

Do we need to change some things?

Open questions posed by the upcoming post-quantum migration to existing standards and deployments

Panos Kampanakis[✉] and Tancrède Lepoint^[0000-0003-3796-042X]

Amazon Web Services, USA
{kpanos,tlepoint}@amazon.com

Abstract. Cryptographic algorithms are vital components ensuring the privacy and security of computer systems. They have constantly improved and evolved over the years following new developments, attacks, breaks, and lessons learned. A recent example is that of quantum-resistant cryptography, which has gained a lot of attention in the last decade and is leading to new algorithms being standardized today. These algorithms, however, present a real challenge: they come with strikingly different size and performance characteristics than their classical counterparts. At the same time, common foundational aspects of our transport protocols have lagged behind as the Internet remains a very diverse space in which different use-cases and parts of the world have different needs.

This vision paper motivates more research and possible standards updates related to the upcoming quantum-resistant cryptography migration. It stresses the importance of amplification reflection attacks and congestion control concerns in transport protocols and presents research and standardization takeaways for assessing the impact and the efficacy of potential countermeasures. It emphasizes the need to go beyond the standardization of key encapsulation mechanisms in order to address the numerous protocols and deployments of public-key encryption while avoiding pitfalls. Finally, it motivates the critical need for research in anonymous credentials and blind signatures at the core of numerous deployments and standardization efforts aimed at providing privacy-preserving trust signals.

Keywords: Post-quantum · Amplification Protection · Congestion Control · Public-key Encryption · Anonymous Authentication

Note: This pre-print has not undergone any post-submission improvements or corrections. It will be presented in the Security Standardisation Research (SSR) Conference 2023.

1 Introduction

Rapid advances in quantum computing [49] have motivated the need to replace current cryptographic schemes based on the believed hardness of traditional

number-theoretic problems, such as integer factorization and discrete logarithms. Since 2016, the National Institute of Standards and Technology (NIST) is running an open Post-Quantum Cryptography standardization process [73] to standardize quantum-resistant key encapsulation mechanisms (KEMs) and digital signatures. In July 2022, NIST completed the third round of the process, and selected Kyber [87] as KEM, and Dilithium [64], Falcon [78], and SPHINCS+ [42] as digital signatures, to become the first NIST post-quantum standards (expected 2024).¹ While these primitives provide the same core functionalities as their classical counterparts, they feature strikingly different size and performance characteristics. Table 1 summarizes the public key, ciphertext, and signature sizes of NIST future post-quantum standards, as well as current NIST standards ECDH P-384, ECDSA P-384, and RSA-3072.

Algorithm	Quantum-safe	Public Key	Ciphertext / Signature
ECDH P-384	✗	48	48
Kyber-512	✓	800	768
Kyber-768	✓	1184	1088
ECDSA P-384	✗	48	96
RSA-3072	✗	387	384
Falcon-512	✓	897	666
Dilithium-2	✓	1312	2420
Falcon-1024	✓	1793	1280
Dilithium-3	✓	1952	3293
SPHINCS+-128s	✓	32	7856
SPHINCS+-192s	✓	48	16224

Table 1. Classical and post-quantum cryptographic schemes selected by NIST for standardization, ordered by ciphertext/signature size.

Other standardization organizations have also been working on introducing post-quantum algorithms to existing protocols and standards [94,99] and focusing on post-quantum migration challenges and solutions.² Additionally, a few Internet Engineering Task Force (IETF) RFC drafts are already introducing these algorithms in IETF standards [68,95,100,69,41,39]. In February 2023, the IETF created PQUIP³, a working group focused on the use of post-quantum cryptography in protocols. Similarly, the European Telecommunications Standards Institute has formed a Quantum-Safe Working Group [28] that aims to make assessments and recommendations on the various proposals from industry and academia regarding real-world deployments of quantum-safe cryptography.

¹ The standardization process continues with a fourth round for alternates key encapsulation mechanisms (BIKE [9], Classic McEliece [4], HQC [3], and SIKE [48]), and a new call for proposal for digital signatures.

² <https://www.nccoe.nist.gov/crypto-agility-considerations-migrating-post-quantum-cryptographic-algorithms>

³ <https://datatracker.ietf.org/wg/pquip/about/>

The integration of post-quantum key encapsulation mechanisms and signatures will affect existing protocols due to their size. For example, using post-quantum cryptography could increase the number of packets which in turns could increase the loss probability in constrained conditions [74]. Post-quantum certificate chains could exceed any certificate chain size that our applications see today. Post-quantum signatures in TLS could lead to connection establishment slowdowns [11,56,91,92]. QUIC would also see challenges with post-quantum signatures related to its amplification protection feature [55]. Most of the transport protocols used today were designed decades ago under different network conditions with different sizes in mind.

Our Contributions. In this work, we emphasize several areas and gaps requiring research and standardization to make the post-quantum migration successful.

1. In Section 2, we discuss how post-quantum authentication will increase the amplification reflection attack risk for UDP-based secure transport protocols, and survey potential amplification protection trade-offs for QUIC, DTLS, and others (Section 2.1). We propose further investigations to identify current protocol use-case behavior in order to find the best option for standardization.
2. We focus on congestion control in Section 2.2; we point out that the initial congestion window value in common secure transport protocols may be already too small for today and could introduce connection slowdowns in a post-quantum world. We suggest a re-evaluation of the initial value as done by Chi et al. a decade ago for RFC 6928. We also stress the potential impact of such a change on various parts of the world due the heavier additional post-quantum handshake data.
3. In Section 3, we identify and discuss the need for standardization of quantum-safe public-key encryption, and in particular that of hybrid public key encryption and key wrapping.
4. We identify the need for research in areas like quantum-safe password-authenticated key exchange (Section 4), oblivious pseudorandom functions, blind signatures (Section 5), and other cryptographic standards which do not have the community or the forum to introduce quantum-safe algorithms (Section 6).

2 Transport Protocol Implications

This section discusses implications of the post-quantum migration on transport protocols, and, in particular, talks about the increased concerns around amplification attacks (Section 2.1) and congestion control (Section 2.2).

2.1 Amplification Attacks

Amplification attacks are distributed denial of service (DDoS) attacks exploiting a disparity in resource consumption between an attacker and a target system.

When a small request triggers a large response, an attacker performing many such requests can disrupt a target system which receives high volumes of response data. To further increase the impact, attackers can perform amplification reflection attacks that leverage a reflector entity which inadvertently serves in amplifying multiple requests spoofed from a victim’s address. Fig. 1 shows a typical amplification reflection.




Fig. 1. Attackers spoofing multiple small requests can trigger big responses by the reflector and deploy reflection amplification DDoS to the victim.

In [67], Majkowski presents historical amplification reflection attacks and explains how attackers managed to trigger Gbps of traffic over multiple protocols. These attacks were successful with protocols where the source address is not validated and thus can be spoofed. Such protocols are usually UDP-based, like DNS or NTP. CISA created an alert for UDP-based amplification⁴ describing the amplification factors for each protocol of concern. MITRE maintains a Common Enumeration Weakness about them⁵ referencing common vulnerabilities created in relation to previous DDoS amplification attacks (e.g., CVE-1999-1379, CVE-1999-0513, CVE-2000-0041, CVE-1999-1066, CVE-2013-5211).

Such attacks were observed widely a decade ago with DNS⁶. In one of the largest amplification attacks, attackers employed open DNS recursors which responded to anyone on the Internet to initiate very large amounts of data towards spamhaus.org [79]. The attacks peaked at 90Gbps on the victim’s system and at 300Gbps on the Tier 1 provider. The amplification factor was 100× by using 21.7 million mis-configured recursors and a UDP protocol.

The impact of an amplification attack grows as the amplification factor increases, i.e., the more data the attacker can trigger in the response with a small, cheap request, the more damage it can cause the victim. In [84], Rossow proposes to quantify the amplification by introducing the notions of bandwidth amplification factor (BAF) and packet amplification factor (PAF). They also investigate the efficiency of discovering amplifiers, and the most amplifying protocols (which include common UDP protocols).

⁴ <https://www.cisa.gov/uscert/ncas/alerts/TA14-017A>

⁵ <https://cwe.mitre.org/data/definitions/406.html>

⁶ <https://www.cisa.gov/uscert/ncas/alerts/TA13-088A>




Fig. 2. Amplification Protection where the reflector limits the size of the response from a not-yet-validated source and waits for an **Acknowledgement** to limit the impact of a potential reflection amplification attack.

Some protocols offer anti-amplification mechanisms which mandate the validation of the source before sending the whole large response. For example, Fig. 2 shows how the reflector can send part of the data and wait for an **Acknowledgement** before sending the rest of it. These mechanisms introduce a round-trip while the sender is waiting for the **ACK** which slows down the connection. Thus, service providers sometimes do not honor them [72]. Additional countermeasures include address spoofing protections deployed at the network level, which validate the source address based on routing information or routing firewall rules, but these are not always available or effective. Furthermore, some of the most commonly exploited amplifying protocols, such as DNS, have seen hardened deployments over time which improved their security and made amplification reflection harder.

Below we discuss specific protocols and the potential impact quantum-safe algorithms would have on their amplification potential.

QUIC. QUIC is a UDP-based encrypted transport protocol, at the core of HTTP/3, designed for performance. It is specified in a set of IETF standards ratified in May 2021 [98]. QUIC protects its UDP datagrams by using encryption and authentication keys established in a TLS 1.3 handshake carried over QUIC transport. QUIC also offers amplification protection by mandating that a sender can send up to $3\times$ the size of the request for non-validated addresses (which could be spoofed). Fig. 3 shows QUIC’s amplification protection.

In typical QUIC deployments, $3\times$ the request size amounts to approximately $\approx 4\text{kB}$. Unfortunately, it is not uncommon for a certificate chain along with the leaf certificate on the Web to exceed that limit, which would trigger QUIC’s amplification protection and add a round-trip to common HTTP/3 connections. As a consequence, many QUIC servers do not honor the amplification window. Nawrocki et al. [72] investigated common servers like Cloudflare and Google which exceed the window up to an amplification BAF factor of 10. Meta’s servers exceed it by an even more significant amount.




Fig. 3. Amplification Protection in QUIC where the server returns $3\times$ the size of the QUIC client `Init` packet and waits for a response before sending more data.

In a post-quantum world, when using the general purpose Dilithium signature scheme, the `ServerHello`, `Certificate` and `CertificateVerify` messages could add up to 15-17kB for the lowest security level. Depending on the digital signature algorithm, parameters, and length of the certificate chain, the post-quantum migration could easily make the BAF ranging from 5 to 20 in typical settings. In [55], Kampanakis et al. discuss how the post-quantum signature schemes selected by NIST will all exceed the amplification window and incur at least a round-trip. Fig. 4 experimentally shows this for classical RSA-2048 and post-quantum (PQ) certificate chains, illustrating further the impact post-quantum authentication has on QUIC amplification protection. If QUIC implementers do not honor the amplification protection window today to prevent the extra round-trip [72], will they do the same in a post-quantum world where the amount of data returned from the server is significantly higher? If the servers not honoring the amplification protection introduce some amplification reflection attack risk today [72], that risk will multiply with post-quantum signatures in the future. In typical high profile attacks of the past, millions of mis-configured open DNS resolvers ended up generating Gbps of traffic with amplification factors between 50-100. Locking down these servers was relatively easy. In a future world, thousands or more of QUIC servers could theoretically serve as amplification reflectors with sizeable amplification factors. Locking them down by identifying illegitimate traffic may not be trivial as they could be serving clients from all over the world.

There are different approaches to addressing the $3\times$ amplification window issue with post-quantum authentication in QUIC. Each of them has advantages and disadvantages. The options are:

- Increase the $3\times$ amplification window to a level that will not incur round-trips for the post-quantum case. That could mean $>15\times$ amplification windows which certainly increase the amplification factor.
- Trim the “authentication data” sent in the handshake. That can include using session resumption [98, Sec. 4.5], [80, Sec. 2.2] or suppressing intermediate CA certificates by caching them [55,97]. These could alleviate the size of the server response and the amplification factor. A similar option could be

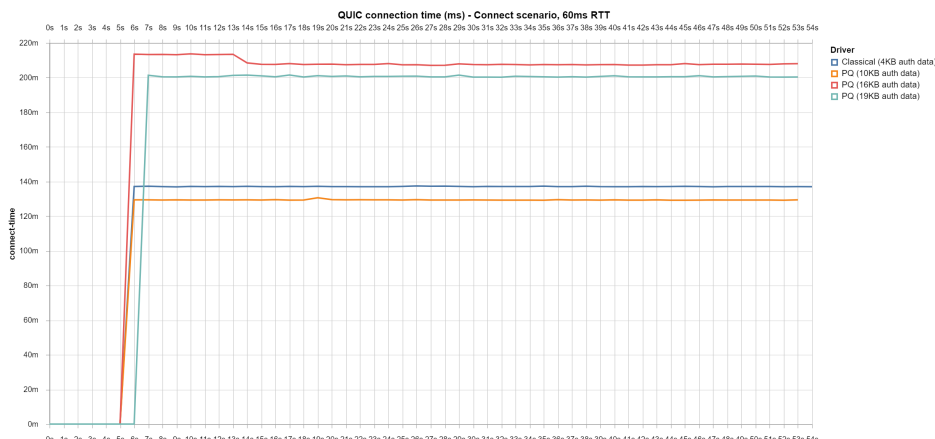


Fig. 4. QUIC connection time (ms) for classical ≈ 4 kB RSA-2048, 10kB PQ certificate chains, 18kB Dilithium-2, and 22kB Dilithium-3 certificate chains. In these experiments the client was creating 1,000 sequential connections. The client-server round-trip was 60ms. The measurements were collected using `s2n-quic`'s `netbench` benchmarking tool. `s2n-quic` is AWS' QUIC library. We can see that all experiments include one round-trip due to QUIC's amplification protection. The 18 and 22kB certificate chain connections include one additional round-trip due to QUIC's initial congestion window (see Section 2.2). Note that captures showed that the 10kB chain connection times are a fraction of 1ms slower than the classical RSA 4kB ones, whereas network condition variability made them measure a few milliseconds faster in the experiment.

certificate compression [34], but that would not improve the size much as the certificate bloat comes from random data (public key, signature).

- Use address validation tokens for every post-quantum authentication. QUIC validation tokens are opaque values which enable the server to confirm it has seen the client address. Tokens add a round-trip and may negatively affect performance unless the client revisits the server for the lifetime of the token. We are not aware of any work that studies the use of QUIC tokens on the Internet and how frequently clients revisit servers. Another option could be for the client to include a new extension requesting a token only when supporting post-quantum signatures and for servers that it knows it visits often. As a less aggressive measure, servers could start sending tokens only when suspecting amplification reflection attacks.
- Artificially increase the `ClientHello` size so that the the quantum-safe authentication data fits in the amplification window. Increasing the `ClientHello` means we will be unnecessarily wasting resources or bandwidth.

In short, post-quantum authentication introduces performance complications with QUIC amplification. We propose the following research investigations and standardization takeaways:

Takeway 1. Study the use of address validation tokens in QUIC and identify if server implementations use them and when. Evaluate their overall effectiveness by quantifying the frequency of revisiting servers for the lifetime of a token. Optimize the lifetime of tokens based on total traffic and client revisits.

Takeway 2. Evaluate the amplification protection options and standardize the best one for post-quantum signature sizes. Trimming the authentication data and leveraging tokens seem to be the most natural options. As [72] demonstrated, $3\times$ is not enough even for classical signatures, thus a more realistic amplification protection limit should also be identified.

DTLS 1.2, 1.3. DTLS 1.3 is TLS’ counterpart over UDP. It was recently standardized and is susceptible to amplification attacks because the server authentication data is disproportionately larger than the `ClientHello`. The DTLS 1.3 standard [83] addresses amplification similarly to QUIC. It recommends, but does not mandate, that “a server *SHOULD* limit the amount of data it sends toward a client address to $3\times$ the amount of data sent by the client before it verifies that the client is able to receive data at that address”. Moreover, it recommends, not mandates, the use of a cookie to validate the source address. We are not aware of any work that investigates if DTLS 1.3 implementations honor these optional amplification protections discussed in the standard.

Like TLS 1.3 and QUIC, DTLS 1.3 post-quantum authentication data size will significantly increase, which will have implications to its amplification potential. DTLS 1.3 servers could theoretically serve as amplification reflectors with sizeable amplification factors.

DTLS 1.2 [81] would suffer the same challenges with amplification and post-quantum certificates. So far, only (D)TLS 1.3 is being planned for post-quantum standardization in the IETF’s TLS WG. Thus, we consider DTLS 1.2 out of scope for this work at this time.

To address some of the open questions for amplifications in post-quantum DTLS 1.3, we propose the following research investigations and standardization takeaways:

Takeway 3. Study how server DTLS 1.3 implementations behave when the certificate chain exceeds $3\times$ the client request (like [72] did for QUIC). Also consider if an amplification protection mechanism should become mandatory in DTLS 1.3.

Takeway 4. Investigate the DTLS 1.3 use-cases that would benefit from key exchange cookies when post-quantum certificates exceed the amplification window. These will depend on the frequency of clients communicating with the same server over the lifetime of the cookie.

Takeway 5. Evaluate the amplification protection options and standardize the best one for post-quantum signatures in DTLS 1.3. These options resemble the ones in QUIC and include increasing the window, using exchange cookies when post-quantum certificates are used and limiting the authentication data in the handshake. A more realistic amplification protection window should also be identified.

DNSSEC. DNS is a protocol widely used for name resolution. It runs over UDP. As we already discussed, it has traditionally been a good candidate for amplification reflection attacks. In a DNS amplification, the EDNS0 extension has been used to include large messages. A DNS request of about 50B could elicit response sizes of up to 4kB which results to a $80\times$ BAF [67]. Although widely effective in the past [79], DNS amplification reflections have become less common as there are fewer open resolvers and protections against them are well established [67].

DNSSEC is a standard for authenticating DNS responses to prevent DNS spoofing attacks. Unfortunately, the post-quantum signatures selected by NIST would not fit in a single packet, and fragmentation is not an option for DNSSEC [71]. To address the issue, Goertzen et al. [35] propose ARRF, a method of fragmenting DNS resource records at the application layer based on client acknowledgments. Sequentially sending each fragment after an acknowledgement would slow down getting the full response. That is why [35] describes a mode for requesting the fragments in parallel after acknowledging the first fragment. Although the BAF is high due to the size of the signed record, such an approach would not increase the PAF per DNSSEC RRSIG record. Requiring the server to keep track of the fragments until they are received by the requester adds some burden on DNS servers today which are stateless. Alternatively, Fregly et al. [29] recently proposed a Merkle-tree structure which can shrink these signatures to manageable sizes.

To address some of the amplification concerns for post-quantum DNSSEC, we propose the following standardization takeaway:

Takeway 6. Identify, evaluate and standardize an amplification-resistant solution for post-quantum DNSSEC in IETF’s DNSOP working group.

TLS 1.2, 1.3. TLS is the status quo for encrypting communications on the Internet. TLS is not susceptible to the same amplification reflection risks because TCP validates the source address before the TLS negotiation. However, post-quantum TLS still induces an amplification risk as the data sent by the client and server are unbalanced in size.

For example, in a TLS 1.3 negotiation, a client could send a `ClientHello` to negotiate quantum-safe key exchanges without including a post-quantum `key_share`. If the server supported post-quantum key exchange, it would respond with a `HelloRetryRequest` and a post-quantum `key_share` which could add up to a few kB depending on the post-quantum KEM. Thus, with a very

small `ClientHello`, the client could trigger a relatively big response ($5-8 \times$ BAF). This response would not be nearly as big as the post-quantum certificate chain from the server which has a higher amplification factor. We consider the amplification concern for TLS 1.3 to only be a denial of service risk for the server and client’s edge capacity, not a reflection attack against a victim. Such attacks have not been seen widely in the past because they affect both the attacker and the victim and they are easy to identify and protect against.

In TLS 1.2, the client could trigger a big quantum-safe certificate chain and `Server Key Exchange` from the server. Fig. 5 shows the `Server Key Exchange`, `Certificate` and `CertificateVerify` messages which could add to 20kB with a very small `ClientHello` ($100 \times$ BAF) which is significant. So far, only TLS 1.3 is being planned for post-quantum standardization in the IETF’s TLS WG. Thus, we consider TLS 1.2 amplifications out of scope for this work at this time.




Fig. 5. Amplification in TLS 1.2 where the server returns a potentially large post-quantum certificate chain and ephemeral public key while the client has only sent a small `ClientHello`. The client address is still validated by the TCP handshake.

IKEv2. IKEv2 [60] is another protocol running over UDP. It is used to negotiate keys, authenticate and establish Security Associations between IPsec VPN peers. CERT/CC published vulnerability VU#419128⁷ about how IKEv2 implementations could amplify requests by $9 \times$ based on a whitepaper by Chad Seaman from Akamai in 2016. The work had discovered IKEv2 implementations where the responder kept re-sending `IKE_SA_INIT` messages after not receiving an `IKE_AUTH` by the initiator. The study found thousands of servers replied 21 times or more per initiator `IKE_SA_INIT` message, and some servers responded thousands of times. IKEv2 specifies that the “*responder MUST never retransmit a response unless it receives a retransmission of the request*”. Thus, the behavior of these responders was violating the standard.

Post-quantum IKEv2 may need to fragment `IKE_SA_INIT` messages as the post-quantum KEM public key or ciphertext may not fit in one packet. RFC 9242 [93] specifies the `IKE_INTERMEDIATE` exchange which can carry and fragment post-quantum key exchanges after performing a classical key exchange with `IKE_SA_INIT`. The quantum-safe response by the server could be significantly bigger than the classical `IKE_SA_INIT` but given that the `IKE_INTERMEDIATE`

⁷ <https://www.kb.cert.org/vuls/id/419128>

exchange occurs after `IKE_SA_INIT` there is no additional amplification reflection possibility. Moreover, Kyber offers a balance between the public key and ciphertext which means there is no additional amplification factor introduced. Even if implementers haven't addressed the IKEv2 amplification vulnerability from 2016, post-quantum algorithms do not bring any additional amplification concerns to IKEv2.

2.2 Congestion Control

Congestion control algorithms aim to transfer data fast without overwhelming the network. Generally, they do that by offering a initial congestion window value which is the starting point. The sender ramps up sending data as long as it is successful and slows down after it sees failures (which could be because of congestion).

TCP. TCP congestion control defines an initial congestion window (`initcwnd`) size which is the limit of data the sender can send without receiving an ACK from the receiver. The congestion window starts in the slow-start phase from `initcwnd` and increases exponentially when data is received successfully until a packet loss occurs, at which point the TCP connection enters the congestion avoidance phase. Fig. 6 shows the sender filling up a window with 10 segments and waiting for an ACK before continuing the slow-start phase and doubling the congestion window.




Fig. 6. After filling the `initcwnd`, the sender waits for a full round-trip for an Acknowledgement before resuming sending data.

The TCP `initcwnd` has historically been changing over time as networks evolve. It originally started at 1 Maximum TCP Segment Size (MSS) [96], and then increased to 2 MSS [76] and to 4 MSS [75]. A decade ago, after careful evaluation and research [25], RFC 6928 [19] updated `initcwnd` to 10 MSS. There was also a proposal [101] to introduce a mechanism to dynamically update the initial congestion window by tracking connections. This idea never got adopted.

Asia Pacific Network Information Centre's (APNIC) 2018 survey [86] showed that 80% of the Alexa Top 1M servers used `initcwnd` = 10 MSS, but there were

thousands that exceeded it. Today’s Content Delivery Network (CDN) providers often use `initcwnd` much higher than 10 MSS to maximize their content performance [77]. CDNs increase their starting window based on performance optimizations and content measurements of their networks, but it is unlikely they use broad Internet data to measure the effect on the rest of the Internet.

Post-quantum authentication will interfere with TCP’s `initcwnd`. A certificate chain with quantum-safe signatures and public keys will exceed 10 MSS. Sikeridis et al. investigated the issue and showed that post-quantum signatures would introduce a round-trip [92]. Westerbaan simulated post-quantum authentication in TLS with `initcwnd` =30 MSS [11], and showed that tweaking the `initcwnd` value could prevent a round-trip but would not address constrained networks, buggy servers or middle-boxes which are affected by the overall size of post-quantum signatures. A discussion⁸ in NIST’s pqc-forum mailing list brought up the topic of increasing the window to alleviate the issue given that CDNs already do it.

Significantly increasing the default TCP `initcwnd` value in every Internet server that wants to enable post-quantum TLS authentication should not be done lightly as it could have adverse effects on TCP Congestion Control. We consider it as a potential improvement that should be carefully studied at a large scale before deployment. For example, the last increase to 10 MSS received thorough analysis and investigation; RFC 6928 [19, Appendix A] discusses concerns which include fairness and effects on slow networks and developing regions.

To address some of these open questions, we propose the following research investigations and standardization takeaway:

Takeway 7. Consider increasing the TCP initial congestion window to `initcwnd` =20-25MSS. A similar exercise to [86] should reveal what values are commonly used today and how much the industry has tweaked the standard. We could also continuously monitor `initcwnd` trends by probing random addresses on the Internet like in [86]. After carefully researching the `initcwnd` optimal values for the Internet like in [25] to ensure the concerns in [19, Appendix A] are addressed, we could standardize a new `initcwnd` value.

QUIC congestion control. QUIC has its own congestion control which is similar to TCP’s (RFC 6928 [19]). [46] states that “*endpoints SHOULD use an initial congestion window of ten times the maximum size (`max_datagram_size`), while limiting the window to the larger of 14,720 bytes or twice the maximum datagram size*”. It is clear that QUIC will suffer from the same connection slow-down with quantum-safe authentication as TCP due to the initial congestion window. RFC 6928 also says that “*Though the anti-amplification limit can prevent the congestion window from being fully utilized and therefore slow down the increase in congestion window, it does not directly affect the congestion window*”. QUIC can be slowed down by both the amplification window (Section 2.1) and

⁸ https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/2ak2U_MxyrQ/m/L-kQ-SubBwAJ

the initial congestion window, whichever is the lowest. Fig. 4 experimentally shows this for classical RSA-2048 and post-quantum (PQ) certificate chains. Other than the standard, alternate congestion control algorithms have also been proposed for QUIC such as CUBIC (RFC 8312).

To address the congestion window issue with post-quantum QUIC, we propose the following research investigation and standardization takeaways:

Takeway 8. Evaluate if QUIC implementations and deployments honor the initial congestion window. If they do not honor it, that may mean that addressing the amplification protection issue (Section 2.1) will be enough to enable the transition to post-quantum authentication.

Takeway 9. Consider updating the QUIC standard to support bigger initial congestion window values after carefully investigating its impact as with TCP’s `initcwnd` above.

DTLS congestion Control. DTLS does not specify a congestion control algorithm and leaves it to the transport protocol. [83] states that “*some transports provide congestion control for traffic carried over them. If the congestion window is sufficiently narrow, DTLS handshake retransmissions may be held rather than transmitted immediately, potentially leading to timeouts and spurious retransmission. When DTLS is used over such transports, care should be taken not to overrun the likely congestion window.*”. It also acknowledges that large certificate chains can lead to congestion and recommends sending part of the chain and waiting for a response packet. It then proposes extensions that will alleviate the data size sent.

While DTLS calls out potential congestion issues with large server responses, it does not mandate any specific countermeasure. It is uncertain if implementations honor the suggestions. If they honor them, a transition to post-quantum certificates could slow down DTLS connections. To address some of these open questions, we propose the following research investigation takeaway:

Takeway 10. Research if DTLS implementations and deployments honor the congestion avoidance suggestions in [83, Sec. 5.8.3]. If these introduce a round-trip to the post-quantum certificate scenario, investigate what value would be optimal and how it would interfere with the DTLS amplification protection investigation takeaway in Section 2.1.

IKEv2. IKEv2 communications consist of pairs of messages where the initiator initiates an exchange and the responder responds. The specification [60] does not define any congestion control. It recognizes that retransmissions could affect congestion and mandates exponential backoffs and Explicit Congestion Notification (ECN) support to alleviate it.

Post-quantum IKEv2 would include large `IKE_INTERMEDIATE` (RFC 9242 [93]) and `IKE_AUTH` messages which could exacerbate the problem. It would be interesting to investigate how broad exponential backoff and ECN support is in IKEv2

implementations, but given that IKEv2 negotiations happen only when bringing up a tunnel and constitute a very small percentage of Internet traffic, we do not believe post-quantum IKEv2 could affect congestion on the Internet.

3 Public Key Encryption and Key Wrapping

In its post-quantum standardization call for proposals, NIST asked for public-key encryption (PKE), KEM, and digital signatures schemes, and said “*As the KEM and public-key encryption functionalities can generally be interconverted, unless the submitter specifies otherwise, NIST will apply standard conversion techniques to convert between schemes if necessary.*” Indeed, almost all of the KEM candidates submitted first apply some variant of the Fujisaki–Okamoto (FO) transform [30,31,24] on a weakly-secure “base” PKE scheme to construct the resulting KEM. This blueprint holds in particular for Kyber which defines an IND-CPA secure PKE [87, Sec. 1.2], and uses a tweaked FO transform to create the IND-CCA2-secure KEM [87, Sec. 1.3].

The generic reverse construction follows the KEM-DEM paradigm [89] (a.k.a. hybrid encryption). The recipient public key is used as input of the encapsulation function of the KEM to create a ciphertext and shared key, and the data is encrypted using the shared key in a data-encapsulation mechanism (DEM) — often instantiated with an AEAD scheme such as AES-GCM. Albeit these generic transformations give a nice theoretical framework to work with, we believe that standardizing public-key encryption will necessitate additional efforts and care to avoid the pitfalls encountered in the past.

Indeed, numerous standards have been created over the years for public-key encryption (and hybrid public-key encryption), including ANSI X9.63 (ECIES) [8], ANSI X9.44 (RSA-KEM) [6], IEEE P1363a [44], ISO/IEC 18033-2 [90], and SECG SEC 1 [88], and no less than four RFCs for PKCS [53,54,51,52]. Martínez et al. [33] provide a thorough comparison of the elliptic-curve public-key encryption standards, highlighting that the differences between them prevented ECIES from being fully interoperable. The lack of clear PKE standard has further led to inconsistent support across libraries; e.g., NaCL and BouncyCastle implement their own versions of hybrid encryption.

(Hybrid) Public-Key Encryption. In 2022, Hybrid Public Key Encryption (HPKE) was published in RFC 9180 [10]. HPKE aims at addressing interoperability issues with ECIES. It has been designed to be generic, with simplicity and modularity in mind. It offers a “base” mode which only encrypts data, and “authenticate” and “psk” modes which authenticate the sender and encrypt data. HPKE has seen immediate adoption by Internet protocols such as TLS Encrypted Client Hello [82], Oblivious DNS-over-HTTPS (RFC 9230 [61]), Message Layer Security⁹, and Privacy Preserving Measurement¹⁰. Since its ciphersuites

⁹ <https://datatracker.ietf.org/wg/mls/about/>

¹⁰ <https://datatracker.ietf.org/wg/ppm/about/>

consist of specifying a KEM, a Key Derivation Function (KDF), and an AEAD, HPKE naturally supports the addition of new ciphersuites in the future [10, Sec. 9.2]. At ICMC 2022, Anastassova et al. [7] presented the first implementation of a post-quantum HPKE (PQ-HPKE) using Kyber as KEM¹¹, and a combination of Kyber and DHKEM [10] in “PQ-hybrid” mode. While they conclude that the performance is acceptable (especially for larger messages), they also emphasize that Kyber does not allow to use HPKE in “authenticated” mode since it does not provide a direct API for authenticated encapsulation and decapsulation. As such, there is no known construction of post-quantum HPKE resisting key impersonation attacks. Similar discussions were held on the CFRG mailing list¹², further emphasizing the need of authenticated KEM in protocols, such as the Certificate Management Protocol (CMP) [70].

To address some of these open issues with quantum-safe PKE, we propose the following research investigation and standardization takeaways:

Takeway 11. Standardize an HPKE variant that provides post-quantum security in the near future, and addresses the security proof gaps with PQ-hybrid HPKE identified in [7, Sec. IV.C].

Takeway 12. Research and standardize post-quantum authenticated KEMs to enable the migration of protocols requiring such a primitive.

Key wrapping. While the KEM-DEM paradigm above (or HPKE) enable us to construct an IND-CCA secure public key encryption, the encryptor does not fully control the value of the shared key for the data-encapsulation mechanism. In some applications, this is not desirable. For example, when sending the same message to many recipients, a natural approach is to encrypt the message with a fresh AES key, and then encrypt the AES key to each recipient. Henceforth, the data to encrypt with the KEM-DEM (or HPKE) would be the fresh AES key. While any secure authenticated encryption scheme can be used to encrypt the fresh key, there exist specialized symmetric key-wrapping algorithms that are more compact like those defined in NIST SP 800-38F, IETF RFC 3394, RFC 5297, RFC 5649. Some of these symmetric key wrapping algorithms have been combined with public key encryption in RFC 6637 [50] and deployed by major vendors^{13,14}. In many other deployments, the fresh AES key is directly

¹¹ They also instantiated the KEM using SIKE [47], which subsequently suffered fatal attacks [16,66] and should no longer be used [48], so we do not discuss this further.

¹² <https://mailarchive.ietf.org/arch/msg/cfrg/zTnaLh05N7ipvPyJ8lmV7Iic9RU/>

¹³ <https://opensource.apple.com/source/Security/Security-59754.80.3/keychain/SecureObjectSync/SOSECWrapUnwrap.c>

¹⁴ <https://cloud.google.com/kms/docs/key-wrapping>

encrypted (using padding) using the PKE scheme, in particular using RSA-OAEP (as in AWS KMS¹⁵, AWS CloudHSM¹⁶, or Apple¹⁷).

It is clear there is a need for a standard to wrap data using a KEM public key. It is worth asking ourselves whether the few bytes that were gained by using key wrapping (rather than an AEAD) would be as impactful when using a post-quantum primitive for which the ciphertext is at least an order of magnitude larger than the expected gain (Table 1). A potential avenue (which requires proper evaluation) may be to use symmetric key wrapping as the data encapsulation mechanism in HPKE. Thus, we propose the following research investigation and standardization takeaway:

Takeway 13. Research and standardize a public-key wrapping method, either as a direct application of a (post-quantum) HPKE, or as a specialized construction.

4 Password-Authenticated Key Exchange

Asymmetric Password Authenticated Key Exchange (PAKE) protocols allow password authentication and mutually authenticated key exchange without disclosing passwords to servers. The most widely deployed PKI-free asymmetric PAKE is the Secure Remote Password (SRP) protocol [102,43,45]. SRP continues, as of today, to be the default asymmetric PAKE in many settings such as authentication in applications (e.g., AWS Cognito, keychains¹⁸, mail authentication¹⁹).

In 2019, the IETF Crypto Forum Research Group held a PAKE selection process, with the goal of recommending a symmetric and an asymmetric PAKE for usage in IETF protocols. It respectively selected CPace [1] and OPAQUE [15] in 2020. Although those state-of-the-art PAKEs now feature proofs in the Universal Composability model and good performance, they rely on primitives which fail to provide quantum resistance. OPAQUE was designed with modularity in mind by combining an oblivious pseudorandom function (OPRF [20]) and an authenticated key exchange. [15, Appendix B] explicitly mentions that a post-quantum AKE can be used in OPAQUE (further strengthening Takeway 12), but a fully-fledged post-quantum OPAQUE would also necessitate the OPRF to be quantum-resistant. Unfortunately, state-of-the-art post-quantum OPRFs are orders of magnitude away from being practical (Section 5). Few papers have looked at constructing post-quantum PAKEs directly [32].

¹⁵ <https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-get-public-key-and-token.html>

¹⁶ https://docs.aws.amazon.com/cloudhsm/latest/userguide/key_mgmt_util-wrapKey.html

¹⁷ <https://support.apple.com/guide/security/how-imessage-sends-and-receives-messages-sec70e68c949/1/web/1>

¹⁸ <https://blog.ipassword.com/developers-how-we-use-srp-and-you-can-too/>

¹⁹ <https://proton.me/blog/encrypted-email-authentication>

During the PAKE selection process, the notion of a PAKE being “quantum annoying” was proposed²⁰, and this property was later formalized and proved to hold for CPace [27]. Informally, a scheme is said to be quantum annoying if being able to solve discrete logarithms does not immediately provide the ability to compromise a system but rather only allows to eliminate a single possible password guess. In the absence of a post-quantum PAKE, such a property becomes very appealing as considerable quantum resources would be needed to compromise a single well-constructed password. While CPace can serve as a quantum-annoying balanced PAKE where possible, we propose the following research investigation and standardization takeaway:

Takeway 14. Research and standardize a post-quantum PAKE (and authenticated KEM; Takeway 12).

5 OPRF, Privacy Pass, and Blind Signatures

Besides OPAQUE, (verifiable) OPRFs are used to construct anonymous tokens [21,62,26], a form of lightweight anonymous credentials used as a trust signal by major vendors^{21,22,23}. These anonymous tokens aim at providing a private-key alternative to blind signatures, and are being developed and standardized in the IETF Privacy Pass working group²⁴.

As mentioned in Section 4, there exists no efficient post-quantum OPRF to date. Boneh et al. proposed two constructions based on isogenies in [14]. The first one was based on SIDH and is therefore insecure [16,66,48], and was recently improved by Basso in [12]. The second was based on CSIDH (a relatively novel hardness assumption) and had communication cost around 500kB per evaluation (no computation cost was provided). Albrecht et al. proposed a lattice-based construction [5] which is “*practically instantiable [but] far less efficient [than its classical counterpart]*”. They suggested that one may want to “*accept, for now, that VOPRFs are less appealing building blocks in a post-quantum world*” and to propose post-quantum alternatives on a per application basis instead. To address this gap, we propose the following research takeaway:

Takeway 15. Construct an efficient post-quantum OPRF for use in anonymous authentication schemes or propose post-quantum anonymous credential primitives using general-purpose zero-knowledge proofs.

The past two years have also seen a significant interest renewal for blind signatures, including constructions [37,38,59,17,22,58,36,18], attacks [13], specification [23], and deployments by major vendors^{22,25}. State-of-the-art post-quantum

²⁰ <https://mailarchive.ietf.org/arch/msg/cfrg/dtf91cmavpzT47U3AVxrVGNB5UM/>

²¹ <https://support.cloudflare.com/hc/en-us/articles/115001992652-Using-Privacy-Pass-with-Cloudflare>

²² <https://blog.cloudflare.com/eliminating-captchas-on-iphones-and-macs-using-new-standard/>

²³ <https://web.dev/trust-tokens/>

²⁴ <https://datatracker.ietf.org/wg/privacypass/about/>

²⁵ <https://one.google.com/about/vpn/howitworks>

blind signatures were initially proposed in 2010 by Rückert [85], and the latest state-of-the-art protocols [2,65,22] yield signatures of the order of 50kB, i.e., two orders of magnitude larger than an RSA-based blind signature as defined in [23]. Recent deployments and standardization of RSA-based blind signatures motivates the following research takeaway:

Takeway 16. Construct efficient quantum-safe blind signatures for privacy-preserving and authentication use-cases.

6 Conclusion

In this paper, we proposed 16 research and standardization open questions posed by the upcoming post-quantum migration. We emphasized the impact of post-quantum authentication in transport protocols: the size increase may exacerbate amplification attacks and congestion and requires new research and standards. Additionally, we pointed out the need of going beyond KEM standardization as done by NIST: many protocols would benefit from an authenticated KEM, as the generic construction of public-key encryption from KEM does not capture the versatility of public-key encryption use-cases today. Finally, we briefly discussed the state of the art for post-quantum password-authenticated key exchange and anonymous authentication to motivate future research.

Beyond what is discussed above, we invite the reader to consider standards and protocols which would benefit from increased public attention. One example is SSH. Although SSH carries huge amounts of proprietary data today, the IETF working group responsible for it has concluded. The harvest-now-decrypt-later concern is an important one, so that the IETF TLS working group embarked on a journey of updating TLS with new quantum-safe hybrid key exchanges. There is no SSH group to introduce post-quantum algorithms to SSH, and although there have been some initial efforts to address this [57], it will not be addressed by the PQUIP working group. Other than SSH, there are important cryptographic standards that will need to embark on a post-quantum journey as well. These include Trusted Platform Modules (ISO/IEC 11889), UEFI Secure Boot, OASIS Key Management Interoperability Protocol (KMIP) and PKCS#11 and more.

Finally, one should always ponder the impact of a post-quantum transition in proper context. Investigations so far [11,56,91,92,63] have been considering the time-to-first-byte at the 90-95th-percentile as an indication of overall impact for post-quantum connections. The tail-ends of the 90-95th-percentile may be overestimating this impact. At the time of this writing, web clients perform ≈ 13 connections per page to fetch $\approx 2\text{MB}$ of total data [40] on average. Connections at the tail-ends of the 90-95th-percentile that suffer significantly with 10-20kB of additional data are probably already suffering with $\approx 150\text{kB}$ per connection. Henceforth, even though one should aim for fairness of impact for upcoming changes, we stress that the post-quantum transition should avoid rendering poor connections with low time-to-last-byte much poorer than before.

References

1. Abdalla, M., Haase, B., Hesse, J.: CPace, a balanced composable PAKE. Tech. rep., Internet Research Task Force (2022), <https://datatracker.ietf.org/doc/draft-irtf-cfrg-pace/> 16
2. Agrawal, S., Kirshanova, E., Stehle, D., Yadav, A.: Can round-optimal lattice-based blind signatures be practical? Cryptology ePrint Archive, Report 2021/1565 (2021), <https://ia.cr/2021/1565> 18
3. Aguilar Melchor, C., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Persichetti, E., Zémor, G., Bos, J., Dion, A., Lacan, J., Robert, J.M., Veron, P.: HQC. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions> 2
4. Albrecht, M.R., Bernstein, D.J., Chou, T., Cid, C., Gilcher, J., Lange, T., Maram, V., von Maurich, I., Misoczki, R., Niederhagen, R., Paterson, K.G., Persichetti, E., Peters, C., Schwabe, P., Sendrier, N., Szefer, J., Tjhai, C.J., Tomlinson, M., Wang, W.: Classic McEliece. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-4-submissions> 2
5. Albrecht, M.R., Davidson, A., Deo, A., Smart, N.P.: Round-optimal verifiable oblivious pseudorandom functions from ideal lattices. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 261–289. Springer, Heidelberg, Germany, Virtual Event (May 10–13, 2021). https://doi.org/10.1007/978-3-030-75248-4_10 17
6. American National Standards Institute, Inc.: ANSI X9.44-2007 key establishment using integer factorization cryptography (2007), <https://webstore.ansi.org/standards/ascx9/ansix9442007r2017> 14
7. Anastasova, M., Kampanakis, P., Massimo, J.: PQ-HPKE: Post-quantum hybrid public key encryption. Intl Cryptographic Module Conference 2022 (2022), <https://ia.cr/2022/414> 15
8. American National Standards Institute (ANSI) X9.F1 subcommittee. ANSI X9.63 Public key cryptography for the Financial Services Industry: Elliptic curve key agreement and key transport schemes (Jul 5, 1998), working draft version 2.0 14
9. Aragon, N., Barreto, P., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Gueron, S., Guneyasu, T., Aguilar Melchor, C., Misoczki, R., Persichetti, E., Sendrier, N., Tillich, J.P., Zémor, G., Vasseur, V., Ghosh, S., Richter-Brokmann, J.: BIKE. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions> 2
10. Barnes, R., Bhargavan, K., Lipp, B., Wood, C.A.: Hybrid public key encryption. RFC, Internet Engineering Task Force (2022), <https://www.rfc-editor.org/rfc/rfc9180> 14, 15
11. Bas Westerbaan, C.: Sizing Up Post-Quantum Signatures (Nov 2021), <https://blog.cloudflare.com/sizing-up-post-quantum-signatures/> 3, 12, 18
12. Basso, A.: A post-quantum round-optimal oblivious PRF from isogenies. Cryptology ePrint Archive, Paper 2023/225 (2023), <https://eprint.iacr.org/2023/225> 17
13. Benhamouda, F., Lepoint, T., Loss, J., Orrù, M., Raykova, M.: On the (in)security of ROS. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 33–53. Springer, Heidelberg, Germany, Zagreb, Croatia (Oct 17–21, 2021). https://doi.org/10.1007/978-3-030-77870-5_2 17

14. Boneh, D., Kogan, D., Woo, K.: Oblivious pseudorandom functions from isogenies. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 520–550. Springer, Heidelberg, Germany, Daejeon, South Korea (Dec 7–11, 2020). https://doi.org/10.1007/978-3-030-64834-3_18 17
15. Bourdrez, D., Krawczyk, D.H., Lewi, K., Wood, C.A.: The OPAQUE asymmetric PAKE protocol. Tech. rep., Internet Research Task Force (2022), <https://datatracker.ietf.org/doc/draft-irtf-cfrg-opaque/> 16
16. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH (preliminary version). Cryptology ePrint Archive, Report 2022/975 (2022), <https://ia.cr/2022/975> 15, 17
17. Chairattana-Apirom, R., Hanzlik, L., Loss, J., Lysyanskaya, A., Wagner, B.: Piccut-choo and friends: Compact blind signatures via parallel instance cut-and-choose and more. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part III. LNCS, vol. 13509, pp. 3–31. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_1 17
18. Chairattana-Apirom, R., Lysyanskaya, A.: Compact cut-and-choose: Boosting the security of blind signature schemes, compactly. Cryptology ePrint Archive, Report 2022/003 (2022), <https://ia.cr/2022/003> 17
19. Chu, J., Dukkupati, N., Cheng, Y., Mathis, M.: Increasing TCP’s Initial Window. RFC 6928 (Apr 2013), <https://www.rfc-editor.org/info/rfc6928> 11, 12
20. Davidson, A., Faz-Hernandez, A., Sullivan, N., Wood, C.A.: Oblivious pseudorandom functions (OPRFs) using prime-order groups. Tech. rep., Internet Research Task Force (2022), <https://datatracker.ietf.org/doc/draft-irtf-cfrg-voprff/> 16
21. Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., Valsorda, F.: Privacy pass: Bypassing internet challenges anonymously. PoPETs 2018(3), 164–180 (Jul 2018). <https://doi.org/10.1515/popets-2018-0026> 17
22. del Pino, R., Katsumata, S.: A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 306–336. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15979-4_11 17, 18
23. Denis, F., Jacobs, F., Wood, C.A.: RSA blind signatures. Tech. rep., Internet Research Task Force (2022), <https://datatracker.ietf.org/doc/draft-irtf-cfrg-rsa-blind-signatures/> 17, 18
24. Dent, A.W.: A designer’s guide to KEMs. In: Paterson, K.G. (ed.) 9th IMA International Conference on Cryptography and Coding. LNCS, vol. 2898, pp. 133–151. Springer, Heidelberg, Germany, Cirencester, UK (Dec 16–18, 2003) 14
25. Dukkupati, N., Refice, T., Cheng, Y., Chu, J., Herbert, T., Agarwal, A., Jain, A., Sutin, N.: An argument for increasing tcp’s initial congestion window. SIGCOMM Comput. Commun. Rev. 40(3), 26–33 (jun 2010). <https://doi.org/10.1145/1823844.1823848> 11, 12
26. Durak, F.B., Vaudenay, S., Chase, M.: Anonymous tokens with hidden metadata bit from algebraic macs. Cryptology ePrint Archive, Paper 2022/1622 (2022), <https://ia.cr/2022/1622> 17
27. Eaton, E., Stebila, D.: The “quantum annoying” property of password-authenticated key exchange protocols. In: Cheon, J.H., Tillich, J.P. (eds.) Post-Quantum Cryptography - 12th International Workshop, PQCrypto 2021. pp. 154–173. Springer, Heidelberg, Germany, Daejeon, South Korea (Jul 20–22, 2021). https://doi.org/10.1007/978-3-030-81293-5_9 17

28. ETSI: ETSI TC Cyber Working Group for Quantum-Safe Cryptography (2017), <https://portal.etsi.org/TBSiteMap/CYBER/CYBERQSCToR.aspx>, Web page. Accessed 2019-07-25. 2
29. Fregly, A., Harvey, J., Jr., B.S.K., Sheth, S.: Merkle tree ladder mode: Reducing the size impact of NIST PQC signature algorithms in practice. Cryptology ePrint Archive, Paper 2022/1730 (2022), <https://ia.cr/2022/1730> 9
30. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999). https://doi.org/10.1007/3-540-48405-1_34 14
31. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. Journal of Cryptology 26(1), 80–101 (Jan 2013). <https://doi.org/10.1007/s00145-011-9114-1> 14
32. Gao, X., Ding, J., Liu, J., Li, L.: Post-quantum secure remote password protocol from RLWE problem. Cryptology ePrint Archive, Report 2017/1196 (2017), <https://ia.cr/2017/1196> 16
33. Gayoso Martínez, V., Hernández Á lvarez, F., Hernández Encinas, L., Sánchez Á vila, C.: A comparison of the standardized versions of ECIES. In: 2010 Sixth International Conference on Information Assurance and Security. pp. 1–4 (2010). <https://doi.org/10.1109/ISIAS.2010.5604194> 14
34. Ghedini, A., Vasiliev, V.: TLS Certificate Compression. RFC 8879 (Dec 2020). <https://doi.org/10.17487/RFC8879>, <https://www.rfc-editor.org/info/rfc8879> 7
35. Goertzen, J., Stebila, D.: Post-quantum signatures in DNSSEC via request-based fragmentation. CoRR abs/2211.14196 (2022). <https://doi.org/10.48550/arXiv.2211.14196> 9
36. Hanzlik, L., Loss, J., Wagner, B.: Rai-choo! Evolving blind signatures to the next level. Cryptology ePrint Archive, Report 2022/1350 (2022), <https://ia.cr/2022/1350> 17
37. Hauck, E., Kiltz, E., Loss, J.: A modular treatment of blind signatures from identification schemes. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 345–375. Springer, Heidelberg, Germany, Darmstadt, Germany (May 19–23, 2019). https://doi.org/10.1007/978-3-030-17659-4_12 17
38. Hauck, E., Kiltz, E., Loss, J., Nguyen, N.K.: Lattice-based blind signatures, revisited. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 500–529. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2020). https://doi.org/10.1007/978-3-030-56880-1_18 17
39. Housley, R.: Use of the HSS/LMS Hash-Based Signature Algorithm in the Cryptographic Message Syntax (CMS). RFC 8708 (Feb 2020), <https://www.rfc-editor.org/info/rfc8708> 2
40. http archive: Report: State of the Web, <http://httparchive.org/trends.php> 18
41. Huelsing, A., Butin, D., Gazdag, S.L., Rijneveld, J., Mohaisen, A.: XMSS: eXtended Merkle Signature Scheme. RFC 8391 (May 2018), <https://rfc-editor.org/rfc/rfc8391> 2
42. Hulsing, A., Bernstein, D.J., Dobraunig, C., Eichlseder, M., Fluhrer, S., Gazdag, S.L., Kampanakis, P., Kolbl, S., Lange, T., Lauridsen, M.M., Mendel, F., Niederhagen, R., Rechberger, C., Rijneveld, J., Schwabe, P., Aumasson, J.P., Westermann, B., Beullens, W.: SPHINCS+. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022> 2

43. IEEE draft standard P1363.2. Password-based public key cryptography. <http://grouper.ieee.org/groups/1363/passwdPK> (May 2004), draft Version 15 16
44. IEEE P1363a Committee. IEEE P1363a / D9 — standard specifications for public key cryptography: Additional techniques. <http://grouper.ieee.org/groups/1363/index.html/> (Jun 2001), draft Version 9 14
45. ISO: Information technology — security techniques — key management — part 4: Mechanisms based on weak secrets. ISO/IEC, International Organization for Standardization (2017), <https://www.iso.org/standard/67933.html> 16
46. Iyengar, J., Swett, I.: QUIC Loss Detection and Congestion Control. RFC 9002 (May 2021), <https://www.rfc-editor.org/info/rfc9002> 12
47. Jao, D., Azarderakhsh, R., Campagna, M., Costello, C., De Feo, L., Hess, B., Jalali, A., Koziel, B., LaMacchia, B., Longa, P., Naehrig, M., Renes, J., Soukharev, V., Urbanik, D., Pereira, G., Karabina, K., Hutchinson, A.: SIKE. Tech. rep., National Institute of Standards and Technology (2020), available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions> 15
48. Jao, D., Azarderakhsh, R., Campagna, M., Costello, C., De Feo, L., Hess, B., Jalali, A., Koziel, B., LaMacchia, B., Longa, P., Naehrig, M., Renes, J., Soukharev, V., Urbanik, D., Pereira, G., Karabina, K., Hutchinson, A.: SIKE. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions> 2, 15, 17
49. Jaques, S.: Landscape of quantum computing in 2022 (2022), https://sam-jaques.appspot.com/quantum_landscape_2022 1
50. Jivsov, A.: Elliptic curve cryptography (ECC) in OpenPGP. RFC, Internet Engineering Task Force (2016), <https://www.rfc-editor.org/rfc/rfc6637> 15
51. Jonsson, J., Kaliski, B.: Public-key cryptography standards (PKCS) #1: RSA cryptography specifications version 2.1. RFC, Internet Engineering Task Force (2003), <https://www.rfc-editor.org/rfc/rfc3447> 14
52. K. Moriarty, E., Kaliski, B., Jonsson, J., Rusch, A.: PKCS #1: RSA cryptography specifications version 2.2. RFC, Internet Engineering Task Force (2012), <https://www.rfc-editor.org/rfc/rfc8017> 14
53. Kaliski, B.: PKCS #1: RSA encryption version 1.5. RFC, Internet Engineering Task Force (1998), <https://www.rfc-editor.org/rfc/rfc2313> 14
54. Kaliski, B., Jonsson, J.: PKCS #1: RSA cryptography specifications version 2.0. RFC, Internet Engineering Task Force (1998), <https://www.rfc-editor.org/rfc/rfc2437> 14
55. Kampanakis, P., Kallitsis, M.: Faster post-quantum TLS handshakes without intermediate CA certificates. In: Dolev, S., Katz, J., Meisels, A. (eds.) *Cyber Security, Cryptology, and Machine Learning*. pp. 337–355. Springer International Publishing, Cham (2022) 3, 6
56. Kampanakis, P., Sikeridis, D.: Two PQ signature use-cases: Non-issues, challenges and potential solutions. *Cryptology ePrint Archive*, Report 2019/1276 (2019), <https://ia.cr/2019/1276> 3, 18
57. Kampanakis, P., Stebila, D., Hansen, T.: Post-quantum Hybrid Key Exchange in SSH. Internet-Draft draft-kampanakis-curdle-ssh-pq-ke-00, Internet Engineering Task Force (Nov 2022), <https://datatracker.ietf.org/doc/draft-kampanakis-curdle-ssh-pq-ke/00/>, work in Progress 18
58. Kastner, J., Loss, J., Xu, J.: The abe-okamoto partially blind signature scheme revisited. *Cryptology ePrint Archive*, Report 2022/1232 (2022), <https://ia.cr/2022/1232> 17

59. Katz, J., Loss, J., Rosenberg, M.: Boosting the security of blind signature schemes. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part IV. LNCS, vol. 13093, pp. 468–492. Springer, Heidelberg, Germany, Singapore (Dec 6–10, 2021). https://doi.org/10.1007/978-3-030-92068-5_16 17
60. Kaufman, C., Hoffman, P.E., Nir, Y., Eronen, P., Kivinen, T.: Internet Key Exchange Protocol Version 2 (IKEv2). RFC 7296 (Oct 2014), <https://www.rfc-editor.org/info/rfc7296> 10, 13
61. Kinnear, E., McManus, P., Pauly, T., Verma, T., Wood, C.A.: Oblivious DNS over HTTPS. RFC, Internet Engineering Task Force (2022), <https://www.rfc-editor.org/rfc/rfc9230> 14
62. Kreuter, B., Lepoint, T., Orrù, M., Raykova, M.: Anonymous tokens with private metadata bit. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part I. LNCS, vol. 12170, pp. 308–336. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2020). https://doi.org/10.1007/978-3-030-56784-2_11 17
63. Kris Kwiatkowski, L.V.: The TLS Post-Quantum Experiment (Oct 2020), <https://blog.cloudflare.com/the-tls-post-quantum-experiment/> 18
64. Lyubashevsky, V., Ducas, L., Kiltz, E., Lepoint, T., Schwabe, P., Seiler, G., Stehlé, D., Bai, S.: CRYSTALS-DILITHIUM. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022> 2
65. Lyubashevsky, V., Nguyen, N.K., Plancon, M.: Efficient lattice-based blind signatures via gaussian one-time signatures. Cryptology ePrint Archive, Report 2022/006 (2022), <https://ia.cr/2022/006> 18
66. Maino, L., Martindale, C.: An attack on SIDH with arbitrary starting curve. Cryptology ePrint Archive, Report 2022/1026 (2022), <https://ia.cr/2022/1026> 15, 17
67. Majkowski, M.: Reflections on reflection (attacks) (May 2017), <https://blog.cloudflare.com/reflections-on-reflections/> 4, 9
68. Massimo, J., Kampanakis, P., Turner, S., Westerbaan, B.: Internet X.509 Public Key Infrastructure: Algorithm Identifiers for Dilithium. Internet-Draft draft-ietf-lamps-dilithium-certificates-00, Internet Engineering Task Force (Sep 2022), <https://datatracker.ietf.org/doc/draft-ietf-lamps-dilithium-certificates/00/>, work in Progress 2
69. McGrew, D., Curcio, M., Fluhrer, S.: Leighton-Micali Hash-Based Signatures. RFC 8554 (Apr 2019), <https://rfc-editor.org/rfc/rfc8554> 2
70. Mononen, T., Kause, T., Farrell, S., Adams, D.C.: Internet X.509 public key infrastructure certificate management protocol (CMP). RFC, Internet Engineering Task Force (2005), <https://www.rfc-editor.org/rfc/rfc4210> 15
71. 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. SIGCOMM Comput. Commun. Rev. **50**(4), 49–57 (oct 2020). <https://doi.org/10.1145/3431832.3431838> 9
72. Nawrocki, M., Tehrani, P.F., Hiesgen, R., Mücke, J., Schmidt, T.C., Wählisch, M.: On the interplay between TLS certificates and QUIC performance. In: Proceedings of the 18th International Conference on emerging Networking EXperiments and Technologies. ACM (nov 2022). <https://doi.org/10.1145/3555050.3569123> 5, 6, 8
73. NIST: NIST PQ project (Feb 2022), <https://csrc.nist.gov/projects/post-quantum-cryptography> 2

74. Paquin, C., Stebila, D., Tamvada, G.: Benchmarking post-quantum cryptography in tls. In: Ding, J., Tillich, J.P. (eds.) *Post-Quantum Cryptography*. pp. 72–91. Springer International Publishing, Cham (2020) [3](#)
75. Partridge, D.C., Allman, M., Floyd, S.: Increasing TCP’s Initial Window. RFC 3390 (Nov 2002), <https://www.rfc-editor.org/info/rfc3390> [11](#)
76. Paxson, D.V., Allman, M., Stevens, W.R.: TCP Congestion Control. RFC 2581 (Apr 1999), <https://www.rfc-editor.org/info/rfc2581> [11](#)
77. Planet, C.: Initcwnd settings of major CDN providers (Feb 2017), <https://www.cdnplanet.com/blog/initcwnd-settings-major-cdn-providers/> [12](#)
78. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022> [2](#)
79. Prince, M.: The DDoS That Almost Broke the Internet (May 2017), <https://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/> [4](#), [9](#)
80. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Aug 2018), <https://rfc-editor.org/rfc/rfc8446> [6](#)
81. Rescorla, E., Modadugu, N.: Datagram Transport Layer Security Version 1.2. RFC 6347 (Jan 2012), <https://www.rfc-editor.org/info/rfc6347> [8](#)
82. Rescorla, E., Oku, K., Sullivan, N., Wood, C.A.: TLS encrypted client hello. Tech. rep., Internet Engineering Task Force (2022), <https://datatracker.ietf.org/doc/draft-ietf-tls-esni/> [14](#)
83. Rescorla, E., Tschofenig, H., Modadugu, N.: The Datagram Transport Layer Security (DTLS) Protocol Version 1.3. RFC 9147 (Apr 2022), <https://www.rfc-editor.org/info/rfc9147> [8](#), [13](#)
84. Rossow, C.: Amplification hell: Revisiting network protocols for ddos abuse (01 2014). <https://doi.org/10.14722/ndss.2014.23233> [4](#)
85. Rückert, M.: Lattice-based blind signatures. In: Abe, M. (ed.) *ASIACRYPT 2010*. LNCS, vol. 6477, pp. 413–430. Springer, Heidelberg, Germany, Singapore (Dec 5–9, 2010). https://doi.org/10.1007/978-3-642-17373-8_24 [18](#)
86. Rühth, J., Bormann, C., Hohlfeld, O.: Large-scale scanning of tcp’s initial window. In: *Proceedings of the 2017 Internet Measurement Conference*. p. 304–310. IMC ’17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3131365.3131370> [11](#), [12](#)
87. Schwabe, P., Avanzi, R., Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Seiler, G., Stehlé, D., Ding, J.: CRYSTALS-KYBER. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022> [2](#), [14](#)
88. Certicom research, standards for efficient cryptography group (SECG) — sec 1: Elliptic curve cryptography. http://www.secg.org/secg_docs.htm (Sep 20, 2000), version 1.0 [14](#)
89. Shoup, V.: A proposal for an ISO standard for public key encryption. *Cryptology ePrint Archive*, Report 2001/112 (2001), <https://ia.cr/2001/112> [14](#)
90. Shoup, V.: ISO 18033-2: An emerging standard for public-key encryption. <https://shoup.net/iso/std6.pdf> (Dec 2004), final Committee Draft [14](#)
91. Sikeridis, D., Kampanakis, P., Devetsikiotis, M.: Assessing the overhead of post-quantum cryptography in TLS 1.3 and SSH. In: *Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies*. p. 149–156. CoNEXT ’20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3386367.3431305> [3](#), [18](#)

92. Sikeridis, D., Kampanakis, P., Devetsikiotis, M.: Post-quantum Authentication in TLS 1.3: A performance study. In: 27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23-26, 2020. The Internet Society (2020), <https://www.ndss-symposium.org/ndss-paper/post-quantum-authentication-in-tls-1-3-a-performance-study/> 3, 12, 18
93. Smyslov, V.: Intermediate Exchange in the Internet Key Exchange Protocol Version 2 (IKEv2). RFC 9242 (May 2022), <https://www.rfc-editor.org/info/rfc9242> 10, 13
94. Stebila, D., Fluhner, S., Gueron, S.: Hybrid key exchange in TLS 1.3. Internet-Draft draft-ietf-tls-hybrid-design-04, Internet Engineering Task Force (Jan 2022), <https://datatracker.ietf.org/doc/html/draft-ietf-tls-hybrid-design-04>, work in Progress 2
95. Stebila, D., Fluhner, S., Gueron, S.: Hybrid key exchange in TLS 1.3. Internet-Draft draft-ietf-tls-hybrid-design-05, Internet Engineering Task Force (Aug 2022), <https://datatracker.ietf.org/doc/draft-ietf-tls-hybrid-design-05/>, work in Progress 2
96. Stevens, W.R.: TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms. RFC 2001 (Jan 1997), <https://www.rfc-editor.org/info/rfc2001> 11
97. Thomson, M., Kampanakis, P., Bytheway, C., Westerbaan, B.: Suppressing CA Certificates in TLS 1.3. Internet-Draft draft-kampanakis-tls-scas-latest-02, Internet Engineering Task Force (Jul 2022), <https://datatracker.ietf.org/doc/draft-kampanakis-tls-scas-latest-02/>, work in Progress 6
98. Thomson, M., Turner, S.: Using TLS to Secure QUIC. RFC 9001 (May 2021), <https://www.rfc-editor.org/info/rfc9001> 5, 6
99. Tjhai, C., Tomlinson, M., Bartlett, G., Fluhner, S., Geest, D.V., Garcia-Morchon, O., Smyslov, V.: Multiple Key Exchanges in IKEv2. Internet-Draft draft-ietf-ipsecme-ikev2-multiple-ke-04, Internet Engineering Task Force (Sep 2021), <https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-ikev2-multiple-ke-04>, work in Progress 2
100. Tjhai, C., Tomlinson, M., Bartlett, G., Fluhner, S., Geest, D.V., Garcia-Morchon, O., Smyslov, V.: Multiple Key Exchanges in IKEv2. Internet-Draft draft-ietf-ipsecme-ikev2-multiple-ke-12, Internet Engineering Task Force (Dec 2022), <https://datatracker.ietf.org/doc/draft-ietf-ipsecme-ikev2-multiple-ke-12/>, work in Progress 2
101. Touch, D.J.D.: Automating the Initial Window in TCP. Internet-Draft draft-touch-tcpm-automatic-iw-03, Internet Engineering Task Force (Jul 2012), <https://datatracker.ietf.org/doc/draft-touch-tcpm-automatic-iw-03/>, work in Progress 11
102. Wu, T.: The SRP authentication and key exchange system. RFC, Internet Engineering Task Force (2000), <https://www.rfc-editor.org/rfc/rfc2945> 16