
Stream: Internet Engineering Task Force (IETF)
RFC: [9958](#)
Category: Informational
Published: May 2026
ISSN: 2070-1721
Authors: A. Banerjee T. Reddy.K D. Schoinianakis T. Hollebeek M. Ounsworth
 Nokia *Nokia* *Nokia* *DigiCert* *Entrust*

RFC 9958

Post-Quantum Cryptography for Engineers

Abstract

The advent of a cryptographically relevant quantum computer (CRQC) would render state-of-the-art, traditional public key algorithms deployed today obsolete, as the mathematical assumptions underpinning their security would no longer hold. To address this, protocols and infrastructure must transition to post-quantum algorithms, which are designed to resist both traditional and quantum attacks. This document explains why engineers need to be aware of and understand post-quantum cryptography (PQC), and it details the impact of CRQCs on existing systems and the challenges involved in transitioning to post-quantum algorithms. Unlike previous cryptographic updates, this shift may require significant protocol redesign due to the unique properties of post-quantum algorithms.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9958>.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	4
2. Terminology	6
3. Threat of CRQCs on Cryptography	7
3.1. Symmetric Cryptography	7
3.2. Asymmetric Cryptography	8
3.3. Quantum Side-Channel Attacks	9
4. Traditional Cryptographic Primitives That Could Be Replaced by PQC	9
5. NIST PQC Algorithms	10
5.1. NIST Candidates Selected for Standardization	10
5.1.1. PQC Key Encapsulation Mechanisms (KEMs)	10
5.1.2. PQC Signatures	10
6. ISO Candidates Selected for Standardization	11
6.1. PQC Key Encapsulation Mechanisms (KEMs)	11
7. Timeline for Transition	11
8. PQC Categories	13
8.1. Lattice-Based Public Key Cryptography	13
8.2. Hash-Based Public Key Cryptography	14
8.3. Code-Based Public Key Cryptography	15
9. KEMs	15
9.1. Authenticated Key Exchange	16
9.2. Security Properties of KEMs	20
9.2.1. IND-CCA2	20
9.2.2. Binding	20
9.3. HPKE	21

10. PQC Signatures	21
10.1. Security Properties of PQC Signatures	21
10.1.1. EUF-CMA and SUF-CMA	21
10.2. Details of FN-DSA, ML-DSA, and SLH-DSA	22
10.3. Details of XMSS and LMS	23
10.3.1. LMS Key and Signature Sizes	23
10.4. Hash-then-Sign	24
11. NIST Recommendations for Security and Performance Trade-offs	25
12. Comparing PQC KEMs/Signatures and Traditional KEMs/Signatures	27
13. Post-Quantum and Traditional (PQ/T) Hybrid Schemes	29
13.1. PQ/T Hybrid Confidentiality	29
13.2. PQ/T Hybrid Authentication	30
13.3. Hybrid Cryptographic Algorithm Combinations: Considerations and Approaches	30
13.3.1. Hybrid Cryptographic Combinations	30
13.3.2. Composite Keys in Hybrid Schemes	30
13.3.3. Key Reuse in Hybrid Schemes	31
13.3.4. Future Directions and Ongoing Research	32
14. Impact on Constrained Devices and Networks	32
15. Security Considerations	32
15.1. Cryptanalysis	32
15.2. Cryptographic Agility	33
15.3. Jurisdictional Fragmentation	33
15.4. Hybrid Key Exchange and Signatures: Bridging the Gap Between PQ/T Cryptography	34
15.5. Caution: Ciphertext Commitment in KEM vs. DH	34
16. IANA Considerations	35
17. Further Reading and Resources	35
18. References	35
18.1. Normative References	35
18.2. Informative References	36
Acknowledgements	42

1. Introduction

Quantum computing is no longer just a theoretical concept in computational science and physics; it is now an active area of research with practical implications. Considerable research efforts and enormous corporate and government funding for the development of practical quantum computing systems are currently being invested. At the time this document is published, cryptographically relevant quantum computers (CRQCs) that can break widely used asymmetric algorithms (also known as public key algorithms) are not yet available. However, there is ongoing research and development in the field of quantum computing, with the goal of building more powerful and scalable quantum computers.

One common myth is that quantum computers are faster than conventional CPUs and GPUs in all areas. This is not the case; much as GPUs outperform general-purpose CPUs only on specific types of problems, quantum computers also have a niche set of problems on which they excel. Unfortunately for cryptographers, integer factorization and discrete logarithms, the mathematical problems underpinning much of classical public key cryptography, happen to fall within the niche in which quantum computers are expected to excel. As quantum technology advances, there is the potential for future quantum computers to have a significant impact on current cryptographic systems. Predicting the date of emergence of a CRQC is a challenging task, and there is ongoing uncertainty regarding when they will become practically feasible [[CRQCThreat](#)].

Extensive research has produced several post-quantum cryptographic algorithms that offer the potential to ensure cryptography's survival in the quantum computing era. However, transitioning to a post-quantum infrastructure is not a straightforward task, and there are numerous challenges to overcome. It requires a combination of engineering efforts, proactive assessment and evaluation of available technologies, and a careful approach to product development and deployment.

PQC is sometimes referred to as "quantum-proof", "quantum-safe", or "quantum-resistant". It is the development of cryptographic algorithms designed to secure communication and data in a world where quantum computers are powerful enough to break traditional cryptographic systems, such as RSA (Rivest-Shamir-Adleman) and ECC (Elliptic Curve Cryptography). PQC algorithms are intended to be resistant to attacks by quantum computers, which use quantum-mechanical phenomena to solve mathematical problems that are infeasible for classical computers.

As the threat of CRQCs draws nearer, engineers responsible for designing, maintaining, and securing cryptographic systems must prepare for the significant changes that the existence of CRQCs will bring. Engineers need to understand how to implement post-quantum algorithms in applications, how to evaluate the trade-offs between security and performance, and how to ensure backward compatibility with current systems where needed. This is not merely a one-for-

one replacement of algorithms; in many cases, the shift to PQC will involve redesigning protocols and infrastructure to accommodate the significant differences in resource utilization and key sizes between traditional and PQC algorithms. Due to the wide-ranging nature of these impacts, discussions of protocol changes are integrated throughout this document rather than being confined to a single section.

This document aims to provide general guidance to engineers working on cryptographic libraries, network security, and infrastructure development, where long-term security planning is crucial. The document covers topics such as selecting appropriate PQC algorithms and understanding the differences between PQC Key Encapsulation Mechanisms (KEMs) and traditional Diffie-Hellman (DH) and RSA-style key exchanges, and it provides insights into expected differences in keys, ciphertext, signature sizes, and processing times between PQC and traditional algorithms. Additionally, it discusses the potential threat to symmetric cryptography and hash functions from CRQCs.

It is important to remember that asymmetric algorithms (also known as public key algorithms) are largely used for secure communications between organizations or endpoints that may not have previously interacted, so a significant amount of coordination between organizations, and within and between ecosystems, needs to be taken into account. Such transitions are some of the most complicated in the tech industry and will require staged migrations in which upgraded agents need to coexist and communicate with non-upgraded agents at a scale never before undertaken.

The National Security Agency (NSA) of the United States released an article on future PQC algorithm requirements for US national security systems [CNSA2-0] based on the need to protect against deployments of CRQCs in the future. The German Federal Office for Information Security (BSI) has also released a PQC migration and recommendations document [BSI-PQC] that largely aligns with United States National Institute of Standards and Technology (NIST) and NSA guidance but differs in aspects such as specific PQC algorithm profiles.

CRQCs pose a threat to both symmetric and asymmetric cryptographic schemes. However, the threat to asymmetric cryptography is significantly greater due to Shor's algorithm [Shors], which can break widely used public key schemes like RSA and ECC. Symmetric cryptography and hash functions face a lower risk from Grover's algorithm [Grovers], although the impact is less severe and can typically be mitigated by doubling key and digest lengths where the risk applies. It is crucial for the reader to understand that when "PQC" is mentioned in the document, it means asymmetric cryptography (or public key cryptography) and not any symmetric algorithms based on stream ciphers, block ciphers, hash functions, Message Authentication Codes (MACs), etc., which are less vulnerable to quantum computers. This document does not cover topics such as when traditional algorithms might become vulnerable (for that, see documents such as [QC-DNS] and others).

This document does not cover unrelated technologies like quantum key distribution (QKD) or quantum key generation, which use quantum hardware to exploit quantum effects to protect communications and generate keys, respectively. PQC is based on conventional math (not on quantum mechanics) and software, and it can be run on any general-purpose computer.

This document does not go into the deep mathematics or technical specification of the PQC algorithms but rather provides an overview to engineers on the current threat landscape and the relevant algorithms designed to help prevent those threats. Also, the cryptographic and algorithmic guidance given in this document should be taken as non-authoritative if it conflicts with emerging and evolving guidance from the IRTF's Crypto Forum Research Group (CFRG).

2. Terminology

Quantum computer: A computer that performs computations using quantum-mechanical phenomena such as superposition and entanglement.

Physical qubit: The basic physical unit in a quantum computer, which is prone to noise and errors.

Logical qubit: A fault-tolerant qubit constructed from multiple physical qubits using quantum error correction; it is the effective unit for reliable quantum computation.

Post-Quantum Cryptography (PQC): Cryptographic algorithms designed to be secure against quantum and classical attacks.

Cryptographically Relevant Quantum Computer (CRQC): A quantum computer with sufficient logical qubits to break traditional asymmetric cryptographic algorithms (e.g., RSA or ECC) within a practical timeframe.

Public Key Cryptography (also called Asymmetric Cryptography): A class of cryptographic algorithms in which separate keys are used for encryption and decryption or for signing and verification. Throughout this document, the terms Public Key Cryptography and Asymmetric Cryptography are used interchangeably.

There is ongoing discussion about whether to use the term "post-quantum", "quantum ready", "quantum resistant", or "quantum secure" to describe algorithms that resist CRQCs, and a consensus has not yet been reached. NIST has coined the term "post-quantum" to refer to the algorithms that participated in its competition-like selection process; in this context, the term can be interpreted to mean "the set of algorithms that are designed to still be relevant after quantum computers exist" and not a statement about their security. "Quantum resistant" or "quantum secure" is obviously the goal of these algorithms; however, some people have raised concerns that labeling a class of algorithms as "quantum resistant" or "quantum secure" could lead to confusion if one or more of those algorithms are later found to be insecure or to not resist quantum computers as much as theory predicted. "Quantum ready" is often used to refer to a solution -- device, appliance, or software stack -- that has reached maturity with regard to integration of these new cryptographic algorithms. That said, the authors recognize that there is great variability in how these terms are used. This document uses these terms interchangeably to refer to such algorithms.

In this document, the terms "current", "state-of-the-art", and "ongoing" refer to work, research, investigations, deployments, or developments that are applicable at the time of publication.

3. Threat of CRQCs on Cryptography

When considering the security risks associated with the ability of a quantum computer to attack traditional cryptography, it is important to distinguish between the impact on symmetric algorithms and public key ones. Dr. Peter Shor and Dr. Lov Grover developed two algorithms that changed the way the world thinks of security under the presence of a CRQC.

Quantum computers are, by their nature, hybrids of classical and quantum computational units. For example, Shor's algorithm consists of a combination of quantum and classical computational steps. Thus, the term "quantum adversary" should be thought of as "quantum-enhanced adversary", meaning they have access to both classical and quantum computational techniques.

Although large-scale quantum computers do not yet exist to experiment on, the theoretical properties of quantum computation are very well understood. This allows engineers and researchers to reason about the upper limits of quantum-enhanced computation and to design cryptographic algorithms that are resistant to any conceivable form of quantum cryptanalysis.

3.1. Symmetric Cryptography

For unstructured data such as symmetric encrypted data or cryptographic hashes, although CRQCs can search for specific solutions across all possible input combinations (e.g., Grover's algorithm), no quantum algorithm is known to break the underlying security properties of these classes of algorithms. Symmetric-key cryptography, which includes keyed primitives such as block ciphers (e.g., AES) and message authentication mechanisms (e.g., HMAC-SHA256), relies on secret keys shared between the sender and receiver and remains secure even in a post-quantum world. Symmetric cryptography also includes hash functions (e.g., SHA-256) that are used for secure message digesting without any shared key material. Hashed Message Authentication Code (HMAC) is a specific construction that utilizes a cryptographic hash function and a secret key shared between the sender and receiver to produce a message authentication code.

Grover's algorithm is a quantum search algorithm that provides a theoretical quadratic speedup for searching an unstructured database, compared to traditional search algorithms. This has led to the common misconception that symmetric key lengths need to be doubled for quantum security. When you consider the mapping of hash values to their corresponding hash inputs (also known as pre-image) or of ciphertext blocks to the corresponding plaintext blocks as an unstructured database, then Grover's algorithm theoretically requires doubling the key sizes of the symmetric algorithms that are currently deployed at the time of publication to counter the quadratic speedup and maintain the current security level. This is because Grover's algorithm reduces the amount of operations to break 128-bit symmetric cryptography to 2^{64} quantum operations, which might sound computationally feasible. However, quantum operations are fundamentally different from classical ones, as 2^{64} classical operations can be efficiently parallelized but 2^{64} quantum operations must be performed serially, making them infeasible on practical quantum computers.

Grover's algorithm is highly non-parallelizable and even if one deploys 2^c computational units in parallel to brute-force a key using Grover's algorithm, it will complete in time proportional to $2^{\{(128-c)/2\}}$, or, put simply, using 256 quantum computers will only reduce runtime by a factor of 16, 1024 quantum computers will only reduce runtime by a factor of 32, and so forth (see [NIST] and [Cloudflare]). Due to this inherent limitation, the general expert consensus is that AES-128 remains secure in practice and key sizes do not necessarily need to be doubled.

It would be natural to ask whether future research will develop a superior algorithm that could outperform Grover's algorithm in the general case. However, Christof Zalka has shown that Grover's algorithm achieves the best possible complexity for this type of search, meaning no significantly faster quantum approach is expected [Grover-Search].

Finally, in their evaluation criteria for PQC, NIST is assessing the security levels of proposed post-quantum algorithms by comparing them against the equivalent traditional and quantum security of AES-128, AES-192, and AES-256. This indicates that NIST is confident in the stable security properties of AES, even in the presence of both traditional and quantum attacks. As a result, 128-bit algorithms can be considered quantum-safe for the foreseeable future. However, for compliance purposes, some organizations, such as the French National Agency for the Security of Information Systems (ANSSI) [ANSSI] and the National Security Agency (NSA) (CNSA 2.0) [CNSA2-0], recommend the use of AES-256.

3.2. Asymmetric Cryptography

"Shor's algorithm" efficiently solves the integer factorization problem (and the related discrete logarithm problem), which underpin the foundations of the vast majority of public key cryptography that the world uses today. This implies that, if a CRQC is developed, today's public key algorithms (e.g., RSA, Diffie-Hellman, and ECC, as well as less commonly used variants such as ElGamal [RFC6090] and Schnorr signatures [RFC8235]) and protocols would need to be replaced by algorithms and protocols that can offer cryptanalytic resistance against CRQCs. Note that Shor's algorithm cannot run solely on a classical computer; it requires a CRQC.

For example, studies show that, if a CRQC existed, it could break RSA-2048 in hours or even seconds depending on assumptions about error correction [RSAShor] [RSA8HRS] [RSA10SC]. While such machines are purely theoretical at the time of writing, this illustrates the eventual vulnerability of RSA to CRQCs.

For structured data such as public keys and signatures, CRQCs can fully solve the underlying hard problems used in traditional cryptography (see Shor's algorithm). Because an increase in the size of the key pair would not provide a secure solution (short of RSA keys that are many gigabytes in size [PQRSA]), a complete replacement of the algorithm is needed. Therefore, post-quantum public key cryptography must rely on problems that are different from the ones used in traditional public key cryptography (i.e., the integer factorization problem, the finite-field discrete logarithm problem, and the elliptic-curve discrete logarithm problem).

3.3. Quantum Side-Channel Attacks

Cryptographic side-channel attacks exploit physical implementations (such as timing, power consumption, or electromagnetic leakage) to recover secret keys.

The field of cryptographic side-channel attacks potentially stands to gain a boost in attacker power once cryptanalytic techniques can be enhanced with quantum computation techniques [[QuantSide](#)]. While a full discussion of quantum side-channel techniques is beyond the scope of this document, implementers of cryptographic hardware should be aware that current best practices for side-channel resistance may not be sufficient against quantum adversaries.

4. Traditional Cryptographic Primitives That Could Be Replaced by PQC

Any asymmetric cryptographic algorithm based on integer factorization, finite field discrete logarithms, or elliptic-curve discrete logarithms will be vulnerable to attacks using Shor's algorithm on a CRQC. This document focuses on the principal functions of asymmetric cryptography:

Key agreement and key transport: Key agreement schemes, typically referred to as Diffie-Hellman (DH) or Elliptic Curve Diffie-Hellman (ECDH), as well as key transport, typically using RSA encryption, are used to establish a shared cryptographic key for secure communication. They are one of the mechanisms that can be replaced by PQC, as they are based on existing public key cryptography and are therefore vulnerable to Shor's algorithm. A CRQC can employ Shor's algorithm to efficiently find the prime factors of a large public key (in the case of RSA), which, in turn, can be exploited to derive the private key. In the case of DH, a CRQC has the potential to calculate the discrete logarithm of the (short- or long-term) DH public key. This, in turn, would reveal the secret required to derive the symmetric encryption key.

Digital signatures: Digital signature schemes are used to authenticate the identity of a sender, detect unauthorized modifications to data, and underpin trust in a system. Similar to key agreement, signatures also depend on a public-private key pair based on the same mathematics as for key agreement and key transport. Because of this, a break in existing public key cryptography will also affect traditional digital signatures, hence the importance of developing post-quantum digital signatures.

Boneh-Boyen-Shacham (BBS) signatures: BBS signatures are a privacy-preserving signature scheme that offers zero-knowledge proof-like properties by allowing selective disclosure of specific signed attributes without revealing the entire set of signed data. The security of BBS signatures relies on the hardness of the discrete logarithm problem, making them vulnerable to Shor's algorithm. A CRQC can break the data authenticity security property of BBS but not the data confidentiality ([Section 6.9](#) of [[BBS-SIG-SCHEME](#)]).

Content encryption: Content encryption typically refers to the encryption of the data using symmetric key algorithms, such as AES, to ensure confidentiality. The threat to symmetric cryptography is discussed in [Section 3.1](#).

5. NIST PQC Algorithms

At the time of writing, NIST has standardized three PQC algorithms, with more expected to be standardized in the future (see [\[NISTFINAL\]](#)). These algorithms are not necessarily drop-in replacements for traditional asymmetric cryptographic algorithms. For instance, RSA [\[RSA\]](#) and ECC [\[RFC6090\]](#) can be used as both a KEM and a signature scheme, whereas there is currently no post-quantum algorithm that can perform both functions. When upgrading protocols, it is important to replace the existing use of traditional algorithms with either a PQC KEM or a PQC signature method, depending on how the traditional algorithm was previously being used. Additionally, KEMs, as described in [Section 9](#), present a different API than either key agreement or key transport primitives. As a result, they may require protocol-level or application-level changes in order to be incorporated.

5.1. NIST Candidates Selected for Standardization

5.1.1. PQC Key Encapsulation Mechanisms (KEMs)

ML-KEM: Module-Lattice-Based Key-Encapsulation Mechanism. See FIPS 203 [\[ML-KEM\]](#).

HQC: Hamming Quasi-Cyclic. See [\[HQC\]](#). The coding algorithm based on the hardness of the syndrome decoding problem for quasi-cyclic concatenated Reed-Muller and Reed-Solomon (RMRS) codes in the Hamming metric. Reed-Muller (RM) codes are a class of block error-correcting codes commonly used in wireless and deep-space communications, while Reed-Solomon (RS) codes are widely used to detect and correct multiple-bit errors. HQC has been selected as part of the NIST post-quantum cryptography project but has not yet been standardized.

5.1.2. PQC Signatures

ML-DSA: Module-Lattice-Based Digital Signature Algorithm. See FIPS 204 [\[ML-DSA\]](#).

SLH-DSA: Stateless Hash-Based Digital Signature Algorithm. See FIPS 205 [\[SLH-DSA\]](#).

FN-DSA: Fast-Fourier Transform over NTRU-Lattice-Based Digital Signature Algorithm. See [\[FN-DSA\]](#); note that NIST has named this algorithm "FN-DSA" and assigned "FIPS 206" for its specification, but at the time of this document's publication, it has not yet been released.

For more information about these, see [Sections 8.1](#), [8.2](#), and [10.2](#).

6. ISO Candidates Selected for Standardization

At the time of writing, ISO has selected three PQC KEM algorithms as candidates for standardization; these are mentioned in the following subsection.

6.1. PQC Key Encapsulation Mechanisms (KEMs)

FrodoKEM: KEM based on the hardness of learning with errors in algebraically unstructured lattices. See [[FrodoKEM](#)].

ClassicMcEliece: KEM based on the hardness of syndrome decoding of Goppa codes. Goppa codes are a class of error-correcting codes that can correct a certain number of errors in a transmitted message. The decoding problem involves recovering the original message from the received noisy codeword. See [[ClassicMcEliece](#)].

NTRU: KEM based on the "N-th degree Truncated polynomial Ring Units" (NTRU) lattices. Variants include Streamlined NTRU Prime (sntrup761), which is leveraged for use in SSH [[RFC9941](#)]. See [[NTRU](#)].

7. Timeline for Transition

The timeline and driving motivation for transition differ slightly between data confidentiality (e.g., encryption) and data authentication (e.g., signature) use cases.

For data confidentiality, one is concerned with the so-called "harvest now, decrypt later" (HNDL) attack where a malicious actor with adequate resources can launch an attack to store sensitive encrypted data today that they hope to decrypt once a CRQC is available. This implies that, every day, sensitive encrypted data is susceptible to the attack by not implementing quantum-safe strategies, as it corresponds to data possibly being deciphered in the future.

For authentication, it is often the case that signatures have a very short lifetime between signing and verifying (such as during a TLS handshake), but some authentication use cases do require long lifetimes, such as signing firmware or software that will be active for decades, signing legal documents, or signing certificates that will be embedded into hardware devices such as smart cards. Even for short-lived signature use cases, the infrastructure often relies on long-lived root keys, which can be difficult to update or replace on in-field devices.



Figure 1: Mosca Model

These challenges are illustrated nicely by the so-called Mosca model discussed in [Threat-Report]. In Figure 1, "x" denotes the time that systems and data need to remain secure, "y" the number of years to fully migrate to a PQC infrastructure, and "z" the time until a CRQC that can break current cryptography is available. The model assumes either that encrypted data can be intercepted and stored before the migration is completed in "y" years, or that signatures will still be relied upon for "x" years after their creation. This data remains vulnerable for the complete "x" years of their lifetime; thus, the sum "x+y" gives us an estimate of the full timeframe that data remains insecure. The model essentially asks how one is preparing IT systems during those "y" years (in other words, how one can minimize those "y" years) to minimize the transition phase to a PQC infrastructure and hence minimize the risks of data being exposed in the future.

Finally, other factors that could accelerate the introduction of a CRQC should not be underestimated, for example, faster-than-expected advances in quantum computing and more efficient versions of Shor's algorithm requiring fewer qubits. Innovation often comes in waves, so it is to the industry's benefit to remain vigilant and prepare as early as possible. Also, bear in mind that while the industry tracks advances from public research institutions such as universities and companies that publish their results, there is also a great deal of large-budget quantum research being conducted privately by various national interests. Therefore, the true state of quantum computer advancement is likely several years ahead of the publicly available research at the date this document is published.

Organizations should also carefully and honestly consider what their migration timeline "y" actually is. If you only think of the time between receiving a patch from your technology vendor and rolling that patch out, then "y" might seem as short as a few weeks. However, this represents the minority of migration cases; more often, a PQC migration will involve at least some amount of hardware replacement. For example, performance-sensitive applications will need CPUs with PQC hardware acceleration. Security-sensitive applications will need PQC TPMs, Trusted Execution Environments (TEEs), secure enclaves, and other cryptographic co-processors. Smart card applications will require replacement of the cards and readers. The readers can come in many form factors: tap-for-entry door and turnstile readers, PIN pad machines, laptops with built-in smart card readers, and many others.

Included in "y" is not only the deployment time but also the preparation time: integration, testing, auditing, and recertification of cryptographic environments. Also consider upstream effects that contribute to "y", including lead times for vendors to produce PQC-ready products,

which may itself include auditing and certification delays, time for regulating bodies to adopt PQC policies, time for auditors to become familiar with the new requirements, etc. If you measure the full migration time "y" from when your vendors begin implementing PQC functionality to when you switch off your last non-PQC-capable device, then "y" can be quite long, likely measured in years for even most moderately sized organizations. This long tail should not discourage early action.

Organizations responsible for protecting long-lived sensitive data or operating critical infrastructure will need to begin transitioning immediately, particularly in scenarios where data is vulnerable to HNDL attacks. Post-quantum and traditional (PQ/T) [Section 13](#) or PQ key exchange is relatively self-contained, typically requiring changes only to the cryptographic library (e.g., OpenSSL). In contrast, migrating to post-quantum or PQ/T digital signatures involves broader ecosystem changes, including updates to certificates, certificate authorities (CAs), Certificate Management Protocols, HSMs, and trust anchors. Starting early with hybrid key exchange deployments allows organizations to gain operational experience, while prototyping and planning for PQ/T or PQ digital signature integration helps identify ecosystem-wide impacts early. This phased approach reduces long-term migration risks and ensures readiness for more complex updates.

8. PQC Categories

The post-quantum cryptographic schemes standardized by NIST can be categorized into three main groups: lattice-based, hash-based, and code-based. Other approaches, such as isogeny-based, multivariate-based, and MPC-in-the-Head-based cryptography, are also being explored in research and standardization efforts. In addition, NIST issued a call for additional digital signature proposals to expand the set of post-quantum signatures under evaluation [[AddSig](#)].

8.1. Lattice-Based Public Key Cryptography

Lattice-based public key cryptography leverages the simple construction of lattices (i.e., a regular collection of points in a Euclidean space that are evenly spaced) to create "trapdoor" problems. These problems are efficient to compute if you possess the secret information but challenging to compute otherwise. Examples of such problems include the shortest vector, closest vector, short integer solution, learning with errors, module learning with errors, and learning with rounding problems. All of these problems feature strong proofs for worst-to-average case reduction, effectively relating the hardness of the average case to the worst case.

Lattice-based public keys and signatures are larger than those of classical schemes such as RSA or ECC, but typically by less than an order of magnitude for public keys (about 6-10x) and by roughly one to two orders of magnitude for signatures (about 10-100x) rather than by several orders of magnitude, making them the best available candidates for general-purpose use, such as replacing the use of RSA in PKIX certificates.

Examples of this class of algorithms include ML-KEM, FN-DSA, ML-DSA, and FrodoKEM.

It is noteworthy that lattice-based encryption schemes require a rounding step during decryption, which has a non-zero probability of "rounding the wrong way" and leading to a decryption failure, meaning that valid encryptions are decrypted incorrectly. However, the parameters of NIST PQC candidates are carefully chosen so that the probability of such a failure is cryptographically negligible, far lower than the probability of random transmission errors and implementation bugs. In practical terms, these rare decryption failures can be treated the same way as any fatal transport error: Both sides simply perform a fresh KEM operation, generating a new ciphertext and shared secret.

In cryptanalysis, an oracle refers to a system that an attacker can query to learn whether decryption succeeded or failed. If such an oracle exists, an attacker could significantly reduce the security of lattice-based schemes that have a relatively high failure rate. However, for most of the NIST PQC proposals, the number of required oracle queries to force a decryption failure is above practical limits, as shown in [LattFail1]. More recent works have improved upon the results in [LattFail1], showing that the cost of searching for additional failing ciphertexts after one or more have already been found can be sped up dramatically [LattFail2]. Nevertheless, at the time this document is published, the PQC candidates by NIST are considered secure under these attacks, and constant monitoring as cryptanalysis research is ongoing.

8.2. Hash-Based Public Key Cryptography

Hash-based Public Key Cryptography (PKC) has been around since the 1970s, when it was developed by Lamport and Merkle. It is used to create digital signature algorithms, and its security is based on the security of the underlying cryptographic hash function. Many variants of hash-based signatures (HBSs) have been developed since the 1970s, including the recent XMSS [RFC8391], HSS/LMS [RFC8554], or BPQS [BPQS] schemes. Unlike many other digital signature techniques, most hash-based signature schemes are stateful, which means that signing necessitates the update and careful tracking of the state of the secret key. Producing multiple signatures using the same secret key state results in loss of security and may ultimately enable signature forgery attacks against that key.

Stateful hash-based signatures with long service lifetimes require additional operational complexity compared to other signature types. For example, consider a 20-year root key; there is an expectation that 20 years is longer than the expected lifetime of the hardware that key is stored on, so the key will need to be migrated to new hardware at some point. Disaster-recovery scenarios where the primary node fails without warning can be similarly tricky. This requires careful operational and compliance consideration to ensure that no private key state can be reused across the migration or disaster recovery event. One approach for avoiding these issues is to only use stateful HBSs for short-term use cases that do not require horizontal scaling, for example, signing a batch of firmware images and then retiring the signing key.

The SLH-DSA algorithm, which was standardized by NIST, leverages the HORST (Hash to Obtain Random Subset with Trees) technique and remains the only standardized hash based signature scheme that is stateless, thus avoiding the complexities associated with state management. SLH-DSA is an advancement on SPHINCS that reduces the signature sizes in SPHINCS and makes it more compact.

8.3. Code-Based Public Key Cryptography

This area of cryptography started in the 1970s and 1980s and was based on the seminal work of McEliece and Niederreiter, which focuses on the study of cryptosystems based on error-correcting codes. Some popular error-correcting codes include Goppa codes (used in McEliece cryptosystems), encoding and decoding syndrome codes used in HQC, or quasi-cyclic moderate density parity check (QC-MDPC) codes.

Examples include all the unbroken NIST Round 4 finalists: Classic McEliece, HQC (selected by NIST for standardization), and Bit Flipping Key Encapsulation (BIKE) [BIKE].

9. KEMs

A Key Encapsulation Mechanism (KEM) is a cryptographic technique used for securely exchanging symmetric key material between two parties over an insecure channel. It is commonly used in hybrid encryption schemes where a combination of asymmetric (public key) and symmetric encryption is employed. The encapsulation operation of a KEM results in a fixed-length symmetric key that can be used with a symmetric algorithm, typically a block cipher, in one of two different ways:

- To derive a data encryption key (DEK) to encrypt the data
- To derive a key encryption key (KEK) used to wrap a DEK

These techniques are often referred to as the Hybrid Public Key Encryption (HPKE) [RFC9180] mechanism.

The term "encapsulation" is chosen intentionally to indicate that KEM algorithms behave differently at the API level from the key agreement or key encipherment and key transport mechanisms that are in use today. Key agreement schemes imply that both parties contribute a public-private key pair to the exchange, while key encipherment and key transport schemes imply that the symmetric key material is chosen by one party and "encrypted" or "wrapped" for the other party. KEMs, on the other hand, behave according to the following API primitives [PQCAPI]:

- `def kemKeyGen() -> (pk, sk)`
- `def kemEncaps(pk) -> (ss, ct)`
- `def kemDecaps(ct, sk) -> ss`

where `pk` is the public key, `sk` is the secret key, `ct` is the ciphertext representing an encapsulated key, and `ss` is the shared secret. The following figure illustrates a sample flow of a KEM-based key exchange:

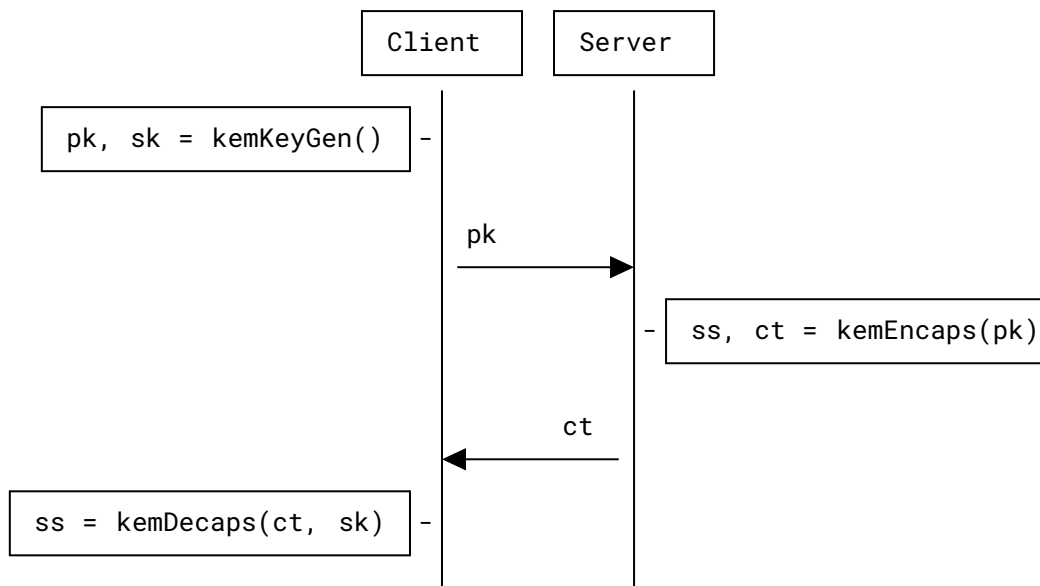


Figure 2: KEM-Based Key Exchange

9.1. Authenticated Key Exchange

Authenticated Key Exchange (AKE) with KEMs where both parties contribute a KEM public key to the overall session key is interactive as described in [Section 9.4](#) of [\[RFC9528\]](#). However, a single-sided KEM, such as when one peer has a KEM key in a certificate and the other peer wants to encrypt for it (as in S/MIME or OpenPGP email), can be achieved using non-interactive HPKE [\[RFC9180\]](#). The following figure illustrates the DH Key exchange:

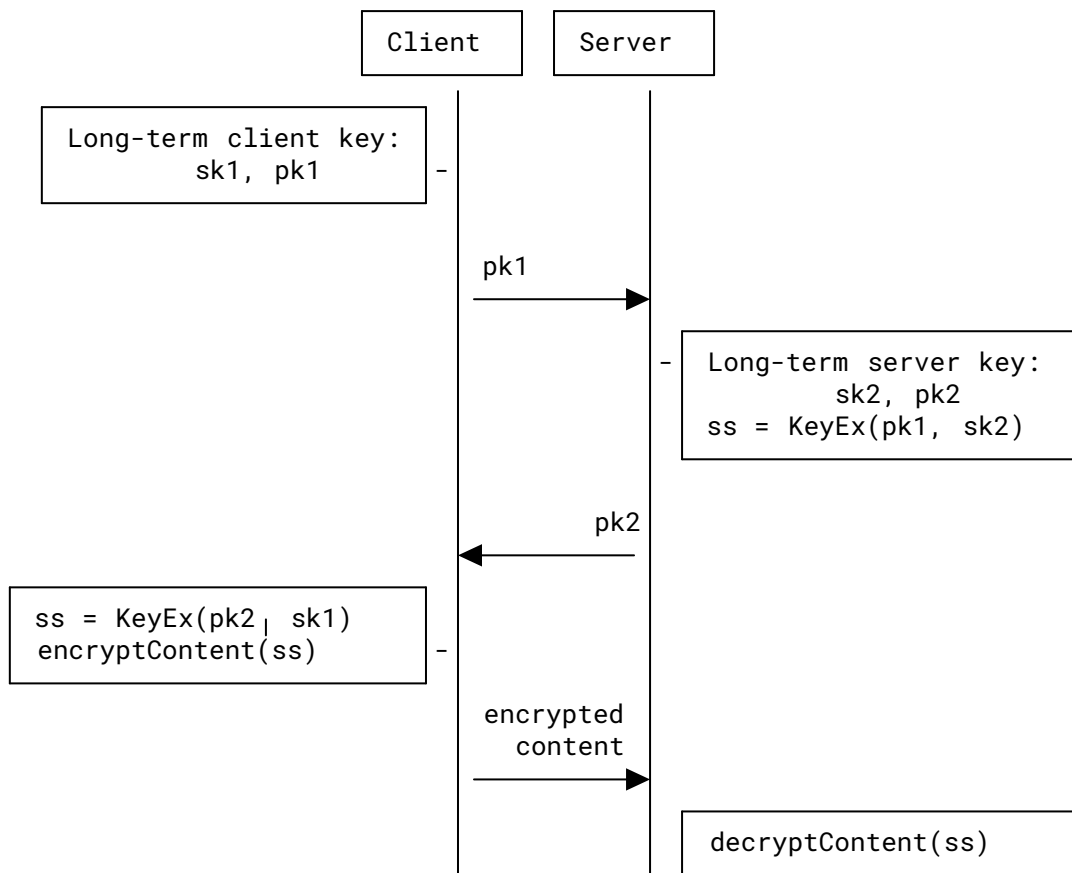


Figure 3: DH-Based AKE

In the sample flow above, it is important to note that the shared secret ss is derived using key material from both the client and the server, which classifies it as an AKE. There is another property of a key exchange, called Non-Interactive Key Exchange (NIKE), that refers to whether the sender can compute the shared secret ss and encrypt content without requiring active interaction (an exchange of network messages) with the recipient. Figure 3 shows a DH key exchange, which is an AKE since both parties are using long-term keys that can have established trust (for example, via certificates), but it is not a NIKE since the client needs to wait for the network interaction to receive the receiver's public key $pk2$ before it can compute the shared secret ss and begin content encryption. However, a DH key exchange can be an AKE and a NIKE at the same time if the receiver's public key is known to the sender in advance (see Figure 4), and many Internet protocols rely on this property of DH-based key exchanges.

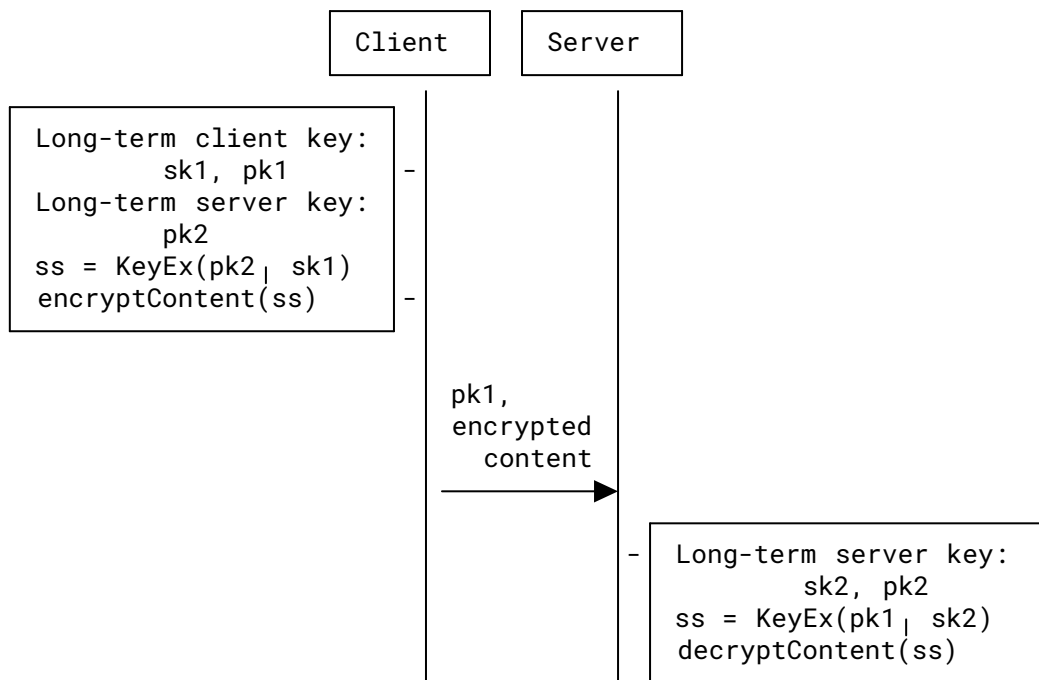


Figure 4: Simultaneous DH-Based AKE and NIKÉ

The complication with KEMs is that a KEM `Encaps()` is non-deterministic; it involves randomness chosen by the sender of that message. Therefore, in order to perform an AKE, the client must wait for the server to generate the needed randomness and perform `Encaps()` against the client key, which necessarily requires a network round-trip. Therefore, a KEM-based protocol can either be an AKE or a NIKÉ, but it cannot be both at the same time. Consequently, certain Internet protocols will necessitate a redesign to accommodate this distinction, either by introducing extra network round trips or by making trade-offs in security properties.

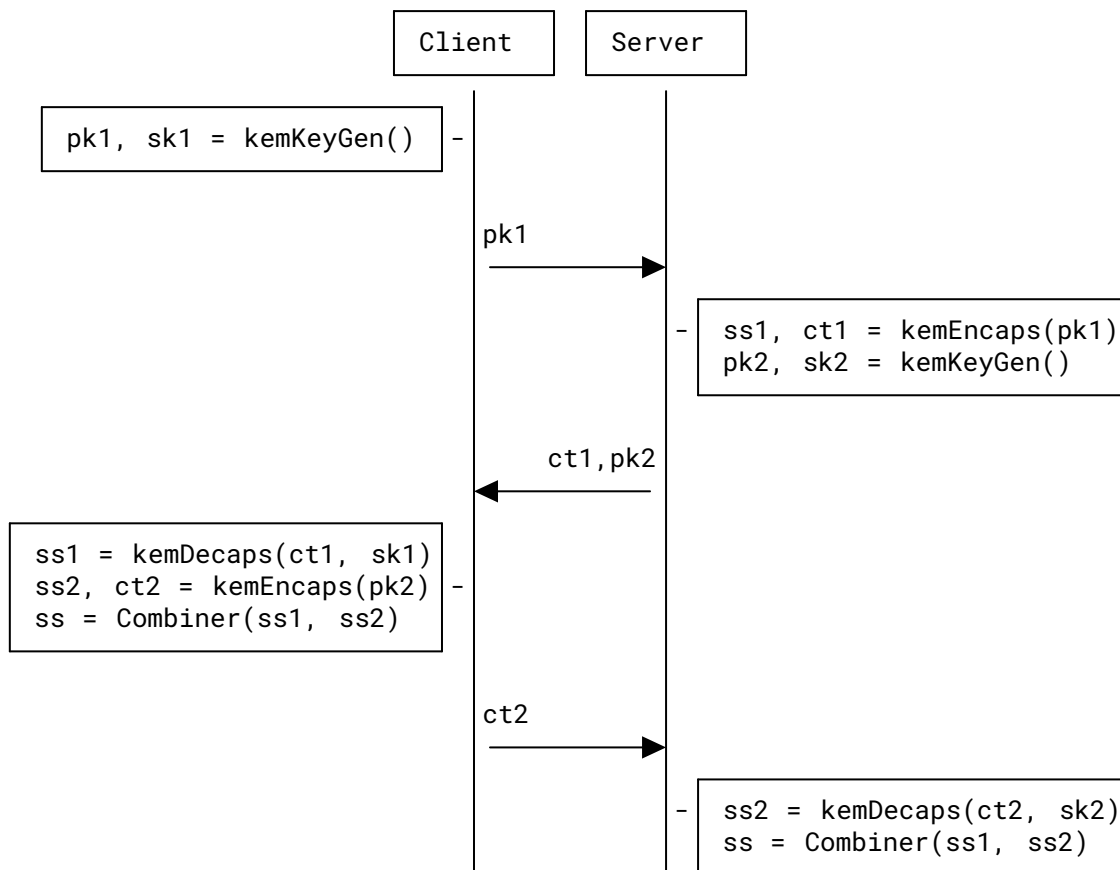


Figure 5: KEM-Based AKE

In the figure above, $\text{Combiner}(ss1, ss2)$, often referred to as a KEM combiner, is a cryptographic construction that takes in two shared secrets and returns a single combined shared secret. The simplest combiner is concatenation $ss1 \parallel ss2$, but combiners can vary in complexity depending on the cryptographic properties required. For example, if the combination should preserve IND-CCA2 (see [Section 9.2.1](#)) of either input, even if the other is chosen maliciously, then a more complex construct is required. Another consideration for combiner design is the so-called "binding properties" introduced in [\[KEEPINGUP\]](#), which may require the ciphertexts and recipient public keys to be included in the combiner. KEM combiner security analysis becomes more complicated in hybrid settings where the two KEMs represent different algorithms, for example, where one is ML-KEM and the other is ECDH. For a more thorough discussion of KEM combiners, see [\[KEEPINGUP\]](#), [\[KEM-COMBINER\]](#), and [\[PQ-KEM\]](#).

9.2. Security Properties of KEMs

The security properties described in this section (IND-CCA2 and binding) are not an exhaustive list of all possible KEM security considerations. They were selected because they are fundamental to evaluating KEM suitability in protocol design and are commonly discussed in current PQC work.

9.2.1. IND-CCA2

IND-CCA2 (INDistinguishability under adaptive Chosen-Ciphertext Attack) is an advanced security notion for encryption schemes. It ensures the confidentiality of the plaintext and resistance against chosen-ciphertext attacks. An appropriate definition of IND-CCA2 security for KEMs can be found in [CS01] and [BHK09]. ML-KEM [ML-KEM] and Classic McEliece provide IND-CCA2 security.

Understanding IND-CCA2 security is essential for individuals involved in designing or implementing cryptographic systems and protocols in order to evaluate the strength of the algorithm, assess its suitability for specific use cases, and ensure that data confidentiality and security requirements are met. Understanding IND-CCA2 security is generally not necessary for developers migrating to using an IETF-vetted KEM within a given protocol or flow. IND-CCA2 is a widely accepted security notion for public key encryption mechanisms, making it suitable for a broad range of applications. When an IETF specification defines a new KEM, its security considerations should fully describe the relevant cryptographic properties, including IND-CCA2.

9.2.2. Binding

KEMs also have an orthogonal set of properties to consider when designing protocols around them: binding [KEEPINGUP]. This can be "ciphertext binding", "public key binding", "context binding", or any other property that is important to not be substituted between KEM invocations. In general, a KEM is considered to bind a certain value if substitution of that value by an attacker will necessarily result in a different shared secret being derived. As an example, if an attacker can construct two different ciphertexts that will decapsulate to the same shared secret, can construct a ciphertext that will decapsulate to the same shared secret under two different public keys, or can substitute whole KEM exchanges from one session into another, then the construction is not ciphertext binding, public key binding, or context binding, respectively. Similarly, protocol designers may wish to bind protocol state information such as a transaction ID or nonce so that attempts to replay ciphertexts from one session inside a different session will be blocked at the cryptographic level because the server derives a different shared secret and is thus unable to decrypt the content.

The solution to binding is generally achieved at the protocol design level: It is recommended to avoid using the KEM output shared secret directly without integrating it into an appropriate protocol. While KEM algorithms provide key secrecy, they do not inherently ensure source authenticity, protect against replay attacks, or guarantee freshness. These security properties should be addressed by incorporating the KEM into a protocol that has been analyzed for such protections. Even though modern KEMs such as ML-KEM produce full-entropy shared secrets, it

is still advisable for binding reasons to pass the shared secret through a key derivation function (KDF) and also include all values that you wish to bind; finally, you will have a shared secret that is safe to use at the protocol level.

9.3. HPKE

Modern cryptography has long used the notion of "hybrid encryption" where an asymmetric algorithm is used to establish a key and then a symmetric algorithm is used for bulk content encryption. The previous sections explained important security properties of KEMs, such as IND-CCA2 security and binding, and emphasized that these properties must be supported by proper protocol design. One widely deployed scheme that achieves this is Hybrid Public Key Encryption (HPKE) [RFC9180].

HPKE [RFC9180] works with a combination of KEMs, KDFs, and Authenticated Encryption with Associated Data (AEAD) schemes. HPKE includes three authenticated variants, including one that authenticates possession of a pre-shared key and two optional ones that authenticate possession of a KEM private key. HPKE can be extended to support hybrid post-quantum KEM [PQ-HPKE]. ML-KEM does not support the static-ephemeral key exchange that allows HPKE that is based on DH-based KEMs and its optional authenticated modes as discussed in Section 1.5 of [X-WING].

10. PQC Signatures

Any digital signature scheme that provides a construction defining security under a post-quantum setting falls under this category of PQC signatures.

10.1. Security Properties of PQC Signatures

10.1.1. EUF-CMA and SUF-CMA

EUF-CMA (existential unforgeability under chosen message attack) [GMR88] is a security notion for digital signature schemes. It guarantees that an adversary, even with access to a signing oracle, cannot forge a valid signature for an arbitrary message. EUF-CMA provides strong protection against forgery attacks, ensuring the integrity and authenticity of digital signatures by preventing unauthorized modifications or fraudulent signatures. ML-DSA, FN-DSA, and SLH-DSA provide EUF-CMA security.

SUF-CMA (strong unforgeability under chosen message attack) builds upon EUF-CMA by requiring that an adversary cannot produce a different valid signature for a message that has already been signed by the signing oracle. Like EUF-CMA, SUF-CMA provides robust assurances for digital signature schemes, further enhancing their security posture. ML-DSA, FN-DSA, and SLH-DSA also achieve SUF-CMA security.

Understanding EUF-CMA and SUF-CMA security is essential for designing or implementing cryptographic systems in order to ensure the security, reliability, and robustness of digital signature schemes. These notions allow for informed decision making, vulnerability analysis, compliance with standards, and designing systems that provide strong protection against

forgery attacks. For developers migrating to an IETF-vetted PQC signature scheme within a given protocol or flow, a deep understanding of EUF-CMA and SUF-CMA security may not be necessary, as the schemes vetted by IETF adhere to these stringent security standards.

EUF-CMA and SUF-CMA are considered strong security benchmarks for public key signature algorithms, making them suitable for most applications. Authors of IETF specifications should include all security concerns in the "Security Considerations" section of the relevant RFC and should not assume that implementers are experts in cryptographic theory.

10.2. Details of FN-DSA, ML-DSA, and SLH-DSA

ML-DSA [ML-DSA] is a digital signature algorithm based on the hardness of lattice problems over module lattices (i.e., the Module Learning with Errors (MLWE) problem). The design of the algorithm is based on the "Fiat-Shamir with Aborts" [Lyu09] framework introduced by Lyubashevsky that leverages rejection sampling to render lattice-based Fiat-Shamir (FS) schemes compact and secure. ML-DSA uses uniformly distributed random number sampling over small integers to compute coefficients in error vectors, which makes the scheme easier to implement compared to FN-DSA [FN-DSA], which uses Gaussian-distributed numbers, necessitating the need to use floating-point arithmetic during signature generation.

ML-DSA offers both deterministic and randomized signing and is instantiated with three parameter sets providing different security levels. Security properties of ML-DSA are discussed in Section 9 of [RFC9881].

FN-DSA [FN-DSA] is based on the GPV hash-and-sign lattice-based signature framework introduced by Gentry, Peikert, and Vaikuntanathan, which is a framework that requires a certain class of lattices and a trapdoor sampler technique.

The main design principle of FN-DSA is compactness, i.e., it was designed in a way that achieves minimal total memory bandwidth requirement (the sum of the signature size plus the public key size). This is possible due to the compactness of NTRU lattices. FN-DSA also offers very efficient signing and verification procedures. The main potential downsides of FN-DSA refer to the non-triviality of its algorithms and the need for floating-point arithmetic support in order to support Gaussian-distributed random number sampling where the other lattice schemes use the less efficient but easier to support uniformly distributed random number sampling.

Implementers of FN-DSA need to be aware that FN-DSA signing is highly susceptible to side-channel attacks unless constant-time 64-bit floating-point operations are used. This requirement is extremely platform-dependent, as noted in NIST's report [NIST].

The performance characteristics of ML-DSA and FN-DSA may differ based on the specific implementation and hardware platform. Generally, ML-DSA is known for its relatively fast signature generation, while FN-DSA can provide more efficient signature verification. The choice may depend on whether the application requires more frequent signature generation or signature verification (see [LIBOQS]). For further clarity on the sizes and security levels, please refer to the tables in Sections 11 and 12.

SLH-DSA [SLH-DSA] utilizes the concept of stateless hash-based signatures, where each signature is unique and unrelated to any previous signature (as discussed in Section 8.2). This property eliminates the need for maintaining state information during the signing process. SLH-DSA was designed to sign up to 2^{64} messages under a given key pair, and it offers three security levels. The parameters for each of the security levels were chosen to provide 128 bits of security, 192 bits of security, and 256 bits of security. SLH-DSA offers smaller public key sizes, larger signature sizes, slower signature generation, and slower verification when compared to ML-DSA and FN-DSA. SLH-DSA does not introduce a new hardness assumption beyond those inherent to the underlying hash functions. It builds upon established foundations in cryptography, making it a reliable and robust digital signature scheme for a post-quantum world.

All of these algorithms (ML-DSA, FN-DSA, and SLH-DSA) include two signature modes: pure mode, where the entire content is signed directly, and pre-hash mode, where a digest of the content is signed.

10.3. Details of XMSS and LMS

The eXtended Merkle Signature Scheme (XMSS) [RFC8391] and Hierarchical Signature Scheme (HSS) / Leighton-Micali Signature (LMS) [RFC8554] are stateful hash-based signature schemes, where the secret key state changes over time. In both schemes, reusing a secret key state compromises cryptographic security guarantees.

XMSS and LMS can be used for signing a potentially large but fixed number of messages, and the number of signing operations depends upon the size of the tree. XMSS and LMS provide cryptographic digital signatures without relying on the conjectured hardness of mathematical problems, instead leveraging the properties of cryptographic hash functions. Multi-tree XMSS and LMS (i.e., XMSS-MT and HSS, respectively) use a hyper-tree-based hierarchical approach with a Merkle tree at each level of the hierarchy. [RFC8391] describes both single-tree and multi-tree variants of XMSS, while [RFC8554] describes the Leighton-Micali One-Time Signature (LM-OTS) system as well as the LMS and HSS N-time signature systems. Comparison of XMSS and LMS is discussed in Section 10 of [RFC8554].

The number of tree layers in multi-tree XMSS and HSS provides a trade-off between signature size on the one side and key generation and signing speed on the other side. Increasing the number of layers reduces key generation time exponentially and signing time linearly at the cost of increasing the signature size linearly. HSS allows for customization of each subtree, whereas XMSS-MT does not, electing instead to use the same structure for each subtree.

Due to the complexities described above, XMSS and LMS are not suitable replacements for traditional signature schemes like RSA or ECDSA. Applications that expect a long lifetime of a signature, like firmware update or secure boot, are typical use cases where those schemes can be successfully applied.

10.3.1. LMS Key and Signature Sizes

The LMS scheme is characterized by four distinct parameter sets: the underlying hash function (SHA2-256 or SHAKE-256), the length of the digest (24 or 32 bytes), the LMS tree height parameter that controls a maximal number of signatures that the private key can produce, and

the width of the Winternitz coefficients (see [RFC8554], Section 4.1) that can be used to trade-off signing time for signature size. Parameters can be mixed, providing 80 possible parameterizations of the scheme.

The public (PK) and private (SK) key size depends on the length of the digest (M). The signature size depends on the digest, the Winternitz parameter (W), the LMS tree height (H), and the length of the digest. The table below provides key and signature sizes for parameterization with the digest size M=32 of the scheme.

PK	SK	W	H=5	H=10	H=15	H=20	H=25
56	52	1	8684	8844	9004	9164	9324
56	52	2	4460	4620	4780	4940	5100
56	52	4	2348	2508	2668	2828	2988
56	52	8	1292	1452	1612	1772	1932

Table 1

10.4. Hash-then-Sign

Within the hash-then-sign paradigm, the message is hashed before signing it. By pre-hashing, the onus of resistance to existential forgeries becomes heavily reliant on the collision-resistance of the hash function in use. The hash-then-sign paradigm has the ability to improve application performance by reducing the size of signed messages that need to be transmitted between application and cryptographic module and making the signature size predictable and manageable. As a corollary, hashing remains mandatory even for short messages and assigns a further computational requirement onto the verifier. This makes the performance of hash-then-sign schemes more consistent, but not necessarily more efficient.

Using a hash function to produce a fixed-size digest of a message ensures that the signature is compatible with a wide range of systems and protocols, regardless of the specific message size or format. Crucially for hardware security modules, hash-then-sign also significantly reduces the amount of data that needs to be transmitted and processed by a Hardware Security Module (HSM). Consider scenarios such as a networked HSM located in a different data center from the calling application or a smart card connected over a USB interface. In these cases, streaming a message that is megabytes or gigabytes long can result in notable network latency, on-device signing delays, or even depletion of available on-device memory.

Note that the vast majority of Internet protocols that sign large messages already perform some form of content hashing at the protocol level, so this tends to be more of a concern with proprietary cryptographic protocols and protocols from non-IETF standards bodies. Protocols like TLS 1.3 and DNSSEC use the hash-then-sign paradigm. In TLS 1.3 [RFC8446] CertificateVerify messages, the content that is covered under the signature includes the transcript hash output

(Section 4.4.1 of [RFC8446]) while DNSSEC [RFC4034] uses it to provide origin authentication and integrity assurance services for DNS data. Similarly, the Cryptographic Message Syntax (CMS) [RFC5652] includes a mandatory message digest step before invoking the signature algorithm.

In the case of ML-DSA, it internally incorporates the necessary hash operations as part of its signing algorithm. ML-DSA directly takes the original message, applies a hash function internally, and then uses the resulting hash value for the signature generation process. In the case of SLH-DSA, it internally performs randomized message compression using a keyed hash function that can process arbitrary length messages. In the case of FN-DSA, the SHAKE-256 hash function is used as part of the signature process to derive a digest of the message being signed.

Therefore, ML-DSA, FN-DSA, and SLH-DSA offer enhanced security over the traditional hash-then-sign paradigm because, by incorporating dynamic key material into the message digest, a pre-computed hash collision on the message to be signed no longer yields a signature forgery. Applications requiring the performance and bandwidth benefits of hash-then-sign may still pre-hash at the protocol level prior to invoking ML-DSA, FN-DSA, or SLH-DSA, but protocol designers should be aware that doing so reintroduces the weakness that hash collisions directly yield signature forgeries. Signing the full un-digested message is recommended where applications can tolerate it.

11. NIST Recommendations for Security and Performance Trade-offs

This information is a reprint of information provided in the NIST PQC project [NIST] as of the time this document is published. Table 2 denotes the five security levels provided by NIST for PQC algorithms. Neither NIST nor the IETF makes any specific recommendations about which security level to use. In general, protocols will include algorithm choices at multiple levels so that users can choose the level appropriate to their policies and data classification, similar to how organizations today choose which size of RSA key to use. The security levels are defined as requiring computational resources comparable to or greater than an attack on AES (128, 192, and 256) and SHA2/SHA3 algorithms, i.e., exhaustive key recovery for AES and optimal collision search for SHA2/SHA3.

PQ Security Level	AES/SHA(2/3) hardness	PQC Algorithm
1	AES-128 (exhaustive key recovery)	ML-KEM-512, FN-DSA-512, SLH-DSA-SHA2/SHAKE-128f/s
2	SHA-256/SHA3-256 (collision search)	ML-DSA-44
3	AES-192 (exhaustive key recovery)	ML-KEM-768, ML-DSA-65, SLH-DSA-SHA2/SHAKE-192f/s

PQ Security Level	AES/SHA(2/3) hardness	PQC Algorithm
4	SHA-384/SHA3-384 (collision search)	No algorithm tested at this level
5	AES-256 (exhaustive key recovery)	ML-KEM-1024, FN-DSA-1024, ML-DSA-87, SLH-DSA-SHA2/SHAKE-256f/s

Table 2

The SLH-DSA-x-yf/s "f/s" in the above table denotes whether SLH-DSA is using SHAKE or SHA-2 as an underlying hash function "x" and whether it is the fast (f) or small (s) version for "y" bit AES security level. Refer to [RFC9814] for further details on SLH-DSA algorithms.

The following table compares the signature sizes for different SLH-DSA algorithm categories at equivalent security levels using the "simple" version. The categories include "f" for fast signature generation and "s" for smaller signature size and faster verification, although with slower signature generation. Both SHA-256 and SHAKE-256 parameterizations produce the same signature sizes and are therefore included together in the table.

PQ Security Level	Algorithm	Public key size (in bytes)	Private key size (in bytes)	Signature size (in bytes)
1	SLH-DSA-{SHA2,SHAKE}-128f	32	64	17088
1	SLH-DSA-{SHA2,SHAKE}-128s	32	64	7856
3	SLH-DSA-{SHA2,SHAKE}-192f	48	96	35664
3	SLH-DSA-{SHA2,SHAKE}-192s	48	96	16224
5	SLH-DSA-{SHA2,SHAKE}-256f	64	128	49856
5	SLH-DSA-{SHA2,SHAKE}-256s	64	128	29792

Table 3

The following table illustrates the impact of performance on different security levels in terms of private key sizes, public key sizes, and ciphertext/signature sizes.

PQ Security Level	Algorithm	Public key size (in bytes)	Private key size (in bytes)	Ciphertext/signature size (in bytes)
1	ML-KEM-512	800	1632	768
1	FN-DSA-512	897	1281	666
2	ML-DSA-44	1312	2560	2420
3	ML-KEM-768	1184	2400	1088
3	ML-DSA-65	1952	4032	3309
5	FN-DSA-1024	1793	2305	1280
5	ML-KEM-1024	1568	3168	1588
5	ML-DSA-87	2592	4896	4627

Table 4

12. Comparing PQC KEMs/Signatures and Traditional KEMs/Signatures

This section provides two tables for comparison of different KEMs and signatures, respectively, in the traditional and post-quantum scenarios. These tables focus on the secret key sizes, public key sizes, and ciphertext/signature sizes for the PQC algorithms and their traditional counterparts of similar security levels.

The first table compares traditional and PQC KEMs in terms of security, public and private key sizes, and ciphertext sizes.

PQ Security Level	Algorithm	Public key size (in bytes)	Private key size (in bytes)	Ciphertext size (in bytes)
Traditional	P256_HKDF_SHA-256	65	32	65
Traditional	P521_HKDF_SHA-512	133	66	133
Traditional	X25519_HKDF_SHA-256	32	32	32
1	ML-KEM-512	800	1632	768

PQ Security Level	Algorithm	Public key size (in bytes)	Private key size (in bytes)	Ciphertext size (in bytes)
3	ML-KEM-768	1184	2400	1088
5	ML-KEM-1024	1568	3168	1568

Table 5

The next table compares traditional and PQC signature schemes in terms of security, public, private key sizes, and signature sizes.

PQ Security Level	Algorithm	Public key size (in bytes)	Private key size (in bytes)	Signature size (in bytes)
Traditional	RSA2048	256	256	256
Traditional	ECDSA-P256	64	32	64
1	FN-DSA-512	897	1281	666
2	ML-DSA-44	1312	2560	2420
3	ML-DSA-65	1952	4032	3309
5	FN-DSA-1024	1793	2305	1280
5	ML-DSA-87	2592	4896	4627

Table 6

As is clear from the above table, PQC KEMs and signature schemes typically have significantly larger keys and ciphertexts/signatures than their traditional counterparts. These increased key and signature sizes could introduce problems in protocols. As an example, the Internet Key Exchange Protocol Version 2 (IKEv2) uses UDP as the transport protocol for its messages. One challenge with integrating a PQC KEM into IKEv2 is that IKE fragmentation cannot be utilized in the initial IKE_SA_INIT exchange. To address this issue, [RFC9242] introduces a solution by defining a new exchange called the "Intermediate Exchange", which can be fragmented using the IKE fragmentation mechanism. [RFC9370] then uses this Intermediate Exchange to carry out the PQC key exchange after the initial IKEv2 exchange and before the IKE_AUTH exchange. Another example from Section 6.3.3 of [SP-1800-38C] shows that increased key and signature sizes cause protocol key exchange messages to span more network packets, which results in a higher total loss probability per packet. In lossy network conditions, this may increase the latency of the key exchange.

13. Post-Quantum and Traditional (PQ/T) Hybrid Schemes

The migration to PQC is unique in the history of modern digital cryptography in that neither the traditional algorithms nor the post-quantum algorithms are fully trusted to protect data for the required lifetimes. The traditional algorithms, such as RSA and ECDH, will fall to quantum cryptanalysis, while the post-quantum algorithms face uncertainty about the underlying mathematics, compliance issues, unknown vulnerabilities, and hardware and software implementations that have not had sufficient maturing time to rule out traditional cryptanalytic attacks and implementation bugs.

During the transition from traditional to post-quantum algorithms, there may be a desire or a requirement for protocols that use both algorithm types. [RFC9794] defines the terminology for PQ/T hybrid schemes.

13.1. PQ/T Hybrid Confidentiality

The PQ/T Hybrid Confidentiality property can be used to mitigate both "harvest now, decrypt now" and HNDL attacks described in Section 7. If the PQ portion were to have a flaw, the traditional (T) algorithm, which is secure against today's attackers, prevents immediate decryption ("harvest now, decrypt now"). If the T algorithm is broken in the future by CRQCs, the PQ portion, assuming it remains secure, prevents later decryption (i.e., HNDL). A hybrid construction therefore provides confidentiality as long as at least one component remains secure. Two types of hybrid key agreement schemes are discussed below.

Concatenated hybrid key agreement scheme: The final shared secret that will be used as an input of the key derivation function is the result of the concatenation of the secrets established with each key agreement scheme. For example, in [TLS-HYB-KEY-EXCH], the client uses the TLS supported groups extension to advertise support for a PQ/T hybrid scheme, and the server can select this group if it supports the scheme. The hybrid-aware client and server establish a hybrid secret by concatenating the two shared secrets, which is used as the shared secret in the existing TLS 1.3 key schedule.

Cascaded hybrid key agreement scheme: The final shared secret is computed by applying as many iterations of the key derivation function as the number of key agreement schemes composing the hybrid key agreement scheme. For example, [RFC9370] extends IKEv2 to allow one or more PQC algorithms in addition to the traditional algorithm to derive the final IKE Security Association (SA) keys using the cascade method as explained in Section 2.2.2 of [RFC9370].

Various instantiations of these two types of hybrid key agreement schemes have been explored. One must be careful when selecting which hybrid scheme to use. The chosen scheme for protocols like TLS 1.3 [TLS-HYB-KEY-EXCH] has IND-CCA2 robustness. That is, IND-CCA2 security is guaranteed for the scheme as long as at least one of the component algorithms is IND-CCA2 secure.

13.2. PQ/T Hybrid Authentication

The PQ/T hybrid authentication property provides resilience against catastrophic breaks or unforeseen vulnerabilities in PQC algorithms, allowing systems additional time to stabilize before migrating fully to pure PQ deployments.

This property ensures authentication using a PQ/T hybrid scheme as long as at least one component algorithm remains secure. For example, a PQ/T hybrid certificate [ML-DSA-X.509] can be employed to facilitate a PQ/T hybrid authentication protocol. However, a PQ/T hybrid authentication protocol does not need to use a PQ/T hybrid certificate; separate certificates could be used for individual component algorithms [RFC9763]. When separate certificates are used, it may be possible for attackers to take them apart or put them together in unexpected ways, including enabling cross-protocol attacks. The exact risks this presents are highly dependent on the protocol and use case, so a full security analysis is needed. Best practices for ensuring that pairs of certificates are only used as intended are discussed in more detail in Sections 13.3.2 and 13.3.3 of this document.

The frequency and duration of system upgrades and the time when CRQCs will become widely available need to be weighed to determine whether and when to support the PQ/T Hybrid Authentication property.

13.3. Hybrid Cryptographic Algorithm Combinations: Considerations and Approaches

13.3.1. Hybrid Cryptographic Combinations

It is also possible to use more than two algorithms together in a hybrid scheme, with various methods for combining them. For post-quantum transition purposes, the combination of a post-quantum algorithm with a traditional algorithm is the most straightforward and recommended. The use of multiple post-quantum algorithms with different mathematical bases has also been considered. Combining algorithms in a way that requires both to be used together ensures stronger security, while combinations that do not require both will sacrifice security but offer other benefits like backwards compatibility and crypto agility. Including a traditional key alongside a post-quantum key often has minimal bandwidth impact.

13.3.2. Composite Keys in Hybrid Schemes

When combining keys in an "and" mode, it may make more sense to consider them to be a single composite key instead of two keys. This generally requires fewer changes to various components of PKI ecosystems, many of which are not prepared to deal with two keys or dual signatures. To those protocol- or application-layer parsers, a "composite" algorithm composed of two "component" algorithms is simply a new algorithm, and support for adding new algorithms generally already exists. Treating multiple "component" keys as a single "composite" key also has security advantages, such as preventing cross-protocol reuse of the individual component keys and guarantees about revoking or retiring all component keys together at the same time, especially if the composite is treated as a single object all the way down into the cryptographic module.

All that needs to be done is to standardize the formats of how the two keys from the two algorithms are combined into a single data structure and how the two resulting signatures or KEMs are combined into a single signature or KEM. The answer can be as simple as concatenation if the lengths are fixed or easily determined. At the time this document is published, security research is ongoing as to the security properties of concatenation-based composite signatures and KEMs versus more sophisticated signature and KEM combiners and protocol contexts in which those simpler combiners are sufficient.

One last consideration is the specific pairs of algorithms that can be combined. A recent trend in protocols is to only allow a small number of "known good" configurations that make sense, often referred to in cryptography as a "ciphersuite", instead of allowing arbitrary combinations of individual configuration choices that may interact in dangerous ways. The current consensus is that the same approach should be followed for combining cryptographic algorithms and that "known good" pairs should be explicitly listed ("explicit composite") instead of just allowing arbitrary combinations of any two cryptographic algorithms ("generic composite").

The same considerations apply when using multiple certificates to transport a pair of related keys for the same subject. Exactly how two certificates should be managed in order to avoid some of the pitfalls mentioned above is still an active area of investigation. Using two certificates keeps the certificate tooling simple and straightforward, but in the end, this simply moves problems (i.e., problems with the requirement that both certificates be used as a pair, that two signatures that must be carried separately, and that both validate) to the certificate management layer, where addressing these concerns in a robust way can be difficult.

At least one scheme has been proposed that allows the pair of certificates to exist as a single certificate when being issued and managed but dynamically split into individual certificates when needed (see [\[ENC-PAIR-CERTS\]](#)).

13.3.3. Key Reuse in Hybrid Schemes

An important security note, particularly when using hybrid signature keys, but also to a lesser extent hybrid KEM keys, is key reuse. In traditional cryptography, problems can occur with so-called "cross-protocol attacks" when the same key can be used for multiple protocols; for example, signing TLS handshakes and signing S/MIME emails. While it is not best practice to reuse keys within the same protocol, e.g., using the same key for multiple S/MIME certificates for the same user, it is not generally catastrophic for security. However, key reuse becomes a large security problem within hybrid schemes.

Consider an {RSA, ML-DSA} hybrid key where the RSA key also appears within a single-algorithm certificate. In this case, an attacker could perform a "stripping attack" where they take some piece of data signed with the {RSA, ML-DSA} key, remove the ML-DSA signature, and present the data as if it was intended for the RSA only certificate. This leads to a set of security definitions called "non-separability properties", which refers to how well the signature scheme resists various complexities of downgrade/stripping attacks [\[HYBRID-SIG-SPECT\]](#). Therefore, it is recommended that implementers either reuse the entire hybrid key as a whole or perform fresh key generation of all component keys per usage, and must not take an existing key and reuse it as a component of a hybrid key.

13.3.4. Future Directions and Ongoing Research

Many aspects of hybrid cryptography are still under investigation. The LAMPS Working Group at IETF is actively exploring the security properties of these combinations, and future standards will reflect the evolving consensus on these issues.

14. Impact on Constrained Devices and Networks

PQC algorithms generally have larger keys, ciphertext, and signature sizes than traditional public key algorithms. This has particular impact on constrained devices that operate with limited data rates. In the IoT space, these constraints have historically driven significant optimization efforts in the IETF (e.g., in the LAKE and CoRE Working Groups) to adapt security protocols to resource-constrained environments.

As the transition to PQC progresses, these environments will face similar challenges. Larger message sizes can increase handshake latency, raise energy consumption, and require fragmentation logic. Work is ongoing in the IETF to study how PQC can be deployed in constrained devices (see [\[CONSTRAIN-DEV-PCQ\]](#)).

15. Security Considerations

15.1. Cryptanalysis

Traditional cryptanalysis exploits weaknesses in algorithm design, mathematical vulnerabilities, or implementation flaws that are exploitable with classical (i.e., non-quantum) hardware, whereas quantum cryptanalysis harnesses the power of CRQCs to solve specific mathematical problems more efficiently. Quantum side-channel attacks are another form of quantum cryptanalysis. In such attacks, a device under threat is directly connected to a quantum computer, which then injects entangled or superimposed data streams to exploit hardware that lacks protection against quantum side channels. Both pose threats to the security of cryptographic algorithms, including those used in PQC. It is crucial to develop and adopt new cryptographic algorithms resilient against these threats to ensure long-term security in the face of advancing cryptanalysis techniques.

Recent attacks on the side-channel implementations using deep learning-based power analysis have also shown that one needs to be cautious while implementing the required PQC algorithms in hardware. Two of the most recent works include one attack on ML-KEM [\[KyberSide\]](#) and one attack on Saber [\[SaberSide\]](#). An evolving threat landscape points to the fact that lattice-based cryptography is indeed more vulnerable to side-channel attacks as in [\[SideCh\]](#) and [\[LatticeSide\]](#). Consequently, some mitigation techniques for side-channel attacks have been proposed; see [\[Mitigate1\]](#), [\[Mitigate2\]](#), and [\[Mitigate3\]](#).

15.2. Cryptographic Agility

Cryptographic agility is recommended for both traditional and quantum cryptanalysis as it enables organizations to adapt to emerging threats, adopt stronger algorithms, comply with standards, and plan for long-term security in the face of evolving cryptanalytic techniques and the advent of CRQCs.

Several PQC schemes are available that need to be tested; cryptography experts around the world are pushing for the best possible solutions, and the first standards that will ease the introduction of PQC are being prepared. This is of paramount importance and is a call for imminent action for organizations, bodies, and enterprises to start evaluating their cryptographic agility, assess the complexity of implementing PQC into their products, processes, and systems, and develop a migration plan that achieves their security goals to the best possible extent.

An important and often overlooked step in achieving cryptographic agility is maintaining a cryptographic inventory. Modern software stacks incorporate cryptography in numerous places, making it challenging to identify all instances. Therefore, cryptographic agility and inventory management take two major forms. First, application developers responsible for software maintenance should actively search for instances of hard-coded cryptographic algorithms within applications. When possible, they should design the choice of algorithm to be dynamic, based on application configuration. Second, administrators, policy officers, and compliance teams should take note of any instances where an application exposes cryptographic configurations. These instances should be managed through either organization-wide written cryptographic policies or automated cryptographic policy systems.

Numerous commercial solutions are available for detecting hard-coded cryptographic algorithms in source code and compiled binaries, as well as providing cryptographic policy management control planes for enterprise and production environments.

15.3. Jurisdictional Fragmentation

Another potential application of hybrid schemes bears mentioning, even though it is not directly related to PQC: using hybrids to navigate inter-jurisdictional cryptographic connections. Traditional cryptography is already fragmented by jurisdiction. Consider that while most jurisdictions support ECDH, those in the United States will prefer the NIST curves while those in Germany will prefer the Brainpool curves. China, Russia, and other jurisdictions have their own national cryptography standards. This situation of fragmented global cryptography standards is unlikely to improve with PQC. If "and" mode hybrid schemes become standardized for the reasons mentioned above, then one could imagine leveraging them to create ciphersuites in which a single cryptographic operation simultaneously satisfies the cryptographic requirements of both endpoints.

15.4. Hybrid Key Exchange and Signatures: Bridging the Gap Between PQ/T Cryptography

Post-quantum algorithms selected for standardization are relatively new and have not been subject to the same depth of study as traditional algorithms. PQC implementations will also be new and therefore more likely to contain implementation bugs than the battle-tested crypto implementations that are relied on today. In addition, certain deployments may need to retain traditional algorithms due to regulatory constraints, e.g., FIPS [SP-800-56C] or Payment Card Industry (PCI) compliance [PCI]. Hybrid key exchange is recommended to enhance security against the HNDL attack. Additionally, hybrid signatures provide for time to react in the case of the announcement of a devastating attack against any one algorithm, while not fully abandoning traditional cryptosystems.

Hybrid key exchange performs both a classical and a post-quantum key exchange in parallel. It provides security redundancy against potential weaknesses in PQC algorithms, allows for a gradual transition of trust in PQC algorithms, and, in backward-compatible designs, enables gradual adoption without breaking compatibility with existing systems. For instance, in TLS 1.3, a hybrid key exchange can combine a widely supported classical algorithm, such as X25519, with a post-quantum algorithm like ML-KEM. This allows legacy clients to continue using the classical algorithm while enabling upgraded clients to proceed with hybrid key exchange. In contrast, overhead-spreading hybrid designs focus on reducing the PQ overhead. For example, approaches like those described in [PQ-MLS] amortize PQ costs by selectively applying PQ updates in key exchange processes, allowing systems to balance security and efficiency. This strategy ensures a post-quantum secure channel while keeping the overhead manageable, making it particularly suitable for constrained environments.

While some hybrid key exchange options introduce additional computational and bandwidth overhead, the impact of traditional key exchange algorithms (e.g., key size) is typically small, helping to keep the overall increase in resource usage manageable for most systems. In highly constrained environments, however, those hybrid key exchange protocols may be impractical due to their higher resource requirements compared to pure post-quantum or traditional key exchange approaches. However, some hybrid key exchange designs distribute the PQC overhead, making them more suitable for constrained environments. The choice of hybrid key exchange design depends on the specific system requirements and use case, so the appropriate approach may vary.

15.5. Caution: Ciphertext Commitment in KEM vs. DH

The ciphertext generated by a KEM is not necessarily directly linked to the shared secret it produces. KEMs allow for multiple ciphertexts to encapsulate the same shared secret, which enables flexibility in key management without enforcing a strict one-to-one correspondence between ciphertexts and shared secrets. This allows for secret reuse across different recipients, sessions, or operational contexts without the need for new secrets for each use, simplifying key

distribution and reducing computational overhead. In contrast, cryptographic schemes like Diffie-Hellman inherently link the public key to the derived shared secret, meaning any change in the public key results in a different shared secret.

16. IANA Considerations

This document has no IANA actions.

17. Further Reading and Resources

A good book on modern cryptography is "Serious Cryptography, 2nd Edition" by Jean-Philippe Aumasson [[Serious-Crypt](#)].

The Open Quantum Safe (OQS) Project [[OQS](#)] is an open-source project that aims to support the transition to quantum-resistant cryptography.

The IETF's PQUIP Working Group [[PQUIP-WG](#)] maintains a list of PQC-related protocol work within the IETF.

18. References

18.1. Normative References

[[ClassicMcEliece](#)] "Classic McEliece", <<https://classic.mceliece.org/>>.

[[FN-DSA](#)] "FALCON: Fast Fourier lattice-based compact signatures over NTRU", <<https://falcon-sign.info/>>.

[[FrodoKEM](#)] "FrodoKEM", <<https://frodokem.org/>>.

[[Grovers](#)] Grover, L. K., "A fast quantum mechanical algorithm for database search", STOC '96: Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing, pp. 212-219, DOI 10.1145/237814.237866, 1 July 1996, <<https://dl.acm.org/doi/10.1145/237814.237866>>.

[[ML-DSA](#)] NIST, "Module-Lattice-Based Digital Signature Standard", NIST FIPS 204, DOI 10.6028/NIST.FIPS.204, August 2024, <<https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.204.pdf>>.

[[ML-KEM](#)] NIST, "Module-Lattice-Based Key-Encapsulation Mechanism Standard", NIST FIPS 203, DOI 10.6028/nist.fips.203, August 2024, <<https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.203.pdf>>.

[[NTRU](#)] "NTRU", <<https://ntru.org/index.shtml>>.

[[RFC4034](#)] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <<https://www.rfc-editor.org/info/rfc4034>>.

-
- [RFC6090] McGrew, D., Igoe, K., and M. Salter, "Fundamental Elliptic Curve Cryptography Algorithms", RFC 6090, DOI 10.17487/RFC6090, February 2011, <<https://www.rfc-editor.org/info/rfc6090>>.
- [RFC8235] Hao, F., Ed., "Schnorr Non-interactive Zero-Knowledge Proof", RFC 8235, DOI 10.17487/RFC8235, September 2017, <<https://www.rfc-editor.org/info/rfc8235>>.
- [RFC8391] Huelsing, A., Butin, D., Gazdag, S., Rijneveld, J., and A. Mohaisen, "XMSS: eXtended Merkle Signature Scheme", RFC 8391, DOI 10.17487/RFC8391, May 2018, <<https://www.rfc-editor.org/info/rfc8391>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8554] McGrew, D., Curcio, M., and S. Fluhrer, "Leighton-Micali Hash-Based Signatures", RFC 8554, DOI 10.17487/RFC8554, April 2019, <<https://www.rfc-editor.org/info/rfc8554>>.
- [RFC9180] Barnes, R., Bhargavan, K., Lipp, B., and C. Wood, "Hybrid Public Key Encryption", RFC 9180, DOI 10.17487/RFC9180, February 2022, <<https://www.rfc-editor.org/info/rfc9180>>.
- [RFC9242] Smyslov, V., "Intermediate Exchange in the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 9242, DOI 10.17487/RFC9242, May 2022, <<https://www.rfc-editor.org/info/rfc9242>>.
- [RFC9370] Tjhai, C.J., Tomlinson, M., Bartlett, G., Fluhrer, S., Van Geest, D., Garcia-Morchon, O., and V. Smyslov, "Multiple Key Exchanges in the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 9370, DOI 10.17487/RFC9370, May 2023, <<https://www.rfc-editor.org/info/rfc9370>>.
- [RFC9881] Massimo, J., Kampanakis, P., Turner, S., and B. E. Westerbaan, "Internet X.509 Public Key Infrastructure -- Algorithm Identifiers for the Module-Lattice-Based Digital Signature Algorithm (ML-DSA)", RFC 9881, DOI 10.17487/RFC9881, October 2025, <<https://www.rfc-editor.org/info/rfc9881>>.
- [RSA] Rivest, R. L., Shamir, A., and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems", Communications of the ACM, vol. 21, no. 2, pp. 120-126, DOI 10.1145/359340.359342, February 1978, <<https://dl.acm.org/doi/pdf/10.1145/359340.359342>>.
- [Shors] Shor, P., "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer", arXiv:quant-ph/9508027v2, 25 January 1996, <<https://arxiv.org/pdf/quant-ph/9508027>>.
- [SLH-DSA] NIST, "Stateless Hash-Based Digital Signature Standard", NIST FIPS 205, DOI 10.6028/NIST.FIPS.205, August 2024, <<https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.205.pdf>>.

18.2. Informative References

- [AddSig]** NIST, "Post-Quantum Cryptography: Additional Digital Signature Schemes", <<https://csrc.nist.gov/Projects/pqc-dig-sig/standardization>>.
- [ANSSI]** ANSSI, "ANSSI views on the Post-Quantum Cryptography transition (2023 follow up)", 21 December 2023, <https://cyber.gouv.fr/sites/default/files/document/follow_up_position_paper_on_post_quantum_cryptography.pdf>.
- [BBS-SIG-SCHEME]** Looker, T., Kalos, V., Whitehead, A., and M. Lodder, "The BBS Signature Scheme", Work in Progress, Internet-Draft, draft-irtf-cfrg-bbs-signatures-10, 8 January 2026, <<https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-bbs-signatures-10>>.
- [BHK09]** Bellare, M., Hofheinz, D., and E. Kiltz, "Subtleties in the Definition of IND-CCA: When and How Should Challenge-Decryption be Disallowed?", Cryptology ePrint Archive, Paper 2009/418, 2009, <<https://eprint.iacr.org/2009/418>>.
- [BIKE]** "BIKE", <<https://bikesuite.org/>>.
- [BPQS]** Chalkias, K., Brown, J., Hearn, M., Lillehagen, T., Nitto, I., and T. Schroeter, "Blockchained Post-Quantum Signatures", Cryptology ePrint Archive, Paper 2018/658, n.d., <<https://eprint.iacr.org/2018/658>>.
- [BSI-PQC]** BSI, "Quantum-safe cryptography - fundamentals, current developments and recommendations", 18 May 2022, <<https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Brochure/quantum-safe-cryptography.html?nn=916626>>.
- [Cloudflare]** Westerbaan, B., "NIST's pleasant post-quantum surprise", Cloudflare Blog, 8 July 2022, <<https://blog.cloudflare.com/nist-post-quantum-surprise/>>.
- [CNSA2-0]** NSA, "Announcing the Commercial National Security Algorithm Suite 2.0", September 2022, <https://media.defense.gov/2025/May/30/2003728741/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS.PDF>.
- [CONSTRAIN-DEV-PCQ]** Reddy, K. T., Wing, D., S. B., and K. Kwiatkowski, "Adapting Constrained Devices for Post-Quantum Cryptography", Work in Progress, Internet-Draft, draft-ietf-pquip-pqc-hsm-constrained-05, 1 April 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-pquip-pqc-hsm-constrained-05>>.
- [CRQCThreat]** Jaques, S., "Landscape of Quantum Computing", <https://sam-jaques.appspot.com/quantum_landscape_2024>.
- [CS01]** Cramer, R. and V. Shoup, "Design and Analysis of Practical Public-Key Encryption Schemes Secure against Adaptive Chosen Ciphertext Attack", Cryptology ePrint Archive, Paper 2001/108, 2001, <<https://eprint.iacr.org/2001/108>>.

- [ENC-PAIR-CERTS]** Bonnell, C., Gray, J., Hook, D., Okubo, T., and M. Ounsworth, "A Mechanism for Encoding Differences in Paired Certificates", Work in Progress, Internet-Draft, draft-bonnell-lamps-chameleon-certs-07, 18 October 2025, <<https://datatracker.ietf.org/doc/html/draft-bonnell-lamps-chameleon-certs-07>>.
- [GMR88]** Goldwasser, S., Micali, S., and R. L. Rivest, "A digital signature scheme secure against adaptive chosen-message attacks", *SIAM Journal on Computing*, vol. 17, no. 2, pp. 281-308, DOI 10.1137/0217017, April 1988, <https://people.csail.mit.edu/silvio/Selected%20Scientific%20Papers/Digital%20Signatures/A_Digital_Signature_Scheme_Secure_Against_Adaptive_Chosen-Message_Attack.pdf>.
- [Grover-Search]** Zalka, C., "Grover's quantum searching algorithm is optimal", *Physical Review A*, vol. 60, no. 4, pp. 2746-2751, DOI 10.1103/PhysRevA.60.2746, October 1999, <<https://link.aps.org/doi/10.1103/PhysRevA.60.2746>>.
- [HQC]** "HQC", <<http://pqc-hqc.org/>>.
- [HYBRID-SIG-SPECT]** Bindel, N., Hale, B., Connolly, D., and F. D., "Hybrid signature spectrums", Work in Progress, Internet-Draft, draft-ietf-pquip-hybrid-signature-spectrums-07, 20 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pquip-hybrid-signature-spectrums-07>>.
- [KEEPINGUP]** Cremers, C., Dax, A., and N. Medinger, "Keeping Up with the KEMs: Stronger Security Notions for KEMs and automated analysis of KEM-based protocols", *Cryptology ePrint Archive*, Paper 2023/1933, 2023, <<https://eprint.iacr.org/2023/1933>>.
- [KEM-COMBINER]** Ounsworth, M., Wussler, A., and S. Kousidis, "Combiner function for hybrid key encapsulation mechanisms (Hybrid KEMs)", Work in Progress, Internet-Draft, draft-ounsworth-cfrg-kem-combiners-05, 31 January 2024, <<https://datatracker.ietf.org/doc/html/draft-ounsworth-cfrg-kem-combiners-05>>.
- [KyberSide]** Ji, Y., Wang, R., Ngo, K., Dubrova, E., and L. Backlund, "A Side-Channel Attack on a Hardware Implementation of CRYSTALS-Kyber", *Cryptology ePrint Archive*, Paper 2022/1452, 2022, <<https://eprint.iacr.org/2022/1452>>.
- [LattFail1]** D'Anvers, J., Guo, Q., Johansson, T., Nilsson, A., Vercauteren, F., and I. Verbauwhede, "Decryption Failure Attacks on IND-CCA Secure Lattice-Based Schemes", *Public-Key Cryptography - PKC 2019, Lecture Notes in Computer Science*, vol. 11443, pp. 565-598, DOI 10.1007/978-3-030-17259-6_19, 6 April 2019, <https://link.springer.com/chapter/10.1007/978-3-030-17259-6_19>.
- [LattFail2]** D'Anvers, J., Rossi, M., and F. Virdia, "(One) Failure Is Not an Option: Bootstrapping the Search for Failures in Lattice-Based Encryption Schemes", *Advances in Cryptology - EUROCRYPT 2020, Lecture Notes in Computer Science*, vol. 12107, pp. 3-33, DOI 10.1007/978-3-030-45727-3_1, 1 May 2020, <https://link.springer.com/chapter/10.1007/978-3-030-45727-3_1>.

-
- [LatticeSide]** Ravi, P., Roy, S. S., Chattopadhyay, A., and S. Bhasin, "Generic Side-channel attacks on CCA-secure lattice-based PKE and KEM schemes", Cryptology ePrint Archive, Paper 2019/948, 2019, <<https://eprint.iacr.org/2019/948>>.
- [LIBOQS]** "LibOQS - Open Quantum Safe", commit 97f6b86, November 2025, <<https://github.com/open-quantum-safe/liboqs>>.
- [Lyu09]** Lyubashevsky, V., "Fiat-Shamir With Aborts: Applications to Lattice and Factoring-Based Signatures", ASIACRYPT 2009, <<https://www.iacr.org/archive/asiacrypt2009/59120596/59120596.pdf>>.
- [Mitigate1]** Hoffmann, C., Libert, B., Momin, C., Peters, T., and F. Standaert, "POLKA: Towards Leakage-Resistant Post-Quantum CCA-Secure Public Key Encryption", Cryptology ePrint Archive, Paper 2022/873, 2022, <<https://eprint.iacr.org/2022/873>>.
- [Mitigate2]** Tsai, T., Huang, S., Tseng, Y., Chuang, Y., and Y. Hung, "Leakage-Resilient Certificate-Based Authenticated Key Exchange Protocol", IEEE Open Journal of the Computer Society, vol. 3, pp. 137-148, DOI 10.1109/OJCS.2022.3198073, 2022, <<https://ieeexplore.ieee.org/document/9855226>>.
- [Mitigate3]** Azouaoui, M., Kuzovkova, Y., Schneider, T., and C. V. Vredendaal, "Post-Quantum Authenticated Encryption against Chosen-Ciphertext Side-Channel Attacks", Cryptology ePrint Archive, Paper 2022/916, 2022, <<https://eprint.iacr.org/2022/916>>.
- [ML-DSA-X.509]** Ounsworth, M., Gray, J., Pala, M., Klaussner, J., and S. Fluhrer, "Composite Module-Lattice-Based Digital Signature Algorithm (ML-DSA) for use in X.509 Public Key Infrastructure", Work in Progress, Internet-Draft, draft-ietf-lamps-pq-composite-sigs-19, 21 April 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-lamps-pq-composite-sigs-19>>.
- [NIST]** NIST, "Post-Quantum Cryptography Standardization", <<https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization>>.
- [NISTFINAL]** NIST, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards", 13 August 2024, <<https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>>.
- [OQS]** "Open Quantum Safe Project", <<https://openquantumsafe.org/>>.
- [PCI]** PCI Security Standards Council, "Payment Card Industry Data Security Standard", PCI DSS: v4.0.1, <https://docs-prv.pcisecuritystandards.org/PCI%20DSS/Standard/PCI-DSS-v4_0_1.pdf>.
- [PQ-HPKE]** Barnes, R. and D. Connolly, "Post-Quantum and Post-Quantum/Traditional Hybrid Algorithms for HPKE", Work in Progress, Internet-Draft, draft-ietf-hpke-pq-04, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-hpke-pq-04>>.

-
- [PQ-KEM]** Connolly, D., Barnes, R., and P. Grubbs, "Hybrid PQ/T Key Encapsulation Mechanisms", Work in Progress, Internet-Draft, draft-irtf-cfrg-hybrid-kems-10, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-hybrid-kems-10>>.
- [PQ-MLS]** Joël, Hale, B., Mularczyk, M., and X. Tian, "Flexible Hybrid PQ MLS Combiner", Work in Progress, Internet-Draft, draft-hale-mls-combiner-01, 26 September 2024, <<https://datatracker.ietf.org/doc/html/draft-hale-mls-combiner-01>>.
- [PQCAPI]** NIST, "PQC - API notes", <<https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/example-files/api-notes.pdf>>.
- [PQRSA]** Bernstein, D. J., Heninger, N., Lou, P., and L. Valenta, "Post-quantum RSA", 19 April 2017, <<https://cr.yp.to/papers/pqrsa-20170419.pdf>>.
- [PQUIP-WG]** IETF, "Post-Quantum Use In Protocols (pquip)", <<https://datatracker.ietf.org/group/pquip/documents/>>.
- [QC-DNS]** Hoffman, P., "Quantum Computing and the DNS", ICANN Office of the Chief Technology Officer, OCTO-031v2, 22 April 2024, <<https://www.icann.org/octo-031-en.pdf>>.
- [QuantSide]** Xu, C., Erata, F., and J. Szefer, "Exploration of Quantum Computer Power Side-Channels", arXiv:2304.03315v2, 9 May 2023, <<https://arxiv.org/pdf/2304.03315>>.
- [RFC5652]** Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC9528]** Selander, G., Preuß Mattsson, J., and F. Palombini, "Ephemeral Diffie-Hellman Over COSE (EDHOC)", RFC 9528, DOI 10.17487/RFC9528, March 2024, <<https://www.rfc-editor.org/info/rfc9528>>.
- [RFC9763]** Becker, A., Guthrie, R., and M. Jenkins, "Related Certificates for Use in Multiple Authentications within a Protocol", RFC 9763, DOI 10.17487/RFC9763, June 2025, <<https://www.rfc-editor.org/info/rfc9763>>.
- [RFC9794]** Driscoll, F., Parsons, M., and B. Hale, "Terminology for Post-Quantum Traditional Hybrid Schemes", RFC 9794, DOI 10.17487/RFC9794, June 2025, <<https://www.rfc-editor.org/info/rfc9794>>.
- [RFC9814]** Housley, R., Fluhrer, S., Kampanakis, P., and B. Westerbaan, "Use of the SLH-DSA Signature Algorithm in the Cryptographic Message Syntax (CMS)", RFC 9814, DOI 10.17487/RFC9814, July 2025, <<https://www.rfc-editor.org/info/rfc9814>>.
- [RFC9941]** Friedl, M., Mojzis, J., and S. Josefsson, "Secure Shell (SSH) Key Exchange Method Using Hybrid Streamlined NTRU Prime sntrup761 and X25519 with SHA-512: sntrup761x25519-sha512", RFC 9941, DOI 10.17487/RFC9941, April 2026, <<https://www.rfc-editor.org/info/rfc9941>>.

-
- [RSA10SC]** QuintessenceLabs, "Breaking RSA Encryption - an Update on the State-of-the-Art", 13 June 2019, <<https://www.quintessencelabs.com/blog/breaking-rsa-encryption-update-state-art>>.
- [RSA8HRS]** Gidney, C. and M. Eker, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits", arXiv:1905.09749v3, 13 April 2021, <<https://arxiv.org/abs/1905.09749>>.
- [RSAShor]** Beauregard, S., "Circuit for Shor's algorithm using $2n+3$ qubits", arXiv:quant-ph/0205095v3, 21 February 2003, <<https://arxiv.org/pdf/quant-ph/0205095.pdf>>.
- [SaberSide]** Ngo, K., Dubrova, E., and T. Johansson, "A side-channel attack on a masked and shuffled software implementation of Saber", Journal of Cryptographic Engineering, vol. 13, pp. 443-460, DOI 10.1007/s13389-023-00315-3, 25 April 2023, <<https://link.springer.com/article/10.1007/s13389-023-00315-3>>.
- [Serious-Crypt]** Aumasson, J., "Serious Cryptography, 2nd Edition", ISBN 9781718503847, August 2024.
- [SideCh]** Ngo, K., Wang, R., Dubrova, E., and N. Paulsruud, "Side-Channel Attacks on Lattice-Based KEMs Are Not Prevented by Higher-Order Masking", Cryptology ePrint Archive, Paper 2022/919, 2022, <<https://eprint.iacr.org/2022/919>>.
- [SP-1800-38C]** Newhouse, W., Souppaya, M., Barke, W., Brown, C., Kampanakis, P., Goodman, J., Prat, J., Larrieu, R., Gray, J., Ounsworth, M., Viana, C., Gong, H. L. V., Kwiatkowski, K., Hu, A., Burns, R., Paquin, C., Gilbert, J., Scinta, G., Kim, E., and V. Krumme, "Migration to Post-Quantum Cryptography Quantum Readiness: Testing Draft Standards, Volume C: Quantum-Resistant Cryptography Technology Interoperability and Performance Report", Preliminary Draft, NIST SP 1800-38C, December 2023, <<https://www.nccoe.nist.gov/sites/default/files/2023-12/pqc-migration-nist-sp-1800-38c-preliminary-draft.pdf>>.
- [SP-800-56C]** Barker, E., Chen, L., and R. Davis, "Recommendation for Key-Derivation Methods in Key-Establishment Schemes", NIST SP 800-56Cr2, DOI 10.6028/NIST.SP.800-56Cr2, August 2020, <<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-56Cr2.pdf>>.
- [Threat-Report]** Mosca, M. and M. Piani, "Quantum Threat Timeline Report 2020", Global Risk Institute, 27 January 2021, <<https://globalriskinstitute.org/publications/quantum-threat-timeline-report-2020/>>.
- [TLS-HYB-KEY-EXCH]** Stebila, D., Fluhrer, S., and S. Gueron, "Hybrid key exchange in TLS 1.3", Work in Progress, Internet-Draft, draft-ietf-tls-hybrid-design-16, 7 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tls-hybrid-design-16>>.
- [X-WING]** Connolly, D., Schwabe, P., and B. Westerbaan, "X-Wing: general-purpose hybrid post-quantum KEM", Work in Progress, Internet-Draft, draft-connolly-cfrg-xwing-kem-10, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-connolly-cfrg-xwing-kem-10>>.

Acknowledgements

This document leverages text from an earlier Internet-Draft by Paul Hoffman. Thanks to Dan Wing, Florence D, Thom Wiggers, Sophia Grundner-Culemann, Panos Kampanakis, Ben S, Sofia Celi, Melchior Aelmans, Falko Strenzke, Deirdre Connolly, Hani Ezzadeen, Britta Hale, Scott Rose, Hilarie Orman, Thomas Fossati, Roman Danyliw, Mike Bishop, Mališa Vučinić, Éric Vyncke, Deb Cooley, Dirk Von Hugo, and Daniel Van Geest for the discussion, review and comments.

In particular, the authors would like to acknowledge the contributions to this document by Kris Kwiatkowski.

Authors' Addresses

Aritra Banerjee

Nokia
London
United Kingdom
Email: aritra.banerjee@nokia.com

Tirumaleswar Reddy.K

Nokia
Bangalore
Karnataka
India
Email: k.tirumaleswar_reddy@nokia.com

Dimitrios Schoinianakis

Nokia
Athens
Greece
Email: dimitrios.schoinianakis@nokia-bell-labs.com

Timothy Hollebeek

DigiCert
Pittsburgh, PA
United States of America
Email: tim.hollebeek@digicert.com

Mike Ounsworth

Entrust Limited
2500 Solandt Road, Suite 100
Ottawa, Ontario K2K 3G5
Canada
Email: mike@ounsworth.ca