not require crypt-analytic knowledge, but when PQC is used, its attack surface grows since more algorithm IDs are added that unpatched resolvers do not recognize.

Adversary on-path (A_{ON}): A_{ON} controls a network node that is situated between a recursive resolver and an authoritative nameserver or between a stub resolver and a recursive resolver. DNS communications can be seen, modified, injected, and suppressed by A_{ON} . A_{ON} can perform man-in-the-middle attacks on unencrypted DNS interactions, selectively remove RRSIG records from dual-signature answers (Section 6.2.3), and control IP fragmentation to change specific fragments [15]. A_{ON} can take advantage of flaws in the protocol and implementation, but it cannot compromise the cryptographic primitives used by A_C .

Adversary that is not concurrent (A_{OFF}): Although A_{OFF} is not in charge of a node on the resolution path, it has the ability to introduce bogus DNS answers into the network. The adversarial architecture for both the off-path downgrade variation described in [15] and traditional DNS cache poisoning attacks is represented by A_{OFF} [4,5]. A_{OFF} presents replies that a target resolver recognizes as authoritative via taking use of IP fragmentation injection, port randomization flaws, or transaction ID prediction. Although the DNSSEC chain-of-trust technique aims to address A_{OFF} , the algorithm agility vulnerability allows resolvers to accept answers without signature verification, and so partially reopening A'_{OFF} s attack surface. Resolve manipulation adversary (A_{RM}): A_{RM} manipulates a malevolent authoritative nameserver and takes advantage of resolver-side protocol compliance mandates. A_{RM} is the adversary model for the KeyTrap vulnerability (CVE-2023-50387, Section 6.3.1) [40], wherein a zone filled with an excessive quantity of DNSKEY and RRSIG records forces resolvers to execute an infinite amount of cryptographic operations in reaction to a single query. A_{RM} serves as the adversarial framework for the NSEC3 CPU exhaustion attack (CVE-2023-50868, Section 6.3.2) [41] and the MTL Mode resource depletion attack (Section 6.3.3). The threat severity of A_{RM} escalates with the implementation of PQC, as the per-operation verification cost for ML-DSA-44 is around 2.8 times greater than that of RSA-2048 [16], hence exacerbating the computational load for each DNSKEY-RRSIG combination with a single-testbed estimate; see Section 6.8.

Parameters of examination: The security analyses in Sections 4–6 utilize the following adversary models: A_C and A_{DG} serve as the principal adversaries for Section 4 (classical DNSSEC vulnerabilities); A_Q and A_{HNDL} are the main adversaries for the PQC motivation analysis in Section 2.4 and the forward secrecy analysis in Section 6.2.4; A_{DG} and A_{ON} are the key adversaries for Section 6.2 (downgrade attacks); A_{RM} is the primary adversary for Section 6.3 (denial-of-service resilience). In the analysis of hybrid proposals (Sections 5.5 and 6.2.3), it is presumed that the adversary is A_{ON} or A_{DG} , targeting one component of a dual-signature pair preferentially. Side-channel adversaries, those able to time, power-profile, or otherwise physically observe signature generation [42], lie outside the network-capability model of the adversary classes defined above and are not treated as a general threat class in this review; they become relevant only for the online-signing configurations and signer-hardware choices analyzed in Sections 6.4.1 and 7.3 and tested by hypothesis H5.

3. Methodology

3.1. Research Questions

Three research questions guided the systematic search, screening, and synthesis process. These questions were designed to cover all aspects of the post-quantum DNSSEC transition difficulty, from high-level ecosystem and standards hurdles to low-level transport constraints.

RQ1 (Transport and Protocol limits): In terms of UDP transport limits, message size, TCP fallback behavior, and denial-of-service resilience, what technological difficulties arise when post-quantum cryptographic primitives are included into DNSSEC?

RQ2 (Proposals and Their Effectiveness): What solutions have been put out in the literature to deal with these issues, and how well do they all concurrently handle the noted operational, security, and transportation constraints?

RQ3 (Residual Gaps): What outstanding issues and research gaps exist in the areas of forward secrecy, operational deployment, standards readiness, and cryptographic agility?

Each research question corresponds to one or more analytical sections of the review: RQ1 to Section 4, RQ2 to Section 5, and RQ3 to Section 6.

3.2. Information Sources and Search Strategy

Four main databases, IEEE Xplore, the ACM Digital Library, Google Scholar, and the IACR Cryptology ePrint Archive, were used to conduct a thorough literature search. Additional searches were conducted on the IETF Datatracker and arXiv (cs.CR and cs.NI subsections), focusing on RFC publications, Internet-Drafts, and working group documents pertaining to DNSSEC and post-quantum cryptography. The examination covered the period from January 2020 to December 2025. The commencement date was set in January 2020 due to the fact that Müller et al. [12] conducted the first systematic empirical examination of PQC integration into DNSSEC, thereby establishing the foundational benchmark for all subsequent research. Previous works have addressed issues that have since been significantly refined or surpassed. The primary Boolean search string, applied across all four databases, was as follows:

("DNSSEC" OR "DNS Security Extensions") AND ("post-quantum" OR "quantum-resistant" OR "lattice-based" OR "hash-based" OR "ML-DSA" OR "SLH-DSA" OR "Falcon" OR "CRYSTALS" OR "SPHINCS+" OR "Dilithium" OR "FN-DSA").

Secondary strings were applied to capture the literature on specific challenge dimensions identified in the research questions:

("DNSSEC" AND "fragmentation" AND ("UDP" OR "TCP fallback" OR "EDNS");
 ("DNSSEC" AND ("downgrade" OR "algorithm agility" OR "cipher negotiation");
 ("DNSSEC" AND ("DDoS" OR "amplification" OR "KeyTrap" OR "denial of service");
 ("DNSSEC" AND ("forward secrecy" OR "KEM" OR "MAC" OR "signatureless").

Backward and forward citation chaining were applied to all included works in order to reduce the possibility that pertinent works would be overlooked by keyword search alone. Backward chaining looked at the reference lists of each included paper, while forward chaining used Google Scholar's "Cited by" function to find works published during the review period that cited an included paper. Drafts submitted to the DNSOP and CFRG working groups that mentioned post-quantum cryptography or DNSSEC were manually searched through the IETF Datatracker. A total of 18 more entries that were not found by the initial database searches were found by this additional procedure.

3.3. Inclusion and Exclusion Criteria

Prior to starting the search, inclusion and exclusion criteria were established in advance and implemented uniformly throughout the screening process. Table 1 outlines the whole eligibility structure along with the justification for each requirement.

Table 1. Eligibility criteria.

Criterion	Inclusion	Exclusion	Rationale
Publication type	Peer-reviewed conference papers, journal articles, IETF RFC/standards-track documents	Opinions, editorials, and gray literature without peer review	Ensures a minimum evidentiary standard across the corpus
Primary focus	DNSSEC security, PQC in DNS, or both; original empirical, analytical, or protocol-design contribution relevant to RQ1–RQ3	PQC works addressing exclusively non-DNS contexts (TLS only, PKI only, IPsec)	Maintains direct relevance to all three research questions
Language	English	Non-English	Ensures comprehensive content extraction by the review authors
Publication date	January 2020–December 2025	Pre-2020, except defined foundational exceptions	Captures the post-PQC-standardization research landscape; exceptions defined below
Contribution type	Original empirical measurement, analytical characterization, or protocol proposal	Works reproducing prior results without novel contribution	Ensures each included work contributes independently to the synthesis
Preprints and drafts	arXiv and IACR ePrint preprints with broad community recognition (≥ 50 citations or formally adopted by an IETF working group); IETF Internet-Drafts with active working group engagement	Preprints with no peer-review track record and no community adoption signal	Balances coverage of rapidly emerging work against minimum quality control

Pre-2020 exceptions: Since they set the protocol standards by which post-quantum proposals are judged, foundational RFC specifications such as RFC 4033–4035 [7] and RFC 5155 [26] as well as groundbreaking DNSSEC measurement studies necessary for contextual framing were included regardless of publication date. These works are not included in the 27 papers that make up the main analytical corpus and are specifically mentioned in the reference list.

3.4. Study Selection Process

Study selection followed the four-phase PRISMA 2020 flow illustrated in Figure 1. Screening was performed independently by two reviewers at each phase; disagreements were resolved through structured discussion. Where consensus could not be reached through discussion, a third reviewer applied the pre-specified eligibility criteria and ren-

dered a binding determination. Phase 1 (Identification): Database searches returned 312 records in total IEEE Xplore ($n = 84$), ACM Digital Library ($n = 71$), Google Scholar ($n = 103$), and IACR ePrint Archive ($n = 54$). Supplementary manual searches of the IETF Datatracker and citation chaining yielded 18 additional records, for a combined total of 330 records.




Figure 1. PRISMA 2020 flow diagram.

Phase 2 (Screening): After removing 41 duplicate records, 289 unique records were screened at title and abstract levels. Records clearly outside the scope of DNSSEC or post-quantum cryptography including general DNS measurement studies without PQC relevance, PQC applications to TLS or PKI without DNS-specific analysis, and quantum computing hardware surveys were excluded ($n = 187$), yielding 102 records for full-text evaluation.

Phase 3 (Eligibility): Full-text assessment of the 102 records resulted in the exclusion of 75 works on the following grounds: non-DNS PQC context ($n = 34$), specifically works addressing PQC exclusively in the context of TLS handshake performance, certificate chains, or IPsec without DNS-specific analysis; insufficient methodological contribution ($n = 19$), comprising works that reproduced prior results without a novel empirical, analytical, or protocol-design contribution; pre-2020 publication outside the defined foundational exceptions ($n = 14$); and non-English language ($n = 8$). The complete list of all 75 excluded studies, organized by exclusion category with the specific rationale for each exclusion, is provided in the Supplementary Materials (File S1).

Phase 4 (Inclusion): The final corpus comprises 27 works distributed as follows: 14 peer-reviewed conference papers (ACM CCS, USENIX Security, ACM ASIA CCS, ACM IMC, ACM SIGCOMM, ACM ANRW, PQCrypto, IEEE/IFIP TMA, IEEE MSWiM, IEEE ICC, USENIX WOOT, SSR); 6 journal articles (ACM SIGCOMM CCR, IEEE/ACM Transactions on Networking, IEEE Transactions on Network and Service Management, Cryptography, Advances in Web Development Journal); 3 IETF RFC specifications; 2 arXiv preprints

satisfying the community-recognition threshold; and 2 IETF Internet-Drafts with active DNSOP engagement. The complete PRISMA 2020 flow diagram is provided in Figure 1.

3.5. Data Extraction

Data extraction was performed using a structured instrument developed a priori and piloted on three randomly selected included works before full application. For each of the 27 included works, the following variables were extracted across five dimensions.

Dimension 1: Threat model and adversarial assumptions of whether the work models a classical adversary, a quantum adversary, or both; whether the HNDL threat is explicitly addressed; and the assumed network position of the adversary (on-path, off-path, or unspecified).

Dimension 2: Proposed solution or primary finding for the specific contribution category (algorithm evaluation, protocol proposal, empirical measurement, standardization analysis, or security analysis), the primary claim advanced, and any explicit limitations acknowledged by the authors.

Dimension 3: Experimental methodology and evaluation conditions include the DNS software and version used, the imposed network conditions (bandwidth, latency, and packet loss rate), the evaluation environment (testbed, simulation, field measurement, or analytical), and the evaluation scale (number of queries, resolver probes, or zone records processed).

Dimension 4: Quantitative results, where reported, signature size in bytes; public key size in bytes; signing and validation throughput in operations per second; DNS resolution latency in milliseconds; Bandwidth Amplification Factor (BAF); and DNS resolution failure rates under various transport conditions.

Dimension 5: Identified limitations and open problems, limitations explicitly acknowledged by the authors, and limitations identified during the extraction process that were not acknowledged by the authors, classified by the five challenge dimensions established in Section 1.

Where quantitative values were presented graphically rather than in tabular form, values were extracted from figures using digital measurement tools and cross-checked against text descriptions in the same work. Where multiple values were reported across different experimental conditions, the full range is recorded and the associated conditions are specified in Sections 4 and 5.

3.6. Quality Assessment

Each included work was assessed on four criteria using a two-point scale (0 = criterion not met, 1 = criterion met), yielding a maximum quality score of 4. The criteria were (1) publication in a peer-reviewed venue or satisfaction of the preprint community-recognition threshold defined in Section 3.3; (2) reproducibility of experimental results, defined as sufficient reporting of experimental conditions to permit independent replication or verification; (3) explicit statement of the threat model and adversarial assumptions; and (4) provision of quantitative evaluation results, as opposed to exclusively qualitative analysis.

All 27 included works scored 3 or 4 on this instrument, satisfying the minimum threshold of 3 established a priori for inclusion. In practice, this threshold did not actuate as a binding filter at the appraisal stage: no full-text-eligible study reaching quality assessment scored below 3, because works lacking evaluable methodology had already been removed at full-text screening under the “insufficient methodological contribution” criterion ($n = 19$; Figure 1), whose exclusion taxonomy accordingly contains no quality-score category. The threshold therefore characterizes the appraised corpus rather than having trimmed it.

This distinction is procedural rather than rhetorical, and it turns entirely on the sequence in which the two methodological gates were applied. Methodological adequacy is assessed at two separate points in the PRISMA workflow, under two different instruments, with two different functions. The first gate operates at full-text screening (Phase 3, Figure 1), where the “insufficient methodological contribution” criterion removed nineteen works’ position statements, high-level overviews, and implementation notes that restated prior results without an evaluable empirical, analytical, or protocol-design contribution. This gate is eliminative: its purpose is to decide inclusion, and it acts before any numerical score is assigned. The second gate is the four-criterion quality instrument of Section 3.6, which is scored only for the works that have already passed the first gate. By the time the quality instrument is applied, every remaining study necessarily possesses a methodology of some evaluable kind, because the studies that lacked one have already been withdrawn upstream. The two gates therefore cannot both bind: whatever the quality instrument would have excluded for want of evaluable method has, by construction, already been excluded on substantive grounds a phase earlier.

The empirical signature of this ordering is visible in Figure 1 itself. The exclusion taxonomy records categories such as non-DNS PQC context, pre-2020 publication outside the foundational exceptions, and non-English language, but it contains no category for a failing quality score, because no work was removed by the 3/4 cut-off. Every study that reached quality assessment scored 3 or 4, and the score distribution reported in Section 3.9 has nine works at 4/4 and eighteen at 3/4, with reproducibility (criterion 2) the single most frequently unmet criterion, which is a description of the surviving set, not a record of attrition. Had the threshold acted as a live filter, one would expect Figure 1 to itemize works excluded for scores of 1 or 2, and the yield at Phase 4 to differ from the count entering quality assessment; neither is the case. The threshold thus functions as a reported property of the corpus (“all included works satisfy at least three of four quality criteria”) rather than as an operator that selected the corpus. This is the precise sense in which it characterizes rather than trims.

Reading the threshold correctly also matters for the survivorship analysis in Section 3.10, and the ordering bounds that concern in a specific way. Because the works removed at the first gate were removed for absence of evaluable methodology and not for the direction of their findings, and because none of those nineteen documented a negative or failed evaluation of a post-quantum DNSSEC approach, the pre-quality screening cannot have preferentially discarded adverse results. The quality threshold, applied only afterward, could not have compounded a survivorship bias that the earlier gate did not introduce. The a priori cut-off is retained in the reporting for transparency and reproducibility of the appraisal procedure, not because it exercised independent selective force on the evidence base.

Works providing foundational analytical or standardization contributions without experimental evaluation of a category that is well represented in a field at the early stage of protocol transition were assessed on criteria appropriate to their contribution type: the reproducibility criterion was interpreted as the sufficiency of the analytical framework description to permit independent verification, rather than requiring replicable experimental conditions. The rationale for this interpretation is that applying a single, experiment-derived reading of the reproducibility criterion uniformly across the corpus would systematically misgrade a category of contribution that is not marginal but central at the present stage of the protocol transition, the stage at which the requirements, the threat characterization, and the security arguments are still being fixed, and before large-scale deployment evidence can exist. Two kinds of non-experimental work carry disproportionate weight in this corpus, and each demands a criterion matched to the claim

it actually makes. The first is standardization and specification material: the foundational documents RFC 4033–4035 [7] and RFC 5155 [26], the algorithm-presence requirement of RFC 4035 [37], and the consolidated deployment guidance of Rose et al. [18]. These works do not report measurements; they define the protocol surface, the chain-of-trust semantics, the denial-of-existence mechanism, the validation rules, and the operational constraints against which every empirical proposal in Sections 4 and 5 is subsequently evaluated. The second is analytical work whose conclusions follow deductively from stated premises rather than from observation: the combiner unforgeability guarantee of Bindel et al. [38], which establishes that a dual-signature composition remains secure so long as one component does, and the Mosca-theorem urgency calculus of Nal and Niemiec [9], which derives a migration-timing condition from parameter estimates rather than measuring one.

For contributions of these kinds, criterion 2 was operationalized as the sufficiency of the specification or analytical description to let an independent reader re-derive the result or verify the argument end to end of the analytical analog of replication rather than as the presence of a network testbed whose conditions could be reproduced. Demanding experimental replicability of a document that makes no experimental claim would not measure its quality; it would penalize it for a property it never purported to have. Worse, it would invert the epistemic order of the field, recording the specifications and proofs that supply the evaluative criteria for the empirical studies as though they were deficient empirical studies themselves. The interpretation adopted here is the same type-relative logic later made explicit in the certainty-of-evidence framework of Section 6.8, where a claim that follows deductively from protocol premises is graded on the rigor and explicitness of those premises and on the absence of unwarranted extrapolation, while a claim established by measurement is graded on replication and on the distance between testbed and production conditions. A foundational RFC and a single-testbed latency measurement are both admissible evidence, but they are admissible in different currencies, and scoring them on one scale would either overstate the measurement or understate the specification. To keep this transparent rather than hidden inside an averaged number, the quality tier of the works supporting any particular claim is surfaced at the point of use in Sections 4 and 5, so that a reader can see whether a given conclusion rests on replicated measurement, on a deductive argument from protocol structure, or on a directly inspectable fact, and can weigh it accordingly.

Quality tier is noted in Sections 4 and 5 where it is relevant to contextualizing the strength of evidence for a specific claim.

3.7. Synthesis Approach

The five difficulty factors listed in Section 1, namely transport restrictions, cryptographic agility, denial-of-service resilience, operational deployment, and standards readiness, were used to organize the narrative analysis of the 27 included works. Values are tabulated to enable direct comparison when several studies provide quantitative measurements on the same variable under similar experimental settings. Values are presented with their conditions clearly stated and direct numerical comparison is avoided when experimental conditions differ significantly between works.

The heterogeneity of experimental conditions across the included works encompassing varying network latency profiles (10 ms to 100 ms RTT), DNS software versions (BIND 9.16 through 9.19, Unbound 1.13 through 1.17, PowerDNS Authoritative 4.5 through 4.8), zone file structures ranging from synthetic test zones to production ccTLD datasets, and measurement methodologies spanning controlled testbeds to large-scale RIPE ATLAS field deployments precludes pooled statistical estimation. Because post-quantum DNSSEC

evaluation lacks the established benchmarking techniques that quantitative synthesis would necessitate, this variation is intrinsic to the current state of the field.

3.8. Limitations of the Review

This systematic review has three limitations that need to be clearly acknowledged. First, the search was limited to English-language publications, which might have left out pertinent contributions written in other languages, especially from Asia-Pacific academic communities engaged in DNS measurement and infrastructure operation. Second, due to the post-quantum DNSSEC field's rapid evolution, some of the gaps identified in Section 6 may be filled by works published or significantly revised between the search cut-off of December 2025 and the submission date of this paper. As a result, the roadmap presented in Section 7 is a summary of the field as of the review period. Third, the gray literature inclusion threshold defined as 50 or more citations or formal IETF working group adoption may have excluded emerging work of relevance that had not yet accumulated sufficient citation counts at the time of the search; this risk was partially mitigated by the manual IETF Datatracker search and the forward and backward citation chaining described in Section 3.2.

3.9. Protocol Deviations and Methodological Clarifications

This document discloses deviations from the pre-registered review methodology and provides methodological clarifications in compliance with PRISMA 2020 reporting standards.

Deviation 1: Status of OSF registration. The review methodology was submitted to the Open Science Framework (OSF) before the data extraction was completed, as noted in the pre-registration record. The DOI placeholder in Section 3 indicates a registration submitted before the search was conducted; the registration record details the process as outlined in Sections 3.1–3.4. The review researchers did not modify the eligibility criteria, research questions, or synthesis approach after conducting the search.

Graphical data extraction is the second deviation. Quantitative results (signature sizes, resolution latencies, BAF values, failure rates) were only displayed graphically in 14 of the 27 experiments that were examined; no corresponding data tables were included. WebPlotDigitizer 4.6 was used to obtain quantitative data for these studies from high-resolution PDF versions of the original graphs. Each graphical source was extracted by two independent reviewers; if the absolute difference between reviewers was less than 2% of the figure axis range, the extracted values were considered acceptable; if not, a third reviewer decided on differences. The data extraction tool records all visibly obtained results and the associated uncertainty range.

Deviation 2: Absence of meta-analysis and statistical heterogeneity of DNS software versions (BIND 9.16 to 9.19, Unbound 1.13 to 1.17, PowerDNS Authoritative 4.5 to 4.8, CoreDNS 1.10 to 1.12), network conditions (10 ms to 100 ms RTT, 1 Mbps to 100 Mbps bandwidth), zone file configurations (synthetic test zones to production ccTLD datasets comprising 59 million queries), and measurement techniques (controlled testbeds, Docker environments, RIPE Atlas field deployments using up to 10,000 probes). There is not a single benchmarking framework for evaluating PQC DNSSEC. There was no pooled statistical estimation. There was no statistical meta-analysis done. This is clearly acknowledged as a limitation of the existing body of evidence rather than a methodological decision made by the review committee. Quality assessment score distribution: The four-criteria test described in Section 3.6 yielded ratings of 3 or 4 for each of the 27 listed works. Score distribution: Nine studies received a score of 4/4, including four conference papers at ACM CCS/USENIX Security, two journal articles, two RFC specifications, and one IETF Internet-Draft with active working group involvement; eighteen studies got 3/4. The

most frequently unmet criterion was reproducibility (criterion 2): 16 of the 18 studies scoring 3/4 received a score of 0 in reproducibility, mainly due to the inadequate detail in the description of experimental conditions, which hindered independent replication without access to the original testbed configuration. All studies achieved a score of at least 3/4 following the pre-screening evaluation outlined in Section 3.3.

Deviation 3: Substitution of “PRISMA-guided” phrase. In accordance with PRISMA 2020 criteria, this review is referred to as “a systematic review following PRISMA 2020 guidelines” instead of “a PRISMA-guided systematic review” consistently. The differentiation is not only terminological: PRISMA 2020 constitutes a reporting standard, rather than a methodological framework; the methodological framework is delineated in Sections 3.1–3.8.

Deviation 4: Inter-rater reliability is not quantified. Although screening was performed by two reviewers with third-reviewer adjudication as described in Section 3.4, formal inter-rater reliability statistics (Cohen’s kappa) were not computed, and per-phase disagreement counts were not logged in a form permitting retrospective calculation. Agreement was managed through structured discussion at each phase, but the absence of quantified reliability is a reporting deviation from PRISMA 2020 Item 8 and is acknowledged as a limitation: the reproducibility of the screening decisions cannot be independently assessed from the reported data. Future updates of this review will log per-record dual decisions to permit kappa computation.

3.10. Constraints of Review-Based Evidence in Relation to Deployment-Scale Validation

A thorough assessment of the literature about post-quantum DNSSEC offers a distinct and more constrained type of evidence compared to empirical validation at the deployment scale. The distinction is epistemic before it is practical, and it follows directly from what a systematic review is capable of producing. A review is a second-order instrument: it aggregates, weighs, and reconciles evidence that other studies have already generated, but it originates no observations of its own. Its conclusions are therefore bounded above by the scope and quality of the primary corpus, and where that corpus has a systematic gap, the review inherits the gap and can at most characterize it but cannot fill it by synthesis. In the post-quantum DNSSEC literature, the decisive gap is the absence of production-scale measurement, and the point is sharpest precisely at the corpus’s strongest empirical study.

The most extensive field experiment available, the RIPE Atlas campaign of Goertzen et al. [43] across roughly 10,000 global probes, was conducted on experimental PQC-signed zones served from OQS-patched authoritative software rather than on production DNS operated at TLD scale under production-registered IANA algorithm numbers, and this distinction is substantive on two counts. First, no PQC signature method currently has an allocated number in the IANA DNSSEC Algorithm Numbers registry, so experimental zones sign under provisional OID values that break interoperability both between independently developed implementations and between experimental and production deployments; a measurement taken under provisional numbers cannot report how the deployed resolver population would treat production-registered algorithms, and this matters directly because an unrecognized algorithm identifier is itself the trigger for the validation-skip and downgrade behavior analyzed in Section 6.2, so the very numbering under which production would run is the variable most likely to change the observed failure profile. Second, the OQS-patched PowerDNS and BIND9 builds and the CoreDNS plugin of Suela et al. [44], the only PQC-capable DNS software identified by Schutijser et al. [45], are research builds rather than the hardened, high-volume nameservers and recursive resolvers that carry TLD-scale traffic, so their fragmentation handling, EDNS(0) buffer-negotiation, and TCP fallback behavior need not reproduce production infrastructure. The three quantities a

deployment decision most depends on, namely production-scale resolver failure rates, EDNS(0) negotiation behavior, and TCP fallback reliability under real PQC response sizes, have therefore not been generated by any included study and are not recoverable by combining studies that did not generate them; this is why Section 7 places production-scale characterization at Phase I as a prerequisite rather than as an optional refinement.

Three further properties, elaborated in the remainder of this section, compound the constraint and are worth naming here because together they explain why the evidence is “more constrained in type” rather than merely thinner. First, the primary studies were run under heterogeneous conditions, namely different DNS software and versions, different zone structures, and network settings ranging from controlled testbeds to large field deployments, so their quantitative results cannot be pooled and can be compared only indicatively; the contrast between the fragmentation failure figures of Van den Broek et al. [46] and the testbed measurements of Rawat and Jhanwar [16] is illustrative, since differing setups may be surfacing different underlying phenomena rather than replicating one. Second, the corpus is exposed to a publication bias that favors novel proposals and positive outcomes over documented failures, so the review’s picture of what “works” is drawn from a sample that underrepresents what did not. Third, the findings concentrate within a small number of co-authorship clusters, which means that agreement across several publications does not automatically constitute independent corroboration, because shared authorship can propagate common assumptions, toolchains, and measurement conventions.

Deployment-scale empirical validation is a different order of evidence precisely because it would relax all four constraints simultaneously: it would generate direct, production-conditioned observations, under a common methodology, across operators who are independent of one another and of the original research groups, on the very infrastructure the decision concerns. A review can identify each of these limitations, estimate its direction, and flag where a conclusion is single-cluster or testbed-bound and this review does so explicitly through the evidence-grading of Section 6.8, but it cannot convert a bounded, indirect, partially dependent evidence base into a direct one. That is the structural reason the roadmap of Section 7 does not treat further synthesis as a substitute for measurement: it places production-scale characterization at Phase I, as the prerequisite on which the Phase II and Phase III recommendations are explicitly gated, rather than as an optional refinement of conclusions the review could otherwise reach on its own.

The subsequent limitations are exclusive to review-based synthesis and are separate from the constraints of the individual studies included, as explained in Section 3.8.

Lack of production-scale veracity: The most extensive field measurement in the analyzed corpus—the Goertzen et al. [43] RIPE Atlas study—utilized experimental PQC-signed zones operating on OQS-patched authoritative servers, rather than on production DNS infrastructure. The performance of the worldwide resolver population under PQC DNSSEC load from production-registered IANA algorithm numbers, across production nameservers at the TLD scale, has not been assessed. Review-based synthesis cannot replace this measurement as it can only evaluate existing evidence, and the essential evidence—production-scale failure rates, EDNS(0) negotiation behavior, and TCP fallback reliability under PQC response sizes—has not been generated at the necessary scale.

Integration of varied experimental conditions: Sections 4–6 present a narrative synthesis that integrates findings from studies utilizing different DNS software versions, varied zone architectures, and unique network conditions, as detailed in Section 3.9. The comparison of quantitative values from different studies (for example, the 10% fragmentation failure rate reported by Van den Broek et al. [46] versus that of Rawat and Jhanwar [16]) does not necessarily indicate replication under uniform conditions; indepen-

dent measurements performed under varied experimental setups may uncover divergent underlying phenomena.

Publication bias favoring favorable results and novel proposals: The studied corpus likely underrepresents adverse outcomes studies evaluating a PQC DNSSEC approach and considering it impractical since such results are less likely to be submitted to or accepted by the venues included in the search strategy. The discovery that several methodologies lack empirical validation (MTL Mode [47], hybrid dual-signature security proofs [48]) may suggest a genuine shortcoming in published evaluations rather than an absence of assessment endeavors. This constraint is partially mitigated by the inclusion of IETF Internet-Drafts and arXiv preprints, which have less stringent acceptance standards; however, it cannot be completely eliminated within a review-based system.

The a priori quality threshold of 3/4 (Section 3.6) could, in principle, compound this survivorship tendency, since less polished reports are disproportionately likely to describe unsuccessful approaches. Two observations bound this risk. First, none of the nineteen works excluded at full-text screening for insufficient methodological contribution constituted a documented negative evaluation of a PQC-DNSSEC approach; they comprised position statements, high-level overviews, and implementation notes without evaluable methodology, and had any excluded work documented a failed approach, it would have been discussed narratively for contrastive purposes without entering the evidence synthesis. Second, in this field documented failures predominantly appear inside the published corpus rather than as standalone negative-result papers: the transport-layer infeasibility of large-signature schemes, the latency penalties of TCP-only operation, and the limitations that led fragmentation designs to be revised are reported within the included studies themselves ([12,49–51]), and the expiration histories of IETF Internet-Drafts function as the field's de facto registry of abandoned directions. Two mechanisms make this registry function concretely. First, the fragmentation designs were not abandoned wholesale but revised in place in response to failure modes documented inside the corpus itself. ARRF [49] introduced request-based application-layer fragmentation using a custom RRFRag pseudo-resource-record, and its three limitations were reported by its own authors: RRFRag carries no IANA allocation and is dropped, truncated, or left unanswered by strict-parsing firewalls, load balancers, and DNS-aware middleboxes [52], precisely on the large-response paths where fragmentation is invoked; request-based reassembly obliges the resolver to hold per-record state, opening an asymmetric memory-exhaustion vector in which an adversary announces a large fragment map and never delivers the remaining fragments; and because each fragment is a separate exchange, the probability that at least one of an n -fragment response is lost grows with n , forcing extra round-trips in exactly the degraded conditions where they are most costly. Raavi et al. [50] added a fourth failure mode, fragment mis-association across concurrent resolutions. QBF [53] is the direct revision that answers these: by encoding fragment identifiers in standard QNAME labels and storing signature bytes in ordinary RRSIG and DNSKEY records, it removes the middlebox-drop and memory-exhaustion exposures and reassembles in a single round-trip. Because the corpus that reports the ARRF limitations ([49,50]) also reports the QBF revision [53], the revision history is legible entirely from within the included studies. Second, because no DNSOP working-group item has been chartered for PQC-DNSSEC and coordination occurs on an unofficial pq-dnssec mailing list [48], the primary durable trace a discontinued direction leaves is the IETF Internet-Draft: a draft lapses and is removed from active status if it is not refreshed or adopted within its validity period, so an expired, never-adopted draft is a dated, publicly retained record that a direction was proposed and then not carried forward. The manual IETF Datatracker search described in Section 3.2 therefore surfaces abandoned directions that never reach peer-reviewed publication, and in a field with no

formal negative-results venue the expiration timeline itself substitutes for a registry of directions the community stopped pursuing. The residual risk that methodologically weak but substantively informative failure reports were never published at all is a manifestation of the publication bias acknowledged above and cannot be corrected at the review level; it is retained as a limitation.

Author-group concentration and epistemic independence: A co-authorship analysis of the included corpus reveals four principal research clusters: the Twente/SIDN network centered on van Rijswijk-Deij and Müller ([12,17,29,45,46,54], extending into the cross-institutional research agenda [48]); the ATHENE group of Heftrig, Shulman, and Waidner ([13,14,40,41]); Rawat and Jhanwar ([16,43,53]); and the Goertzen-Stebila-Thomassen line ([49]). Consequently, concordance among multiple publications does not by itself establish independent corroboration, since shared authorship may propagate common assumptions, toolchains (notably OQS-patched software), and measurement conventions. To mitigate this, the consistency criterion in Section 6.8 is applied at the level of co-authorship-disjoint author clusters rather than publication counts: each High-confidence rating was re-examined for cross-cluster support; one finding was downgraded on this basis, and two others are retained as High on the explicitly stated grounds of protocol-arithmetic or registry verifiability rather than replication count. Findings resting on a single cluster are flagged accordingly in Section 6.8, and independent replication outside the originating groups is identified as a research priority in Section 6.7.

3.11. Treatment of Equations and Quantitative Results

Treatment of equations and quantitative results: As this work is a systematic review rather than an original empirical study, none of the equations reported in this paper were derived or experimentally evaluated by the present authors. Equation (1), Mosca's inequality, is reproduced from [9] and is applied analytically in Sections 2.4 and 7.2 by substituting the parameter estimates reported in the cited literature ($t_{protect} \approx 20$ years [9], $t_{migrate} \approx 7$ years [10,11], $t_{quantum} \approx 15$ years [9]); no independent estimation of these parameters was performed. Equations (2)–(7), specifying the SL-DNSSEC KEM-MAC construction, are reproduced from Sections 3.2–3.4 of [16] and describe protocol logic rather than measured quantities; their correctness rests on the security proofs given in [16]. Equation (8), the Bandwidth Amplification Factor, is reproduced from Section 2 of [12], and all numerical BAF values reported in this paper (Sections 5.7 and 6.3.3) are taken directly from the empirical measurements published in the cited primary studies, principally Tables 11, 12, 16, and 17 of [16], obtained on the hardware and network configurations documented therein. Where a primary study does not fully specify its evaluation methodology, this is noted at the point of citation, and the absence of a unified cross-study benchmarking framework is itself identified as a research gap in Section 6.

4. Systemic Analysis of the Identified DNSSEC Problems

Two decades of experience in DNSSEC utilization and implementation reveal that the protocol's limitations exist at three interrelated levels: the protocol architecture level, the cryptographic algorithm level, and the ecosystem implementation level. The cumulative impact of these three tiers results in a contradiction wherein, despite over twenty years of implementation, a protocol, regardless of its technical clarity, continues to leave over 60% of worldwide internet traffic vulnerable [55]. This section provides a systematic study of concerns identified in the relevant literature, positing that these issues are interconnected systemic challenges necessitating a comprehensive solution rather than isolated technical flaws.

4.1. Packet Size Limitations, Fragmentation and DDoS Amplification

The Extension Mechanisms for DNS (EDNS0) (RFC 2671) later raised this limit to between 1232 and 4096 bytes, but DNS Flag Day 2020 established 1232 bytes as the final maximum. The Domain Name System (DNS) was initially created for User Datagram Protocol (UDP) transmission with a 512-byte response limit (RFC 1035). By including DNSKEY, RRSIG, and NSEC/NSEC3 entries, DNSSEC responses regularly exceed this level. There are three consequences. IP fragmentation or TCP fallback is triggered by answers larger than 1232 bytes; Van den Broek et al. [56] note that approximately 10% of resolvers fail to properly reassemble fragmented responses, leading to resolution failure rather than reduced performance. Second, because each fragment lacks independent verification, fragmented responses expand the attack surface for cache poisoning [56]. Third, augmented responses exacerbate DNS-based DDoS attacks: van Rijswijk-Deij et al. [29] report an average tenfold increase in amplification factors for DNSSEC-signed replies, while Rawat and Jhanwar [16] provide further quantification for PQC techniques. While RSA-2048 responses vary from 2000 to 4000 bytes [18], ECDSA P-256 replies (64-byte signatures) in conventional DNSSEC are closer to the UDP limit [12], requiring TCP fallback in 11% of nameserver interactions.

The literature [56] examines this constraint along three axes: a response beyond the path MTU triggers IP fragmentation or TCP fallback, which in the best case raises round-trip time and in the worst case causes total connection failure; empirically, as many as 10% of resolvers cannot correctly reassemble fragmented packets [46,56]; and, most consequentially, fragmented answers open a cache-poisoning vector because the individual fragments are not independently verified [56]. This last point warrants full elaboration, because on its face it appears paradoxical: if verifying fragments separately “complicates validation,” how does fragmentation aid an attacker rather than merely burden the defender? The resolution is that DNSSEC performs no per-fragment verification at all, and the poisoning vector lives in exactly that gap. A DNSSEC signature is computed over a complete, canonically ordered resource-record set and is checked once, against the reassembled RRset, only after the IP layer has already stitched the fragments back together; validation therefore sits above reassembly in the stack and runs strictly after it. Reassembly, by contrast, is a pre-cryptographic operation: the IP layer binds fragments using only the 16-bit IP identification field together with the fragment offset and the more-fragments flag, and it never consults the DNS transaction ID, the source port, or any signature material, because for the trailing fragments none of that information is present.

This layering is what Herzberg and Shulman [56] show an attacker can exploit. The two randomized fields on which classical DNS anti-spoofing depends, the 16-bit source port and the 16-bit transaction ID, together a search space on the order of 2^{32} , are carried only in the first fragment of a UDP datagram, since only the first fragment holds the UDP and DNS headers; every subsequent fragment carries an IP header and raw payload and nothing else. An off-path adversary, with the off-path class defined in Section 2.5, which cannot observe the query and therefore cannot guess the port or transaction ID, is consequently not required to defeat those defenses at all. It needs only to fabricate a trailing fragment that matches the far weaker IP-ID and the expected offset and race it into the victim’s reassembly buffer ahead of, or in place of, the authentic fragment. The entropy the attacker must overcome has therefore collapsed from the combined port-and-transaction-ID space guarding an unfragmented answer to the single, much smaller IP-ID space guarding reassembly [46,56]; this is a change in category rather than of degree, since an injection that is computationally impractical against a whole datagram becomes routinely winnable against a fragmented one. If the attacker wins the race, the forged bytes are spliced into the

response before the resolver ever parses the message, and the resolver proceeds to validate a datagram that is part genuine and part attacker-authored.

Two outcomes follow at the validation step, and the presence of RRSIG records forecloses neither. In the first, the injected bytes fall in response to content that the resolver may cache without covering signature additional-section records or glue so that validation of the signed portion succeeds while the unsigned, now-forged portion is cached and served as authentic; this is cache poisoning in its literal form, the very result signing was introduced to prevent [4,5]. In the second, the injected bytes corrupt signed data, the reassembled RRset fails verification, and the resolver returns SERVFAIL, but a bare validation failure is not a safe state, because under the algorithm agility and downgrade behavior analyzed in Sections 4.5 and 6.2 such a failure can be steered into abandonment of DNSSEC protection rather than refusal, and the off-path fragment-tag-modification attack of Heftrig et al. [15] shows that the identical fragment-rewriting primitive that enables poisoning also drives off-path downgrade. Because verification is applied once to the reassembled whole and never to the fragments as they arrive, a spoofed fragment that survives reassembly is, at the point of injection, indistinguishable from legitimate data; the signature check that could expose the tampering runs too late to keep the forgery out of the reassembly buffer or the cache, and operates too coarsely to identify which fragment was altered. Fragmentation thereby converts what begins as a transport-size limitation into a security vulnerability, and it is this conversion, not the size increase in isolation, that makes fragmentation avoidance, rather than mere accommodation, a security objective for post-quantum DNSSEC, as developed in Section 6.1 [46].

The supplementary entries introduced by DNSSEC enhance the efficacy of DDoS amplification attacks. Research [29] indicates that DNSSEC-enhanced DNS answers amplify the factor by an average of tenfold, since a simple query comprising DNSKEY, RRSIG, and NSEC records produces significantly bigger results. The mechanism that turns this increase in size into an attack vector is reflection over stateless transport. Because UDP requires no handshake, a DNS server cannot verify the source address of a query and will return its answer to whatever address the query claims to originate from; an adversary exploits this by forging the victim's IP address as the source of a small query, typically under 100 bytes, directed at open resolvers or authoritative nameservers, which then deliver the much larger response to the victim [57]. DNS is an especially effective reflector for this pattern: the infrastructure is publicly reachable by design, open resolvers exist in large numbers across the internet, and a single small query can elicit a response many times its size, so the attacker both conceals the true origin of the flood behind third-party servers and multiplies the traffic directed at the target. The multiplication is quantified by the Bandwidth Amplification Factor, the ratio of response size to query size (formalized in Section 6.3.3), and any protocol feature that enlarges responses raises this ratio directly. DNSSEC contributes to the yield in three compounding ways. First, the protocol mandates the inclusion of RRSIG records in every signed response and returns large DNSKEY RRsets when the keys themselves are queried, so signed answers are intrinsically larger than their unsigned counterparts; measurements of signed versus unsigned zones place the resulting amplification for DNSSEC-enabled responses at roughly an order of magnitude above equivalent unsigned queries [29]. Second, DNSSEC transport requires EDNS0 with enlarged UDP buffer sizes, the very mechanism that permits these large responses to be delivered over UDP in the first place rather than being truncated and forced to TCP, whose three-way handshake would defeat the source-address spoofing on which reflection depends. Third, the reflected responses consist of cryptographically valid signed data, indistinguishable at the packet level from legitimate DNSSEC traffic, and therefore resistant to content-based filtering that might otherwise discard malformed

reflection payloads [29]. The compounding effect is that a single attacker sustains a flood far exceeding its own upstream bandwidth, aimed at a victim that never issued a query, using packets that a resolver cannot refuse to answer and a firewall cannot easily distinguish from valid responses. The cryptography DNSSEC adds provides no defense here, and in fact deepens the problem: signatures authenticate the origin and integrity of records but do nothing to limit the volume or destination of responses, so the same signed payloads that guarantee authenticity to a legitimate resolver serve equally well as amplified, unfilterable ordnance against a spoofed victim a security mechanism converted, by its own size, into attack material. This result is operationally important: DNSSEC, intended to safeguard DNS, concurrently enhances the efficacy of DNS-based amplification attacks because of the protocol's mandate for signature inclusion, as empirically demonstrated in [30] and quantitatively assessed for PQC schemes by Rawat and Jhanwar [16].

Müller et al. [12] report that 11% of nameservers lacked TCP capability during the evaluation period, indicating that TCP fallback is useless for a significant portion of the resolver population.

During algorithm migration phases, dual-signature zones simultaneously contain two RRSIG sets, hence increasing response size. This problem is somewhat mitigated by algorithm selection, but it cannot be fixed within the current protocol architecture; Section 5 evaluates the evidence supporting each mitigation technique.

As a result, packet size continues to be a crucial constraint during the duration of DNSSEC.

4.2. Comparative Analysis of Signature Sizes

A signature that exceeds UDP transmission restrictions cannot be validated, regardless of computational efficiency; hence, Müller et al. [12] identify signature size as the primary criterion for PQC-DNSSEC compatibility, with validation throughput as a secondary factor. Table 2 compares the signature sizes of the classical and post-quantum algorithms relevant to DNSSEC.

Table 2. Comparison of classical and post-quantum signature algorithms in the context of DNSSEC.

Algorithm	Signature (Bytes)	Public Key	Type	Validation
RSA-2048	256	300	Public Key (Classical)	49,367
Ed25519	64	32	Public Key (Classical)	7954
ECDSA P-256	64	64	Public Key (Classical)	13,078
ML-DSA-44 (Dilithium)	2420	1312	Signature (Post-Quantum)	8956
FN-DSA-512 (Falcon)	666	897	Signature (Post-Quantum)	20,228
SLH-DSA-128s (SPHINCS+)	7856	32	Signature (Post-Quantum)	4

No currently standardized PQC algorithm satisfies the UDP size constraint for non-minimal DNSSEC responses without protocol-level modification [12,44].

This confirms that format compactness does not resolve the transport problem; the binding constraint is the inherent signature size, not encoding overhead [54].

Computational throughput is not the binding constraint: Falcon-512 achieves 3307 signatures/s and 20,228 validations/s against a DNSSEC operational requirement of approximately 1000 validations/s [12]. The binding constraint is transport size. DNSSEC's weaknesses extend beyond packet size; the denial-of-existence mechanism introduces additional structural challenges addressed in Section 4.3.

4.3. NSEC, Zone Enumeration, and Temporal Inconsistency of the Attestation Chain

DNSSEC's authenticated denial-of-existence mechanism reveals a fundamental conflict between security and efficiency, originating from the original design decisions of the protocol. The NSEC records serve the purpose of cryptographic authentication of non-existent names and, aiming at this, list all names in a zone in lexicographic order. Finite-state analysis [25] has confirmed that this structure simply enables complete zone enumeration, the so-called "zone walking", which can bring disclosure of an organization's internal infrastructure, server names and other confidential information [18].

Aiming at addressing this problem, NSEC3 (RFC 5155 [26]) was introduced; it uses hashed names to complicate enumeration. According to operational guidelines, NSEC3 cannot guarantee full required protection. Empirical research confirmed that the salt string, whose value is publicly accessible via DNS query, cannot prevent post-query dictionary attacks. Moreover, NSEC3 usage ensues a significant computational overhead on authoritative servers [18,25]. The only supported hash algorithm for NSEC3 is SHA-1. Research [58] has practically demonstrated the first full SHA-1 collision, confirming that SHA-1 shows cryptographic weaknesses and its continued use contains risks in the long-term perspective. Although a direct exploit of a SHA-1 collision in the NSEC3 context has not been demonstrated so far, the cryptographic community considers replacing SHA-1 to be a necessity, while specification of a new hash algorithm for NSEC3 is still ongoing [18].

The NSEC3 opt-out option is another critically vulnerable area. Here too, research [25] has practically demonstrated (in a laboratory environment) that the opt-out bit, while being a useful configuration option for TLD operators, constitutes a critical security vulnerability in enterprise-level zones: an attacker can insert any A record whose hash falls between two hashed names in an opt-out NSEC3 record in the zone. A browser cookie theft attack via DNS name insertion was successfully performed without cryptanalysis [25], confirming that this vulnerability can be exploited practically.

The temporal inconsistency of the attestation chain is another fundamentally significant problem in DNSSEC that is not given due review and discussion. Running of the finite-state model [25] revealed that the DNSSEC protocol specifies a temporal relationship only between the TTL and its directly corresponding RRSIG, while no such relationship is specified between other elements in the attestation chain. As a result, a DNS record may remain valid in the cache after the corresponding higher-level signature in the attestation chain is expired. Specifically, in the event of key compromise, an attacker can use the compromised key to sign records and create an RRSIG record with a long validity period, which remains valid in the cache even after expiration of other elements of the attestation chain [25]. This problem creates a classic downgrade attack vector within DNSSEC and its resolution requires changes at the resolver software level, specifically if TTL expiration depends on the validity of all elements in the attestation chain.

NSEC/NSEC3 presents a three-way conflict between confidentiality (zone enumeration), performance (NSEC3 CPU overhead), and security (SHA-1 weakness [58]) that remains unresolved. Section 6.3 addresses how PQC deployment compounds the NSEC3 CPU exhaustion surface.

4.4. Key Management Complexity

Key lifecycle management is the least analyzed yet practically one of the most significant dimensions of DNSSEC's operational complexity. According to the operational guidelines [18], DNSSEC signing keys are divided into two categories, namely Zone Signing Key (ZSK) and Key Signing Key (KSK); each of these requires a different rollover procedure. KSK rollover is particularly challenging, as it requires coordination with the parent zone, which is often under the administrative control of a different organization.

Eight years after it was first introduced, the first root DNS KSK rollover occurred in 2018. Up to 2% of validating resolvers were not ready to update the root trust anchor, posing a serious risk of widespread DNS resolution failure, according to a thorough empirical investigation of the procedure [10]. The study [10] also verified that the rollover procedure took 11 months, during which ICANN had to make every effort to coordinate. The operational complexity of algorithm migration [12] is also well-documented based on this experience; from standardization to full deployment, the EdDSA-Ed25519 integration into DNSSEC took more than three years.

RRSIG validity periods of 5–7 days require continuous automated resigning [18], increasing the operational burden and the probability of human error. DNS resolution fails entirely for DNSSEC-validating resolvers if a key expires or is compromised without immediate replacement. Hardware security modules are operationally required to protect the KSK [18], raising infrastructure costs that disproportionately affect small and medium-sized operators.

4.5. Typology of DNS Attacks: From Cache Poisoning to Downgrade Attacks

Fragmented responses represent a significant cache-poisoning vector since individual pieces cannot be independently verified [56]; the correlation between packet size (Section 4.1) and injection assaults is fundamentally interrelated rather than incidental. Cache poisoning exemplifies a quintessential DNS attack; yet, the threat model spans a wider range. Alharbi et al. [4] classify man-in-the-middle attacks and DNS hijacking as distinct vectors; the 2020 SADDNS attack [5] demonstrated that side-channel vulnerabilities can circumvent port randomization, confirming that the DNS attack surface continues to evolve post-DNSSEC deployment. Downgrade attacks compromise DNSSEC without requiring any cryptanalysis. Heftrig et al. [14] demonstrate an off-path attack that undermines a DNSSEC-validated connection and reverts it to unsigned DNS through transaction-ID prediction and UDP spoofing, and Kambourakis et al. [59] find that six major browsers were vulnerable to at least one downgrade variant while presenting no visible indication to the user. The reason a fully signed zone does not by itself defeat these attacks is that DNSSEC authenticates record content along the signer-to-resolver path but neither encrypts nor authenticates the stub-to-recursive hop, and the AD bit on which a stub relies to learn that a response was validated is a single cleartext flag that any element on that last hop can clear or forge. Consequently, however completely a zone is DNSSEC-signed, an adversary positioned on the stub-to-recursive link of the on-path class of Section 2.5 can strip the AD bit, coerce a fallback to unsigned, or present a downgraded response, and the stub has no cryptographic means to detect it; a partially implemented DNSSEC therefore provides strictly weaker guarantees than the specification, because unvalidated paths remain open to the cache-poisoning and man-in-the-middle vectors of Alharbi et al. [4] and the SADDNS side-channel [5]. End-to-end downgrade resistance thus requires an encrypted and authenticated stub-to-recursive channel DoT or DoH regardless of the extent of DNSSEC implementation. This is precisely why the absence of an integrated analysis across DNSSEC, DoT, and DoH identified by Lee et al. [60] is a security gap rather than a mere performance question, and the cipher-suite-negotiation gap of Shrishak and Shulman [61] compounds it, since without negotiation the recursive cannot even signal to the stub which algorithms it will honor on that hop.

4.6. Low Global Adoption: A System of Self-Reinforcing Barriers

A global validity rate of about 40% suggests that adoption constraints are structural rather than technical, since DNSSEC regulations were approved in 2005.

Operational (Barrier I): The complexity of key management is the main obstacle to adoption for small and medium-sized businesses (Section 4.4) [12,62]. The coordination costs that hinder deployment were highlighted by the 2018 Root KSK rollover, which took 11 months and put 2% of resolvers at risk of resolution failure [10].

Infrastructure (Barrier II): Operational guidelines [18] state that most stub resolvers rely on an upstream recursive resolver and are unable to independently perform DNSSEC validation, indicating that full security necessitates an entire infrastructure chain. The efficacy of the entire system is compromised if DNSSEC is absent at any point in this chain.

Ecosystem (Barrier III): The study [63], which methodically examined the issue of trust chain disruption brought on by inadequate DNSSEC implementation, was the first to formally define the idea of islands of security. Current research [64] reveals that the issue is still unsolved over twenty years after this investigation. Only 11 of the top 20 registrars support DNSSEC, and three of them rely on a subpar email-based approach for DS record submission, according to study [27], the most thorough empirical investigation of registrar adoption. Consequently, just 0.31% of .org domains exhibit complete DNSSEC compliance [64].

Deployment Rate (Barrier IV): According to empirical data [12], the DNS environment takes a long time to integrate new algorithms; even two years after EdDSA-Ed25519 was standardized, 30% of resolvers were still not supported by this algorithm. More than three years elapsed between the standardization of EdDSA-Ed25519 and its complete deployment, according to a thorough investigation of the algorithm's lifecycle [11]. According to research, it will take ten years to fully deploy PQC algorithms at the current rate [65].

The DoH Paradox (Barrier V): Research [59] has shown that while DNS traffic encryption safeguards confidentiality, it also compromises the infrastructure monitoring capabilities required for the successful deployment of DNSSEC. According to research [66], the DNS security framework is further complicated by the detection of malicious traffic through DoH, which increases fragmentation.

The five obstacles fall into two related categories. A self-reinforcing cycle is created by Barriers I–III (operational, infrastructural, and ecosystem): complexity hinders acceptance, low adoption lowers the incentive to address Island of Security fragmentation, and fragmentation undermines DNSSEC's efficacy even for its supporters. This cycle is externally intensified by Barriers IV–V (deployment velocity, DoH dilemma). Dual-signature zones during the transition will incur additional response-size penalties on an ecosystem that has not completely implemented standard DNSSEC, and therefore the PQC migration must take existing dynamics into account.

4.7. Synthesis of the Systemic Nature of DNSSEC's Problems

The analysis presented above reveals that DNSSEC's problems are not isolated technical flaws, but structurally interconnected challenges. Table 3 presents a synthetic analysis of these problems.

Table 3 supports three conclusions. First, PQC migration will not resolve the four “unresolved” problems (downgrade attacks, islands of security, low adoption, attestation chain temporal inconsistency); these must be addressed in parallel rather than sequentially. Second, problems with “partial” PQC impact—notably NSEC zone walking—remain equally significant after PQC adoption. Third, the problem hierarchy is circular: protocol-level constraints (packet size, DDoS amplification, attestation chain) create the operational barriers that suppress adoption, and low adoption reduces the ecosystem incentive to address protocol-level constraints. This circularity explains the two-decade deployment gap and will equally constrain PQC adoption.

Table 3. Systemic classification of DNSSEC Problems.

Problem	Resolution Status	Impact on PQC	Level
Packet size and fragmentation	Partially unsolved	Critical—10–25x increase	Protocol
NSEC zone walking	Partially unsolved		Cryptography
DDoS amplification	Partially unsolved	Worsens with size increase	Protocol
Downgrade attacks	Unsolved		Protocol
Islands of security	Unsolved	Critically worsens	Ecosystem
Low adoption rate	Unsolved	Critically worsens	Ecosystem
Key management complexity	Partially unsolved	Critically worsens	Operational
Attestation chain temporal inconsistency	Unsolved	Worsens with key rollover	Protocol

Second, issues with “not directly worsened” PQC impact—such as downgrade assaults and NSEC zone walking—do not negate their significance. Conversely, it indicates that switching to PQC does not shield against these issues, and they will continue to be equally significant difficulties even after PQC.

Third, Table 3 shows that the DNSSEC problem system’s hierarchical protocol-level issues like packet size, DDoS amplification, and attestation chain serve as the basis for ecosystem issues, while ecosystem issues like low adoption rate and islands of security lessen the motivation to address protocol-level issues. Because of this circular dependency, DNSSEC has not been widely deployed during the past 20 years, and switching to PQC is both a technical and an ecosystem difficulty. The ideas put forth in the literature for the systematic resolution of these issues are examined in the section that follows.

5. Post-Quantum DNSSEC Proposals

The systemic complexity of the protocol environment and the restriction imposed by signature size are two obstacles to incorporating post-quantum cryptography (PQC) into DNSSEC, as the systematic analysis in Section 4 demonstrates. The relevant literature states that the three primary approaches that can be used to get around these problems are algorithm optimization, fragmentation-based methods, and signature-free authentication. This section offers a methodical assessment of the data underpinning each methodology, analyzes the comparative efficacy of each strategy, and identifies the challenges requiring resolution, which are addressed in Section 6.

5.1. Quantitative Transport Overhead and Fragmentation Analysis

This section provides a comprehensive comparison of transport overhead and fragmentation for four DNSSEC signature schemes, namely ECDSA P-256, FN-DSA-512 (Falcon-512), ML-DSA-44, and SLH-DSA-128s, based only on empirically measured response sizes documented in the original literature. No formula for response size has been established; all data in Table 4 are sourced directly from production measurements and testbed studies conducted by Fabrizio et al. [54], Rawat and Jhanwar [16], and Müller et al. [12]. In order to give a direct evaluation against the UDP fragmentation threshold, TCP Slow Start window, and DDoS amplification risk, this section offers a unified cross-scheme tabulation of previously distinct measurements using a consistent set of comparison dimensions. It must be emphasized that this tabulation harmonizes the presentation of these measurements, not their experimental conditions: the underlying studies differ in network configuration, DNS software, and zone structure, so the cross-scheme comparison is indicative rather

than a replication under uniform conditions, consistent with the methodological limitation discussed in Section 3.9.

Table 4. DNSSEC response size comparison.

Scheme	Sig. (Bytes)	Pub. Key (Bytes)	R_NOERROR (Bytes)	R_NXDOMAIN (Bytes)	UDP Frags	UDP Status	TCP Slow Start Window Exceeded
ECDSA P-256 [12]	64	64	~512 [12]	~680 [12]	1	PASS	No
FN-DSA-512 (Falcon) [54]	666	897	2100–2950 [54]	2269–3767 [54]	3	FAIL	No
ML-DSA-44 [16]	2420	1312	~9800 [16]	~14,500 [16]	10	FAIL	No
SLH-DSA-128s [16]	7856	32	~24,800 [16]	~38,000 [16]	22	FAIL	No (SLH-DSA-256s exceeds 65,535)

Sources and breadth of measurement: The response size metrics for ECDSA P-256 are derived from Müller et al. [12], who indicate a median NOERROR response size of 512 bytes and an NXDOMAIN response size of 680 bytes for a typical authoritative zone. The response sizes for FN-DSA-512 (Falcon-512) are derived from the measurements conducted by Fabrizio et al. [54], who assessed 59 million live queries. The authoritative infrastructure for the nl ccTLD, with OQS-patched BIND 9.16, produced NOERROR replies varying from 2100 to 2950 bytes, with a median of 2420 bytes. NXDOMAIN responses, containing two to three NSEC3 records each accompanied by a complete RRSIG, varied from 2269 to 3767 bytes. The response sizes for ML-DSA-44 and SLH-DSA-128 are based on the testbed measurements conducted by Rawat and Jhanwar [16], who assessed both schemes on a 100 Mbps Docker-based DNS testbed utilizing OQS-patched PowerDNS Authoritative 4.7: ML-DSA-44 generated NOERROR responses of roughly 9800 bytes and NXDOMAIN responses of approximately 14,500 bytes; SLH-DSA-128 yielded NOERROR responses of about 24,800 bytes and NXDOMAIN responses of around 38,000 bytes. UDP fragment counts are determined by the formula $\text{ceil}(R/1232)$, with 1232 bytes representing the maximum DNS response size via UDP as defined by the DNS Flag Day 2020 operational consensus [59].

Table 4 illustrates the empirically measured response sizes for all four schemes, derived from sources [12,16,54], within the framework of the NOERROR and NXDOMAIN scenarios, alongside the UDP fragmentation threshold (1232 bytes), the TCP Slow Start initial window (12,200 bytes as specified by RFC 6928 with $\text{initcwnd} = 10$), and the maximum DNS message size limit (65,535 bytes as determined by the 16-bit DNS message length field). The fragment count is determined by taking the ceiling of R divided by 1232 for IP-layer fragmentation.

Comparison of transport overhead and fragmentation for four DNSSEC signature techniques: Response sizes were obtained directly from original sources: ECDSA P-256 as reported by Müller et al. [12]; FN-DSA-512 as detailed by Fabrizio et al. [54] (.nl ccTLD, 59 million queries, OQS-BIND 9.16); ML-DSA-44 and SLH-DSA-128s as presented by Rawat and Jhanwar [16] (100 Mbps Docker testbed, OQS-PowerDNS 4.7). Fragment count equals the ceiling of R divided by 1232. UDP Status: PASS = conforms to the 1232-byte DNS Flag Day 2020 threshold; FAIL = threshold surpassed. These values were obtained under heterogeneous experimental conditions, a production ccTLD measurement campaign ([54]: .nl authoritative infrastructure, OQS-patched BIND 9.16, 59 million live queries), a controlled 100 Mbps Docker testbed ([16]: OQS-patched PowerDNS Authoritative 4.7, synthetic zones), and classical-baseline measurements on conventional production zones ([12]) with differing network settings, software versions, and zone structures. The tabulated figures are therefore not direct replications under uniform conditions and should be interpreted as indicative of the response-size regime of each scheme rather than as strictly equivalent measurements. Cross-scheme differences in the magnitude reported here (roughly $2\times$ to

$30\times$ the 1232-byte UDP threshold) are robust to such variation; small absolute differences between adjacent rows should not be over-interpreted.

Table 4 consolidates three observations that were present in the primary literature but had not been placed side by side under a uniform comparison framework. First, the UDP fragmentation threshold is crossed by all four evaluated PQC schemes at NOERROR response size. Even FN-DSA-512, the most compact PQC candidate, produces a median NOERROR response of approximately 2420 bytes [54]—approximately 1.97 times the 1232-byte UDP limit—and NXDOMAIN responses ranging from 2269 to 3767 bytes [54], confirming that fragmentation is unavoidable under any realistic zone configuration using Falcon-512.

Second, TCP Slow Start behavior (RFC 6928 `initcwnd` = 10 segments of 1220 bytes = 12,200 bytes initial congestion window) does not affect ECDSA or FN-DSA-512 NOERROR responses, but is relevant for ML-DSA-44 responses at NXDOMAIN scale (~14,500 bytes) and SLH-DSA-128s responses at both NOERROR and NXDOMAIN scales. Rawat and Jhanwar [51] quantify the resulting TCP resolution penalty at 200 ms and additional round-trip time per congestion window overflow at 100 ms RTT, with SPHINCS+ (SLH-DSA-128s) responses requiring more than two additional round-trips to complete transmission. This penalty is independent of implementation choice; it follows from the mandatory protocol event sequence specified by RFC 5681 and the response sizes produced by the cryptographic schemes.

Third, SLH-DSA-256s presents a qualitatively distinct incompatibility. Three RRSIG records at 29,792 bytes each yield a minimum three-RRSIG contribution of 89,376 bytes, which exceeds the absolute DNS message ceiling of 65,535 bytes imposed by the 16-bit length field of the DNS message format. SLH-DSA-256s cannot be deployed in any DNSSEC configuration carrying three or more RRSIG records without a fundamental redesign of the DNS message format itself; this is not a transport-layer constraint but a protocol-format constraint.

DDoS amplification under fragmentation: Rawat and Jhanwar [16] measure a Bandwidth Amplification Factor (BAF) of 10.8 for QBF (Falcon-512, application-layer fragmentation) and 0.48 for SL-DNSSEC Level I (ML-KEM-512 + HMAC-SHA-256) under identical query conditions on a 100 Mbps testbed. Under IP-layer fragmentation—the mechanism that applies to all signature-based PQC proposals—the effective amplification factor is higher than application-layer BAF because the entire response is transmitted to the victim without per-fragment confirmation. Rawat and Jhanwar [16] do not report IP-layer BAF values for ML-DSA-44 or SLH-DSA-128s directly; the application-layer BAF values of 10.8 (QBF) and 13.3 (ARRF) [16] are therefore the most conservative available empirical estimates for the amplification risk of fragmented PQC DNSSEC responses.

These findings have direct implications for the research roadmap in Section 7.3. IANA algorithm registration for any scheme whose NXDOMAIN responses exceed the TCP Slow Start initial congestion window (12,200 bytes under RFC 6928 `initcwnd` = 10)—specifically ML-DSA-44 at approximately 14,500 bytes [16] and SLH-DSA-128s at approximately 38,000 bytes [16]—requires simultaneous specification of resolver-side TCP connection management recommendations, not merely algorithm parameter tables. The fragmentation failure analysis further motivates the Phase I hypothesis H1 (Section 6.6): the 10% resolver fragmentation failure rate reported by Van den Broek et al. [46] was measured under classical DNSSEC response sizes; whether this rate holds, increases, or decreases under FN-DSA-512 NXDOMAIN responses of 2269–3767 bytes [54] has not been empirically established.

PQC signature dimensions, transit restrictions, security flaws, compounding factors, and implementation implications are all part of the causal chain; references [9–11,13,14,44,45,53,58,63] are cited.

The illustration in Figure 2 is a depiction of systemic interactions for the post-quantum DNSSEC transition. It includes PQC signature size, transport limitations (UDP constraints, IP fragmentation, TCP fallback), security vulnerabilities (DDoS amplification, off-path downgrade [15], resolver failures [46]), compounding factors (KeyTrap $\times$ PQC cost [40], algorithm agility lock-in [61], silent validation failure [43]), and deployment ramifications (Mosca urgency gap [9], deployment deadlock [44,45], lack of forward secrecy [16]). No single proposal takes care of everything at once. The specified primary sources are the source of quantitative labels. BAF data are obtained from [16]; the fragmentation failure rate is obtained from [46]; the KeyTrap CPU effect is referenced in [40]; the resolver false-positive rate is extracted from [43].




Figure 2. PQC-DNSSEC transition systemic interaction diagram: from PQC signature size through transport constraints and security risks to deployment barriers. Sources: [8–11,13,14,40,44,45,53,58,63].

Five families: Transport-layer, fragmentation-based, signatureless, hybrid, and protocol-redesign. Properties annotated from Tables 5A, 5B, and 6. Sources: [12,16,18,45,47–49,53,54,67].

Figure 3: Classification of PQC-DNSSEC mitigation strategies. It includes five categories of proposals: transport-layer methods (TCP fallback [12], out-of-band key distribution [12]); fragmentation-based methods (ARRF [49], QBF [53], compact denial-of-existence [18]); signatureless authentication (SL-DNSSEC [16]); hybrid methods (Double-Signed Fragmented DNSSEC [67], MTL Mode [47]); and protocol-redesign methods (KEM-MAC architecture [16]). Every family is accompanied by its essential operational characteristics extracted from Tables 6 and 7. Proposals can belong to several families; SL-DNSSEC fulfills both signature-free authentication and protocol-revision requirements.




Figure 3. Taxonomy of PQC-DNSSEC mitigation approaches, classifying eight reviewed proposals into five families with their operational characteristics and standardization readiness. Colors denote the proposal families—blue: transport-layer; orange: fragmentation-based; green: signatureless authentication; purple: hybrid; red: protocol-redesign—and the navy header box denotes the taxonomy root. Sources: [12,16,18,45,47–49,53,54,67].

5.2. Algorithm Evaluation and Selection

The first systematic evaluation of NIST Round 3 PQC candidate algorithms in the DNSSEC operational context was carried out by Müller et al. [12], who established three main selection criteria that are either implicitly or explicitly calibrated against all subsequent algorithm assessments in the reviewed literature. The first criterion deals with the UDP payload size threshold, the total size of any DNS answer that must be transmitted without TCP fallback, or fragmentation is limited by the 1232-byte maximum operational limit, which was calculated in Section 2.4 [12]. The second criterion specifies validation throughput: a medium-sized resolver requires a minimum of 1000 signature validations per second to sustain production query loads [12]. The third criterion addresses signing throughput: a minimum of 100 signatures per second is required for large zones to support continuous re-signing and key rollover operations [12].

Applying these criteria to seven evaluated candidates, Müller et al. [12] established that only three, Falcon-512, Rainbow-Ia, and RedGeMSS128, satisfy the size constraint. Rainbow-Ia and RedGeMSS128 were subsequently broken by classical cryptanalysis and did not advance beyond Round 3 [12]; Falcon-512 therefore remains the sole candidate satisfying all three criteria among the algorithms evaluated. Müller et al. [12] also identified TCP fallback, out-of-band key distribution via HTTP or FTP, and TTL extension as potential protocol-level adaptation pathways to accommodate algorithms that exceed the UDP limit proposals evaluated in Section 5.4.

Current NIST standardization status: Three PQC signature and key encapsulation schemes, ML-KEM (FIPS 203, originally Kyber), ML-DSA (FIPS 204, formerly Dilithium), and SLH-DSA (FIPS 205, formerly SPHINCS+), were fully standardized by NIST in August

2024. Falcon is being standardized under the FN-DSA designation [44]. According to Schutijser et al. [45] and Rose et al. [18], the formal IETF specification needed to implement any of these algorithms within DNSSEC is still lacking as of this writing. The remaining standardization barrier examined in Section 6.5 is the lack of a finalized IETF specification and registered IANA algorithm identifiers. NIST standardization eliminates the cryptographic barrier to PQC integration but does not establish operational authorization.

Stateful hash-based signatures: Stateful hash-based signature schemes, specifically LMS (Leighton-Micali Signatures) and XMSS (eXtended Merkle Signature Scheme), both standardized in RFC 8391 and NIST SP 800-208 [18], offer security guarantees grounded solely in well-understood hash function security assumptions. However, Müller et al. [12] explicitly exclude stateful schemes from consideration for DNSSEC on the basis of a fundamental operational incompatibility: these schemes require the signer to maintain a counter tracking its position in a Merkle tree across all signing operations, and reuse of any counter value whether through implementation error or concurrent signing by independent signers renders the entire scheme insecure under its formal security definition [18]. Various authoritative nameservers regularly sign zone records independently and concurrently in operational DNSSEC implementations; without coordination mechanisms that are unprecedented in DNS architecture, it is impossible to guarantee the uniqueness of counters among distributed signers during network partition scenarios. As a result, the only realistic options for DNSSEC are the stateless PQC signature techniques Falcon-512, ML-DSA, and SLH-DSA, a conclusion supported by all four studies reviewed in this section [12].

Performance evaluation in four different investigations: The implementation of Falcon-512, CRYSTALS-Dilithium, and SPHINCS+ in CoreDNS was examined by Suela et al. [44], who confirmed that all assessed post-quantum cryptography methods regularly exceed the 1232-byte UDP payload limit. Since the ML-DSA signature at NIST Security Level I is 263% stronger than the Falcon-512 signature [44], any ML-DSA implementation must use TCP fallback or fragmentation because it is technically impossible to transmit it in a single UDP datagram. Raavi et al. [68] assessed PQC algorithm performance in the context of TLS and PKI integration, providing complementary throughput data applicable to DNSSEC resolver contexts. Falcon-512 validation throughput was measured at 20,228 operations per second [68], substantially exceeding the 1000 sig/s minimum established by Müller et al. [12], confirming that computational performance does not constitute a barrier to DNSSEC integration for Falcon-512 and that the size constraint, not the speed constraint, is the binding limitation.

Schutijser et al. [45] evaluated Falcon-512 and MAYO-2 a NIST additional signature candidate producing 186-byte signatures with a 4912-byte public key from the perspective of TLD operators, using the .nl, .se, and .nu ccTLD zone files. Non-parametric statistical methods (Kruskal–Wallis H-test and Dunn’s pairwise post hoc test) were applied, as the sample size of 16 observations precludes assumption of a normal distribution. Performance ratios across the three TLDs were not statistically significantly different for the NSEC signing case ($p > 0.05$), indicating consistent behavior across zone file characteristics. On x86-64-v3 processors with AVX2 instruction set support, Falcon-512 signing performance is comparable to RSA-1280 (986 ± 32.5 s versus 1046 ± 37.1 s) [45]. On x86-64-v2 processors without AVX2 support, Falcon-512 is 20.3 times slower than ECDSA-P256 (9216 ± 233.3 s versus 458 ± 9.2 s) [45], a critical deployment limitation for the substantial fraction of DNS infrastructure that does not support AVX2 acceleration. MAYO-2’s 4912-byte public key independently exceeds the 1232-byte UDP threshold, requiring DNSKEY records to be served over TCP regardless of the RRSIG size [45].

The collective finding across these four studies is that no currently available PQC algorithm provides a fully adequate drop-in replacement for DNSSEC without protocol-

level modification; Falcon-512 is the only candidate satisfying the size constraint under minimal response conditions, yet even Falcon-512 requires fragmentation or TCP fallback for NXDOMAIN responses and multi-RRSIG queries, as demonstrated empirically in Section 5.3. This conclusion motivates the alternative approaches reviewed in Sections 5.4–5.7.

5.3. Empirical Analysis of Falcon Signature Formats

Fabrizio et al. [54] conducted an empirical comparison of the padded (666 bytes, fixed) and compressed formats of Falcon-512 signatures, drawing on 59,312,128 NOERROR and 113,348,878 NXDOMAIN queries recorded in the .nl ccTLD. As reported directly in Figure 1 of [54], the median size of the compressed format is 655 bytes, with a standard deviation of 2.55 bytes. Among 9,912,711 signatures, only a single 667-byte instance was observed [54], indicating that in practice the size overhead of the padded format is negligible. As provided in Figure 3 of [54], NXDOMAIN responses containing two NSEC3 records range between 2269 and 2420 bytes, while those containing three NSEC3 records range between 3075 and 3767 bytes, in all cases substantially exceeding the 1232-byte threshold. The study also demonstrated that domain name length variation is comparable in magnitude to the size difference between the padded and compressed formats; at the same time, it was established that format-level optimization is insufficient to resolve the NXDOMAIN response size problem [54]. Accordingly, the authors [54] recommend the padded format for standardization purposes, as its fixed, predictable size reduces the probability of implementation errors.

5.4. Protocol-Level Modification Proposals

The most operationally conservative response to the UDP size incompatibility established in Section 4.1 is to adapt existing DNS protocol mechanisms rather than introduce new ones. Three such adaptations appear in the reviewed literature: TCP fallback, out-of-band key distribution, and compact denial-of-existence responses. While each addresses a specific aspect of the size problem, none resolves the structural incompatibility between PQC signature sizes and the DNS transport layer.

TCP fallback: Müller et al. [12] proposed mandatory TCP fallback support as the primary protocol-level adaptation pathway, building on the observation established in Section 4.1 that a significant fraction of authoritative nameservers lacked TCP support at the time of their evaluation. Even in cases where TCP is available, Rawat and Jhanwar [16] confirm that fallback results in a minimum 50% latency penalty because the PQC signature size structurally exceeds the 1232-byte EDNS buffer, wasting the first UDP round-trip on a truncated response before the TCP connection can be established. This penalty is independent of network conditions and cannot be eliminated by implementation optimization, because it arises from the mandatory protocol event sequence specified by DNS rather than from any implementation choice. The full latency impact of TCP fallback under PQC response sizes, including the interaction with the TCP Slow Start algorithm for SLH-DSA-128s responses, is quantified in Section 6.1.2.

Out-of-band key distribution: As noted in Section 5.2, Müller et al. [12] additionally proposed distributing large PQC public keys through out-of-band channels such as HTTP or FTP to reduce DNSKEY record sizes within the DNS protocol. Rawat and Jhanwar [16] evaluate this approach and confirm that it introduces approximately 30% additional resolution overhead from the secondary round-trip and connection establishment cost. More fundamentally, this approach leaves the signature size problem unresolved; RRSIG records constitute the dominant contribution to per-response size under PQC schemes and must be transmitted within the DNS response regardless of how public keys are distributed. Out-of-

band key distribution therefore reduces but does not resolve the transport incompatibility and introduces an operational dependency on secondary distribution infrastructure with no precedent in the DNS architecture.

Compact denial-of-existence: Rose et al. [18] note that an IETF draft for compact denial-of-existence responses aims to reduce NXDOMAIN response sizes by minimizing the number of NSEC or NSEC3 records required per non-existence proof. As established in Section 5.3, NXDOMAIN responses represent the most size-critical response category under PQC schemes [54]; however, the compact denial-of-existence mechanism addresses the NSEC record contribution to response size rather than the PQC signature contribution, which is the structurally dominant term. Standardization of this mechanism remains in progress at the time of writing [18], and its interaction with PQC RRSIG overhead has not been evaluated.

Taken together, these three protocol-level modifications address peripheral contributions to the DNS response size problem while leaving its structural core unresolved: ML-DSA-44 signatures at 2420 bytes per RRSIG record exceed the 1232-byte UDP limit independently of all other response components, and no combination of TCP fallback, key distribution optimization, and denial-of-existence compaction can bring a multi-RRSIG PQC response within that threshold. This ceiling motivates the application-layer fragmentation schemes reviewed in Section 5.5, which address the signature size directly by distributing it across multiple DNS exchanges.

5.5. Fragmentation-Based Approaches

5.5.1. ARRF—A Resource Record Fragmentation

Goertzen and Stebila [49] proposed ARRF, a request-based application-layer fragmentation mechanism in which all fragments are transmitted only upon explicit request, using custom RRfrag pseudo-resource records. The resolver uses the initial response as a map for subsequent fragment requests [49]. The empirical evaluation was conducted on Amazon EC2 (c5.2xlarge 8-core Intel Xeon Platinum 8124 M, 16 GB RAM) and assessed both Sequential and Parallel ARRF variants across three network scenarios; the results are presented graphically in Figures 3–6 of [49].

The study established three principal conclusions [49]. First, Parallel ARRF outperforms TCP fallback in all evaluated scenarios, as parallel fragment requests constitute a single additional round-trip. Second, Sequential ARRF outperforms TCP fallback only for Falcon-512; for Dilithium-2-2 and SPHINCS $\pm$ 128s under 100 ms latency conditions, Sequential ARRF produces substantially worse resolution times, approximately 1604 ms and 4207 ms, respectively, owing to the latency-amplifying effect of sequential round-trips [49]. Third, under constrained bandwidth conditions (128 KB/s), Parallel ARRF outperforms Standard DNS for all evaluated algorithms [49]. As stated directly in the abstract of [49], total data transmission is lower under ARRF than under Standard DNS for both Falcon-512 and Dilithium2; specifically for Falcon-512, ARRF transmits 2557 bytes compared to 3112 bytes under Standard DNS [49]. Three principal limitations are identified [49]. First, RRfrag is a non-standard pseudo-resource-record type that carries no IANA allocation and is unknown to any resolver or middlebox not specifically patched to understand it. Operational DNS infrastructure is documented to handle unknown record types poorly: firewalls, load balancers, and DNS-aware inspection devices that enforce strict parsing of known types frequently drop, truncate, or fail to respond to messages containing types they do not recognize, rather than passing them through unmodified [52]. This exposure is particularly acute for ARRF because request-based fragmentation is invoked precisely when responses are large, which is exactly the condition under which oversized or unusual DNS traffic is most likely to encounter middlebox interference so ARRF risks resolution

failure on the very network paths it is intended to traverse, and does so silently, since a dropped RRFragment message is indistinguishable to the resolver from ordinary packet loss. Second, request-based fragmentation obliges the resolver to allocate and retain states for each partially reassembled resource record until all of its fragments have arrived, introducing a stateful reassembly buffer that classical single-response DNS does not require. The cost of this state is asymmetric: an adversary can announce a large fragment map, prompting the resolver to reserve reassembly buffers, and then never deliver the remaining fragments, so that memory is consumed on the resolver side at negligible cost to the attacker, a memory-exhaustion vector that the authors explicitly acknowledge as unresolved [49]. Compounding this, Raavi et al. [50] demonstrate that when multiple resolutions are in flight concurrently, fragments belonging to different responses can be mis-associated during reassembly, allowing an attacker to graft a fragment of its own into a victim resolution and constituting a distinct attack surface specific to fragmented post-quantum DNSSEC. Third, ARRF behavior under packet loss and unreliable network conditions remains under active investigation [49]. Because each fragment is requested and delivered as an individual exchange, the probability that at least one fragment of an n -fragment response is lost grows with n , and the loss of any single fragment leaves the resolver holding an incomplete record set that cannot be validated, and an RRSIG cannot be checked against a partially assembled RRset. Unlike TCP, the DNS application layer provides no retransmission or reordering semantics of its own, so recovery from such a loss depends entirely on the resolver re-requesting the missing fragment, adding further round-trips precisely in the degraded network conditions where round-trips are most expensive. Caching Implications of Fragmentation: An important operational consideration not addressed in [49] or [16] concerns the interaction between application-layer fragmentation and DNS caching. In standard DNSSEC, a resolver cache has complete, validated RRsets. Under ARRF or QBF, however, the resolver must reassemble fragmented records before validation occurs. For the protocol logic it can be concluded that incomplete fragment sets arising from packet loss or middlebox interference leave the resolver in an indeterminate state with respect to cache population. This class of caching failure modes is absent from standard DNSSEC and has not yet been empirically quantified at scale; accordingly, it is identified as a research gap in Section 6.

5.5.2. QBF—QNAME-Based Fragmentation

Rawat and Jhanwar [53] proposed QBF, a fragmentation scheme that addresses the ARRF limitations by employing standard DNS record types. Signature bytes are stored within RRSIG and DNSKEY records, while the fragment number is encoded in the QNAME field using the format `?n?domain -name` [53]. The principal advantages of QBF over ARRF are that [53] the use of standard DNS record types eliminates the risk of middlebox discarding; backward compatibility is automatic; the mechanism proves to be resistant to memory depletion attacks; reassembly requires only a single round-trip. Rawat and Jhanwar [53] confirm that, for Falcon-512, QBF achieves a resolution time comparable to standard ECDSA-based DNSSEC (45 ms versus 44 ms). Under a 1% packet loss rate, Rawat and Jhanwar [53] calculate that the probability of SPHINCS+-128s requiring an additional round-trip is approximately 0.37, derived from the approximate 46 DNS packets in total exchanged during a SPHINCS+-128s QBF session [53]; this implies that statistically one in three sessions requires an additional round-trip.

5.5.3. Double-Signed Fragmented DNSSEC

Pan et al. [67] proposed a hybrid security strategy for the transitional period, combining a pre-quantum and a post-quantum signature within a single DNS response. Drawing

on ENISA recommendations, the authors [67] argue that PQC algorithms have not yet undergone sufficient cryptanalytic scrutiny to justify the exclusive reliance on them. The response is considered valid only if both signatures are legitimate, offering protection against three adversary scenarios: simultaneous breaking of both, which is probably impossible; classical-computer assault on the PQC signature; and quantum-computer attack on the classical signature [67]. The BIND9 resolver must be changed in order to allow validation of both signatures [67].

A FALCON + ECDSA double signature introduces a resolution overhead of only 8.3% compared to FALCON alone (205.9 ms versus 190.1 ms), while FALCON + RSA introduces a 7.6% overhead (204.5 ms), according to an empirical evaluation [67] carried out on Amazon EC2 t2.xlarge (50 Mbps, 10 ms latency). The overhead is considerably smaller for DILITHIUM combinations: DILITHIUM + RSA adds 2.9% (220.7 ms) and DILITHIUM + ECDSA adds 5.2% (225.6 ms against 214.5 ms) [67]. At the fragment level, a double signature adds only one additional fragment for FALCON and DILITHIUM queries of QTYPE A (three and eight fragments, respectively, compared to two and seven for single signatures), while the fragment count for SPHINCS+ remains unchanged at 23 [67].

The analysis of all three fragmentation-based approaches ARRF [49], QBF [53], and Double-Signed DNSSEC [67] shows that each of them creates a potential UDP flooding risk due to parallel packet transmission and fails to resolve the DDoS amplification problem, as the response size exceeds the query size in all cases [16]. These limitations motivate the fundamentally different approach discussed in Section 5.7.

5.6. Hybrid Approaches

The IETF draft specification for SLH-DSA in Merkle Tree Ladder (MTL) Mode for DNSSEC [47] proposed a strategy in which a conservative algorithm signs Merkle tree ladders and individual messages are verified using authentication paths. Currently there is no empirical evaluation of this approach within an operational DNSSEC environment [47]; the proposal remains at the stage of theoretical analysis, a gap being identified as a research limitation, which is discussed in Section 6.

MTL Mode Operational Complexity: The main reason for the empirical evaluation absence concerns the operational complexity of Merkle tree management in dynamic DNS zones. In a static zone, the set of signed records is fixed and a Merkle tree can be pre-computed; however, DNS zones are frequently updated as records are added, modified and removed, which requires continuous recomputations of tree nodes and ladder structures. This constraint is particularly significant for large TLDs, where zone updates occur at high frequency. This restriction is acknowledged in the IETF draft specification [47], which also points out that MTL Mode works best in zones with infrequent updates, where the amortization of signature size over several records offers the biggest advantage. Additionally, the pre-computation approach that underpins MTL Mode is incompatible with the online signing necessary for NSEC3 white lies and black lies countermeasures against zone wandering [47]. Despite its potential benefits in lowering the effective signature size, MTL Mode has not yet been proven in a real-world DNSSEC setting due to these operational limitations.

Additionally, the mechanism requires modifications to the DNSSEC protocol and does not constitute a drop-in replacement [47]. Müller et al. [12] identified a fundamental limitation of hybrid models more broadly: the simultaneous use of classical and PQC algorithms increases DNS message size, as both signatures and both public keys must be transmitted concurrently [12]. As demonstrated by Google and Cloudflare installations, Scrivano [69] affirms that X25519 + Kyber hybrid schemes are practically deployable in

TLS and observes that this experience can be applied to the DNSSEC context during the transitional phase [69].

5.7. SL-DNSSEC Signature-Free Post-Quantum DNSSEC

All approaches discussed in Sections 5.4–5.6 offer incremental mitigation of the signature size problem but cannot achieve a fundamental resolution. Rawat and Jhanwar [16] proposed the SL-DNSSEC protocol, which performs DNSSEC validation using a KEM and MAC combination, without signatures, within a single UDP query/response exchange.

Protocol structure is specified directly in Sections 3.2–3.4 and Figure 2 of [16].

Phase 1: The zone operator generates a KEM key pair; the public key ZKK_{pk} is stored in a DNSKEY record with flags = 258 and re-signed by the KSK [16].

Phase 2: The resolver performs encapsulation [16]:

$$(ss, ct) \xleftarrow{\$} KEM.encap(ZKK_{pk}), \quad (2)$$

The ciphertext ct is appended to the additional section of the DNS query.

Phase 3: The name server performs decapsulation [16]:

$$ss \leftarrow KEM.decrap(ZKK_{sk}, ct), \quad (3)$$

$$k \leftarrow KDF(ss) \quad (4)$$

For every $RRSIG_i$ in response R, the name server constructs

$$msg := RRSIG_i \rightarrow RDATA \| RR_i(1) \| RR_i(2) \| \dots, \quad (5)$$

where $RDATA$ excludes the signature field and $RR_i(j)$ denotes the j -th resource record in $RRset_i$. The MAC tag is computed and stored in the signature field.

$$\begin{aligned} \omega_i &\leftarrow MAC_k(msg) \\ RRSIG_i &\rightarrow RDATA \rightarrow Signature = \omega_i' \end{aligned} \quad (6)$$

Phase 4: The resolver derives $k \leftarrow KDF(ss)$ from the stored shared secret and validates each $RRSIG_i$:

$$RRSIG_i \rightarrow RDATA \rightarrow Signature = \omega_i, \quad (7)$$

If all RRSIGs are verified, the response R is marked as secure.

Security Model [16]: Three attack scenarios are considered. First, DNSKEY modification is prevented by the EUF-CMA-secure signature over the RRSIG record. Second, ciphertext modification is prevented by the IND-CCA-secure KEM, as a modified ct yields a distinct shared secret, resulting in MAC verification failure. Third, response modification is prevented by the EUF-CMA-secure MAC [16].

Bandwidth Amplification Factor is defined directly in Section 2 of [12].

$$BAF = \frac{\text{Number of bytes received (RX)}}{\text{Number of bytes sent (TX)}}, \quad (8)$$

The empirical results, reproduced directly from Tables 11 and 12 of [16], are presented in Table 5 below. No re-execution or independent replication of these measurements was performed in the present review; the values are cited as reported by the original authors, combined with their described experimental configuration (AWS EC2 instances; NIST Level I parameter sets).

Table 5. SL-DNSSEC bandwidth and resolution results—NIST Level I.

Method	TX (Bytes)	RX (Bytes)	BAF	Total (Bytes)	100 Mbps/10 ms (ms)	1 Mbps/100 ms (ms)
ECDSA-SD	84	512	6.1	596	44	407
RSA-SD	84	1088	13	1172	44	408
Falcon-SD	598	2700	-	3298	89	811
Dilithium-SD	928	8292	-	9220	89	817
SPHINCS+-SD	1654	25,389	-	27,043	111	1025
Falcon-QBF	258	2788	10.8	3046	45	410
Dilithium-QBF	693	9225	13.3	9918	46	415
Kyber-HMAC (SL-DNSSEC)	867	416	0.48	1283	44	408

These results demonstrate that SL-DNSSEC (Kyber-HMAC) reduces bandwidth consumption by up to 95% in relation to SPHINCS+-SD, 86% in relation to Dilithium-SD, and 58% in relation to Falcon-SD [16]. Resolution speed is practically equivalent to RSA-2048 (44 ms versus 44 ms).

Crucially, the DDoS amplification attack vector is substantially reduced when $BAF = 0.48$: since the response is smaller than the query, an amplification attack using SL-DNSSEC would actually reduce traffic at the victim rather than amplify it since the answer is less than the query [16].

According to Tables 16 and 17 of [16], the BAF further drops to 0.29 at NIST Level V (Kyber 1024 + HMAC SHA 512), showing higher amplification resistance. Nevertheless, the Level V KEM ciphertext (1568 bytes) is above the 1232-byte threshold, necessitating QBF [16] a restriction that does not exist at Level I. According to Appendix Table A14 of [16], computational performance verifies that Kyber-512 encapsulation runs at 130,503 operations per second and HMAC-SHA-256 operates at 8,398,588 operations per second, both significantly above any realistic DNS server workload.

NIST Security Levels in the DNSSEC Framework: Beyond theoretical security margins, the choice between Level I and Level V has practical implications. NIST Level I is made to withstand attacks by a quantum adversary with the same amount of processing power as an AES-128 key search [70]. As of this writing, NIST's evaluation indicates that there is not a cryptanalytically relevant quantum computer (CRQC) that can compromise Level I security, and this is not expected to change very soon [44]. Because Level V imposes a ciphertext size (1568 bytes) that exceeds the UDP threshold and introduces additional fragmentation complexity without any corresponding near-term security benefit, Rawat and Jhanwar [16] recommend Level I as the appropriate target for current DNSSEC deployments. The suggestion adheres to and validates the idea that rather than being unconditionally maximized, cryptographic security levels should be tuned to realistic threat timelines [16].

Non-Repudiation in the DNSSEC Context: It is important to note that non-repudiation has never been a stated security goal of DNSSEC. RFC 4033 clearly defines the security objectives of the protocol as data origin authentication and data integrity verification and does not include non-repudiation among its security services [7]. In the legal sense of providing irrefutable evidence of signing, non-repudiation is not required for DNS resolution and is not provided by any of the DNSSEC algorithms that are deployed currently. Therefore, absence of non-repudiation in SL-DNSSEC does not represent a regression relative to the protocol's defined security objectives, though it does preclude use cases such as forensic DNS audit trails that may require asymmetric proof of signing [16].

In Appendix A.3 Rawat and Jhanwar [16] confirm that SL-DNSSEC naturally supports online signing, as the MAC is computed on-the-fly—a particular advantage for the NSEC3 zone walking countermeasures (white lies and black lies) that are discussed in

Section 4.3, which require on-the-fly signature generation. The identified limitations [16] include that initial DNSKEY retrieval exceeds 1232 bytes, requiring TCP or fragmentation; HSM deployment is recommended for KEM private key protection; and SL-DNSSEC has not yet been standardized within the IETF.

5.8. Synthesis Comparative Analysis of Post-Quantum DNSSEC Proposals

The proposals discussed in Section 5 present three principal strategies: algorithm-level optimization, application-layer fragmentation, and signature-free authentication. The collective findings of the reviewed literature give four principal conclusions.

First, simultaneous resolution of the UDP size constraint and the DDoS amplification problem is achieved exclusively by SL-DNSSEC, which produces a BAF of 0.48 at NIST Level I and 0.29 at Level V [16]—the only evaluated approach in which the response is smaller than the query. All other approaches leave at least one of these two challenges: unresolved fragmentation-based schemes (ARRF [49], QBF [53], Double-Signed DNSSEC [67]) address the size constraint but retain a response size exceeding the query size, while protocol-level modifications (TCP fallback, out-of-band key distribution [12]) do not address either of the problems at a fundamental level.

Second, SL-DNSSEC does not provide non-repudiation [16], although this limitation does not represent a regression in relation to DNSSEC's defined security objectives of data origin authentication and data integrity [7]. Non-repudiation is not a stated requirement of any currently deployed DNSSEC implementation and becomes relevant only in specific forensic or audit contexts that fall outside the protocol's core operational scope.

Third, the research-stage standardization status of the most innovative approaches, namely SL-DNSSEC [16], MTL Mode [47], and Double-Signed DNSSEC [67], indicates that currently there exists no standardized PQC integration solution for DNSSEC. ARRF [49] and QBF [53] have progressed to IETF draft status, while Falcon-512 under the FN-DSA designation remains in the process of NIST standardization [44]. The absence of a completed IETF specification for any PQC algorithm in DNSSEC [18] we are facing currently means that operational deployment at scale is not yet feasible regardless of the chosen approach.

Fourth, the combination of Falcon-512 with Parallel ARRF [49] or QBF [53], supplemented by the Double-Signed scheme [67] where practicable, represents the most reliable incremental migration pathway at this stage, as these approaches are closest to standardization and have been empirically validated in operational environments. SL-DNSSEC [16] serves as the long-term target, while the DNSKEY size problem resolution and IETF standardization are pending.

The unresolved challenges discussed in Section 6 arise directly from these findings: the DNSKEY size problem in SL-DNSSEC [16]; the ARRF [49] memory exhaustion vulnerability; the caching failure modes introduced by fragmentation-based schemes [43,49]; the requirement for BIND9 modification in Double-Signed DNSSEC [67]; the absence of empirical evaluation of MTL Mode in an operational DNSSEC environment, especially with respect to dynamic zone management [47]; and the lack of current TCP adoption statistics following the DNS Flag Day 2020 initiative [12].

5.9. Formal Evaluation Framework

Cross-Proposal Comparison Across Nine Operational Criteria: Section 5.8 synthesizes four conclusions about the comparative advantages of post-quantum DNSSEC approaches, although it does so without maintaining consistent dimensions throughout the proposals. This part employs a consistent nine-criterion evaluation approach for all proposals examined in part 5, facilitating direct comparison for standardization and deployment decisions. Each criterion pertains to a specific operational or security requirement that any production

DNSSEC implementation must fulfill; these criteria are directly derived from the analytical dimensions outlined in Sections 4 and 6, as well as the research questions articulated in Section 3.1.

The following lists the nine criteria along with the operational definitions that go with them. As required by the DNS Flag Day 2020 operational restriction, the proposal may transmit a full, verified DNSSEC response for a QTYPE A query, including three RRSIG records, in a single 1232-byte UDP datagram. According to Herzberg and Shulman [56] and Van den Broek et al. [46], fragmentation resistance is the extent to which the technique completely avoids IP-layer fragmentation, hence reducing the chances of cache poisoning, middlebox obstruction, and reassembly failures. TCP Contingency Dependency: Classified as LOW (no TCP required), MEDIUM (TCP required only for exceptional response types), or HIGH (TCP critical for all non-trivial responses) depending on how much the proposal needs TCP fallback for standard functionality. According to Rawat and Jhanwar's BAF tests, the DDoS Amplification Exposure approach yields a Bandwidth Amplification Factor (BAF) of less than 1.0 (PASS), a BAF similar to the present ECDSA DNSSEC (HIGH), or an intermediate value (MEDIUM/PARTIAL) [16]. (Downgrade Attack Resistance, anchored as follows: PASS requires that the proposal structurally eliminates the algorithm-selection downgrade vector identified by Heftrig et al. [13–15]; that is, the validation outcome cannot depend on an attacker-influenceable choice among algorithm identifiers, whether through a single mandated algorithm or an authenticated negotiation mechanism; PARTIAL denotes that the proposal inherits the baseline DNSSEC algorithm agility vector unchanged, neither mitigating nor enlarging it, so that downgrade exposure is equivalent to that of currently deployed ECDSA DNSSEC; FAIL denotes that the proposal enlarges the vector relative to this baseline, for example by introducing an additional unauthenticated negotiation surface. Downgrade enlargement mediated specifically by IP fragmentation (the off-path tag-modification variant [15]) is scored under Fragmentation Resistance rather than here, to avoid double-counting a single structural cause across two criteria. Forward Secrecy Support: As covered in Section 6.2, the plan incorporates a per-session key freshness mechanism that reduces the possibility of retroactive forging in the event of a future ZSK compromise. Fourth, deployment complexity—the operational burden placed on zone operators and resolver operators—is classified as LOW (drop-in replacement), MEDIUM (protocol extension required), or HIGH (substantial DNS-software updates required). Resolver Compatibility: the degree to which the proposal functions with unmodified production resolvers, including those that do not support new algorithm identifiers. Standardization Readiness: the proposal's current standing in the IANA algorithm registry and the IETF DNSOP working group, classified as PASS (RFC-track with registered algorithm numbers), PARTIAL (IETF draft or research phase), or FAIL (lack of official standardization activities).

Ratings are divided into three categories: FAIL/HIGH (criterion not satisfied or exposure is heightened, red), PARTIAL (criterion partially satisfied or satisfied under limited conditions, yellow), and PASS (criterion entirely satisfied, green). A LOW rating indicates a PASS in the context of directional criteria (TCP fallback dependency, DDoS amplification exposure, deployment complexity), while a HIGH grade indicates a FAIL. The number of criteria for which the proposal is rated as PASS or LOW is shown by the overall score.

Table 6 summarizes four results that the per-dimension analysis in Sections 5.4–5.7 did not present in a comparative format. The only proposal that simultaneously achieves a PASS rating for DDoS amplification vulnerability and UDP compatibility is SL-DNSSEC [16], which also offers a unique level of forward secrecy (PARTIAL, at the session level). Although SL-DNSSEC has the greatest overall score of 6/9 out of all the proposals that were examined, it falls short of achieving full PASS in terms of downgrade resistance, resolver

compatibility, and standardization readiness. The uniformity of the Downgrade Resistance column—PARTIAL for all eight proposals, including SL-DNSSEC—reflects the fact that no reviewed proposal alters DNSSEC’s algorithm-agility semantics in either direction; as the sensitivity analysis in Section 5.10 makes explicit, a criterion on which all proposals score identically is rank-inactive and cannot have influenced the comparative outcome. The conclusion of Section 5.8 that no current solution concurrently satisfies all operational needs is supported by the fact that no option receives a score of 7/9 or above.

Table 6. Formal framework for cross-proposal assessment of post-quantum DNSSEC proposals using nine operational criteria. Cell shading indicates the degree of criterion satisfaction: green denotes a favorable rating (PASS, or LOW for TCP fallback dependence and deployment complexity); yellow denotes an intermediate rating (PARTIAL or MEDIUM); red denotes an unfavorable rating (FAIL, or HIGH for TCP fallback dependence, DDoS amplification exposure, and deployment complexity). The textual label in each cell carries the authoritative rating; color is applied redundantly to the text so that the table remains fully interpretable for readers with color-vision deficiencies.

Proposal	UDP Compat.	Frag. Resist.	TCP Fallback Dep.	DDoS Amp. Exposure	Downgrade Resist.	Forward Secrecy	Deploy Complexity	Resolver Compat.	Std. Readiness	Overall (Pass/9)
TCP Fallback [12,51]	FAIL	FAIL	HIGH	HIGH	PARTIAL	FAIL	LOW	HIGH	PARTIAL	2/9
EDNS(0)/Out-of-Band Keys [12]	FAIL	FAIL	HIGH	HIGH	PARTIAL	FAIL	MEDIUM	HIGH	PARTIAL	2/9
ARRF [49]	PASS	PARTIAL	MEDIUM	HIGH	PARTIAL	FAIL	MEDIUM	PARTIAL	PARTIAL	3/9
QBF [53]	PASS	PARTIAL	LOW	HIGH	PARTIAL	FAIL	LOW	HIGH	PARTIAL	4/9
Double-Signed DNSSEC [67]	PARTIAL	PARTIAL	MEDIUM	HIGH	PARTIAL	FAIL	MEDIUM	PARTIAL	PARTIAL	3/9
MTL Mode (SLH-DSA) [47]	PARTIAL	PARTIAL	LOW	HIGH	PARTIAL	FAIL	HIGH	PARTIAL	PARTIAL	3/9
SL-DNSSEC [16]	PASS	PASS	LOW	PASS	PARTIAL	PARTIAL	MEDIUM	PARTIAL	PARTIAL	6/9
Falcon-512 (Direct) [45,54]	PARTIAL	PARTIAL	HIGH	HIGH	PARTIAL	FAIL	MEDIUM	HIGH	PARTIAL	3/9

The two conditions on which every proposal fails are upfront secrecy and readiness for standardization. The failure of forward secrecy is inherent to all signature-based systems and is only partially addressed by SL-DNSSEC, as detailed in Section 6.2.4. The failure of standardization preparedness in all proposals indicates the lack of a designated IETF DNSOP working group item and the absence of registered IANA algorithm numbers for any PQC scheme, as outlined in Section 6.5.1. These two universal failures delineate the essential circumstances that any imminent deployment decision must openly acknowledge as residual risks, as no existing strategy mitigates either.

Third, among the proposals that maintain the digital-signature chain of trust (all except SL-DNSSEC), QBF [53] attains the highest overall score (4/9), primarily due to its avoidance of custom resource record types that may be disregarded by middleboxes, and its single-round-trip fragmentation design reduces TCP fallback reliance compared to ARRF and Double-Signed DNSSEC. While clearly acknowledging that its DDoS amplification vulnerability (BAF = 10.8 for Falcon-512-QBF [16]), vulnerability to downgrade attacks, and lack of forward secrecy remain unresolved, QBF is the most practically feasible near-term migration strategy for operators requiring the persistence of the signature-based trust model.

Correspondingly, the identification of SL-DNSSEC as the long-term architectural target is conditional: it presupposes the formal security proof and standardization milestones defined as Phase II deliverables in Section 7.3, where the recommendation is revisited with an explicit contingency path with QBF, as noted above, serving as the default signature-preserving alternative.

As shown in Table 6, the relationship between the fragmentation resistance criterion and the DDoS amplification criterion reveals a structural dependency: any proposal that

lacks fragmentation resistance also fails to lower the risk of DDoS amplification because the mechanisms that cause fragmentation-sending responses to exceed the query also create conditions for amplification. Adjusting the response to be smaller than the query, a task only accomplished by SL-DNSSEC, is the only way to address this collective failure. This dependence, which is only visible when all nine criteria are assessed concurrently, demonstrates that denial-of-service resilience and transport-layer constraints are interrelated issues that cannot be addressed separately, necessitating that any migration strategy take them into account as a single issue rather than as sequential engineering tasks.

5.10. Normalized Quantitative Scoring of Post-Quantum DNSSEC Proposals

A methodical quantitative analysis of the eight evaluated options is shown in Table 7. Table 7 assigns a score of 0 (criterion not met), 1 (criterion partially met), or 2 (criterion fully met) to each of the nine criteria. The criteria are not evenly weighted: because they directly impact a proposal's immediate deployability, UDP compatibility, DDoS amplification resistance, downgrade resistance, and standardization readiness are given double weight (up to 4 points each). Each of the last five criteria has a weight of one, and each criterion may receive up to two points. Thus, $4 \times 4 + 5 \times 2 = 26$ points is the maximum possible score. Section 7.3 roadmap's priority hierarchy is explained by the weighted scoring system: long-term operating characteristics are considered less important than rapid deployability (UDP, DDoS, downgrade, and standardization).

Table 7. DNSSEC proposal scores.

Proposal	UDP Compat. ($\times 2$)	DDoS Amp. ($\times 2$)	Downgrade ($\times 2$)	Std. Ready ($\times 2$)	Frag. Resist.	TCP Dep.	Fwd Secrecy	Deploy.	Resolver Compat.	Total/26	Tier
TCP Fallback	0	0	1	1	0	0	0	2	2	8	III
Out-of-Band Keys	0	0	1	1	0	0	0	1	2	7	III
ARRF [49]	1	0	1	1	1	1	0	1	1	10	II
QBF [53]	1	0	1	1	1	2	0	2	2	13	II
Double-Signed [67]	1	0	1	1	1	1	0	1	1	10	II
MTL Mode [47]	1	0	1	1	1	2	0	0	1	10	II
SL-DNSSEC [16]	2	2	1	1	2	2	1	1	1	19	I
Falcon-512 Direct [54]	1	0	1	1	1	0	0	1	2	10	II

Table 7 presents the normalized quantitative grading of the post-quantum DNSSEC proposals. Scores: 0 indicates that the condition is not met, 1 indicates that it is somewhat met, and 2 indicates that it is fully met. UDP compatibility, DDoS amplification resistance, downgrade resistance, and standardization readiness are the weighted criteria ($\times 2$). Tier I ≥ 17 ; Tier II = 10–16; Tier III ≤ 9 ; maximum score: 26.

The grade clarifies the differences between plans and supports the qualitative evaluation of Section 5.9. The sole proposal to reach Tier I status is SL-DNSSEC (19/26, Tier I); its superiority over the next-ranked proposal (QBF, 13/26) demonstrates its exceptional compliance with both the DDoS amplification double-weighted criterion and UDP compatibility. All of the Tier II proposals (ARRF, QBF, Double-Signed DNSSEC, MTL Mode, Falcon-512 Direct) receive scores ranging from 10 to 13. QBF receives the highest score in this category because of its exceptional TCP dependency score of 2, which indicates that TCP is not necessary for standard NOERROR responses, and its superior resolver compatibility in comparison to ARRF. TCP fallback and out-of-band keys receive Tier III ratings (8/26 and 7/26, respectively), which reflect their limited standardization status and their shortcomings in UDP compatibility, DDoS amplification, and forward secrecy. Since forward secrecy (0 for all signature-based methods) and standardization readiness (maxi-

mum 1 for all proposals given the current IANA/IETF status) impose a theoretical limit on all assessed approaches, with the exception of SL-DNSSEC, the maximum achievable score under current constraints is 20/26 for any signature-based scheme.

Sensitivity of the ranking to the weighting scheme: Because the double weighting of UDP compatibility, DDoS amplification resistance, downgrade resistance, and standardization readiness embeds a judgment about the primacy of immediate deployability, the stability of the resulting ranking was examined under four alternative weighting schemes: (i) an unweighted scheme in which all nine criteria contribute equally (maximum 18); (ii) an inverse scheme in which the five operational criteria receive double weight instead (maximum 28); (iii) a security-emphasis scheme doubling DDoS resistance, downgrade resistance, and forward secrecy (maximum 24); and (iv) a deployment-emphasis scheme doubling standardization readiness, deployment complexity, and resolver compatibility (maximum 24). SL-DNSSEC ranks first and QBF second under all four alternatives as well as under the published scheme, with the narrowest margin (16 versus 15) occurring under deployment emphasis; the Tier I position of SL-DNSSEC is therefore not an artifact of the chosen weights. Two structural observations sharpen this conclusion. First, the downgrade-resistance and standardization-readiness scores are identical across all eight proposals under the current IETF and IANA status, so their double weighting shifts every total equally and cannot influence the rank order at all; of the four double-weighted criteria, only UDP compatibility and DDoS resistance are rank-active. Second, a knockout test removing these two criteria entirely produces a tie between SL-DNSSEC and QBF (9 versus 9 on the remaining unweighted criteria), which makes the source of SL-DNSSEC's advantage explicit: it derives specifically from transport-layer and amplification performance, the constraints that Sections 4 and 6.1 establish as binding, rather than from a diffuse preference encoded in the weights.

6. Remaining Problems and Open Challenges

Despite the considerable research activity conducted since 2020, the literature on post-quantum DNSSEC reveals a set of challenges that remain basically unresolved. The following analysis is organized into five thematic areas. Each section presents a detailed account of the problem, evaluates the proposals provided in the reviewed literature and identifies the specific residual gaps. All statements given in this section are supported by verified, peer-reviewed sources.

6.1. Transport and Message Size

6.1.1. UDP Size Barrier

The 1232-byte UDP payload limit set by DNS Flag Day 2020 serves as the fundamental transport constraint for post-quantum DNSSEC. The maximum DNS payload of 1232 bytes is obtained by subtracting the IPv6 header (40 bytes) and the UDP header (8 bytes) from the IPv6 minimum MTU (1280 bytes). The first thorough examination of this constraint in the PQC context is presented by Müller et al. [12], who show that it is an impassable obstacle for all NIST-standardized post-quantum signature schemes. Their analysis of DNSSEC response composition—Answer, Authority, and Additional sections, each of which potentially contains one or more RRSIG records—shows that even the most compact PQC candidate, Falcon-512, produces a minimum QTYPE A response of approximately 1998 bytes when three RRSIG records are included, exceeding the limit by 766 bytes.

Beernink [71] uses real-world examples to expand this analysis. No existing defined PQC scheme can fit within the 1232-byte restriction without protocol modification, according to nl ccTLD traffic data from SIDN. In particular, SPHINCS+–128s (7856-byte signature) surpasses the limit by a factor of roughly 19 and Dilithium-2 (2420-byte signature) by a

factor of roughly 2.9. A minimal QTYPE A response with three RRSIGs would require about 89,376 bytes, exceeding even the theoretical DNS maximum of 65,535 bytes. The implementation of SPHINCS+–256s in DNSSEC is theoretically impossible without a significant reworking of the DNS message format due to its qualitatively unique problem.

Suela et al. [44] confirm these results in a study on an independent CoreDNS implementation and note that PQC signatures may be 10–50 times bigger than their classical counterparts, leading to proportionate increases in nameserver memory and disk space requirements. Further proof that the problem extends beyond the transport layer comes from Raavi et al. [50], who show that larger PQC response sizes can lead to DNS responses exceeding standard Ethernet frame limitations at the network layer, resulting in additional fragmentation cascades that are unrelated to the DNS Flag Day constraint. The mechanism operates one layer below the DNS Flag Day constraint. The maximum transmission unit of standard Ethernet is 1500 bytes, of which the IPv4 and UDP headers consume 28 bytes (48 bytes under IPv6), leaving approximately 1472 bytes of UDP payload that a single link-layer frame can carry [72]. A single ML-DSA-44 signature of 2420 bytes therefore exceeds the capacity of one Ethernet frame before any DNSKEY material, record data, or additional RRSIG records are counted, and an SLH-DSA-128s signature of 7856 bytes spans approximately six frames. When a UDP datagram exceeds the path MTU, IPv4 splits it into multiple IP fragments, of which only the first carries the UDP header, the property that causes stateless firewalls to drop trailing fragments while IPv6 prohibits in-network fragmentation entirely, obliging the sender to fragment at the source in response to ICMPv6 Packet Too Big signals; where such ICMP messages are filtered, path MTU discovery fails silently and the response is lost [72]. Because reassembly requires every fragment to arrive, the probability of losing the complete response increases with each additional fragment, giving rise to the fragmentation cascades reported by Raavi et al. [50]. Van den Broek et al.’s empirical findings [46] show that fragmentation is both a performance and availability problem. Their measurements show that approximately 10% of resolvers either fail to receive fragmented DNS responses or fail to reassemble them correctly, and that stateless firewalls which inspect only the first fragment of a UDP datagram routinely block subsequent fragments, since only the first fragment carries the UDP header. Rawat and Jhanwar [16] further document that approximately 10% of resolvers fail to handle IP fragments correctly, and that stateless firewalls blocking fragmented packets are commonplace in operational deployments.

The proposals put forward to address this barrier can be categorized into three classes: application-layer fragmentation, signature optimization and signature replacement. Goertzen and Stebila [49] propose Application-layer Request-based Fragmentation (ARRF), which splits large RRSIG records across multiple DNS responses requested by the resolver. Rawat and Jhanwar [53] propose QNAME-Based Fragmentation (QBF), which encodes fragment identifiers in QNAME labels to enable parallel retrieval. Pan et al. [67] propose Double-Signed Fragmented DNSSEC, which combines classical and PQC signatures to provide security during the transition period. For signature optimization, Goertzen and Stebila [49] additionally propose the Merkle Tree Ladder (MTL) mode, which amortizes a single SPHINCS+ signature across multiple DNS records by using a Merkle tree structure. For signature replacement, Rawat and Jhanwar [16] propose SL-DNSSEC, which eliminates signatures entirely in favor of a KEM-MAC construction.

Table 8 presents a systematic comparison of these proposals against the key operational requirements for post-quantum DNSSEC.

Table 8. Comparison of transport-layer proposals.

Proposal	UDP Packets	Round-Trips	BAF	Sig-Based	Backward Compat.	Primary Residual Limitation
QBF [53]	3–46	1	10.8–14.1	✓	✓	37% session failure at 1% packet loss
ARRF [49]	N (parallel)	2	>1	✓		Memory exhaustion; 2 RTT minimum
Double-Signed [67]	N	2	>1	✓	✓	2× signature overhead; mix-and-match risk
MTL Mode [47]	1	1	>1	✓	✓	Resource exhaustion attack unresolved
Out-of-band keys [12]	1	2		✓	✓	+30% resolution overhead; sig size unchanged
SL-DNSSEC [16]	1	1	0.48	✗	✓	Replaces trust model; no signature chain

Table 8 shows that all operational requirements cannot be simultaneously satisfied by any single proposal. SL-DNSSEC achieves the most favorable transport characteristics—one UDP packet per transaction with a BAF below 1—but it replaces the digital-signature-based trust model of DNSSEC with a KEM-MAC construction, modifying the chain-of-trust semantics and requiring careful security analysis of the protocol resulting from it. There has been no formal security proof published for SL-DNSSEC in the context of full DNSSEC chain of trust.

6.1.2. TCP Fallback Latency

When a DNS response exceeds the UDP payload limit, the truncated bit (TC = 1) is set and the resolver is expected to retransmit the query over TCP. In the post-quantum context this fallback mechanism imposes significant latency penalties. Rawat and Jhanwar [51] experimentally quantify these penalties using BIND 9.19.17 in a Docker environment with traffic-shaping via Linux tc. For the 100 Mbps/10 ms latency scenario, in Standard DNS mode (with TCP fallback) Falcon-512 achieves a resolution time equal to 89 ms, compared to 44 ms for ECDSA; for the 1 Mbps/100 ms scenario, resolution time rises to 811 ms, compared to 407 ms for ECDSA. The penalty is most severe for SPHINCS+128s 111 ms in the high-bandwidth scenario and 1025 ms in the low-bandwidth scenario.

The main cause of the SPHINCS+ latency penalty is the TCP Slow Start algorithm (RFC 5681). The initial congestion window (initcwnd) under RFC 6928 is 10 segments of 1220 bytes each, yielding 12,200 bytes. A SPHINCS ± 128s QTYPE A response in Standard DNS mode requires approximately 27,043 bytes, which exceeds the initial congestion window by more than a factor of two. The consequence is that two additional round-trips are required to complete the transmission, adding 200 ms at 100 ms RTT. This cascade does not affect Falcon-512, whose responses are generally within 5200 bytes, thus fitting within the initial congestion window.

Dikshit et al. [73] provide complementary empirical evidence at scale. Having analyzed more than 14 million measurements from 2527 RIPE Atlas probes across 10 public DNS resolvers, they have found that resolvers frequently fail to fall back to TCP when receiving large responses. Three of the ten resolvers that were studied prove EDNS(0) buffer sizes to be substantially larger than the 1232-byte DNS Flag Day recommendation, which increases the likelihood of IP-layer fragmentation taking place. For IPv6, DoTCP failure rates reach 10% and DoUDP failure rates 9.61%, while median TCP response times are approximately twice as high as those of UDP. In a separate large-scale study, Kosek et al. [35] confirm that DNS/TCP is 37% slower than DNS/UDP on average and that TCP failure rates at probe resolvers can reach 75%, underscoring the operational risks of relying on TCP fallback as a general solution for oversized PQC responses.

The Rawat and Jhanwar [51] TurboDNS extension partially addresses the TCP fallback penalty by pre-establishing TCP connections and using connection pooling. Tur-

boDNS substantially reduces resolution times for Falcon-512 and Dilithium-2 but does not eliminate the extra round-trips for SPHINCS $\pm$ 128s, because the problem in that case is not connection establishment overhead, but TCP Slow Start behavior in relation to response size.

Taken together, these measurements show that the latency cost of TCP fallback under PQC response sizes is not a single fixed penalty but scales with signature size through two compounding mechanisms that must be separated to be understood. The first is connection establishment: when a response exceeds the UDP payload limit the truncated bit is set and the resolver re-queries over TCP, whose three-way handshake adds a round-trip that Rawat and Jhanwar [51] measure at roughly a doubling of resolution time for Falcon-51,289 ms versus 44 ms for ECDSA at 100 Mbps/10 ms, rising to 811 ms versus 407 ms at 1 Mbps/100 ms. This component is largely amortizable: the TurboDNS extension [51] pre-establishes TCP connections and pools them, substantially reducing resolution time for Falcon-512 and Dilithium-2. The second mechanism, which pooling cannot remove, is TCP Slow Start (RFC 5681). Under RFC 6928 the initial congestion window is ten segments of 1220 bytes, or 12,200 bytes, whereas a SPHINCS $\pm$ 128s QTYPE A response in standard mode requires approximately 27,043 bytes more than double the initial window, so transmission requires two additional congestion-window-limited round-trips, adding about 200 ms at 100 ms RTT. This cascade is a function of response size rather than of connection setup, which is why it does not affect Falcon-512 (responses generally within 5200 bytes fit the initial window) and why TurboDNS cannot eliminate it for SPHINCS+, whose penalty reaches 111 ms in the high-bandwidth and 1025 ms in the low-bandwidth scenarios. Layered on top of latency is an availability risk: Dikshit et al. [73], analyzing more than 14 million measurements from 2527 RIPE Atlas probes across ten public resolvers, find that resolvers frequently fail to fall back to TCP at all (IPv6 DoTCP failure rates of 10%, median TCP response times about twice those of UDP), and Kosek et al. [35] confirm that DNS/TCP is on average 37% slower than DNS/UDP and that TCP failure rates at probe resolvers can reach 75%. The full impact of TCP fallback is therefore the combination of a size-dependent latency penalty, driven at large signature sizes by Slow Start rather than by handshake overhead, with a non-trivial probability of outright resolution failure, which is why TCP fallback cannot be treated as a general solution for oversized PQC responses.

6.2. Cryptographic Agility, Downgrade Attacks, and Forward Secrecy

6.2.1. Absence of Cipher-Suite Negotiation

Shrishak and Shulman [61] identify a structural gap in DNSSEC that directly impedes PQC deployment: the protocol provides no mechanism for cipher-suite negotiation between nameservers and resolvers. Under the current design, a nameserver has no means to determine which algorithms a resolver supports and therefore cannot safely deprecate legacy algorithms; doing so implies risks of returning SERVFAIL to resolvers that do not recognize the replacement algorithm. This creates a symmetric deadlock in which PQC deployment requires universal resolver support, and resolver support cannot be established without a deployed PQC ecosystem to validate against.

Shrishak and Shulman [61] propose both end-to-end and hop-by-hop negotiation mechanisms as solutions. End-to-end negotiation would allow a resolver to signal algorithm preferences directly to authoritative nameservers; hop-by-hop negotiation would allow intermediate resolvers to communicate algorithm capabilities to upstream nameservers. Either mechanism would break the deployment deadlock and also reduce the IP fragmentation risk, since nameservers could select algorithms, whose signatures fit within the resolver's advertised EDNS(0) buffer size. Despite the practical importance of this

proposal, no corresponding IETF DNSOP working group draft has been chartered, and the mechanism remains unstandardized.

6.2.2. Algorithm Agility as a Downgrade Vector

Heftrig, Shulman and Waidner [13] demonstrate a counterintuitive consequence of DNSSEC's algorithm: agility, the same property that makes PQC migration possible, also introduces a class of downgrade vulnerabilities. When a signed DNS zone contains an RRSIG record using an algorithm that the resolver does not recognize, certain resolver implementations do not return SERVFAIL but silently accept the response without validation instead. Therefore, an adversary which can inject a response containing an unrecognized algorithm tag can cause a DNSSEC-validating resolver to skip signature verification altogether, reducing its security to that of an unsigned zone. Heftrig et al. [13] find that 45% of DNS resolvers used by web clients are vulnerable to this class of attack.

Heftrig, Shulman and Waidner [14] extend this analysis substantially in USENIX Security 2023. Measuring the DNSSEC resolver ecosystem over time, they provide that in 2021, 65% of open resolvers, including the ones operated by Google and Cloudflare, were vulnerable to downgrade attacks exploiting algorithm agility. Following coordinated disclosure and patches to BIND, Unbound, and PowerDNS, by 2022 the vulnerable fraction was reduced to 7.5%. However, the underlying specification ambiguity that enabled the vulnerability—in the form of the absence of a mandatory SERVFAIL requirement for unrecognized algorithms—remains in the standard. The use of new PQC algorithm identifiers that are not recognizable for older resolver versions will bring this vulnerability class forth for any resolver that has not been updated in a way to handle unknown algorithm tags correctly.

To go further, Heftrig, Shulman and Waidner [15] have identified an off-path variant of the attack. By exploiting IP fragmentation, which, as provided in Section 6.1 above, is unavoidable for PQC DNSSEC responses, an adversary without a privileged network position can modify the algorithm tag byte in a fragmented RRSIG record, causing the resolver to encounter an unrecognized algorithm identifier. Out of the 43,000 DNSSEC-signed domains in the Tranco top one million that were examined by the authors, 7700 were found to be vulnerable. In a post-quantum DNSSEC deployment, in which fragmentation is unavoidable for all schemes except SL-DNSSEC structurally, this off-path attack vector becomes systematically exploitable.

6.2.3. Hybrid Signature Mix-And-Match Attacks

Pan et al. [67] offer Double-Signed Fragmented DNSSEC as a transitional security mechanism, following ENISA guidance that recommends running classical and post-quantum algorithms in parallel during the migration period. In order to provide security against both conventional and quantum adversaries, each DNS record set is signed using both a PQC method (like Falcon-512) and a classical approach (like ECDSA P-256). This strategy makes a significant contribution to transitional security and is operationally sound.

However, a class of threats introduced by dual-signature schemes is identified by Fregly et al. [48]. Before a DNS answer reaches the resolver, one of the two RRSIG records can be selectively removed by an on-path attacker. If only the classical RRSIG is retained, the response remains insecure against a quantum adversary; if only the PQC RRSIG is retained and the resolver does not yet support the PQC algorithm, the resolver may entirely skip validation (per the vulnerability described in Section 6.2.2). Speaking more generally, any combination of RRSIG additions and removals that leaves the resolver with a syntactically valid but cryptographically weakened response constitutes a successful mix-and-match attack. No formal security model has been published for hybrid DNSSEC

that would address adversarial manipulation of RRSIG sets; no security proof exists for any dual-signature DNSSEC construction either.

6.2.4. Fundamental Absence of Forward Secrecy

A definitional clarification is required before this property is analyzed. Forward secrecy is classically defined for session-oriented key-exchange protocols: in TLS 1.3, ephemeral Diffie–Hellman key exchange guarantees that compromise of a long-term private key does not expose the confidentiality of previously recorded sessions. DNSSEC, by contrast, is a signature-based authentication protocol: it establishes no sessions, negotiates no keys, and provides no confidentiality, so the TLS notion of forward secrecy does not transfer directly. Throughout this paper, the term is therefore used in its signature-domain sense: retroactive protection of authentication guarantees against future key compromise, formalized in the literature as forward security for digital signatures [74], in which signing keys evolve over time so that compromise of the current key does not render signatures attributed to earlier periods forgeable. It is precisely this property whose absence constitutes the harvest-now–decrypt-later exposure (in the DNSSEC context more accurately “harvest now, forge later”, since the archived material is public and the risk is retroactive forgeability rather than decryption) that DNSSEC structurally lacks. RRSIG records are distributed in plaintext over the internet for their validity period duration, which is typically from 14 to 30 days. The signed material—DNS record data, TTL, and signer name—is public, and the corresponding Zone Signing Key (ZSK) public key is published in the DNSKEY record. In case if, at any future point, the ZSK private key is compromised, all RRSIG records signed with that key are retroactively forgeable, because an adversary who recovers the private key can produce new signatures over arbitrary DNS records that are indistinguishable from legitimate signatures for the same period.

In the “harvest now, decrypt later” threat model [8], an adversary archives RRSIG records, DNSKEY records, and DS records from the DNS today. When a Cryptographically Relevant Quantum Computer (CRQC) becomes available, Shor’s algorithm can be applied to recover the ZSK private key from the public key. All historically archived signatures can then be forged, and the chain of trust can be broken retroactively for any domain whose signing history was captured. This threat is relevant to DNSSEC even under a PQC migration scenario, because classical ZSKs will remain in use for an extended transition period and any signatures produced with RSA or ECDSA during that period are exposed.

Rawat and Jhanwar [16] note that SL-DNSSEC partially mitigates this issue. The KEM ciphertext (ct) included in each SL-DNSSEC query is session-specific for each DNS session; a fresh encapsulation is performed, producing a new shared secret that is used to derive a MAC key solely for that session. Disclosure of a past shared secret does not compromise future sessions, approximating forward secrecy at the session level. However, this property applies only to SL-DNSSEC. For all signature-based DNSSEC deployments—representing the entirety of currently deployed production infrastructure—forward secrecy is architecturally absent. This limitation has not been analyzed as an independent research objective in the reviewed literature, and no proposal in the reviewed literature applies forward-secure signature constructions [74] to the DNSSEC setting.

6.3. Denial-of-Service Resilience

6.3.1. KeyTrap a Structural Flaw in the DNSSEC Standard

Heftrig, Schulmann, Vogel, and Waidner [40] identify and formally characterize KeyTrap (CVE-2023-50387) as a denial-of-service vulnerability rooted not in any implementation defect but in a foundational design principle of the DNSSEC specification itself. The specification, following the principle of liberal acceptance associated with Postel’s

robustness guideline, requires validating resolvers to attempt DNSSEC validation against every available DNSKEY-RRSIG combination present in a response. If a zone publishes k DNSKEY records and m RRSIG records, the resolver must attempt up to $k \times m$ cryptographic validation computations before concluding that validation has failed. An adversary who controls an authoritative nameserver, or who can inject responses into the resolution path, can populate a zone with arbitrarily large numbers of DNSKEY and RRSIG records, compelling a resolver to perform an unbounded number of cryptographic operations in response to a single DNS query of normal size.

An extraordinary empirical harshness was shown by Heftrig et al. [40]. A single 100-byte DNS query packet can cause a 2,000,000-fold increase in CPU instruction count in vulnerable resolver implementations, perhaps rendering some implementations inoperable for up to 16 h. Since the vulnerability affects all major DNS resolver software, such as BIND, Unbound, and PowerDNS, as well as the DNS validation components of critical cloud infrastructure providers, the attack surface includes the resolvers handling the majority of DNS traffic worldwide. According to APNIC global data, over 31% of web clients use DNSSEC-validating resolvers in 2024, indicating a substantial potential impact on a sizable portion of internet users [40].

The basic complexity of the problem is revealed by the corrective action taken by top DNS software suppliers. Heftrig et al. [40] specifically state that the implemented patches intentionally deviate from RFC requirements by restricting the number of DNSKEY-RRSIG combinations a resolver will attempt per query, which is technically in violation of the stated DNSSEC standard. This leads to a conclusion that has a direct impact on post-quantum DNSSEC standards. KeyTrap is a specification flaw rather than an implementation problem, and fixing it completely necessitates updating the DNSSEC RFC rather than just patching individual implementations. The current remediation state, in which the deployed global resolver infrastructure purposefully violates the specification it is intended to implement, is therefore a temporary rather than a permanent solution because the RFC ambiguity that allowed the vulnerability to continue has not been addressed in the standards process. KeyTrap's significance for post-quantum DNSSEC is demonstrated by two separate methods that compound rather than just add. The number of distinct algorithm IDs in a zone increases in direct proportion to the combinatorial difficulty of DNSKEY-RRSIG validation. A zone will be needed to accommodate both classical and post-quantum algorithm signatures at the same time during the PQC transition phase. In contrast to a single-algorithm zone, a zone that uses both RSA and Falcon-512 signing infrastructure will simultaneously publish RSA KSK, RSA ZSK, FN-DSA KSK, and FN-DSA ZSK records, increasing the combinatorial attack surface and enhancing the $k \times m$ product. Second, compared to classical algorithms, PQC signature verification has a substantially higher per-operation cost. According to Rawat and Jhanwar [16], ML-DSA-44 achieves around 37,417 verification procedures per second compared to about 104,064 for RSA-2048, which leads to a 2.8-fold increase in per-operation cost. The epistemic status of these two factors differs and must be kept separate. The combinatorial factor is protocol arithmetic: it follows from the RFC 4035 validation logic exploited by KeyTrap [40] and from the DNSKEY and RRSIG multiplicities of dual algorithm zones, and involves no measurement. The per-operation factor, by contrast, rests on a single testbed estimate ($2.8 \times$ for ML-DSA-44 [16]) graded Low confidence in Table 9. The compounding claim advanced here, and formalized as hypothesis H2, is therefore parametric in form: for any per-operation cost ratio $c > 1$, aggregate validation cost in a dual-algorithm zone scales as the product of the quadratically growing pair count and c , so the super-linear character of the interaction does not depend on the precise value of c . The value 2.8 serves only as an illustrative magnitude and independent hardware-level benchmarks of the NIST lattice signature schemes [75] pro-

vide a non-DNS baseline against which a replicated ratio can be checked and it is not treated as a design constant prior to the independent replication that H2 specifies. The combinatorial growth of DNSKEY-RRSIG pairs and the increase in per-operation costs compound under KeyTrap conditions, creating a compounded CPU consumption scenario whose severity in dual-algorithm transition zones greatly exceeds the already critical baseline set by Heftrig et al. [40]. Although the combined threat landscape resulting from this interaction has been established analytically, it has not yet been experimentally characterized; Section 6.6 designates this deficiency as a research priority, and Section 7.2 summarizes its implications for the overall post-quantum migration strategy.

6.3.2. NSEC3 CPU Exhaustion

Gruza, Heftrig, Jacobsen, Schulmann, Vogel and Waidner [41] identify a separate denial-of-service vulnerability (CVE-2023-50868) affecting the NSEC3 authenticated denial-of-existence mechanism. NSEC3 is used to prove that a queried name does not exist in a signed DNS zone without enumerating the zone's contents. The closest encoder proof computation required for NSEC3 involves iterative hashing of DNS names, and the number of hash iterations is governed by a parameter in the zone configuration. RFC 5155 recommends limiting this parameter in order to reduce CPU consumption, but Gruza et al. [41] demonstrate that even within the RFC-recommended bounds the closest encoder computation can be driven to exhaustion by an adversary.

The attack works by constructing queries that require the resolver to compute the closest encoder proof over a carefully chosen set of NSEC3 records. Experimentally, Gruza et al. [41] report that their attack increases CPU instruction counts by 72 times relative to legitimate NSEC3 processing. Depending on the particular DNS implementation being attacked, the collateral loss rate for benign DNS inquiries varies from 2.7% to 30% at a query rate of 150 malicious NSEC3 entries per second. The variations in how BIND, Unbound, and PowerDNS carry out NSEC3 proof computation are reflected in this variability.

There is currently no research on the relationship between post-quantum DNSSEC and NSEC3 CPU fatigue. Each NSEC3 record in a post-quantum DNSSEC implementation contains a PQC RRSIG, verification of which calls for significantly more processing power than ECDSA verification. The combined burden of PQC RRSIG validation and NSEC3 closest-encoder proof computation results in a compounded CPU consumption situation that could be far worse than either vulnerability alone. There are no assessments of this combined threat surface in the examined literature.

6.3.3. DDoS Amplification

The DNS response size to query size ratio is measured by the Bandwidth Amplification Factor (BAF). When an adversary forges a DNS query source address to match a victim's, the DNS infrastructure can transmit more traffic to the victim than the adversary has supplied, intensifying the attack, if the BAF is greater than 1. This feature has been devastatingly exploited in past DNS-based DDoS incidents; the 2013 attack on Spamhaus used DNS amplification to create almost 300 Gbps of traffic [16].

DDoS amplification is a major issue for post-quantum DNSSEC, according to Kampakis and Lepoint [76], who contend that any successful PQC DNSSEC implementation needs to include amplification-resistant techniques. All of the main PQC DNSSEC systems have experimental BAF values measured by Rawat and Jhanwar [16]: SPHINCS+–QBF reaches BAF = 14.1, Dilithium–QBF BAF = 13.3, and Falcon–QBF BAF = 10.8. Compared to classical DNSSEC (ECDSA BAF $\approx$ 6.1), these numbers are significantly higher. The only method with BAF < 1 is SL-DNSSEC (Kyber-HMAC BAF = 0.48), where amplification is physically impossible because the answer is smaller than the query.

Amplification cannot be avoided for all signature-based PQC DNSSEC schemes; a response that contains a public key, a signature, and DNS record data will always be larger than the query triggering it. Goertzen and Stebila [49] note that ARRF introduces an additional vulnerability to memory exhaustion attacks, in which an adversarial resolver sends continuous fragment requests to occupy the reassembly buffer of a legitimate nameserver. Fregly et al. [48] identify a structurally similar resource exhaustion attack against MTL Mode; a malicious nameserver can reference a new Merkle tree ladder in each response, compelling a resolver to issue repeated follow-up queries to obtain the current ladder. The literature does not provide any facts of either of these residual vulnerabilities being resolved.

6.4. Operational Deployment

6.4.1. Resolver-Side Validation: A Systematic Measurement Gap

The majority of research on PQC DNSSEC performance has been focused on the signing side—the computational cost and output size of producing PQC signatures at authoritative nameservers and TLD operators. The complementary problem of PQC signature verification at resolvers has received significantly less attention. Schutijser et al. [45] identify this asymmetry and state that “these faults/drawbacks need to be studied even better to assess their operational impact on the validation side, e.g., by studying the DNS resolver’s behavior with PQC algorithms.”

Goertzen, Thomassen and Wisiol [43] present the most comprehensive field study to date on PQC DNSSEC resolver behavior, conducted using RIPE ATLAS measurements across approx. 10,000 global probes. The study covered PQC-signed zones using both BIND and PowerDNS implementations of Falcon-512, Dilithium-2, SPHINCS+-128s, and XMSS, and measured DNS response success and failure rates as a function of signing scheme and response size. The results have confirmed that failure rates are primarily a function of response packet size. Out of the tested schemes Falcon-512 performs best, Dilithium-2 follows it and SPHINCS+-128s and XMSS show substantially higher failure rates. An additional finding which should be mentioned is that 8.5% of resolvers incorrectly indicated successful validation for Falcon signatures despite not supporting the algorithm, revealing a class of implementation errors in which resolvers report authenticated responses falsely.

Suela et al. [44] note that this large-scale RIPE ATLAS study was conducted using an experimental DNS infrastructure, rather than a production one. The behavior of production resolvers at major public DNS providers—Cloudflare 1.1.1.1, Google 8.8.8.8, AWS Route53—under PQC DNSSEC load has not been measured. In addition to that, on resource-constrained devices, home routers, mobile devices, and IoT endpoints, energy consumption of PQC signature verification has not been evaluated. In an independent study, Howe and Westerbaan [75] demonstrate that Falcon-512’s floating-point arithmetic produces constant-time irregularities on ARM Cortex M7 and Raspberry Pi 3 devices, concluding that on these devices Falcon is insecure for applications in which signature generation can be timed by an adversary. Because Falcon verification operates exclusively on public inputs (the public key, the message, and the signature), this timing channel concerns signature generation rather than resolver-side verification; single-trace side-channel key recovery against the Falcon Gaussian sampler has since been demonstrated in the hardware literature [42]. The DNSSEC-relevant exposure therefore lies not with the heterogeneous resolver population but with signing infrastructure specifically with online-signing configurations, in which each query can trigger an attacker-timed signature, and with resource-constrained signer hardware lacking a hardened floating-point path. These deployment conditions are formalized as scoping preconditions of the Phase I recommendation in Section 7.3 and as hypothesis H5 in Section 6.6.

6.4.2. Stateful Hash-Based Signatures Operational Incompatibility

XMSS (RFC 8391) [77] and LMS/HSS (RFC 8554) [78] are stateful hash-based signature schemes with their security relying on non-reusing of one-time signing keys derived from a Merkle tree. Each signing operation increments a counter that tracks the position in the Merkle tree; reuse of any counter value, i.e., signing two different messages with the same one-time key, completely breaks the security of the scheme, allowing an adversary who observes both signatures to recover the signing key.

This state management requirement is fundamentally incompatible with the operational DNSSEC model in several common deployment configurations. In multi-signer zones with two or more operators each maintaining a copy of the zone and signing records independently, each operator maintains a separate Merkle tree state. To ensure that two signers never utilize the same counter value simultaneously, there must be network synchronization between operators. This synchronization cannot be ensured under network partition situations, and any race scenario that results in two signers using the same counter value leads to a security breach. The transfer between the old and new signing keys in automatic ZSK rollover scenarios where signature keys are changed on a schedule must maintain the state of both Merkle trees, which increases operational complexity. There are no operational rules for stateful hash-based signatures in DNSSEC multi-signer installations, as noted by Fregly et al. [48].

Goertzen et al. [43] confirm that in their RIPE ATLAS field study XMSS in practice achieves DNS response success rates of approximately 80% over UDP and 90% over TCP. These rates are below the availability thresholds expected of production DNSSEC infrastructure and reflect the fact that XMSS responses are substantially larger than Falcon-512 responses, leading to higher fragmentation and TCP fallback rates.

6.4.3. Root KSK Rollover Timeline Gap and Operational Risk

Fregly et al. [48] state that the Root KSK rollover to PQC algorithms is planned for approximately 2028–2029. Mallick, Kundu and Kompella [8] document that NIST IR 8547 Draft recommends that RSA and ECDSA be deprecated by 2030 and retired by 2035. The gap between the planned PQC Root KSK rollover and the RSA/ECDSA deprecation deadline is one to two years, during which the root zone would be protected by PQC algorithms while legacy algorithms remain in use at lower levels of the hierarchy.

The Root KSK rollover process alone has a significant operational risk. The 2018 ICANN Root KSK Ceremony demonstrated that this approach requires coordinated trust anchor updates across the global resolver ecosystem, extensive telemetry monitoring to track acceptance, and backup preparations for resolvers that cannot be updated. After several years of planning and two delays, the 2018 rollover was the first Root KSK modification since DNSSEC was implemented at the root in 2010. There is extra technological complexity involved in a PQC Root KSK rollover, as the Falcon-512 public key (897 bytes) is 3.4 times larger than the currently used RSA-2048 public key (260 bytes), increasing the DNSKEY RRsets size in the root zone and affecting all resolvers that retrieve and cache the root DNSKEY. Schutijser et al. [45] discuss weekly ZSK rotation as a short-term mitigation strategy for reducing the classical ZSKs exposure window in the period before PQC migration, but acknowledge that this shifts the attack surface to the KSK which rotates much less frequently and does not substitute for a complete PQC migration.

6.5. Standardization and Systemic Gaps

6.5.1. IANA, IETF Process and Absence of Production Deployments

There is presently no algorithm number for any PQC signature method in the IANA DNSSEC Algorithm Numbers registry. Falcon-512 and ML-DSA-44 running in experimen-

tal deployments using provisional OID values currently prevent interoperability between independently developed implementations as well as between experimental and production deployments. The only DNS programs that support PQC, according to Schutijser et al. [45], are OQS-patched versions of PowerDNS and BIND9 and the CoreDNS plugin developed by Suela et al. [44]. No production-grade DNS includes a reliable, interoperable PQC DNSSEC implementation. Fregly et al. [48] state that there is currently no chartered IETF DNSOP working group item on PQC DNSSEC standardization; instead, academics collaborate via an unofficial, standard-track-illegitimate pq-dnssec mailing list.

NIST's second-round signature review, which includes MAYO, SNOVA, and CROSS, has generated candidates that might provide benefits above the established algorithms for DNSSEC use cases, according to Suela et al. [44]. In addition to ML-DSA, Falcon, and SPHINCS+, their own CoreDNS plugin supports MAYO and SNOVA, making it the first DNS implementation to handle these algorithms. Schutijser et al. [45] investigated MAYO-2 at TLD size and discovered that it outperformed RSA-1280 in signing throughput, whereas Falcon-512 achieved equivalent throughput. However, it has not been measured how MAYO-2 and SNOVA behave at the resolver validation side, which is the primary operational concern for high-volume recursive DNS infrastructure.

6.5.2. Absence of Integrated Analysis Across DNSSEC, DoT, and DoH

Lee, Hoque, Aydeger, and Zeydan [60] assert that a thorough, comparative analysis of the cumulative impacts of PQC on DNSSEC, DNS over TLS (DoT), and DNS over HTTPS (DoH) is lacking. Public resolvers that concurrently handle DNSSEC validation, DoT, and DoH are run by Cloudflare, Google, and NextDNS; nevertheless, the interaction between these protocols under PQC signatures is yet unknown, making this gap operationally relevant. Although they do not address fragmentation behavior in the DoT or DoH transport scenarios, Raavi et al. [50] offer a solution to the Fragmentation Mis-Association Threat. Falcon-512 requires TLD operators to quadruple their nameserver disk space and memory, as demonstrated by Schutijser et al. [45], but the overall infrastructure cost of enforcing this requirement throughout the global network of DNSSEC-enabled authoritative nameservers using DoT and DoH simultaneously is still unknown.

6.6. Testable Research Hypotheses and Proposed Methodologies

Below are testable ideas that address the research gaps mentioned in Sections 6.1–6.5. Every hypothesis presents a verifiable claim, a suggested structure for evaluation, and measurable success indicators. By making each study goal operationally clear, this strategy seeks to improve the Phase I–III plan described in Section 7.3.

H1: FN-DSA-512 NXDOMAIN response failure rates at production resolvers exceed 10%. Traditional DNSSEC answer widths of 512–4096 bytes were used to measure the 10% fragmentation failure rate reported by Van den Broek et al. [46]; FN-DSA-512 NXDOMAIN replies range from 2269 to 3767 bytes [54], much exceeding that range. Methodology: On authoritative nameservers running PowerDNS 4.8 or BIND 9.19 with OQS patches, create a DNSSEC-signed zone using FN-DSA-512. Execute QTYPE A and QTYPE NS queries from a minimum of 5000 RIPE Atlas probes directed to the zone, documenting per-probe success/failure rates, truncation indicators, and utilized transport protocols. Control: Similar zone authenticated with ECDSA P-256 queried from the same probe set inside the same measurement frame. The success criterion stipulates that H1 is validated if the FN-DSA-512 NXDOMAIN failure rate surpasses the ECDSA control rate by a statistically significant margin (Mann–Whitney U, $p < 0.05$). H1 is disproven if the rates are statistically indistinguishable, indicating that interim TCP fallback support in modern resolvers is more resilient than indicated in [46]. H2: The KeyTrap DNSKEY-RRSIG com-

binatorial attack surface increases super-linearly during dual-algorithm PQC transition zones. Rationale: Heftrig et al. [40] characterize the $k \times m$ product under single-algorithm zones; dual-algorithm zones (RSA + FN-DSA-512) approximately double both k and m , producing a $4 \times$ combinatorial expansion before accounting for PQC's higher per-operation cost. Methodology: Configure a zone with increasing numbers of DNSKEY records across four conditions: (a) RSA-2048 only, (b) FN-DSA-512 only, (c) RSA-2048 + FN-DSA-512, and (d) RSA-2048 + ML-DSA-44. Measure resolver CPU time per query using a patched Unbound 1.17 with instruction-count profiling for $k = \{1, 2, 4, 8, 16\}$ DNSKEY records per algorithm. Success criterion: H2 is supported if the CPU time per query under condition (c) exceeds the sum of conditions (a) and (b) by more than the per-operation overhead factor measured in situ from the single-algorithm conditions (a) and (b) for which the current literature point estimate is $2.8 \times$ for ML-DSA-44 ([16]; Low confidence, Table 9) indicating compounding rather than additive behavior. Formulating the criterion in terms of the in situ factor ensures that the test of super-linearity does not inherit the uncertainty of the imported estimate.

H3: SL-DNSSEC Level I (ML-KEM-512 + HMAC-SHA-256) attains resolution latency within 10 ms of ECDSA P-256 for the 90th percentile of requests among a geographically varied resolver demographic. Rationale: Rawat and Jhanwar [16] indicate a 44 ms resolution time with a 10 ms RTT on a 100 Mbps testbed, comparable to ECDSA; nevertheless, this singular experimental setting fails to represent tail latency or diverse network conditions. Methodology: Implement SL-DNSSEC on a zone managed by three geographically dispersed authoritative nameservers. Request from a minimum of 3000 RIPE Atlas probes documenting complete resolution timings. Primary metric: 90th-percentile resolution time per probe, compared against ECDSA control. Secondary metric: Rate of queries for which the KEM ciphertext exceeds the EDNS(0) buffer and triggers an additional round-trip. Success criterion: H3 is supported if the 90th-percentile resolution time exceeds the ECDSA control by less than 10 ms for at least 80% of probes.

H4: A quantum-resistant NSEC3 replacement using SHA-3-256 with NSEC5-style verifiable random function (VRF) construction maintains zone-enumeration resistance while reducing closest-encloser proof CPU cost to within $1.5 \times$ of classical NSEC3. Rationale: The NSEC3 CPU exhaustion vulnerability (Section 6.3.2) [41] is amplified by the higher per-record PQC RRSIG verification cost; a VRF-based construction eliminates the iterative hashing that creates the exhaustion surface. Methodology: Implement the NSEC5 draft specification on a patched BIND 9.19 authoritative server. Measure closest-encloser proof computation time per query under the attack conditions of Gruza et al. [41] (150 malicious queries per second against a zone with 1 million NSEC3 records). Success criterion: H4 is supported if the VRF construction reduces the collateral query loss rate to below 1% under the same attack conditions that produce 2.7–30% loss in NSEC3 [41], while providing a formal proof of zone-enumeration resistance equivalent to NSEC3 under the random oracle model.

H5: Timing side-channel attacks against FN-DSA-512 signature generation are not feasible at DNS resolution throughput on x86-64-v3 hardware with AVX2 support, but are feasible on ARM Cortex-M7 without floating-point hardening. Rationale: Howe and Westerman [75] demonstrate constant-time irregularities on ARM Cortex-M7 and Raspberry Pi 3; the question of whether these translate to exploitable timing channels at DNSSEC-relevant query rates on production resolver hardware has not been investigated. Methodology: Measure per-signature timing variance of FN-DSA-512 using a high-resolution timer (RDTSC) across 10^8 signing operations on (a) x86-64-v3 with AVX2, (b) ARM Cortex-A72 (Raspberry Pi 4), and (c) ARM Cortex-M7. Apply timing attack templates using the tlsfuzzer timing analysis framework. Success criterion: H5 is supported if the timing variance on

x86-64-v3 is below the noise floor of a network-based timing attack (estimated at 0.1 ms [75]) while the ARM Cortex-M7 variance exceeds this threshold, confirming the hardware dependency of the vulnerability.

6.7. Future Research Directions

The prior analysis highlights the subsequent research avenues as insufficiently explored in the examined literature.

Extensive resolver assessment: Fregly et al. [48] observe that the influence of PQC on DNSSEC resolver behavior at an internet scale—specifically, EDNS(0) buffer negotiation in relation to PQC response sizes, UDP truncation cascade rates, and DNS resolution failure rates among the Tranco top one million domains—has not been quantified. The approach of Van den Broek et al. [46], which determined the 10% resolver fragmentation failure rate for conventional DNSSEC in 2014, necessitates thorough replication and extension to the post-quantum scenario. The Goertzen et al. [43] RIPE ATLAS investigation, utilizing 10,000 probes, offers the most extensive data presently accessible; nonetheless, it was executed with experimental infrastructure, and there is a lack of production-scale measurement studies.

Formal security demonstration for hybrid and dual-signature DNSSEC: Pan et al. [67] introduce Double-Signed Fragmented DNSSEC, while Fregly et al. [48] examine the attack surface of dual-signature methods; nonetheless, a formal security model delineating adversarial capabilities against hybrid DNSSEC is absent. A conclusive security demonstration for dual-signature DNSSEC, which formally restricts an adversary's ability to arbitrarily remove, reorder, or substitute RRSIG records, is essential before adopting this framework for production deployment.

NSEC3 is resistant to quantum attacks. When exposed to Grover's technique, which decreases the brute-force effort required for SHA-1 preimage calculation from 2^{160} to roughly 2^{80} operations, the SHA-1 hash function used in NSEC3 diminishes its effective security strength. To maintain the verified denial-of-existence characteristics and zone-walking resistance, a switch from NSEC3 to a quantum-resistant hash function, like SHA-3 or BLAKE3, must be planned and assessed. Targeted experimental evaluation is necessary to assess the cumulative danger of CPU depletion resulting from the interaction between NSEC3 closest-encloser proof calculation and PQC RRSIG validation.

Timing side-channel vulnerabilities in Falcon-512 under high-throughput resolver conditions are analyzed. Howe and Westerbaan [75] show that on ARM Cortex M7 and Raspberry Pi 3 platforms, the floating-point computations of Falcon-512 result in consistent-time anomalies. A thorough analysis is necessary to assess the viability of timed attacks on Falcon verification and the factors that make them conceivable in the context of a DNSSEC resolver, where RRSIG verification may take place billions of times every day at a production scale.

The expenses related to the PQC DNSSEC infrastructure are economically analyzed. According to Schutijser et al. [45], TLD operators using Falcon-512 experience a fourfold increase in RAM and disk space requirements. According to Suela et al. [44], the number of papers and signatures increases by a factor of 10 to 50 in comparison to ECDSA. The extensive infrastructure investment needed to implement PQC DNSSEC throughout the global network of TLD operators, ISPs, root server operators, and significant recursive resolver providers has not yet been modeled mathematically. For evidence-based policy decisions on PQC migration schedules, this model is essential.

6.8. Certainty of Evidence for Principal Findings

A systematic evidence-grading framework for the main results of this review is used in the following table. Three criteria are used to classify each finding as High, Moderate, or Low confidence: (a) consistency, defined as the degree to which results from multiple co-authorship-independent research groups (author clusters, Section 3) are concordant, where multiple publications from a single cluster are counted as one line of evidence; (b) methodological quality of the primary studies, which is evaluated using the four-criteria instrument described in Section 3.6; and (c) directness, which is the degree to which the evidence directly relates to the particular claim or necessitates extrapolation from analogous situations. When findings are based on single research, when experimental settings differ greatly from real deployment, or when quantitative values are obtained graphically rather than from primary tabular data, confidence is reduced. Because the principal findings are not epistemically homogeneous, each finding in Table 9 additionally carries an evidence-type tag that makes explicit the basis on which the three criteria are applied. Type E (empirical) denotes claims established by measurement, for which consistency means concordance across co-authorship-independent clusters, quality refers to the methodological instrument of Section 3.6, and directness contrasts production with testbed conditions. Type A (analytical) denotes claims that follow deductively from protocol specifications taken as premises such as the architectural absence of forward secrecy, which is a property of the signature-based validation model defined in RFC 4033-4035 rather than an observed frequency; for these, consistency means independent derivations reaching the same conclusion, quality means the rigor and explicitness of the premises, and directness means the absence of extrapolation beyond what the premises entail. Type V (verifiable) denotes facts checkable by direct inspection of a public registry or corpus, such as the absence of registered IANA algorithm numbers. A single ordinal confidence scale is retained across types because the roadmap of Section 7.3 consumes one prioritization signal; the interpretation of a High grade is, however, type-relative: for E-claims it asserts replicated observation, for A-claims it asserts deductive validity conditional on the cited premises falsifiable by premise revision or counter-derivation, not by additional sampling, and for V-claims it asserts a currently inspectable fact subject to change upon registry or corpus update.

Table 9. Certainty-of-evidence grading for principal findings. Confidence: High = multiple concordant high-quality studies with direct evidence; Moderate = single high-quality study or indirect evidence; Low = single study with limited generalizability. Type: E = empirical measurement; A = analytical, i.e., deductively derived from protocol specifications taken as premises; V = verifiable by direct inspection of a public registry or corpus. The consistency, quality, and directness criteria are interpreted type-relatively, as defined in the accompanying text.

Principal Finding	Supporting Sources	Type	Consistency	Study Quality	Directness	Confidence
Every NIST-standardized PQC scheme exceeds the 1232-byte UDP limit for non-minimal DNSSEC responses	[12,16,44,45]	A+E	High (5 studies; 3 independent clusters; bound fixed by FIPS constants)	High (all Q4)	High (production data)	High
FN-DSA-512 is the only evaluated PQC candidate satisfying the UDP size constraint for minimal NOERROR responses	[12,54]	E	High (2 concordant)	High	Moderate (minimal response only)	Moderate
TCP fallback imposes $\geq 50\%$ latency penalty for Falcon-512 relative to ECDSA under 10 ms RTT	[51]	E	Low (single study)	High	High (experimental)	Moderate

Table 9. Cont.

Principal Finding	Supporting Sources	Type	Consistency	Study Quality	Directness	Confidence
KeyTrap (CVE-2023-50387) causes up to 16 h CPU stall in susceptible resolver implementations	[40]	E	Low (single study)	High	High (lab + production)	Moderate
SL-DNSSEC achieves BAF = 0.48 at NIST Level I (Kyber-512 + HMAC-SHA-256)	[16]	E	Low (single study)	High	Moderate (testbed only)	Moderate
8.5% of resolvers falsely report successful PQC validation despite lacking algorithm support	[43]	E	Low (single study)	High	High (10,000 RIPE Atlas probes)	Moderate
Forward secrecy is architecturally absent from signature-based DNSSEC	[8,9,16]	A	High (3 independent clusters)	High	High (architectural analysis)	High
Algorithm agility mechanism enables downgrade attacks affecting $\geq 45\%$ of resolvers (pre-patch)	[13,14]	E	Moderate (2 publications, 1 author cluster)	High	High (large-scale measurement)	Moderate
No formal security proof exists for hybrid dual-signature DNSSEC constructions	[48,67]	V	High (concordant)	High	High (analytical)	High
PQC ML-DSA-44 validation is approximately $2.8\times$ more expensive per operation than RSA-2048	[16]	E	Low (single study)	High	Moderate (testbed conditions)	Low
NXDOMAIN responses using FN-DSA-512 range from 2269 to 3767 bytes in production	[54]	E	Low (single ccTLD)	High	High (.nl production data)	Moderate
No IANA algorithm number is registered for any PQC signature scheme in DNSSEC	[45,48]	V	High (1 cluster; registry-verifiable)	High	High (registry check)	High

Four findings are rated High confidence: the universal incompatibility of NIST-standardized post-quantum schemes with the 1232-byte UDP limit, the architectural absence of forward secrecy from signature-based DNSSEC, the absence of any formal security proof for hybrid dual-signature constructions, and the absence of IANA algorithm registration for any PQC signature scheme. The first rests on three co-authorship-independent clusters and, more fundamentally, on protocol arithmetic: signature sizes fixed by the FIPS 204, 205, and 206 standards exceed the threshold irrespective of measurement environment. The second is concordant across three independent groups, and the third across two. The fourth, although supported by publications from a single author cluster, is independently verifiable against the public IANA registry. These distinctions are now carried explicitly by the evidence-type tags in Table 9: of the four High-confidence findings, one is empirical–analytical (A+E), one analytical (A), and two are verifiable corpus or registry facts (V), so their High grades assert, respectively, replicated-and-derived status, deductive validity conditional on protocol premises, and inspectable fact rather than a uniform claim of measurement strength. Seven findings are assigned Moderate confidence, mostly owing to reliance on a single study, testbed rather than production conditions, or a single ccTLD dataset; this now includes the downgrade-attack finding, previously rated High on two concordant publications, which is reclassified because both originate from the same author group and therefore constitute a single line of evidence pending replication by an independent group. The KeyTrap severity estimate (16 h stall), the SL-DNSSEC BAF measurement, and the $2.8\times$ ML-DSA-44 validation overhead likewise require independent replication under production conditions before being treated as design constants for protocol specifications. The per-operation cost estimate for ML-DSA-44 retains Low confidence,

resting on a solitary testbed measurement under ambiguous workload conditions without independent corroboration.

7. Conclusions

7.1. Summary of Findings

This systematic review synthesizes 27 peer-reviewed works published between 2020 and 2025, providing the first unified analytical framework for post-quantum DNSSEC research across five dimensions: transport constraints, denial-of-service resilience, cryptographic agility, operational deployment, and forward secrecy. The central analytical result is that these dimensions are structurally coupled rather than independently resolvable; every proposal reviewed in Section 5 that resolves a constraint in one dimension introduces or amplifies a constraint in at least one other. This coupling, not the size of any individual PQC signature, is the defining challenge of the post-quantum DNSSEC transition.

Transport constraints: As established in Section 4.1, every NIST-standardized PQC signature scheme is structurally incompatible with DNS UDP transport in non-minimal DNSSEC responses, and this incompatibility is not addressable through parameter optimization because it arises from the mathematical foundations of lattice-based and hash-based cryptography. The six mitigation strategies evaluated in Section 5 divide along a fundamental axis: fragmentation-based approaches, ARRF [49], QBF [53], and Double-Signed DNSSEC [67], preserve the digital-signature-based chain-of-trust model but cannot eliminate DDoS amplification, since a response containing a public key, a signature, and DNS record data will always exceed the query that triggered it; SL-DNSSEC [16] eliminates amplification and satisfies the UDP constraint but replaces the chain-of-trust model with a KEM-MAC construction whose security within the full DNSSEC authentication hierarchy has not been formally proved [48]. Transport and denial-of-service dimensions are inextricably linked: any proposal that transmits PQC signatures over fragmented UDP inherits the amplification attack surface; the only proposal that avoids fragmentation does so by removing the signatures that are essential to the chain of trust.

Denial-of-service resilience: The KeyTrap vulnerability (CVE-2023-50387), analyzed in Section 6.3.1, is a specification defect confirmed by the fact that all major vendor patches intentionally deviate from RFC requirements [40]. PQC deployment necessitates dual-algorithm transition zones where both classical and post-quantum DNSKEY and RRSIG records coexist. This increases the number of validation combinations required by the DNSSEC specification while also raising the per-operation computational cost of each validation attempt, which is why it is relevant to the post-quantum transition. These two factors compound in a way that has been identified analytically but not yet measured experimentally, as noted in Section 6.3.1. The NSEC3 CPU exhaustion vulnerability (CVE-2023-50868), established in Section 6.3.2, presents a structurally independent compounded threat: its interaction with PQC RRSIG verification overhead has not been characterized, and no evaluation of the joint attack surface exists in the reviewed literature [41].

Downgrade attacks and cryptographic agility: The algorithm agility mechanism that technically enables PQC migration also serves as the delivery channel for downgrade assaults, as shown in Section 6.2. This vulnerability is not accidental; rather, it is a structural result of the DNSSEC specification's failure to require resolver behavior for algorithm IDs that are not recognized. Coordinated vendor patching lessened the impact of this specification ambiguity but did not fix it [14]. The coupling with the transport dimension makes this particularly significant; the off-path downgrade variant documented by Heftrig et al. [15] exploits IP fragmentation, which is structurally unavoidable for all PQC DNSSEC schemes except SL-DNSSEC. The result is that any fragmentation-based PQC deployment inherits both the DDoS amplification risk established above and the off-path downgrade

attack surface—two independent vulnerabilities linked through a single architectural dependency. In the IETF DNSOP working group, the lack of a cipher-suite negotiation mechanism identified by Shrishak and Shulman [61] results in a deployment deadlock with no current resolution path, indicating that the signaling infrastructure needed to manage this risk does not exist and has not been chartered for standardization.

Deployment of operations: Three conclusions are established in Section 6.4, and their combined implications are more important than their separate conclusions. Major public resolver operators have not measured production-scale resolver behavior under PQC DNSSEC load [43]; 8.5% of resolvers falsely report successful PQC validation despite lacking algorithm support [43], indicating that a sizable portion of the resolver population would silently deteriorate the security guarantee DNSSEC provides rather than fail visibly; and no IANA algorithm number is registered for any PQC scheme [44,45], which precludes interoperability between implementations and prevents any production DNS software from implementing stable PQC DNSSEC support. These three gaps interact without IANA registration, and resolver support cannot be tested at scale; without production-scale measurement, the 8.5% silent failure rate cannot be confirmed or bounded; and without confirmed resolver support data, the deployment deadlock identified by Shrishak and Shulman [61] cannot be resolved. The five structural adoption barriers documented in Section 4.6, operational complexity, infrastructural dependency, ecosystem fragmentation through islands of security, slow algorithm deployment, and the DoH monitoring paradox, will amplify rather than merely accompany the PQC transition, because each barrier reduces the fraction of the resolver ecosystem that can be reliably updated within any given timeline, directly affecting the $t_{migrate}$ term in Mosca’s theorem [9].

Forward secrecy and hybrid security: As analyzed in Section 6.2.4, forward secrecy is architecturally absent from signature-based DNSSEC by construction, and this absence cannot be remediated through ZSK rotation policy: the exposure window is bounded by key rotation frequency but cannot be reduced to zero within the signature-based model, because any archived RRSIG record remains permanently forgeable once the corresponding private key is recovered by a CRQC. Hybrid dual-signature schemes, operationally sensible for the transition period as demonstrated by Pan et al. [67], reduce but do not eliminate this exposure. Their principal residual vulnerability, identified by Fregly et al. [48], is a class of mix-and-match attacks in which selective removal of one RRSIG record from a dual-signature response can reduce security to that of the weaker algorithm alone, a threat that has been characterized but not formally bounded, because no security proof exists for any hybrid DNSSEC construction under adversarial RRSIG set manipulation. The obstruction here must be stated precisely, since it operates at the level of protocol semantics rather than mathematical impossibility. At the primitive level no barrier exists: for dual signatures with enforced conjunctive verification, combiner theorems establish unforgeability of the composition whenever at least one component remains unforgeable, including against adversaries that are classical today and quantum later [38]. Under the validation semantics actually deployed in DNSSEC, however, the property to be proven, downgrade resistance, is false rather than unproven: because a resolver accepts any single valid RRSIG, the removal of one signature by A_DG is indistinguishable from legitimate algorithm agility behavior, and no security reduction can exclude conduct that the protocol permits [13,14]. What remains genuinely open, and constitutes the Phase II objective, is the intermediate case: a security proof for conjunctively validated hybrid DNSSEC within the full DNS adversary model on-path and off-path RRSIG manipulation, fragmentation, caching, and the composition, within a single reduction, of security guarantees established in different parameter regimes (classical for RSA, lattice-based for ML-DSA) [48]. This is the absence of a completed proof for a well-posed problem, not an established impossibility.

7.2. Main Takeaway

The review's findings show that the post-quantum DNSSEC transition has urgent operational implications and is really a protocol-redesign issue rather than just an algorithm-substitution one.

The previous DNSSEC algorithm changes from RSA-1024 to RSA-2048, from SHA-1 to SHA-256 in DS records, and the addition of ECDSA were carried out inside the current protocol architecture without changing the transport semantics, chain-of-trust logic, or DNS message format. Even these seemingly simple changes took nearly three years to fully penetrate the ecosystem, according to the operational record [11], and the first root KSK rollover involved two delays and eleven months of coordinated management [10]. The post-quantum transition differs from earlier antecedents in nature rather than degree. It requires simultaneous changes to the DNS message format, the transport model, the chain-of-trust authentication architecture, the NSEC denial-of-existence mechanism, and the DNSSEC specification's validation rules. These changes are interrelated and have an impact on the current deployment landscape that has never been seen before.

A comparison with TLS 1.3 clarifies the fundamental basis of this disparity. Hybrid PQC key exchange in TLS could be implemented at a production scale with little disruption because of TLS's generation of ephemeral per-session keys, explicit cipher-suite negotiation between endpoints, separation of the authentication layer from the key exchange layer, and intrinsic provision of forward secrecy [68]. These characteristics have nothing to do with DNSSEC. It uses long-lasting signature keys, the public versions of which are kept in caches at millions of resolvers worldwide. Because it lacks a negotiation mechanism, nameservers are unable to determine which algorithms a resolver supports and adjust their responses accordingly [61]. The size of the signature directly affects whether a response complies with the limitations of the transport layer because it does not differentiate between authentication and transit. Because it lacks forward secrecy, any signatures created inside the current infrastructure are vulnerable to retroactive forging. DNSSEC lacks the architectural aspects that enable the transition to TLS post-quantum cryptography, and none of the methods discussed in Section 5 have these capabilities without simultaneously changing the chain-of-trust paradigm that defines DNSSEC's architecture. Mosca's theorem [9] provides a formal quantification of this situation's urgency. The necessary protection term t_{protect} for essential DNS infrastructure is roughly 20 years, which accounts for the operational lifespan of enterprise resolver deployments, TLD signing infrastructure, and root zone trust anchors. Based on the empirical record of the EdDSA–Ed25519 transition and the root KSK rollover, the migration timeline t_{migrate} is conservatively estimated at 7 years [10,11], and this estimate does not account for the additional complexity of the PQC transition relative to those precedents. Expert consensus places the quantum horizon t_{quantum} at approximately 15 years [9]. The resulting inequality $27 > 15$ is not a marginal exceedance under pessimistic assumptions but a robust result under representative estimates from the reviewed literature. The risk window defined by Mosca's theorem has already opened, and this calculation precedes the availability of a comprehensive solution: IANA registration is absent, IETF standardization is unchartered, no proposal has been shown in large-scale production infrastructure, and no formal security proof exists for any leading transitional architecture. Migration must therefore proceed on multiple parallel tracks before a solution that addresses all five dimensions simultaneously is available. In order to avoid the unbounded risk of delaying action, one must accept bounded interim risk. Two structural properties make this urgency conclusion insensitive to the production-measurement gap documented in Section 6.5.1. First, none of the three parameters derives from resolver measurements: t_{protect} reflects infrastructure lifespans, t_{migrate} is anchored in the historical record of completed DNS transitions [10,13], and t_{quantum} is an expert-

consensus estimate [9]. Second, the inequality is already binding on the archival term alone $t_{\text{protect}} > t_{\text{quantum}}$ so the harvest-now–forge-later exposure of signatures archived today precedes any migration activity, and the conclusion would survive even an instantaneous migration with $t_{\text{migrate}} = 0$. Production measurements can therefore modulate only t_{migrate} , and monotonically: resolver behavior worse than the experimental findings lengthens t_{migrate} and widens the inequality’s margin, while behavior better than the experimental findings shortens it without approaching the point at which the inequality could reverse.

7.3. Prioritized Research Roadmap

Here, the research directions listed in Section 6.7 are arranged into a three-phase roadmap that is based on dependency ordering as opposed to thematic grouping. Dependency ordering prioritizes tasks whose results are necessary for later ones, and dependency rather than importance determines which tasks must be completed first. Where concurrent work is feasible, phases overlap; the crucial path includes Root KSK operational preparation, formal security analysis, and IANA registration. Phase I (2025–2026) focuses on standardization and establishing the empirical and infrastructural basis. All succeeding stages require a precise empirical assessment of the DNS resolver ecosystem under PQC response sizes in production conditions. The Goertzen et al. field measurement is currently the most thorough one [43]. The RIPE Atlas study was conducted with experimental infrastructure rather than production deployments; the behavior of production resolvers at major public DNS operators under PQC DNSSEC load remains uncharacterized. Before any deployment recommendation can rest on verified operational ground, production-verified failure rates, truncation-cascade rates, and TCP-fallback reliability must be established through a systematic replication of Van den Broek et al. [46], applying their fragmentation-failure methodology extended to PQC response sizes across the Tranco top one million domains. The deployment stalemate found by Shrishak and Shulman [61] cannot be handled with data-informed certainty without this measurement, and the 8.5% silent validation failure rate determined by Goertzen et al. [43] cannot be verified or bounded at production scale. In parallel, Schutijser et al. [45] and Fregly et al. [48] have established that the main interoperability barrier is IANA algorithm number registration for FN-DSA-512 and ML-DSA-44, which is the single most immediately actionable item identified by this research. Without registered identifiers, all experimental deployments documented in Section 5 remain non-interoperable, and the OQS-patched DNS software forks that currently represent the entirety of PQC DNSSEC implementation cannot transition to production-grade status. All further standardization work requires the conversion of the Fregly et al. [48] research agenda into a chartered IETF DNSOP working group item with a clear milestone schedule and RFC target dates. This should be done concurrently with the IANA registration process rather than sequentially.

Operational deployment during this phase should proceed with Falcon-512 combined with QBF [53] or Parallel ARRF [49], supplemented where dual-algorithm security is required by Double-Signed DNSSEC [67]. According to Section 5.8, these methods are the most dependable incremental migration pathway now available, have been empirically proven in controlled conditions, and are closest to IETF draft state. The Falcon-512 component of this recommendation is conditioned on an explicit signing-side threat-model boundary. The timing irregularities documented for FN-DSA-512 arise in signature generation, whose floating-point Gaussian sampling is difficult to render constant-time [75], and single-trace key-recovery attacks against the sampler have been demonstrated in the hardware literature [42]; verification, operating exclusively on public inputs, is unaffected. The recommendation therefore presupposes the offline pre-signing model of classical DNSSEC operation, in which signatures are generated on operator-controlled x86-64 hardware with

AVX2 support (where constant-time behavior is reported [75]) or within hardware security modules, so that no network adversary can trigger or time the signing operations. Two deployment configurations fall outside this boundary and are excluded from the Phase I recommendation pending resolution of hypothesis H5: online (dynamic) signing, including on-the-fly denial-of-existence generation, in which each query elicits an attacker-timeable signature; and signing on resource-constrained or heterogeneous hardware without a verified constant-time floating-point path, such as the ARM Cortex-M7 and Raspberry Pi 3 platforms profiled in [75]. Operators requiring either configuration should employ ML-DSA-44, whose integer-only arithmetic is designed for constant-time implementation, accepting the associated response-size penalty documented in Table 4, or defer PQC signing on the affected infrastructure to Phase II. To prevent the technical debt of entrenching techniques whose residual risks are highlighted in Sections 6.1 and 6.3, they should be clearly positioned as transitory infrastructure with sunset commitments linked to Phase II milestones rather than as target architectures.

Phase II (2026–2028): Root KSK Preparation, Protocol Repair, and Formal Analysis.

Because the second step deals with the formal and protocol-level foundations that production deployment must rely on, it fills in gaps that empirical measurement and incremental deployment are unable to handle. The critical route is defined by three objectives, and their reciprocal dependencies are reflected in their sequence. The first and foundational objective is a formal security proof for hybrid DNSSEC under adversarial RRSIG manipulation. As established in Section 6.2.3, the IND-CMP composition problem between classical and lattice-based security parameter spaces means that no dual-signature DNSSEC construction currently has a security proof that formally bounds an adversary's ability to selectively remove, reorder, or substitute RRSIG records. No hybrid DNSSEC proposal can obtain standards-track RFC status under the IETF's security considerations requirements without such proof, and deployment recommendations continue to rely on informal security arguments rather than verifiable guarantees. This is not a theoretical concern that is incidental to deployment. In contrast to the idealized channel model usually assumed in signature scheme security proofs, the reduction must address the particular adversarial model of DNS in which responses transit untrusted network paths, are susceptible to IP fragmentation, and may be altered by on-path and off-path adversaries.

Experimental characterization of the KeyTrap-PQC relationship across the main DNS resolver implementations is the second goal. The current KeyTrap patches purposefully deviate from RFC requirements [40], and controlled measurement under dual-algorithm zone configurations is necessary to determine whether their caps remain effective when per-operation PQC verification costs are applied. As stated in Section 6.3.1, this interaction has been identified analytically but not measured experimentally. This experiment determines whether further RFC change is required; if the results are positive, they must be reported to the IETF DNSOP working group on a schedule that aligns with Phase III implementation. The third goal is to build and submit a quantum-resistant NSEC3 substitute to the IETF. As stated in Section 6.7, SHA-1, the only hash function specified in RFC 5155 [26], needs to be replaced. The replacement design must simultaneously maintain zone-enumeration resistance, uphold the authenticated denial-of-existence semantics required by validating resolvers, prevent the introduction of the CPU exhaustion vulnerability identified by Gruza et al. [41], and support online signing for NSEC3 white lie and black lie countermeasures. These constraints are jointly non-trivial and require targeted design analysis before submission to the IETF DNSOP working group. Coordinating this submission with the PQC algorithm registration track allows both to progress through the RFC process in a unified protocol revision rather than as sequential amendments that require separate implementation and deployment cycles.

Phase II concludes with the operational preparation required for the Root KSK rollover to PQC algorithms, planned for approximately 2028–2029 per Fregly et al. [48]. The 2018 ICANN Root KSK rollover demonstrated the operational intricacies of this process under conventional circumstances; a PQC rollover introduces more technological complexities due to the greater public key sizes of PQC candidates compared to existing RSA keys, hence augmenting root zone DNS.KEY RRset dimensions worldwide and their interaction with the downgrade vulnerability outlined in Section 6.2 for resolvers facing new algorithm IDs without updated software. Timelines for preparation, contingency planning, and resolver update monitoring infrastructure must be initiated well before the planned rollover date to allow adequate lead time given these compounding factors.

Phase III (2028–2029): Full PQC Architecture and Classical Algorithm Deprecation.

The third phase focuses on architectural constraints that determine the global DNS's long-term security posture and cannot be addressed by incremental solutions. As stated in Section 5.8, SL-DNSSEC [16] is identified as the long-term architectural target because it is the only evaluated method that concurrently satisfies the UDP payload constraint, removes DDoS amplification ($BAF = 0.48$), and approximates session-level forward secrecy through per-session KEM encapsulation, a combination that Table 6 shows no signature-preserving proposal achieves even partially. This designation is, however, explicitly conditional rather than unconditional. Three barriers currently preclude production deployment: the initial DNSKEY retrieval size, the absence of IETF standardization, and, most fundamentally, the absence of a formal security proof for the KEM-MAC construction within the full DNSSEC chain-of-trust model (Sections 5.7 and 6.2.3). The resolution of precisely these barriers constitutes the explicit deliverables of Phase II: the recommendation of SL-DNSSEC as the Phase III target is contingent on the successful completion of the Phase II formal analysis objective, and Phase III therefore proceeds on a foundation of completed standardization and formal guarantees rather than the provisional basis that characterizes Phase I deployment. Should the Phase II security analysis uncover a fundamental flaw in the KEM-MAC chain-of-trust semantics, the contingency path is a continuation of the signature-preserving track, with QBF [53] the highest-scoring signature-based proposal in Table 6 as the default long-term architecture, accepting its residual amplification and forward-secrecy limitations as documented in Section 5.8. The external constraint governing Phase III is the NIST recommendation to deprecate RSA and ECDSA by 2030 and retire them by 2035. The one-to-two-year gap between the planned PQC Root KSK rollover and the 2030 deprecation deadline is a period during which the root zone will carry PQC signatures while classical algorithms remain operationally active at lower levels of the hierarchy, a transitional configuration that inherits the hybrid DNSSEC vulnerabilities established in Section 6.2. Managing this gap requires explicit policy coordination among ICANN, the IETF DNSOP working group, TLD operators, and major resolver operators; no existing institutional forum for this coordination has been established, and its creation is a prerequisite for an orderly Phase III transition.

A quantitative economic model for the aggregate infrastructure cost of full PQC DNSSEC migration constitutes the final deliverable of this phase and a prerequisite for evidence-based policy decisions about migration timelines and the distribution of migration costs across the DNS operator community. The component inputs to this model name-server resource requirements documented by Schutijser et al. [45], signature size increases documented by Suela et al. [44], and resolver fragmentation failure rates documented by Van den Broek et al. [46] are individually established in the reviewed literature but have not been integrated into a systems-level cost model scaled across TLD operators, root server operators, ISPs, enterprise DNS administrators, and major recursive resolver providers. Without this model, regulatory and standards organizations are unable to make

well-informed recommendations regarding the speed of the transition, and the risk of a technically inadequate or financially limited migration which, according to Mosca's theorem [9], would be a worse outcome than an orderly delayed migration remains unquantifiable.

Based on the analysis's findings, the three stages' dependency structure is a logical requirement rather than just an organizational convenience. Phase I provides the empirical measurements and IANA registrations that provide the operational basis for the formal analysis in Phase II, while Phase II generates the security proofs, protocol modifications, and Root KSK preparation necessary to advise the production deployment of SL-DNSSEC in Phase III. According to the requirements of Mosca's theorem [9], each stage must start before the previous one is finished. As the quantitative analysis in Section 7.2 shows, delaying migration until a complete solution is ready ensures that cryptographic compromise will occur before migration is finished. The dependency structure also supplies the roadmap's contingency mechanism for the scenario in which production-scale measurements diverge from the experimental findings on which Phases II and III are provisionally planned. The Phase I deliverables, namely the Tranco-scale fragmentation measurement, the resolver assessment, and hypothesis H1, constitute an explicit decision gate with two branches. If production behavior proves better than the experimental record (H1 falsified: FN-DSA-512 failure rates statistically indistinguishable from the ECDSA control), the urgency schedule is unchanged but mechanism selection simplifies: direct Falcon-512 deployment with TCP fallback becomes viable for a broader operator population, the fragmentation-avoidance mechanisms of Phase I are repositioned from default to optional, and the Phase III architecture decision is re-scored by re-evaluating Tables 6 and 7 with production-measured values, a re-evaluation whose outcome is genuinely open, since the sensitivity analysis of Section 5.10 shows that SL-DNSSEC's ranking advantage derives precisely from the transport and amplification criteria that this branch would relax. If production behavior proves worse, with failure or truncation rates substantially exceeding the 10% fragmentation and 8.5% silent-failure baselines of [43,46], then $t_{migrate}$ lengthens, the Mosca margin widens, and the roadmap responds by advancing the Phase II formal analysis triggers and prioritizing the UDP-fitting architectures. In neither branch is the urgency conclusion of Section 7.2 affected; what the measurements govern is mechanism selection and the architecture decision, both of which are deliberately gated on Phase I outputs rather than committed in advance.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/math14142555/s1>, File S1: Excluded Full-Text Studies—the complete list of 75 studies excluded at the full-text assessment stage, organized by exclusion category with the rationale for each exclusion.

Author Contributions: Conceptualization, M.I. and A.L.; methodology, M.I.; software, N.B.; validation, M.I., A.L. and N.B.; formal analysis, A.L. and N.B.; investigation, M.I. and A.L.; resources, A.L.; data curation, N.B. and M.I.; writing—original draft preparation, M.I., A.L. and N.B.; writing—review and editing, M.I. and A.L.; visualization, N.B.; supervision, M.I.; project administration, M.I.; funding acquisition, M.I. and A.L. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Shota Rustaveli National Science Foundation of Georgia (SRNSFG) N FR-24-15007.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The authors sincerely thank the editors and reviewers for their time, effort, and valuable comments, which helped improve the quality of this manuscript.

Conflicts of Interest: The authors declare no conflicts of interest. The Scientific Cyber Security Association is a non-profit research organization.

References

- Shor, P.W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.* **1999**, *41*, 303–332. [\[CrossRef\]](#)
- Mockapetris, P. *Domain Names-Concepts and Facilities*; RFC 1034; Internet Engineering Task Force (IETF): Fremont, CA, USA, 1987.
- Mockapetris, P. *Domain Names-Implementation and Specification*; RFC 1035; Information Sciences Institute: Marina del Rey, CA, USA, 1987.
- Bator, M.; Przysasz, J.; Serafin, M. Security of the DNSSEC protocol and its impact on online privacy protection. *Adv. Web Dev. J.* **2023**, *1*, 43–63.
- Man, K.; Qian, Z.; Wang, Z.; Zheng, X.; Huang, Y.; Duan, H. Dns cache poisoning attack reloaded Revolutions with side channels. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*; Association for Computing Machinery (ACM): New York, NY, USA, 2020; pp. 1337–1350.
- Liu, B.; Lu, C.; Duan, H.; Liu, Y.; Li, Z.; Hao, S.; Yang, M. Who is answering my queries: Understanding and characterizing interception of the {DNS} resolution path. In *Proceedings of the 27th USENIX Security Symposium (USENIX Security 18)*; USENIX Association: Berkeley, CA, USA, 2018; pp. 1113–1128.
- Arends, R.; Austein, R.; Larson, M.; Massey, D.; Rose, S. *DNS Security Introduction and Requirements*; RFC 4033; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2005.
- Mallick, T.; Kundu, A.; Kompella, R. Study of Post Quantum status of Widely Used Protocols. *arXiv* **2026**, arXiv:2603.28728.
- Nal, T.; Niemiec, M. Analysis and Evaluation of Post-Quantum Cryptography for DNSSEC. *Quantum Inf. Comput.* **2025**, *25*, 438–452. [\[CrossRef\]](#)
- Müller, M.; Thomas, M.; Wessels, D.; Hardaker, W.; Chung, T.; Toorop, W.; Rijswijk-Deij, R.V. Roll, roll, roll your root A comprehensive analysis of the first ever DNSSEC root KSK rollover. In *Proceedings of the Internet Measurement Conference*; Association for Computing Machinery (ACM): New York, NY, USA, 2019; pp. 1–14.
- Müller, M.; Toorop, W.; Chung, T.; Jansen, J.; van Rijswijk-Deij, R. The reality of algorithm agility Studying the DNSSEC algorithm life-cycle. In *Proceedings of the ACM Internet Measurement Conference*; Association for Computing Machinery (ACM): New York, NY, USA, 2020; pp. 295–308.
- Müller, M.; de Jong, J.; van Heesch, M.; Overeinder, B.; van Rijswijk-Deij, R. Retrofitting post-quantum cryptography in internet protocols a case study of DNSSEC. *ACM SIGCOMM Comput. Commun. Rev.* **2020**, *50*, 49–57. [\[CrossRef\]](#)
- Heftrig, E.; Shulman, H.; Waidner, M. Poster the unintended consequences of algorithm agility in DNSSEC. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*; Association for Computing Machinery (ACM): New York, NY, USA, 2022; pp. 3363–3365.
- Heftrig, E.; Shulman, H.; Waidner, M. Downgrading {DNSSEC} How to Exploit Crypto Agility for Hijacking Signed Zones. In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security 23)*; USENIX Association: Berkeley, CA, USA, 2023; pp. 7429–7444.
- Heftrig, E.; Shulman, H.; Waidner, M. Poster Off-path DNSSEC downgrade attacks. In *Proceedings of the ACM SIGCOMM 2023 Conference*; USENIX Association: Berkeley, CA, USA, 2023; pp. 1120–1122.
- Rawat, A.S.; Jhanwar, M.P. Quantum-safe signatureless dnssec. In *Proceedings of the 20th ACM Asia Conference on Computer and Communications Security*; Association for Computing Machinery (ACM): New York, NY, USA, 2025; pp. 267–282.
- Müller, M.; Chung, T.; Mislove, A.; van Rijswijk-Deij, R. Rolling with confidence: Managing the complexity of DNSSEC operations. *IEEE Trans. Netw. Serv. Manag.* **2019**, *16*, 1199–1211. [\[CrossRef\]](#)
- Rose, S.; Liu, C.; Gibson, R. *Secure Domain Name System (DNS) Deployment Guide*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2025.
- Arends, R.; Austein, R.; Larson, M.; Massey, D.; Rose, S.W. *Resource Records for the DNS Security Extensions*; RFC 4034; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2005.
- Chung, T.; van Rijswijk-Deij, R.; Chandrasekaran, B.; Choffnes, D.; Levin, D.; Maggs, B.M.; Mislove, A.; Wilson, C. A Longitudinal, {End-to-End} View of the {DNSSEC} Ecosystem. In *Proceedings of the 26th USENIX Security Symposium (USENIX Security 17)*; USENIX Association: Berkeley, CA, USA, 2017; pp. 1307–1322.
- Gudmundsson, O.; Wouters, P. *Managing DS Records from the Parent via CDS/CDNSKEY*; RFC 8078; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2017.
- Nosyk, Y. The Domain Name System a Tool and a Target for Internet-Wide Measurements. Doctoral Dissertation, Université Grenoble Alpes, Saint-Martin-d'Hères, France, 2024.
- Goldberg, S.; Naor, M.; Papadopoulos, D.; Reyzin, L.; Vasant, S.; Ziv, A. *NSEC5: Provably Preventing DNSSEC Zone Enumeration*; Cryptology ePrint Archive: Rochester, NY, USA, 2014.

24. Gieben, R.; Mekking, W. *Authenticated Denial of Existence in the DNS*; RFC 7129; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2014.
25. Bau, J.; Mitchell, J.C. *A Security Evaluation of DNSSEC with NSEC3*; Cryptology ePrint Archive: Ithaca, NY, USA, 2010.
26. Laurie, B.; Sisson, G.; Arends, R.; Blacka, D. *DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*; RFC 5155; Engineering Task Force (IETF): Fremont, CA, USA, 2008.
27. Chung, T.; van Rijswijk-Deij, R.; Choffnes, D.; Levin, D.; Maggs, B.M.; Mislove, A.; Wilson, C. Understanding the role of registrars in DNSSEC deployment. In *Proceedings of the 2017 Internet Measurement Conference*; Association for Computing Machinery (ACM): New York, NY, USA, 2017; pp. 369–383.
28. Lian, W.; Rescorla, E.; Shacham, H.; Savage, S. Measuring the practical impact of {DNSSEC} deployment. In *Proceedings of the 22nd USENIX Security Symposium (USENIX Security 13)*; USENIX Association: Berkeley, CA, USA, 2013; pp. 573–588.
29. van Rijswijk-Deij, R.; Sperotto, A.; Pras, A. DNSSEC and its potential for DDoS attacks a comprehensive measurement study. In *Proceedings of the 2014 Conference on Internet Measurement Conference*; Association for Computing Machinery (ACM): New York, NY, USA, 2014; pp. 449–460.
30. Shulman, H.; Waidner, M. Towards forensic analysis of attacks with DNSSEC. In *Proceedings of the 2014 IEEE Security and Privacy Workshops*; IEEE: Piscataway, NJ, USA, 2014; pp. 69–76.
31. Lenstra, A.K.; Lenstra, H.W. *The Development of the Number Field Sieve*; Springer Science & Business Media: Berlin, Germany, 1993; Volume 1554.
32. Barker, E. *Recommendation for Key Management: Part 1—General (Nist Special Publication 800-57 Part 1 Revision 5)*; National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2020.
33. Boudot, F.; Gaudry, P.; Guillevic, A.; Heninger, N.; Thomé, E.; Zimmermann, P. Comparing the difficulty of factorization and discrete logarithm: A 240-digit experiment. In *Proceedings of the Annual International Cryptology Conference*; Springer International Publishing: Cham, Switzerland, 2020; pp. 62–91.
34. Hankerson, D.; Vanstone, S.; Menezes, A. *Guide to Elliptic Curve Cryptography*; Springer New York: New York, NY, USA, 2004.
35. Kosek, M.; Doan, T.V.; Huber, S.; Bajpai, V. Measuring DNS over TCP in the era of increasing DNS response sizes a view from the edge. *ACM SIGCOMM Comput. Commun. Rev.* **2022**, *52*, 44–55. [[CrossRef](#)]
36. Grover, L.K. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*; Association for Computing Machinery (ACM): New York, NY, USA, 1996; pp. 212–219.
37. Arends, R.; Austein, R.; Larson, M.; Massey, D.; Rose, S.W. *Protocol Modifications for the DNS Security Extensions*; RFC 4035; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2005.
38. Bindel, N.; Herath, U.; McKague, M.; Stebila, D. Transitioning to a quantum-resistant public key infrastructure. In *Proceedings of the International Workshop on Post-Quantum Cryptography*; Springer International Publishing: Cham, Switzerland, 2017; pp. 384–405.
39. Weiler, S.; Blacka, D. *Clarifications and Implementation Notes for DNS Security (DNSSEC)*; RFC 6840; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2013.
40. Heftrig, E.; Schulmann, H.; Vogel, N.; Waidner, M. The harder you try, the harder you fail The keytrap denial-of-service algorithmic complexity attacks on dnssec. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*; Association for Computing Machinery (ACM): New York, NY, USA, 2024; pp. 497–510.
41. Gruza, O.; Heftrig, E.; Jacobsen, O.; Schulmann, H.; Vogel, N.; Waidner, M. Attacking with Something That Does Not Exist ‘Proof of {Non-Existence} Can Exhaust {DNS} Resolver {CPU}’. In *Proceedings of the 18th USENIX WOOT Conference on Offensive Technologies (WOOT 24)*; USENIX Association: Berkeley, CA, USA, 2024; pp. 45–57.
42. Karabulut, E.; Aysu, A. Falcon down: Breaking falcon post-quantum signature scheme through side-channel attacks. In *Proceedings of the 2021 58th ACM/IEEE Design Automation Conference (DAC)*; IEEE: Piscataway, NJ, USA, 2021; pp. 691–696.
43. Goertzen, J.; Thomassen, P.; Wisiol, N. Field experiments on post-quantum DNSSEC. In *Proceedings of the RIPE 89*; RIPE Network Coordination Centre (RIPE NCC): Amsterdam, The Netherlands, 2024.
44. Suela, J.G.; Blanco-Romero, J.; Mendoza, F.A.; Díaz-Sánchez, D. Implementing and Evaluating Post-Quantum DNSSEC in CoreDNS. In *Proceedings of the 2025 International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*; IEEE: Piscataway, NJ, USA, 2025; pp. 420–427.
45. Schutijser, C.; Koning, R.; Lastdrager, E.; Hesselman, C. Evaluating post-quantum cryptography in dnssec signing for top-level domain operators. In *Proceedings of the 2025 9th Network Traffic Measurement and Analysis Conference (TMA)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–10.
46. Van Den Broek, G.; van Rijswijk-Deij, R.; Sperotto, A.; Pras, A. DNSSEC meets real world dealing with unreachability caused by fragmentation. *IEEE Commun. Mag.* **2014**, *52*, 154–160. [[CrossRef](#)]
47. Fregly, A.; Harvey, J.; Kaliski, B.; Wessels, D. *Stateless Hash Based Signatures in Merkle Tree Ladder Mode (SLH DSA MTL) for DNSSEC*; Internet Engineering Task Force: Fremont, CA, USA, 2025.

48. Fregly, A.; van Rijswijk-Deij, R.; Müller, M.; Thomassen, P.; Schutijser, C.; Chung, T. *Research Agenda for a Post-Quantum DNSSEC*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2024.
49. Goertzen, J.; Stebila, D. Post-quantum signatures in DNSSEC via request-based fragmentation. In *Proceedings of the International Conference on Post-Quantum Cryptography*; Springer Nature: Cham, Switzerland, 2023; pp. 535–564.
50. Raavi, M.; Wuthier, S.; Chang, S.Y. Securing post-quantum dnssec against fragmentation mis-association threat. In *Proceedings of the ICC 2024-IEEE International Conference on Communications*; IEEE: Piscataway, NJ, USA, 2024; pp. 97–102.
51. Rawat, A.S.; Jhanwar, M.P. Post-quantum dnssec with faster tcp fallbacks. In *Proceedings of the International Conference on Cryptology in India*; Springer Nature: Cham, Switzerland, 2024; pp. 212–236.
52. Andrews, M.; Bellis, R. *A Common Operational Problem in DNS Servers: Failure to Communicate*; RFC 89062020; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2020.
53. Rawat, A.S.; Jhanwar, M.P. Post-quantum DNSSEC over UDP via qname-based fragmentation. In *Proceedings of the International Conference on Security, Privacy, and Applied Cryptography Engineering*; Springer Nature: Cham, Switzerland, 2023; pp. 66–85.
54. Fabrizio, G.; Koning, R.; Lastdrager, E.; Schutijser, C.; Sperotto, A.; van Rijswijk-Deij, R. PQC for DNSSEC a format size analysis on Falcon signatures. In *Proceedings of the 2025 Applied Networking Research Workshop*; Association for Computing Machinery (ACM): New York, NY, USA, 2025; pp. 143–149.
55. Nosyk, Y.; Korczyński, M.; Duda, A. Guardians of DNS integrity: A remote method for identifying DNSSEC validators across the internet. In *Proceedings of the 2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*; IEEE: Piscataway, NJ, USA, 2023; pp. 1470–1479.
56. Herzberg, A.; Shulman, H. Fragmentation considered poisonous, or One-domain-to-rule-them-all. org. In *Proceedings of the 2013 IEEE Conference on Communications and Network Security (CNS)*; IEEE: Piscataway, NJ, USA, 2013; pp. 224–232.
57. Rossow, C. Amplification hell: Revisiting network protocols for ddos abuse. In *NDSS*; Internet Society: Reston, VA, USA, 2014; pp. 1–15.
58. Stevens, M.; Bursztein, E.; Karpman, P.; Albertini, A.; Markov, Y. The first collision for full SHA-1. In *Proceedings of the Annual International Cryptology Conference*; Springer International Publishing: Cham, Switzerland, 2017; pp. 570–596.
59. Dawood, M.; Tu, S.; Xiao, C.; Haris, M.; Alasmay, H.; Waqas, M.; Rehman, S.U. The impact of Domain Name Server (DNS) over Hypertext Transfer Protocol Secure (HTTPS) on cyber security limitations, challenges, and detection techniques. *Comput. Mater. Contin.* **2024**, *80*, 4513–4542. [[CrossRef](#)]
60. Lee, J.; Hoque, S.; Aydeger, A.; Zeydan, E. Quantum-Resistant Domain Name System: A Comprehensive System-Level Study. *arXiv* **2025**, arXiv:2506.19943.
61. Shrishak, K.; Shulman, H. Negotiating pqc for dnssec. In *Proceedings of the 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks-Supplemental Volume (DSN-S)*; IEEE: Piscataway, NJ, USA, 2021; pp. 9–10.
62. Yang, H.; Osterweil, E.; Massey, D.; Lu, S.; Zhang, L. Deploying cryptography in Internet-scale systems A case study on DNSSEC. *IEEE Trans. Dependable Secur. Comput.* **2010**, *8*, 656–669. [[CrossRef](#)]
63. Kim, E.; Gupta, A.; Tsendjav, B.; Massey, D. Resolving islands of security problem for dnssec. In *Proceedings of the 2006 International Conference on Wireless Communications and Mobile Computing*; Association for Computing Machinery (ACM): New York, NY, USA, 2006; pp. 1271–1276.
64. Rishith, A.; Kulkarni, A.; Das, T.; Balachandran, V. Overcoming DNSSEC Islands of Security A TLS and IP-Based Certificate Solution. In *Proceedings of the 2024 IEEE Conference on Engineering Informatics (ICEI)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1–7.
65. Armakoon, J.; Paracha, U.; Iqbal, J.; Khan, M.A. Post-Quantum Cryptography From Theory to Practice. *Spectr. Eng. Sci.* **2025**, 892–918.
66. Hynek, K.; Vekshin, D.; Luxemburk, J.; Cejka, T.; Wasicek, A. Summary of DNS over HTTPS abuse. *IEEE Access* **2022**, *10*, 54668–54680. [[CrossRef](#)]
67. Pan, S.W.S.; Nguyen, D.D.N.; Doss, R.; Armstrong, W.; Gauravaram, P. Double-signed fragmented dnssec for countering quantum threat. *arXiv* **2024**, arXiv:2411.07535.
68. Raavi, M.; Khan, Q.; Wuthier, S.; Chandramouli, P.; Balytskyi, Y.; Chang, S.Y. Security and performance analyses of post-quantum digital signature algorithms and their TLS and PKI integrations. *Cryptography* **2025**, *9*, 38. [[CrossRef](#)]
69. Scrivano, A. A comparative study of classical and post-quantum cryptographic algorithms in the era of quantum computing. *arXiv* **2025**, arXiv:2508.00832.
70. Alagic, G.; Apon, D.; Cooper, D.; Dang, Q.; Dang, T.; Kelsey, J.M.; Lichtinger, J.; Liu, Y.K.; Miller, C.A.; Smith-Tone, D. *Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process*; National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2022.
71. Beernink, G. Taking the Quantum Leap Preparing DNSSEC for Post Quantum Cryptography. Master's Thesis, University of Twente, Enschede, The Netherlands, 2022.
72. Bonica, R.; Baker, F.; Huston, G.; Hinden, B.; Troan, O.; Gont, F. *IP Fragmentation Considered Fragile*; RFC 8900; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2020.

73. Dikshit, P.; Kosek, M.; Faulhaber, N.; Sengupta, J.; Bajpai, V. Evaluating DNS Resiliency and Responsiveness With Truncation, Fragmentation & DoTCP Fallback. *IEEE Trans. Netw. Serv. Manag* **2024**, *23*, 1145–1161. [[CrossRef](#)]
74. Bellare, M.; Miner, S.K. A forward-secure digital signature scheme. In *Proceedings of the Annual International Cryptology Conference*; Springer: Berlin/Heidelberg, Germany, 1999; pp. 431–448.
75. Howe, J.; Westerbaan, B. Benchmarking and analysing the NIST PQC lattice-based signature schemes standards on the ARM Cortex M7. In *Proceedings of the International Conference on Cryptology in Africa*; Springer Nature: Cham, Switzerland, 2023; pp. 442–462.
76. Kampanakis, P.; Lepoint, T. Vision paper Do we need to change some things? Open questions posed by the upcoming post-quantum migration to existing standards and deployments. In *Proceedings of the International Conference on Research in Security Standardisation*; Springer Nature: Cham, Switzerland, 2023; pp. 78–102.
77. Hülsing, A.; Butin, D.; Gazdag, S.; Rijneveld, J.; Mohaisen, A. XMSS *eXtended Merkle Signature Scheme*; RFC 8391; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2018.
78. McGrew, D.; Curcio, M.; Fluhrer, S. *Leighton-Micali Hash-Based Signatures*; RFC 8554; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2019.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.