Quantum Information & Computation
ISSN 1533-7146 Volume 25 (2025) pp.438–452
©Cerebration Science Publishing

Article

Analysis and Evaluation of Post-Quantum Cryptography for DNSSEC

Tomasz Nal and Marcin Niemiec *

AGH University of Krakow, Faculty of Computer Science, Electronics and Telecommunications, Mickiewicza 30, 30-059 Krakow, Poland
* Corresponding author: niemiec@agh.edu.pl

Received date: 11 June 2025; Accepted date: 24 August 2025; Published online: 30 September 2025

Abstract: The development of quantum computing poses a threat to many internet protocols because it undermines the security of current asymmetric cryptography. New types of algorithms, collectively known as post-quantum cryptography (PQC), are being developed and tested as potential replacements. Despite being engineered to counter quantum computers using Shor's algorithm, these systems frequently demand larger keys or signatures and substantial computational resources. This complicates the implementation of these algorithms in higher-level protocols, which requires thorough studies of the transition consequences. This paper evaluates the usefulness of one such PQC algorithm, FALCON-512, in the DNSSEC protocol. Using a containerized testbed, simulated DNS traffic is analyzed with a focus on key performance metrics, namely network latency and error rates, as the number of DNS clients in the network increases. The results show that PQC-enabled DNSSEC introduces higher error rates compared to today's algorithms, especially in NSEC responses used to deny the existence of DNS records, which represent a significant increase, resulting in the network's overload. The main contribution of this article is the validation of previous theoretical assumptions on the practical implications of FALCON-512 signature sizes in the implementation to DNSSEC. The differences in latency observed between nameservers in the DNS hierarchy during this study may inform DNS operators during the transition to PQC.

Keywords: post-quantum cryptography; FALCON; DNSSEC; benchmarking; latency; error rate

1. Introduction

The threat of an advanced quantum computer capable of outperforming classical computers in some particular applications (also known as *quantum supremacy*) has sparked a movement toward the adoption of post-quantum cryptography (PQC) algorithms—also known under a more descriptive name: quantum-secure algorithms [1]. These are algorithms that run on classical computers, but their design makes them resistant to cryptanalysis using both classical and quantum computers. In 2016, the US standardization body, the National Institute of Standards and Technology (NIST), launched a competition to determine the best replacement for classical cryptography primitives in the application of digital signatures and public-key encryption encapsulation mechanisms (KEM). After four rounds of competition, NIST has decided to standardize five algorithms: two key-exchange algorithms (Kyber and HQC) and three digital signature algorithms (Dilithium, FALCON, and SPHINCS+). Of those three digital signature algorithms, Dilithium and SPHINCS+ have already been standardized. FALCON is currently awaiting standardization and will be released as a draft FIPS 206 standard [2]. Following this decision is an extensive and ongoing effort to adapt and incorporate newly standardized algorithms into existing protocols. Such a transition to new algorithms can be a burdensome task, as very often it is not only a question of replacing one cryptographic primitive with another. Quantum secure algorithms are, in principle, characterized by a relatively large number of parameters, longer operating times, larger key sizes, and various limiting assumptions about how they can be generated and used securely [3]. At first, it might not be obvious why we should go through all this effort for standardization and post-quantum transition if the scientific community is not yet convinced when and in what form quantum computers will be a real threat to cryptography.

The reason for this is that the effect of replacing current algorithms with quantum-secure ones is not yet very well understood. A lot of time and real-world testing is needed to select the best parameters or adapt existing protocols to balance the security and usability for the end-users. Another threat of the advent of quantum computers analyzed

DOI: 10.2478/qic-2025-0025

by NIST is the so-called *harvest now, decrypt later* attack (also known as *store now, decrypt later*). The rationale behind this cryptographic attack is the observation that an attacker might not have the possibility to obtain the secret key and break the encryption of a given scheme now, but such a possibility might materialize itself in the future—either in the form of a vulnerability in the encryption scheme or by the growing capabilities of computing power. This is exactly the case for quantum computers—as of now, we do not have enough quantum computing power to break classical encryption, but a threat actor might store the data to decrypt it later once the quantum computer's computing power reaches its usable state. For this reason, it is crucial not to wait too long for the standardization and practical implementation of post-quantum cryptography. The problem, as mentioned above, is commonly referred to as Mosca's theorem [4]. The idea behind it can be summarized by the two values representing time—how long the encrypted data must be kept securely and how much time is needed to transition to quantum-secure schemes. By estimating those values and combining them, we get the approximated expectation for the adoption of post-quantum cryptography for a given use case:

- x years—How long must the encrypted data be guarded confidently?
- y years—How long does it take to create new quantum-safe solutions and implement them?
- z years—How long will it take until a sufficient quantum computer is built?
- If $y + x \geq z$, then the secret keys are revealed.

The rationale behind Mosca's theorem should give a transition schedule for a given use case. Special care should be dedicated to protocols that are expected to pose more problems in the process, as adopting post-quantum cryptography into high-layer protocols is not always straightforward—just replacing one primitive with another will not always be enough. One such example is the DNSSEC protocol. It is a security extension of the Domain Name System (DNS) protocol. The DNS is responsible for translating human-readable domain names into IP addresses.

1.1. DNSSEC

The Domain Name System Security Extensions Protocol (DNSSEC) enhances traditional DNS by adding authenticity and integrity to DNS responses [5]. These properties are achieved through the use of public-key cryptography, more specifically, digital signature schemes. When the requester receives a DNS response in return for a query, they can verify a digital signature attached to it using a public key published in the DNS hierarchy. This verification ensures that the client is certain that the response has not been tampered with, and it originates from an authoritative source. In DNS, records such as A and AAAA are used to map domain names to IP addresses. An A (Address) record maps a domain name to an IPv4 address, while an AAAA record maps a domain name to an IPv6 address. These are among the most commonly queried record types and are critical for Internet communication. However, DNSSEC does more than secure positive responses. It also provides authenticated denial of existence (the so-called *proof of non-existence*), which prevents attackers from forging non-existent domain responses. This is achieved through the NSEC (Next Secure) and NSEC3 resource records. An NSEC record explicitly states which names do not exist by indicating the next existing name in the DNS zone. When a resolver queries a non-existent domain name, it receives an NSEC record that proves the queried name falls between two known names, and thus does not exist. These authenticated denials are vital for preventing, for example, cache poisoning attacks, but they also introduce performance considerations, especially with post-quantum digital signature schemes, which often produce larger signatures. Each NSEC or NSEC3 record may be accompanied by multiple signatures (RRSIG). This amplifies the bandwidth and latency impact of DNSSEC, particularly when switching from traditional algorithms like ECDSA or RSA to post-quantum algorithms such as FALCON, characterized by larger signature sizes. The above considerations highlight the importance of integrating post-quantum cryptography into DNSSEC in a way that preserves its core security guarantees without undermining usability or performance. The operations performed by the DNSSEC protocol should be modified with careful consideration for the end user—adding new post-quantum algorithms should not unreasonably increase network latency, the number of errors, network overhead, etc.

1.2. FALCON

Fast Fourier Lattice-based Compact Signature over NTRU (FALCON) is a post-quantum digital signature scheme based on the hardness of problems over NTRU lattices. Designed for efficiency and compactness, FALCON employs fast Fourier sampling techniques to generate short signatures with strong security guarantees. Among the finalists of the NIST Post-Quantum Cryptography Standardization Project, FALCON was selected for standardization due to its favorable balance of performance, security, and particularly its compact signature size—an advantage especially notable when compared to other schemes such as Dilithium [6]. It is currently awaiting finalization in the forthcoming

NIST draft [7]. FALCON supports two security levels [8], with FALCON-512 targeting NIST Level 1 security—defined as resistant to an attack comparable to one requiring computation resources comparable to or greater than those required for a key search on AES-128. The FALCON-512 variant achieves this by using a lattice dimension of 512, resulting in a public key size of 897 bytes and a signature size of 666 bytes. Its compactness and efficiency make it well suited for constrained environments, such as embedded systems and IoT devices, while maintaining robustness against classical and quantum adversaries. This security represents the lowest of the five security levels defined by NIST in its Post-Quantum Cryptography Standardization Project, suitable for applications where moderate security suffices, but efficiency and compactness are critical. The second parameter set used by FALCON-1024, which has a lattice dimensionality of 1024, achieves the highest NIST security level, Level 5, which is comparable to AES-256 security. FALCON-1024 is double the size of FALCON-512, with a public key size of 1,993 bytes and a signature size of 1,280 bytes [8]. However, this makes the second variant incompatible with the minimal theoretical requirements of DNSSEC, since the size of the public key and the digital signature exceed the theoretical limit of 1,223 bytes [9]. For this reason, the paper will focus solely on the FALCON-512 variant.

2. Related Work

Prior to the conclusion of the Post-Quantum Cryptography Standardization Project, Müller et al. [9] showed that the requirements of most NIST candidate algorithms make them unsuitable for use in DNSSEC. The primary reason is the public key or signature size generated by these algorithms. For instance, NIST's chosen digital signature standard, Crystals-Dilithium, has a 2 kB signature size, which exceeds the minimum of 1,232 B established by the paper's authors and is 32 times larger than the 64 B signature size of the IETF-recommended classical algorithm, EdDSA-Ed22519. SPHINCS+ has an even larger signature size of almost 8 kB. Of the digital signature algorithms considered, FALCON-512 was the only one that was, at least in theory, compatible with the minimal requirements. Because implementing post-quantum digital signatures in DNSSEC is complex, other prospects have been studied. The authors proposed changes to the DNSSEC protocol, such as out-of-band key distribution, in which the resolver fetches the key from an HTTP web server. The paper also briefly mentioned the potential problem with NSEC and NSEC3 responses, which require the transmission of up to 2 signed responses and up to 3 signed responses, respectively. This implies that making use of FALCON-512 would not fit in the message size of 1,232 B. It suggests that some protocol-level adjustment would be needed even if using FALCON-512, whose signature in theory fits in the limits. There exist some ideas, like modifications to how NSEC operates, that would require sending only one signature [10]. These findings and conclusions were further expanded and confirmed by Müller in his PHD thesis [11], again highlighting the signature size as the problematic component and suggesting the movement toward the out-of-bound key distribution. He pointed at the FALCON-512 variant to be the best pick as it is the only one with the possibility to be deployed to DNSSEC without protocol modifications, but still having significantly larger keys and signatures than the algorithms currently in use. Two other alternatives he proposed, which were at the time still considered by NIST, were Rainbow-Ia and RedGeMSS128; however, since then, Rainbow-Ia has been broken [12,13], and RedGeMSS128 was rejected by NIST in part due to performance considerations [13]. In his master's thesis, G. Beernink [14] explored the prospects of out-of-band key distribution. From his work, he concluded that out of several tested PQC algorithms (Dilithium and SPHINCS+ included), only FALCON-512 did not exceed the maximum packet size allowed for DNS responses, from which he implied that only that algorithm can be used as an alternative without redesigning the DNSSEC protocol. He also found that regarding the computational overhead on zone servers and resolvers verifying the signature, the algorithm Rainbow-Ia performed best (however, this algorithm was later proven to be not secure against a classical attack [12,13]). The IETF draft "Research Agenda for a Post-Quantum DNSSEC" [15] highlights other ideas, such as fragmenting DNSSEC data and splitting it across multiple Resource Records (RRs). This idea reached an empirical evaluation by Goertzen et al. [16] who proposed a DNS fragmentation scheme, called "A Resource Record Fragmentation mechanism" (ARRF), in which the RRs are fragmented at the application layer to fit into UDP limits. The client retrieves the response for each fragment one at a time. The prospect of response fragmentation was further expanded in Goertzen's master's thesis [17]. More studies are building up on this idea, for example, Raavi et al. [18] have developed a method called commit-and-reveal that, through fragment chaining and blockchain, prevents some known attacks on the DNS response fragmentation.

3. Motivation and Methodology

This section provides the main rationale behind the conducted experiments, explaining their purpose and significance in the broader context of the post-quantum transition. It also outlines the expected results, the hypotheses being tested, and the specific objectives that the experiments aim to achieve. By doing so, it sets the stage for understanding the methodology and the importance of the findings presented later.

3.1. DNSSEC and Post-quantum Migration

According to the findings of the preceding section, FALCON in its 512 lattice dimension variant is the only option that, at least in theory, is compatible with the minimal DNSSEC requirements. The objective of this study is to analyze the impact of transitioning from DNSSEC to FALCON-512 on latency and error rates—to assess the behavior of the system under consideration both on a per-request basis and to verify how the characteristics of a network will change with the increasing number of PQC-enabled nodes in the network. The investigation will also focus on isolating the case of NSEC responses to assess the influence of the mechanism responsible for sending multiple signatures. These signatures cannot fit into a single message and require fragmentation, which is handled unreliably by UDP. In some cases, this may result in a fallback to TCP, although this is not always supported. The objective of this study is to determine whether triggering NSEC responses will have a more significant impact on latencies or error rates than the scenario of regular DNS requests that do not trigger NSEC. The idea for the study was given by the IETF Research Draft [15]. The document describes challenges and expected issues that post-quantum cryptography will present to DNSSEC, various ways to address this impact, and proposals for research activities that should be undertaken to best prepare for a transition to post-quantum digital signatures in DNSSEC. One of the proposed research activities invoked by the authors of the draft is the following:

DNS response time impact for common use cases (browsing, web services, mobile, IoT,...). It will likely take some creative approaches to broadly assess this so that results are reflective of the "real" user experience. [15]

The above draft's proposal was selected as the basis for this project, with the primary aim of evaluating the integration of the FALCON-512 post-quantum algorithm within a DNSSEC implementation. Rather than replicating a fully realistic, production-level scenario, the experiment focused on establishing an initial, simplified setup to assess the algorithm's feasibility and performance, along with the selection of benchmarking tools for a proof-of-concept deployment. This preliminary approach was designed as a foundational step, prioritizing insights into system configuration and response latencies over a comprehensive simulation of real-world traffic conditions. That being said, the configuration and results from this project can later be expanded to better reflect real-world conditions.

3.2. The PATAD Project

The experiment setup lab is based on the PATAD project by SIDN Labs (PATAD—Post-quantum Algorithm Testing and Analysis for the DNS)—a small-scale testbed for empirical investigation of the post-quantum algorithms [19]. The researchers from SIDN Labs shared a fork of the open-source DNS software PowerDNS with support for post-quantum algorithms in the DNSSEC protocol. The project contains the complete setup necessary to run a DNS testbed with support for post-quantum algorithms. Currently, the algorithms supported are FALCON-512, SQIsign1, and Mayo2. The authors have shared an example lab setup with all three algorithms used in a DNS zone signing at different stages. The project and testbed setup is recommended for hands-on testing of the new post-quantum algorithms in a lab environment. The authors have already provided some findings discovered during the testing, mainly consisting of the zone file increases in DNS and validation speed comparison between the classical algorithm and the post-quantum one.

For the purpose of the testing scenarios, in order to compare the results obtained for FALCON-512, the PATAD testbed was modified to also include the currently recommended algorithm for DNSSEC—ECDSA-P256 (Elliptic Curve Digital Signature Algorithm), using the implementation of the P256 elliptic curve [20].

4. Testbed Configuration

The main goal of this project is to benchmark the usage of FALCON-512 in the DNSSEC protocol. The secondary objective was to evaluate the usage of the PATAD environment and adapt it as a benchmarking solution for testing PQC algorithms in DNSSEC. The testbed configuration was modified to include only the relevant signing algorithms on all authoritative nameservers and the recursor. Additionally, some PowerDNS configuration options were modified according to the requirements of the experimental scenario—one of these changes was to limit cache usage as much as possible to ensure consistent results throughout the runs. Changes in configuration settings are presented in Listings 1 and 2.

Listing 1: Configuration settings in pdns.conf file used by the authoritative nameservers to control cache usage.

```
# pdns auth config
cache-ttl=0
query-cache-ttl=0
negquery-cache-ttl=0
```

Listing 2: Configuration settings in pdns.conf file used by the recursor to control cache usage.

```
# pdns recursor config
disable-packetcache=yes
max-cache-ttl=0
max-cache-bogus-ttl=0
max-cache-entries=0
max-negative-ttl=0

packetcache-ttl=0
packetcache-negative-ttl=0
packetcache-servfail-ttl=0
max-packetcache-entries=0

aggressive-nsec-cache-size=0
nothing-below-nxdomain=no
```

It is not feasible to disable the cache in its entirety because PowerDNS depends on it for its functionality. However, the selected parameters are designed to restrict caching to the greatest extent possible. This approach enables the observation of uninfluenced parameters or with limited influence, including actual network latencies, internal server processing time, error rates, packet propagation through the network, retransmissions, and other factors that would otherwise be obscured by caching mechanisms [21]. In basic terms, with caching disabled, each error should be made even more visible, as caching would not mitigate its consequences. Consequently, other studies could concentrate on the optimization and design of caching mechanisms to potentially improve latency parameters. These studies could consider broader factors, such as internal key generation times and signing and validation processes on DNS servers.

To measure DNS query latency, the dnspyre benchmarking tool was used [22]. It is a command-line tool that sends DNS queries to the server and performs various measurements based on the response received. Its variety of configuration options allows for adjusting traffic to specific experiment scenarios, such as simulating a specific number of nodes in the network or setting the required query record type.

Below are the dnspyre parameters used for the experiments, along with the proper explanation.

`--duration 450 s`—It configures the execution time of the benchmark scenario. In this example, the benchmark will run for 7.5 minutes while sending DNS requests in a loop based on the data source provided.

`--dnssec`—It sets the DO bit for all DNS requests to 1, which enables DNSSEC.

`-c 10`—This parameter instructs dnspyre to use the given number of workers when executing. It is used to configure scenarios closer to real-world experience, representing multiple nodes sending DNS queries in the network.

`-t A`—It sets the DNS request type (e.g., A, AAAA).

`--separate-worker-connections`—It instructs dnspyre not to try to share connections between concurrent workers.

The network configuration of the testbed is presented in Table 1. The container images are identified as shared by the authors of the PATAD project—`pqc-resolver-powerdns` for the resolver [23], and `pqc-auth-powerdns` for the authoritative nameserver [24]. Table 2 contains the DNS zone configuration used and the list of hosts in the experiments. In essence, there is a `root` domain, a `.pl` top-level domain, with a `dnsseclab.pl` second-level domain that has a few subdomains. All containers have established support for IPv4 and IPv6 and accept requests on UDP and TCP ports. By default, in the experiments, DNS via UDP is used, but TCP support is guaranteed in the case of a TCP fallback, e.g., in the case when the response is too large to fit within UDP bounds.

Table 1. Testbed networking configuration.

DNS architecture	*Image*	*Port*
.	pqc-auth-powerdns	5302, tcp/udp
pl.	pqc-auth-powerdns	5303, tcp/udp
dnsseclab.pl.	pqc-auth-powerdns	5304, tcp/udp
resolver	pqc-resolver-powerdns	5311, tcp/udp

Table 2. Testbed zone configuration.

Zone	*Host*	*IPv4/IPv6*
.	s.root-servers.net.	10.0.1.2 / fc01::2
pl.	ns1.example.pl.	10.0.1.3 / fc01::3
dnsseclab.pl.	n21.dnsseclab.pl.	10.0.1.4 / fc01::4
resolver	-	10.0.1.10 / fc01::10
resolver (dnssec)	-	10.0.1.11 / fc01::11

The characteristics of the hardware environment are presented in Table 3. The metrics were collected through the Prometheus data endpoint using Prometheus-formatted metrics. Exposure of the metrics by the DNS server requires setting in the `pdns.conf` file of the options: `webserver=yes` and `webserver-port=8081` [25]. The experiment consisted of repeated DNS queries sent for a specified duration, with an increasing number of workers in the network at each iteration. The code in Listing 3 presents the most important fragments of the script used for data collection, which should give the idea of the experiment.

Table 3. Hardware environment—the testbed is running inside Podman containers, on a virtual machine with Ubuntu Linux kernel 6.8.0-54-generic.

Cores	*Number of CPUs*	*RAM*
11th Gen Intel(R) Core(TM) i7-1185G7 @ 3.00GHz	4	8 GB

Listing 3: Key parts of the Bash script showing the overall flow of the experiments and data collection process.

```
#!/bin/bash
PROM_URL="http://localhost:9090"

for concurrent in {10..100..10}; do
    C_VALUE=$concurrent

    for iteration in {1..3}; do
        # Reset the environment state.
        # ...

        # Start the experiment
        dnspyre --duration 450s -c ${C_VALUE} --server [::1]:5311 wiki.
            dnsseclab.pl --dnssec -t AAAA --no-distribution --separate-worker
            -connections


        # PROM queries
        q1='pdns_recursor_noerror_answers{instance="localhost:8091", job="
            powerdns_resolver_dnssec"}'
        q2='pdns_recursor_servfail_answers{job="powerdns_resolver_dnssec"}'
        q3='pdns_recursor_dnssec_queries{job="powerdns_resolver_dnssec"}'
        q4='avg_over_time(pdns_auth_latency{job="powerdns_auth_dnsseclab"}[5m
            ])'
        # ...

        NOERROR=$(curl -sG "${PROM_URL}/api/v1/query" --data-urlencode "query
            =${q1}" | jq -r '.data.result[0].value[1]')
        # ...

        # Append the output to the file.
        # ...
    done
done
```

5. Experiment Design

This section contains detailed descriptions of the characteristics being tested during the experiments, the exact research questions, and the general context behind them. The main goal of the experiments is to expand on theoretical requirements for post-quantum transition research activities [15], theoretical assumptions [9], and some research results that have already been conducted and presented [26].

5.1. Latency

Due to the inherent complexity of post-quantum cryptography, an important parameter to evaluate those algorithms is the latency compared to DNSSEC secured with classical algorithms like ECDSA. The PATAD environment facilitates the evaluation of latency metrics at each level of the DNS architecture. The objective of the present series of experiments is to determine whether there is a difference in the manner in which authoritative nameservers manage traffic and the impact this has on latencies after the implementation of the FALCON-512 algorithm compared to the classical ECDSA. The goal of the experiment will be to measure the latency introduced at the level of PowerDNS Authoritative Servers. In DNS, when a request is recursively passed through multiple nameservers to reach the authoritative nameserver, each step adds latency due to processing time. During the post-quantum transition, it is crucial to anticipate the hardware modifications that may be required. Understanding the latency differences introduced by classical and post-quantum DNSSEC algorithms is a valuable first step in this process. The experiments

will be conducted using identical parameters for both IPv4 and IPv6 deployments, first with FALCON-512 and then with ECDSA. Each experiment will involve 7.5 minutes of querying by a preset number of nodes. After this period, the value of the metric `pdns_auth_latency` is recorded. This metric represents the average number of microseconds a packet spends within PowerDNS [27], and its value is calculated as the average over the last 5 minutes of the experiment (to exclude the initial period before resource saturation). Each experiment is repeated 12 times for each node count in the range of 10 to 100.

5.2. Error Rates

As expanded in previous sections, although the FALCON-512 algorithm is technically a good fit for implementation as it meets the minimal constraints on key and signature sizes, the margin is so small that it is believed to cause problems in practical deployments [9]. This argument will be examined in this work to determine if there is a relevant difference and how the error rate will differ from the classical algorithms. The second experiment will measure the error rates. It is carried out using the same parameters and setup as in the previous experiment, with one exception: in the final step, the metric recorded is the number of occurrences of the `SERVFAIL` error. As before, each experiment is repeated 12 times for each node count, with the number of nodes ranging from 10 to 100.

5.3. NSEC Records

NSEC responses are expected to be a major issue for FALCON-512 as they require the transmission of multiple signatures, which can prevent the reliable functioning of the DNS protocol. In this experiment, it will be assessed how having queries that trigger NSEC responses influences the workings in the network, both in terms of latency and error rates. In this experiment, instead of querying a single domain, four domains are used: three valid domains and one non-existent domain. The non-existent domain triggers the server to return NSEC records. Each experiment iteration runs for 2 minutes, after which the metric counting the number of `SERVFAIL` responses is recorded. This number is then compared with the total number of DNSSEC queries sent, and the result is expressed as an error rate percentage.

6. Results

This section contains a description of the experiments conducted, the parameters used, and the results obtained.

6.1. Experiment 1—Nameserver Latency

The results, after statistical analysis, are presented in Figure 1 and Figure 2 for FALCON-512, and in Figure 3 and Figure 4 for ECDSA. Although the level of noise, represented by the confidence interval, prevents a clear inference regarding the correlation between the number of querying threads and the added latency at the nameserver, a notable trend can still be observed. Specifically, the architecture component most responsible for the added latency differs between the deployments. In the case of FALCON-512, it is the `pl` server (the top-level domain), while in the ECDSA deployment, it is the `root` server (the root domain). For ECDSA the trend is only present when DNS type A queries are used. For DNS type AAAA queries the level of noise is too high to infer the exact differences.

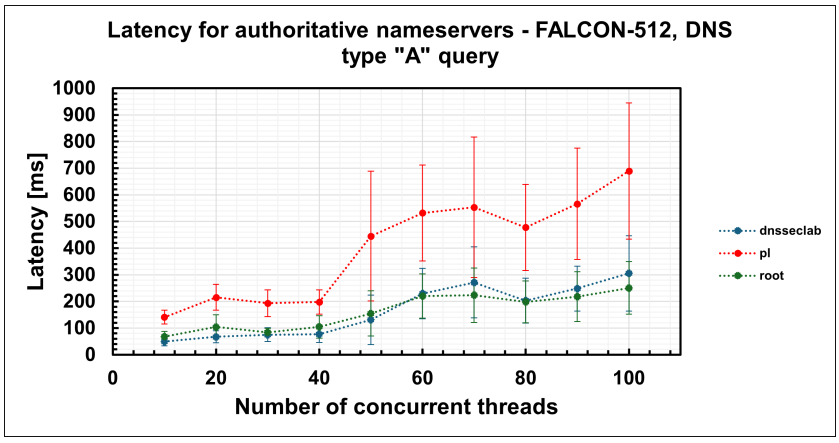

Figure 1. Latency measured at the authoritative nameservers for FALCON-512 with A-type DNS queries.

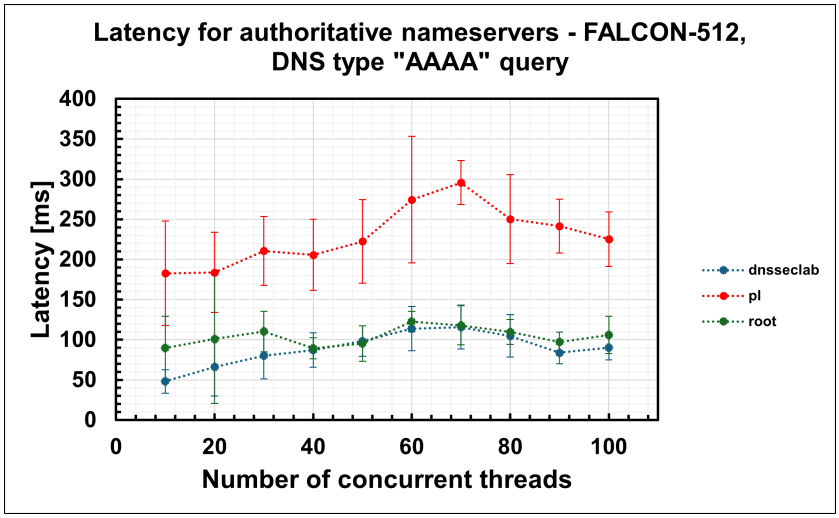

Figure 2. Latency measured at the authoritative nameservers for FALCON-512 with AAAA-type DNS queries.

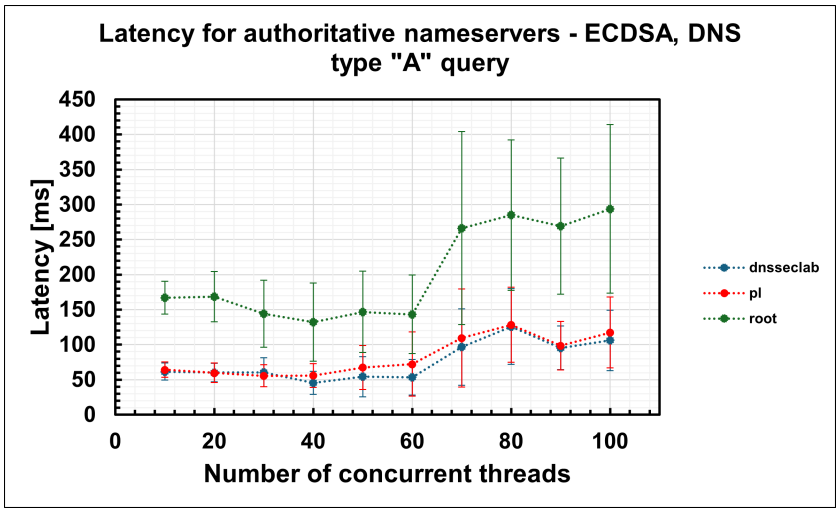

Figure 3. Latency measured at the authoritative nameservers for ECDSA with A-type DNS queries.

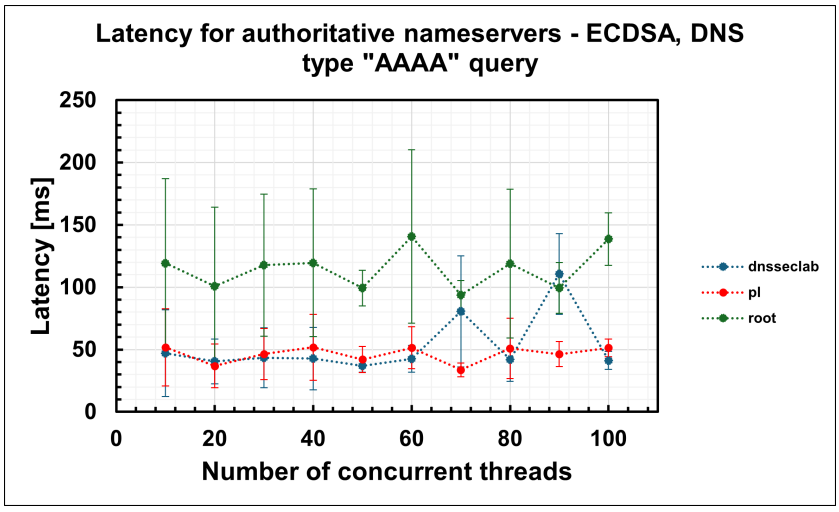

Figure 4. Latency measured at the authoritative nameservers for ECDSA with AAAA-type DNS queries.

6.2. Experiment 2—Error Rates

The results, including statistical analysis, are shown in Figure 5 for the IPv4 configuration and Figure 6 for the IPv6 configuration. For the selected range of concurrent threads, no instances of `SERVFAIL` errors were observed when using DNSSEC with the ECDSA algorithm.

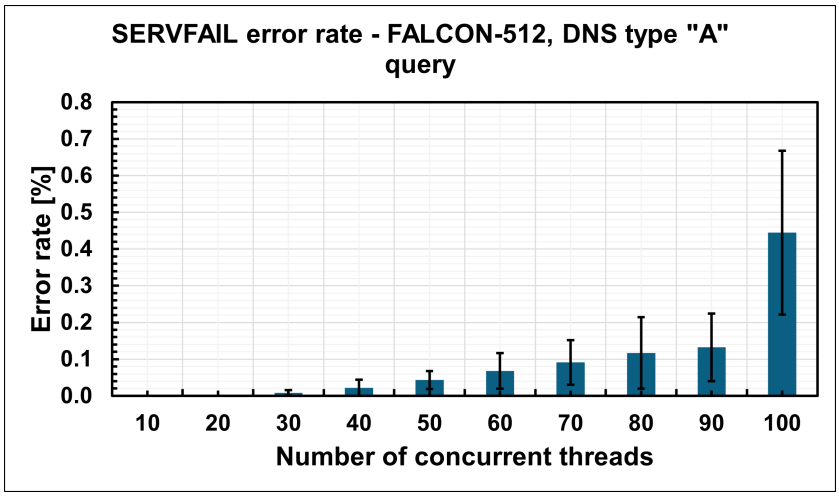

Figure 5. Error rate measured at the authoritative nameservers for FALCON-512 with A-type DNS queries.

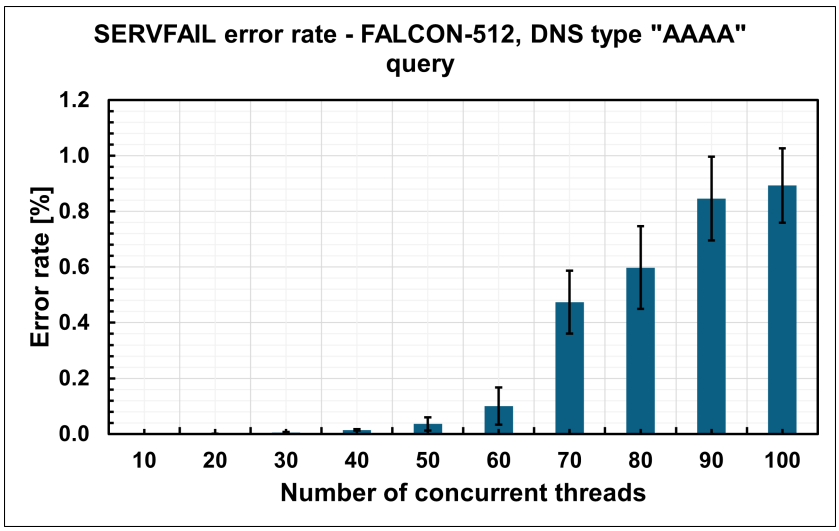

Figure 6. Error rate measured at the authoritative nameservers for FALCON-512 with AAAA-type DNS queries.

6.3. Experiment 3—NSEC

The results are shown in Figure 7 and can be compared to the `SERVFAIL` error rates for standard A-type requests shown in Figure 5. In addition, latency data for each PowerDNS authoritative nameserver was recorded during each experiment run. The latency results for FALCON-512 are presented in Figure 8, and the corresponding results for ECDSA are shown in Figure 9. Unlike the previous experiments, these results indicate that the component contributing the most to latency remains the same across both classical and post-quantum deployments. In both cases, it is the `dnsseclab` server, responsible for the second-level domain. However, a notable difference is that servers that support the FALCON-512 algorithm exhibit both higher latency and greater variability.

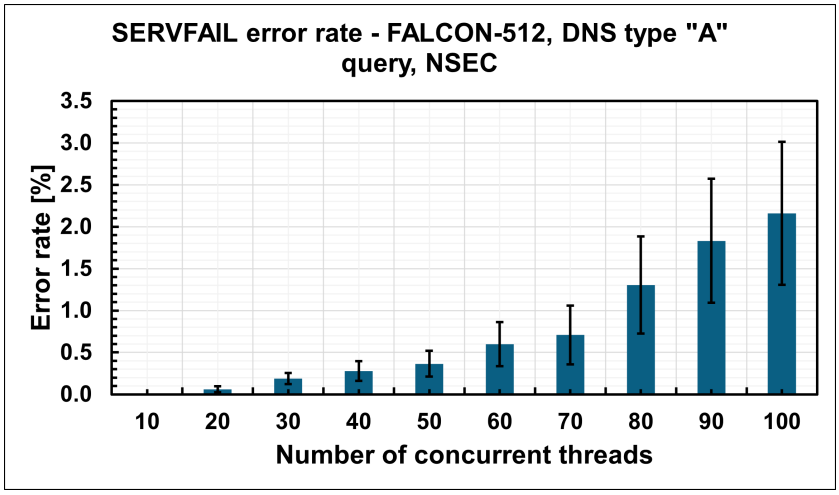

Figure 7. Error rate with queries for non-existent domains.

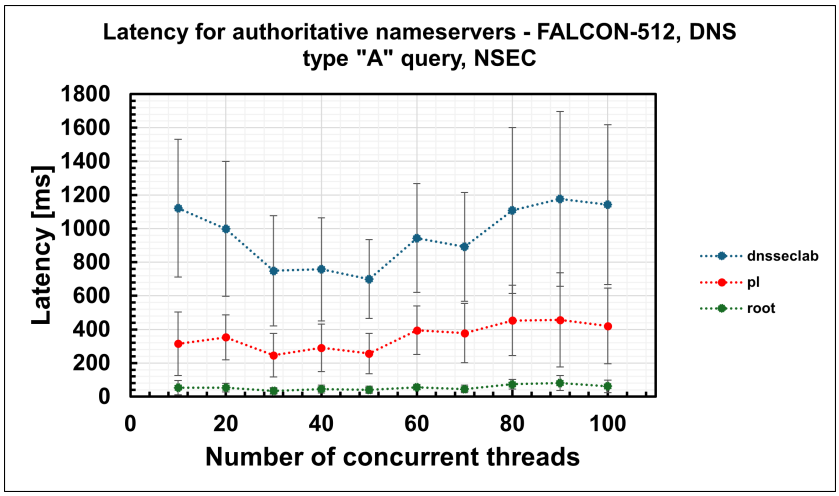

Figure 8. Latency measured at the authoritative nameservers for FALCON-512 with A-type DNS queries for non-existent domains.

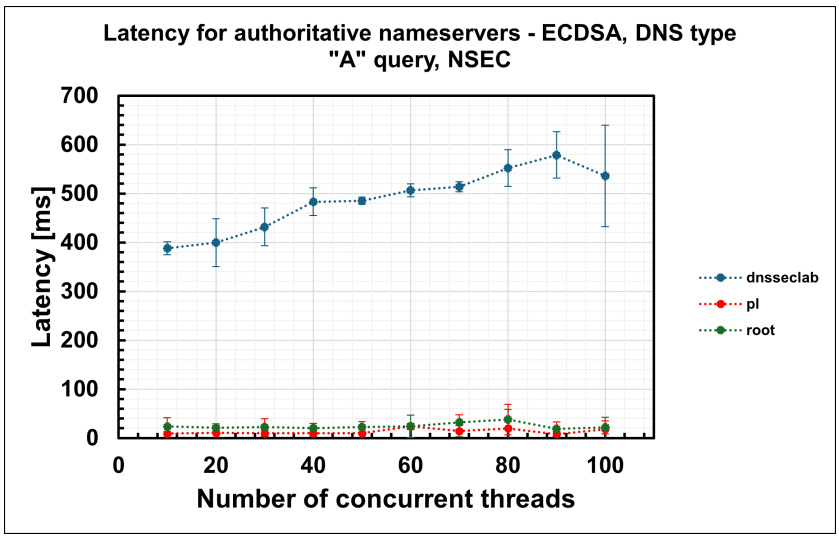

Figure 9. Latency measured at the authoritative nameservers for ECDSA with A-type DNS queries for non-existent domains.

7. Analysis

The result represents valuable insights into the prospects of the implementation of post-quantum cryptography in the DNS protocol. This section analyzes the collected data and allows the implementation of the PQC algorithm in the DNSSEC protocol to be evaluated.

7.1. Latency

A surprising result was that there is no indication of a large increase in latencies when switching to FALCON-512, in the scenario where no NSEC responses are triggered. However, the large noise in the data, represented by the wide confidence intervals, makes it impossible to infer the exact differences. It is clear that in the selected experiment scenarios, the latency differences were not very noticeable. Notable increases in latency could only be observed when including queries for non-existent domains. These queries are closest to a real-world implementation because the NSEC mechanism is an inherent part of DNS. The results indicate a difference in the influence of the authoritative nameserver on overall latency. In the case of ECDSA, the highest latency values were measured in the `root` nameserver (the root domain); however, in the case of FALCON-512, it is the `pl` (the top-level domain) that scored the highest values. It illustrates how, in the event of deployment of this or similar algorithms for DNSSEC hardware requirements and network load projections should be carefully studied, as it may differ significantly from the case of classical algorithms. The scenario that measured the influence of NSEC records has some different observations—in both FALCON-512 and ECDSA experiments, it was the `dnsseclab` nameserver (the second-level domain) that scored the highest latency values.

7.2. Error Rates

In all experiment scenarios, FALCON-512 demonstrated growing error rates—specifically `SERVFAIL` responses—which increased significantly as the number of nodes grew. This trend suggests that FALCON-512 introduces operational overhead or compatibility issues that may not scale with the current protocol configuration. In contrast, ECDSA remained stable throughout the experiments, with no recorded instances of SERVFAIL errors, even under increased node counts. These results raise concerns about the robustness of current FALCON-512 implementations in DNSSEC deployments. Current FALCON-512 performance and reliability issues in larger networks suggest that further optimization and testing are necessary before it can be considered a viable replacement.

7.3. NSEC Records

The scenario of measuring NSEC records proved to be a worrying influence for both metrics considered. When assessing `SERVFAIL` error rates, their contribution was higher for the scenario including NSEC records than in the simple query case. Also, the latency level and latency variability were higher for queries including non-existent domains compared to the first one. Both of those experiment results demonstrate difficulties of DNS transition to FALCON-512 or similar algorithms, as their usage represents a significant increase in both error rates as in latency.

8. Summary

The experimental results presented in this article indicate that FALCON-512 does not introduce significant latency overhead compared to ECDSA in the scenarios tested that do not trigger NSEC responses. It is mainly the NSEC mechanism that is responsible for the largest differences in latencies from the tested scenarios. This suggests that, in terms of response times, FALCON-512 could theoretically be suitable for DNSSEC if some alternative to NSEC is implemented. However, this potential advantage is overshadowed by a considerably higher error rate. In particular, the number of SERVFAIL errors increased rapidly with the number of concurrent workers. Most importantly, NSEC testing revealed high error rates and latency, which would likely make the deployment of FALCON-512 in DNSSEC infeasible under current implementations. These issues point to serious reliability concerns that would need to be addressed before any practical use. Despite this, the experiments helped to identify the circumstances in which FALCON-512 performs well and where it fails, contributing valuable information on the strengths and limitations of this and similar algorithms.

8.1. Limitations

A potential limitation of the study is the set of hardware constraints under which the measurements were performed. The experimental platform relied on limited resources, and during peak testing periods, CPU utilization and I/O load were often close to saturation. Such conditions may have obscured latency differences by introducing

additional, non-cryptographic delays. However, both algorithms were affected, making it possible to perform a comparative analysis. Occasionally, the unusual variability in results rendered certain assessments impossible, such as the direct comparison of latency levels between classical and post-quantum implementations. We were cautious in our analysis to avoid drawing any statistically significant conclusions that might have been affected by the low confidence in the results. The influence of hardware constraints was mitigated through multiple repetitions of the experiments.

A second limitation concerns the interpretation of PowerDNS performance counters. The official documentation provides only short descriptions of many internal metrics, and during the analysis, several inconsistencies became apparent. For example, the metric `auth-latency` did not equal the sum of `auth-send-latency` and `auth-receive-latency`, even though, by definition, it should. A possible explanation is that individual counters are updated with different temporal granularity, so that snapshots taken at a millisecond scale capture some values before others have been committed. Without precise knowledge of the underlying update logic, it is impossible to determine how much of the observed variance is a result of real network delays and how much is a result of a side factor.

8.2. Future Work

A key direction for future work is to understand the source of the high error rates observed, particularly the SERVFAIL responses. This requires a more in-depth examination of the authoritative nameservers' behavior and potentially the internals of the FALCON-512 implementation. Furthermore, the high error rates in scenarios involving NSEC must be thoroughly investigated, as they currently pose a significant obstacle to the adoption of this and similar algorithms in DNSSEC. Solving this problem may require novel approaches or adjustments at the protocol level. Efforts should be made to better understand the exact impact of increasing the signature and key size on query resolution when transitioning to FALCON-512 from ECDSA-P256. In this paper, we have shown that transitioning to FALCON results in higher latency (see Figure 8) and higher error rates (see Figure 7) as predicted by Müller et al. [9,11]. Additional research should be dedicated to empirically verify the exact link between the signature and the increase in key size between ECDSA-P256 and FALCON-512. A packet capture analysis could test how exactly it affects error rates and transmission success.

Another important outcome of this study is the development of a reusable benchmarking methodology. The experimental setup and evaluation framework established here can serve as a foundation for testing other post-quantum algorithms under consideration for DNSSEC. This allows future research to evaluate how alternative candidates compare to FALCON-512 in terms of both performance and stability.

Ultimately, while FALCON-512 appeared promising in theory, particularly in terms of latency, its practical deployment in DNSSEC faces serious challenges. However, the insights gained from evaluating its implementation and behavior under real-world conditions can inform the selection and refinement of more suitable post-quantum alternatives in the future.

Author Contributions
Conceptualization: T.N. and M.N.; conceived and performed the experiments: T.N.; analyzed and interpreted the data: T.N. and M.N.; writing-review and editing: T.N. and M.N.; supervision and funding acquisition: M.N. All authors have read and agreed to the published version of the manuscript.

Funding
This research was partially funded by AGH University of Krakow ("Excellence initiative—research university" program) and the European Commission (Horizon Europe Framework Program) under Grant Agreement no. 101119547 (PQ-REACT).

Conflicts of Interest Statement
The authors declare no conflict of interest.

Data Availability Statement
The original contributions presented in the study are included in the article; further inquiries can be directed to the corresponding author.

Acknowledgments
This research was partly supported by the program "Excellence initiative—research university" for the AGH University of Krakow. This work was also supported by the EU Horizon Europe Framework Program under Grant Agreement no. 101119547 (PQ-REACT).

References

1. P. W. Shor (1999). “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.” *SIAM Review*, 41: 2, 303–332.
2. National Institute of Standards and Technology (2024). *NIST Releases First 3 Finalized Post-Quantum Encryption Standards*, [news release]. Available at: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards (Accessed on 08 June 2025).
3. D. Dziechciarz, M. Niemiec (2025). “E?iciency analysis of NIST-standardized post-quantum cryptographic algorithms for digital signatures in various environments.” *Electronics*, 14: 1, 70, doi: 10.3390/electron-ics14010070
4. NIST Post-Quantum Cryptography Team (2016). *Post-Quantum Cryptography: NIST’s Plan for the Future* [Presentation]. Available at: https://csrc.nist.gov/csrc/media/projects/post-quantum-cryptography/documents/pqcrypto-2016-presentation.pdf
5. S. Rose, M. Larson, D Massey, R. Austein and R. Arends (2005). “DNS security introduction and requirements,” RFC 4033, doi: 10.17487/RFC4033.
6. G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, Y.-K. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, and D. Smith-Tone (2022). *Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process*, National Institute of Standards and Technology Interagency or Internal Report NIST IR 8413-upd1, https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8413-upd1.pdf
7. NIST - Information Technology Laboratory - Computer Security Resource Center. *Post-Quantum Cryptography* [web resource], https://csrc.nist.gov/pqc-standardization
8. P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang, *FALCON* [web resource] (Accessed: 2025, June 08), https://falcon-sign.info/
9. M. Müller, J. de Jong, M. van Heesch, B. Overeinder, and R. van Rijswijk-Deij (2020). “Retrofitting post-quantum cryptography in internet protocols: a case study of DNSSEC.” *SIGCOMM Computer Communication Review*, 50: 4, 49–57, doi: 10.1145/3431832.3431838.
10. S. Huque, C. Elmero, O. Guomundsson (2025). “Compact denial of existence in DNSSEC.” Active Internet-Draft. Available at: https://datatracker.ietf.org/doc/draft-ietf-dnsop-compact-denial-of-existence/07/, expires 31 August 2025.
11. M. Müller (2021). *Making DNSSEC Future Proof*, PhD Thesis - Research UT, graduation UT, University of Twente.
12. W. Beullens (2022). *Breaking Rainbow Takes a Weekend on a Laptop*, Cryptology ePrint Archive, Paper 2022/214
13. G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, Y. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, D. Smith-Tone (2022). “Status report on the third round of the NIST post-quantum cryptography standardization process.” Available at: https://csrc.nist.gov/pubs/ir/8413/upd1/final
14. G. Beernink (2022). “Taking the quantum leap: Preparing dnssec for post-quantum cryptography.” Master’s thesis, University of Twente.
15. A. Fregly, R. van Rijswijk-Deij, M. Müller, P. Thomassen, C. Schutijser, and T. Chung (2024), *Research Agenda for a Post-Quantum DNSSEC*, Expired Internet-Draft, Available: https://datatracker.ietf.org/doc/draft-fregly-research-agenda-for-pqc-dnssec/, expired 26 June 2025.
16. A. Fregly, R. van Rijswijk-Deij, M. Müller, P. Thomassen, C. Schutijser, and T. Chung (2024). “Research agenda for a post-quantum DNSSEC, expired internet-draft.” Available at: https://datatracker.ietf.org/doc/draft-fregly-research-agenda-for-pqc-dnssec/, expired 26 June 2025.
17. J. Goertzen (2022). “Enabling post-quantum signatures in DNSSEC: one ARRF at a time,” Master’s thesis, University of Waterloo.
18. M. Raavi, S. Wuther, S. Chang (2024). “Securing post-quantum DNSSEC against fragmentation mis-association threat”, in *IEEE International Conference on Communications (ICC): Communication and Information System Security Symposium.*
19. *SIDN Labs*, [software]. Available at: https://github.com/SIDN
20. IANA (2024). *Domain Name System Security (DNSSEC) Algorithm Numbers*, [web resource], Last updated: December 05. Available at: https://www.iana.org/assignments/dns-sec-alg-numbers/dns-sec-alg-numbers.xhtml

21. P. Boryło, P. Chołda, J. Domżał, P. Jaglarz, P. Jurkiewicz, A. Lasoñ, M. Niemiec, M. Rzepka, G. Rzym, R. Wój-cik (2017). “SDNRoute: Integrated system supporting routing in software defined networks” in *19th International Conference on Transparent Optical Networks (ICTON*), 2–6 July 2017, doi: 10.1109/ICTON.2017.8024959
22. Tantalor93 (GitHub user). dnspyre—a command-line DNS benchmark tool [web resource]. Available at: https://tantalor93.github.io/dnspyre/
23. SIDN Labs. PowerDNS recursive resolver with support for PQC algorithms [software]. Available at: https://github.com/SIDN/pqc-resolver-powerdns
24. SIDN Labs. PowerDNS authoritative nameserver with PQC algorithms [software]. Available at: https://github.com/SIDN/pqc-auth-powerdns
25. PowerDNS. PowerDNS authoritative server documentation: built-in webserver and HTTP API [web resource]. Available at: https://doc.powerdns.com/authoritative/http-api/index.html
26. J. Goertzen, P. Thomassen, N. Wisiol (2024). Field experiments on post-quantum DNSSEC. Available at: https://ripe89.ripe.net/presentations/57-RIPE-89-Field-Experiments-on-Post-Quantum-DNSSEC.pdf
27. PowerDNS. PowerDNS documentation [web resource]. Available at: https://doc.powerdns.com/ (Accessed on 28 May 2025).