

PQUIP
Internet-Draft
Intended status: Informational
Expires: 30 July 2026

T. Reddy
Nokia
D. Wing
Citrix
B. Salter
UK National Cyber Security Centre
K. Kwiatkowski
PQShield
26 January 2026

Adapting Constrained Devices for Post-Quantum Cryptography
draft-ietf-pquip-pqc-hsm-constrained-03

Abstract

This document provides guidance on integrating Post-Quantum Cryptography (PQC) into resource-constrained devices, such as IoT nodes and lightweight Hardware Security Modules (HSMs). These systems often operate with strict limitations on processing power, RAM, and flash memory, and may even be battery-powered. The document emphasizes the role of hardware security as the basis for secure operations, supporting features such as seed-based key generation to minimize persistent storage, efficient handling of ephemeral keys, and the offloading of cryptographic tasks in low-resource environments. It also explores the implications of PQC on firmware update mechanisms in such constrained systems.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-pquip-pqc-hsm-constrained/>.

Discussion of this document takes place on the pquip Working Group mailing list (<mailto:pqc@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/pqc/>. Subscribe at <https://www.ietf.org/mailman/listinfo/pqc/>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 30 July 2026.

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	3
2. Key Management in Constrained Devices for PQC	4
2.1. Seed Management	4
2.1.1. Seed Storage	5
2.1.2. Efficient Key Derivation	6
2.1.3. Exporting Seeds and Private Keys	7
2.2. Ephemeral Key Management	8
3. Optimizing Memory Footprint in Post-Quantum Signature Schemes	8
3.1. Memory requirements of Lattice-Based Schemes	9
3.1.1. Lazy Expansion as a Memory Optimization Technique	10
3.2. Pre-hashing as a Memory Optimization Technique	11
4. Optimizing Performance in PQC Signature Schemes	11
4.1. Impact of rejection sampling in ML-DSA Signing on performance	12
4.1.1. Practical Implications for Constrained Cryptographic Modules	15
4.1.2. Suggestions for benchmarking ML-DSA Signing Performance	16

5.	Additional Considerations for PQC Use in Constrained Devices	16
5.1.	Key Rotation and Renewal	16
5.2.	Cryptographic Artifact Sizes for Post-Quantum Algorithms	17
6.	Post-quantum Firmware Upgrades for Constrained Devices	19
6.1.	Post-Quantum Firmware Authentication	19
6.2.	Hybrid Signature Approaches	20
7.	Security Considerations	21
7.1.	Side Channel Protection	21
	Acknowledgments	21
	References	21
	Normative References	21
	Informative References	21
	Authors' Addresses	24

1. Introduction

The transition to post-quantum cryptography (PQC) poses significant challenges for resource-constrained devices, such as Internet of Things (IoT) devices, which are often equipped with Trusted Execution Environments (TEEs), secure elements, or other forms of hardware security modules (HSMs). These devices typically operate under strict limitations on processing power, RAM, and flash memory, and in some cases are battery-powered. Adopting PQC algorithms in such environments is difficult due to their substantially larger key sizes and, in some cases, higher computational demands. Consequently, the migration to PQC requires careful planning to ensure secure and efficient key management within constrained platforms.

Constrained devices are often deployed as clients initiating outbound connections, but some also act in server roles or enforce local authentication policies. As a result, designers may need to consider PQ solutions to address confidentiality, both outbound and inbound authentication, and signature verification used in secure boot, firmware updates, and device attestation.

This document provides guidance and best practices for integrating PQC algorithms into constrained devices. It reviews strategies for key storage, ephemeral key management, and performance optimization tailored to low-resource environments. The document also examines ephemeral key generation in protocols such as TLS, along with techniques to optimize PQC signature operations to improve performance within constrained cryptographic modules.

The focus is on PQC in constrained devices, with particular attention to the three algorithms standardized by NIST:

- * Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) [FIPS203],
- * Module-Lattice-Based Digital Signature Algorithm (ML-DSA) [FIPS204], and
- * Stateless Hash-Based Digital Signature Algorithm (SLH-DSA) [FIPS205].

The Hierarchical Signature System/Leighton-Micali Signature (HSS/LMS) [RFC8554] is also considered in the context of firmware signing. Future revisions may extend the scope to additional PQC algorithms, such as the Hamming Quasi-Cyclic (HQC) KEM [HQC] and the Fast Fourier Transform over NTRU-Lattice-Based Digital Signature Algorithm (FN-DSA) [FN-DSA].

This document focuses on device-level adaptations and considerations necessary to implement PQC efficiently on constrained devices. Actual protocol behaviour is defined in other documents.

2. Key Management in Constrained Devices for PQC

The embedded cryptographic components used in constrained devices are designed to securely manage cryptographic keys, often under strict limitations in RAM, flash memory, and computational resources. These limitations are further exhausted by the increased key sizes and computational demands of PQC algorithms.

One mitigation of storage limitations is to store only the seed rather than the full expanded private key, as the seed is far smaller and can derive the expanded private key as necessary. [FIPS204] Section 3.6.3 specifies that the seed sk generated during ML-DSA.KeyGen can be stored for later use with ML-DSA.KeyGen_internal. To reduce storage requirements on constrained devices, private keys for Initial Device Identifiers (IDevIDs), Locally Significant Device Identifiers (LDevIDs), and the optional attestation private key can be stored as seeds instead of expanded key material. This optimization does not apply to device certificates or trust anchors, which must be stored in persistent device storage since they are signed public data structures (see [RFC5280]). The terms IDevIDs and LDevIDs are explained in IEEE Std 802.1AR [IEEE-802.1AR].

2.1. Seed Management

To comply with [FIPS203], [FIPS204], [FIPS205] and [REC-KEM] guidelines:

2.1.1.1. Seed Storage

Several post-quantum algorithms use a seed to generate their private keys (e.g., ML-KEM, ML-DSA, and HQC). Those seeds are smaller than private keys, hence some implementations may choose to retain the seed rather than the full private key to save on storage space. The private key can then be derived from the seed when needed or retained in a cache within the security module.

The seed is a Critical Security Parameter (CSP) as defined in [ISO19790], from which the private key can be derived, hence it must be safeguarded with the same level of protection as a private key. Seeds should be securely stored within a cryptographic module of the device whether hardware or software-based to protect against unauthorized access.

The choice between storing a seed or an expanded private key involves trade-offs between storage efficiency and performance. Some constrained cryptographic modules may store only the seed and derive the expanded private key on demand, whereas others may prefer storing the full expanded key to reduce computational overhead during key usage.

The choice between storing the seed or the expanded private key has direct implications on performance, as key derivation incurs additional computation. The impact of this overhead varies depending on the algorithm. For instance, ML-DSA key generation, which primarily involves polynomial operations using the Number Theoretic Transform (NTT) and hashing, is computationally efficient compared to other post-quantum schemes. In contrast, SLH-DSA key generation requires constructing a Merkle tree and multiple calls to Winternitz One-Time Signature (WOTS+) key generation, making it significantly slower due to the recursive hash computations involved. Designers of constrained systems must carefully balance storage efficiency and computational overhead based on system requirements and operational constraints. While constrained systems employ various key storage strategies, the decision to store full private keys or only seeds depends on design goals, performance considerations, and standards compliance (e.g., PKCS#11).

While vulnerabilities like the "Unbindable Kemmy Schmidt" misbinding attack [BIND] demonstrate the risks of manipulating expanded private keys in environments lacking hardware-backed protections, these attacks generally assume an adversary has some level of control over the expanded key format. However, in a hardware-backed protected environment, where private keys are typically protected from such manipulation, the primary motivation for storing the seed rather than the expanded key is not directly tied to mitigating such misbinding attacks.

If the seed is not securely stored at the time of key generation, it is permanently lost because the process of deriving an expanded key from the seed relies on a one-way cryptographic function. This one-way function derives the private key from the seed, but the reverse operation, deriving the original seed from the expanded key, is computationally infeasible.

A challenge arises when importing an existing private key into a system designed to store only seeds. When a user attempts to import an already expanded private key, there is a mismatch between the key format used internally (seed-based) and the expanded private key. This issue arises because the internal format is designed for efficient key storage by deriving the private key from the seed, while the expanded private key is already fully computed. As NIST has not defined a single private key format for PQC algorithms, this creates a potential gap in interoperability.

2.1.2. Efficient Key Derivation

When storing only the seed in a constrained cryptographic module, it is crucial that the device is capable of deriving the private key efficiently whenever required. However, repeatedly re-deriving the private key for every cryptographic operation may introduce significant performance overhead. In scenarios where performance is a critical consideration, it may be more efficient to store the expanded private key directly (in addition to the seed). Implementations may choose to retain (cache) several recently-used or frequently-used private keys to avoid the computational overhead and delay of deriving private keys from their seeds with each request.

The key derivation process, such as ML-KEM.KeyGen_internal for ML-KEM or similar functions for other PQC algorithms, must be implemented in a way that can securely operate within the resource constraints of the device. If using the seed-only model, the derived private key should only be temporarily held in memory during the cryptographic operation and discarded immediately after use. However, storing the expanded private key may be a more practical solution in time-sensitive applications or for devices that frequently perform cryptographic operations.

2.1.3. Exporting Seeds and Private Keys

Given the potential for hardware failures or the end-of-life of devices containing keys, it is essential to plan for backup and recovery of cryptographic seeds and private keys. Constrained devices should support secure seed- or key-backup mechanisms, leveraging protections such as encrypted storage and ensuring that security measures are in place so that the backup data is protected from unauthorized access.

There are two distinct approaches to exporting private keys or seeds from a constrained device:

2.1.3.1. Direct Transfer Over TLS

In scenarios where the constrained device has sufficient capability to initiate or terminate a mutually-authenticated TLS session, the device can securely transfer encrypted private key material directly to another cryptographic module.

2.1.3.2. Export of Encrypted Seeds and Private Keys

In more common constrained device scenarios for secure exporting of seeds and private keys, a strong symmetric encryption algorithm, such as AES in key-wrap mode ([RFC3394]), should be used to encrypt the seed or private key before export. This ensures that the key remains protected even if the export process is vulnerable to quantum attacks.

Operationally, the exported data and the symmetric key used for encryption must both be protected against unauthorized access or modification.

2.1.3.3. Security Requirements for Export Operations

The encryption and decryption of seeds and private keys must occur entirely within the cryptographic modules to reduce the risk of exposure and ensure compliance to established security standards.

2.2. Ephemeral Key Management

Given the increased size of PQC key material, ephemeral key management will have to be optimized for both security and performance.

For PQC KEMs, ephemeral key pairs are generated from an ephemeral seed, that is used immediately during key generation and then discarded. Furthermore, once the shared secret is derived, the ephemeral private key will have to be deleted. Since the private key resides in the constrained cryptographic module, removing it optimizes memory usage, reducing the footprint of PQC key material in the cryptographic module. This also ensures that that no unnecessary secrets persist beyond their intended use.

Additionally, ephemeral keys, whether from traditional ECDH or PQC KEM algorithms, are intended to be unique for each key exchange instance and kept separate across connections (e.g., TLS). Deleting ephemeral keying material after use helps ensure that key material cannot be reused across connections, which would otherwise introduce security and privacy issues.

Constrained devices implementing PQC ephemeral key management will have to:

- * Generate ephemeral key pairs on-demand from an ephemeral seed stored temporarily within the cryptographic module.
- * Enforce immediate seed erasure after the key pair is generated and the cryptographic operation is completed.
- * Delete the private key after the shared secret is derived.
- * Prevent key reuse across different algorithm suites or sessions.

3. Optimizing Memory Footprint in Post-Quantum Signature Schemes

A key consideration when deploying post-quantum cryptography in cryptographic modules is the amount of memory available. For instance, ML-DSA, unlike traditional signature schemes such as RSA or ECDSA, requires significant memory during signing due to multiple Number Theoretic Transform (NTT) operations, matrix expansions, and rejection sampling loops. These steps involve storing large polynomial vectors and intermediate values, making ML-DSA more memory-intensive.

Some constrained systems, i.e. those battery-operated, may have very limited RAM available for cryptographic operations. In such cases, straightforward implementations of PQ schemes may exceed the available memory, making it infeasible to use without optimizations.

Several post-quantum schemes can be optimized to reduce the memory footprint of the algorithm. For instance, SLH-DSA has two flavours: the "f" variants which are parameterized to run as fast as possible, and the "s" variants which produce shorter signatures. Developers wishing to use SLH-DSA may wish to utilize the "s" variants on devices with insufficient RAM to use the "f" variants. Further optimizations may be possible by running the signature algorithm in a "streaming manner" such that constrained device does not need to hold the entire signature in memory at once, as discussed in [Stream-SPHINCS].

Both the ML-KEM and ML-DSA algorithms were selected for general use. Two optimization techniques that can be applied to make ML-DSA more feasible in constrained cryptographic modules are discussed in Section 3.1.1 and Section 3.2.

3.1. Memory requirements of Lattice-Based Schemes

The dominant source of memory usage in ML-DSA comes from holding the expanded matrix A and the associated polynomial vectors needed to compute the noisy affine transformation $t = A \cdot s_1 + s_2$, where A is a large public matrix derived from a seed, and t , s_1 , s_2 are polynomial vectors involved in the signing process. The elements of those matrices and vectors are polynomials with integer coefficients modulo Q . ML-DSA uses a 23-bit long modulus Q , where in case of ML-KEM it is 12 bits, regardless of security level. Conversely, the sizes of those matrices depend on the security level.

To compute memory requirements, we need to consider the dimensions of the public matrix A and the size of the polynomial vectors. Using ML-KEM-768 as an example, the public matrix A has dimensions 5×5 , with each polynomial having 256 coefficients. Each coefficient is stored on 2 bytes (uint16), leading to a size of $5 * 5 * 256 * 2 = 12,800$ bytes (approximately 12.5 KB) for the matrix A alone. The polynomial vectors t , s_1 , and s_2 also contribute significantly to memory usage, with each vector requiring $5 * 256 * 2 = 2,560$ bytes (approximately 2.5 KB) each. Hence, for straightforward implementation, the minimal amount of memory required for these vectors is $12,800 + 3 * 2,560 = 20,480$ bytes (approximately 20 KB). Similar computation can be easily done for other security levels as well as ML-DSA. The ML-DSA has much higher memory requirements due to larger matrix and polynomial sizes (i.e. ML-DSA-87 requires approximately 79 KB of RAM during signing operations).

It's worth noting that different cryptographic operations may have different memory requirements. For example, during ML-DSA verification, the memory usage is lower since the private key components are not needed.

3.1.1. Lazy Expansion as a Memory Optimization Technique

The lazy expansion technique is an optimization that significantly reduces memory usage by avoiding the need to store the entire expanded matrix A in memory at once. Instead of pre-computing and storing the full matrix, lazy expansion generates parts of it on-the-fly as needed for the process. This approach leverages the fact that not all elements of the matrix are required simultaneously, allowing for a more efficient use of memory.

As an example, we can look at the computation of matrix-vector multiplication $t = A \cdot s_1$. The matrix A is generated from a seed using a PRF, meaning that any element of A can be computed independently when needed. Similarly, the vector s_1 is expanded from random seed and a nonce using a PRF.

The lazy expansion would first generate first element of a vector s_1 ($s_1[0]$) and then iterate over each row of matrix A in a first column. This approach generates partial result, that is a vector t . To finalize the computation of a vector t , the next element of s_1 ($s_1[1]$) is generated, and the process is repeated for each column of A until all elements of s_1 have been processed. This method requires significantly less memory, in case of ML-KEM-768, size of element s_1 (512 bytes) and a vector t (2560 bytes) is $256 * 2 = 512$ bytes, meaning that only 512 bytes + one row of matrix A ($5 * 256 * 2 = 2560$ bytes) + one element of t ($5 * 2 = 10$ bytes) need to be stored in memory at any time, leading to a total of approximately 3 KB of memory usage, compared to the approximately 20 KB required for a straightforward implementation. The savings are even more pronounced for higher security levels, such as ML-DSA-87, where lazy expansion can reduce memory usage from approximately 79 KB to around 12 KB.

With lazy expansion, the implementation differs slightly from the straightforward version. Also, in some cases, lazy expansion may introduce additional computational overhead. Notably, applying it to ML-DSA signing operation may require to recompute vector y (FIPS-204, Algorithm 7, line 11) twice. In this case implementers need to weigh the trade-off between memory savings and additional computation.

Further memory optimizations to ML-DSA can be found in [BosRS22].

3.2. Pre-hashing as a Memory Optimization Technique

To address the memory consumption challenge, algorithms like ML-DSA offer a form of pre-hash using the `ホシ` (message representative) value described in Section 6.2 of [FIPS204]. The `ホシ` value provides an abstraction for pre-hashing by allowing the hash or message representative to be computed outside the cryptographic module. This feature offers additional flexibility by enabling the use of different cryptographic modules for the pre-hashing step, reducing memory consumption within the cryptographic module. The pre-computed `ホシ` value is then supplied to the cryptographic module, eliminating the need to transmit the entire message for signing. [RFC9881] discusses leveraging External`ホシ`-ML-DSA, where the pre-hashing step (External`ホシ`-ML-DSA.Prehash) is performed in a software cryptographic module, and only the pre-hashed message (`ホシ`) is sent to the hardware cryptographic module for signing (External`ホシ`-ML-DSA.Sign). By implementing External`ホシ`-ML-DSA.Prehash in software and External`ホシ`-ML-DSA.Sign in a hardware cryptographic module, the cryptographic workload is efficiently distributed, making it practical for high-volume signing operations even in memory-constrained cryptographic modules.

The main advantage of this method is that, unlike HashML-DSA, the External`ホシ`-ML-DSA approach is interoperable with the standard version of ML-DSA that does not use pre-hashing. This means a message can be signed using ML-DSA.Sign, and the verifier can independently compute `ホシ` and use External`ホシ`-ML-DSA.Verify for verification -- or vice versa. In both cases, the verifier does not need to know whether the signer used internal or external pre-hashing, as the resulting signature and verification process remain the same.

4. Optimizing Performance in PQC Signature Schemes

When implementing PQC signature algorithms in constrained cryptographic modules, performance optimization becomes a critical consideration. Transmitting the entire message to the cryptographic module for signing can lead to significant overhead, especially for large payloads. To address this, implementers can leverage techniques that reduce the data transmitted to the cryptographic module, thereby improving efficiency and scalability.

One effective approach involves sending only a message digest to the cryptographic module for signing. By signing the digest of the content rather than the entire content, the communication between the application and the cryptographic module is minimized, enabling better performance. This method is applicable for any PQC signature algorithm, whether it is ML-DSA, SLH-DSA, or any future signature scheme. For such algorithms, a mechanism is often provided to pre-hash or process the message in a way that avoids sending the entire

raw message for signing. In particular, algorithms like SLH-DSA present challenges due to their construction, which requires multiple passes over the message digest during the signing process. The signer does not retain the entire message or its full digest in memory at once. Instead, different parts of the message digest are processed sequentially during the signing procedure. This differs from traditional algorithms like RSA or ECDSA, which allow for more efficient processing of the message, without requiring multiple passes or intermediate processing of the digest.

4.1. Impact of rejection sampling in ML-DSA Signing on performance

In constrained and battery-powered IoT devices that perform ML-DSA signing, the rejection-sampling loop introduces variability in signing latency and energy consumption due to the probabilistic nature of the signing process. While this results in a variable number of iterations in the signing algorithm, the expected number of retries for the standardized ML-DSA parameter sets is quantified below.

The analysis in this section follows the algorithmic structure and assumptions defined in FIPS-204. Accordingly, the numerical results are analytically derived and characterize the expected behavior of ML-DSA.

The ML-DSA signature scheme uses the Fiat-Shamir with Aborts construction [Lyu09]. As a result, the signature generation algorithm is built around a rejection-sampling loop. This section examines the rejection-sampling behavior of ML-DSA, as rejection sampling is not commonly used as a core mechanism in traditional digital signature schemes.

Rejection sampling is used to ensure that intermediate and output values follow the distributions required by the security proof. In particular, after computing candidate signature components, the signer checks whether certain norm bounds are satisfied. If any of these bounds are violated, the entire signing attempt is discarded and restarted with fresh randomness.

The purpose of rejection sampling is twofold. First, it prevents leakage of information about the secret key through out-of-range values that could otherwise bias the distribution of signatures. Second, it ensures that the distribution of valid signatures is statistically close to the ideal distribution assumed in the security reduction.

The number of rejections during signature generation depends on four factors:

- * the message (i.e., the value of ホシ)
- * the secret key material
- * when hedged signing is used (see [FIPS204], Section 3.4), the random seed
- * the context string (see [FIPS204], Section 5.2)

As a result, some message-key combinations may lead to a higher number of rejection iterations than others.

Using Equation (5) from [Li32] and assuming an RBG as specified in [FIPS204] (Section 3.6.1), the rejection probability during ML-DSA signing can be computed. These probabilities depend on the ML-DSA parameter set and are summarized below.

+=====+	
ML-DSA Variant	Acceptance Probability
+=====+	
ML-DSA-44	0.2350
+-----+	
ML-DSA-65	0.1963
+-----+	
ML-DSA-87	0.2596
+-----+	

Table 1: Acceptance probability - per-attempt probability of successful signing for the given ML-DSA variant.

Each signing attempt can be modeled as an independent Bernoulli trial: an attempt either succeeds or is rejected, with a fixed per-attempt acceptance probability. Under this assumption, the expected number of iterations until a successful signature is generated is the reciprocal of the acceptance probability. Hence, if r denotes the per-iteration rejection probability and $p = 1 - r$ the acceptance probability, then the expected number of signing iterations is $1/p$. Using this model, the expected number of signing attempts for each ML-DSA variant is shown below.

ML-DSA Variant	Expected Number of Attempts
ML-DSA-44	4.255
ML-DSA-65	5.094
ML-DSA-87	3.852

Table 2: Expected Number of Attempts for the given ML-DSA variant.

This model also allows computing the probability that the rejection-sampling loop completes within a given number of iterations. Specifically, the minimum number of iterations n required to achieve a desired completion probability can be computed as: $n \geq \ln(1 - \text{desired_probability}) / \ln(1 - p)$, where p is the per-iteration acceptance probability. For example, achieving a 99% probability of completing the signing process for ML-DSA-65 requires at most 21 iterations of the rejection-sampling loop.

Finally, based on these results, the cumulative distribution function (CDF) can be derived for each ML-DSA variant. The CDF expresses the probability that the signing process completes within at most a given number of iterations.

Iterations	ML-DSA-44	ML-DSA-65	ML-DSA-87
1	0.2350	0.1963	0.2596
2	0.4148	0.3541	0.4518
3	0.5523	0.4809	0.5941
4	0.6575	0.5828	0.6995
5	0.7380	0.6647	0.7775
6	0.7996	0.7305	0.8353
7	0.8467	0.7834	0.8780
8	0.8827	0.8259	0.9097
9	0.9103	0.8601	0.9331
10	0.9314	0.8876	0.9505
11	0.9475	0.9096	0.9634

Table 3: CDF values denote the probability of completing the signing process within the given number of iterations, for each ML-DSA variant.

The table Table 3 shows that while acceptance rate is relatively high for ML-DSA, the probability quickly grows with increasing number of iterations. After 11 iterations, each ML-DSA variant achieves over 90% probability of completing the signing process.

4.1.1. Practical Implications for Constrained Cryptographic Modules

As shown above, the rejection-sampling loop in ML-DSA signing leads to a probabilistic runtime with a geometrically distributed number of iterations. While the expected execution time is small, the tail of the distribution implies that, with low probability, a signing operation may require significantly more iterations than average. This unfavorable tail behavior represents a practical concern for ML-DSA deployments on constrained devices with limited execution capability and may require additional consideration.

This consideration primarily applies to devices that perform ML-DSA signing. Devices that only generate ML-DSA keys or verify signatures are not affected, as those operations do not involve rejection sampling and have deterministic execution times.

4.1.2. Suggestions for benchmarking ML-DSA Signing Performance

When benchmarking ML-DSA signing performance in constrained cryptographic modules, it is important to account for the probabilistic nature of the rejection-sampling loop. Reporting only a single timing measurement or a best-case execution time may lead to misleading conclusions about practical performance.

To provide a more comprehensive assessment of ML-DSA signing performance, benchmarks should report the following two metrics:

1. Single-iteration signing time: The signing time for a signature operation that completes within a single iteration of the rejection-sampling loop. This metric reflects the best-case performance of the signing algorithm and provides insight into the efficiency of the core signing operation without the overhead of repeated iterations.
2. Average signing time: The average signing time measured over a sufficiently large number of signing operations, using independent messages and, where applicable, independent randomness. Alternatively, an implementation MAY report the signing time corresponding to the expected number of iterations (see Table 2). This approach requires identifying a message, key, and randomness combination that results in the expected iteration count.

Libraries implementing ML-DSA should provide a mechanism to report the number of rejection-sampling iterations used during the most recent signing operation. This enables benchmarking tools to accurately compute average signing times across multiple signing operations.

5. Additional Considerations for PQC Use in Constrained Devices

5.1. Key Rotation and Renewal

In constrained devices, managing the lifecycle of cryptographic keys including periodic key rotation and renewal is critical for maintaining long-term security and supporting cryptographic agility. While constrained devices may rely on integrated secure elements or lightweight HSMs for secure key storage and operations, the responsibility for orchestrating key rotation typically resides in

the application layer or external device management infrastructure.

Although the underlying cryptographic module may offer primitives to securely generate new key pairs, store fresh seeds, or delete obsolete keys, these capabilities must be integrated into the device's broader key management framework. This process is especially important in the context of PQC, where evolving research may lead to changes in recommended algorithms, parameters, and key management practices.

The security of PQC schemes continues to evolve, with potential risks arising from advances in post-quantum algorithms, cryptanalytic or implementation vulnerabilities. As a result, constrained devices should be designed to support flexible and updatable key management policies. This includes the ability to:

- * Rotate keys periodically to provide forward-secrecy,
- * Update algorithm choices or key sizes based on emerging security guidance,
- * Reconfigure cryptographic profile of the device via firmware updates.

5.2. Cryptographic Artifact Sizes for Post-Quantum Algorithms

The sizes of keys, ciphertexts, and signatures of post-quantum algorithms are generally larger than those of traditional cryptographic algorithms. This increase in size is a significant consideration for constrained devices, which often have limited memory and storage capacity. For example, the key sizes for ML-DSA and ML-KEM are larger than those of RSA or ECDSA, which can lead to increased memory usage and slower performance in constrained environments.

The following table provides the sizes of cryptographic artifacts associated with instantiations of ML-DSA, SLH-DSA, FN-DSA, and ML-KEM, aiming for "Level 1 security", as defined in [NISTSecurityLevels]. For comparison we also include the sizes of cryptographic artifacts associated with X25519 and Ed25519, which are traditional schemes widely used in constrained environments.

Algorithm	Type	Size (bytes)
ML-DSA-44	Public Key	1312
	Private Key	2560

	Signature	2420
SLH-DSA-SHA2-128s	Public Key	32
	Private Key	64
	Signature	7856
SLH-DSA-SHA2-128f	Public Key	32
	Private Key	64
	Signature	17088
FN-DSA-512	Public Key	897
	Private Key	1281
	Signature	666
ML-KEM-512	Public Key	800
	Private Key	1632
	Ciphertext	768
	Shared Secret	32
X25519	Public Key	32
	Private Key	32
	Shared Secret	32
Ed25519	Public Key	32
	Private Key	32
	Signature	64

Table 4

Corresponding sizes for higher security levels will typically be larger - see [FIPS203], [FIPS204], [FIPS205], and [FN-DSA] for sizes for all parameter sets.

6. Post-quantum Firmware Upgrades for Constrained Devices

Constrained devices deployed in the field require periodic firmware upgrades to patch security vulnerabilities, introduce new cryptographic algorithms, and improve overall functionality. However, the firmware upgrade process itself can become a critical attack vector if not designed to be post-quantum. If an adversary compromises the update mechanism, they could introduce malicious firmware, undermining all other security properties of the cryptographic modules. Therefore, ensuring a post-quantum firmware upgrade process is critical for the security of deployed constrained devices.

CRQCs pose an additional risk by breaking traditional digital signatures (e.g., RSA, ECDSA) used to authenticate firmware updates. If firmware verification relies on traditional signature algorithms, attackers could generate forged signatures in the future and distribute malicious updates.

6.1. Post-Quantum Firmware Authentication

To ensure the integrity and authenticity of firmware updates, constrained devices will have to adopt PQC digital signature schemes for code signing. These algorithms must provide long-term security, operate efficiently in low-resource environments, and be compatible with secure update mechanisms, such as the firmware update architecture for IoT described in [RFC9019].

[I-D.ietf-suit-mti] defines mandatory-to-implement cryptographic algorithms for IoT devices, and recommends use of HSS/LMS [RFC8554] to secure software devices.

Stateful hash-based signature schemes, such as HSS/LMS or the similar XMSS [RFC8391], are good candidates for signing firmware updates. Those schemes offer efficient verification times, making them more practical choices for constrained environments where performance and memory usage are key concerns. Their security is based on the security of the underlying hash function, which is well-understood. A major downside of stateful hash-based signatures is the requirement to keep track of which One-Time Signature (OTS) keys have been reused, since reuse of a single OTS key allows for signature forgeries. However, in the case of firmware updates, the OTS keys will be signing versioned updates, which may make state management easier. [I-D.ietf-pquip-hbs-state] discusses various strategies for a correct state and backup management for stateful hash-based signatures.

Other post-quantum signature algorithms may also be viable for firmware signing:

- * SLH-DSA, a stateless hash-based signature specified in [FIPS205], also has well-understood security based on the security of its underlying hash function, and additionally doesn't have the complexities associated with state management that HSS and XMSS have. However, signature generation and verification are comparatively slow, and signature sizes are generally larger than other post-quantum algorithms. SLH-DSA's suitability as a firmware signing algorithm will depend on the capabilities of the underlying hardware.
- * ML-DSA is a lattice-based signature algorithm specified in [FIPS204]. It is more performant than SLH-DSA, with significantly faster signing and verification times, as well as shorter signatures. This will make it possible to implement on a wider range of constrained devices. The mathematical problem underpinning ML-DSA, Module Learning With Errors (M-LWE), is believed to be a hard problem by the cryptographic community, and hence ML-DSA is believed to be secure. Cryptographers are more confident still in the security of hash-based signatures than M-LWE, so developers may wish to factor that in when choosing a firmware signing algorithm.

6.2. Hybrid Signature Approaches

To enable secure migration from traditional to post-quantum security, hybrid signature methods can be used for firmware authentication. Parallel signatures, where a traditional and a post-quantum signature are generated and attached separately, is simple to implement, requires minimal changes to existing signing, and aligns well with current secure boot and update architectures.

Other hybrid techniques, such as cross-linked signatures (where signatures cover each other's values), composite signatures (which combine multiple signatures into a single structured signature), or counter-signatures (where one signature signs over another) introduce more complexity and are not yet typical in resource-constrained firmware workflows.

7. Security Considerations

The security considerations for key management in constrained devices for PQC focus on the secure storage and handling of cryptographic seeds, which are used to derive private keys. Seeds must be protected with the same security measures as private keys, and key derivation should be efficient and secure within resource-constrained cryptographic module. Secure export and backup mechanisms for seeds are essential to ensure recovery in case of hardware failure, but these processes must be encrypted and protected from unauthorized access.

7.1. Side Channel Protection

Side-channel attacks exploit physical leaks during cryptographic operations, such as timing information, power consumption, electromagnetic emissions, or other physical characteristics, to extract sensitive data like private keys or seeds. Given the sensitivity of the seed and private key in PQC key generation, it is critical to consider side-channel protection in cryptographic module design. While side-channel attacks remain an active research topic, their significance in secure hardware design cannot be understated. Cryptographic modules must incorporate strong countermeasures against side-channel vulnerabilities to prevent attackers from gaining insights into secret data during cryptographic operations.

Acknowledgments

Thanks to Jean-Pierre Fiset, Richard Kettlewell, Mike Ounsworth, and Aritra Banerjee for the detailed review.

References

Normative References

- [RFC3394] Schaad, J. and R. Housley, "Advanced Encryption Standard (AES) Key Wrap Algorithm", RFC 3394, DOI 10.17487/RFC3394, September 2002, <<https://www.rfc-editor.org/rfc/rfc3394>>.
- [RFC9019] Moran, B., Tschofenig, H., Brown, D., and M. Meriac, "A Firmware Update Architecture for Internet of Things", RFC 9019, DOI 10.17487/RFC9019, April 2021, <<https://www.rfc-editor.org/rfc/rfc9019>>.

Informative References

- [BIND] Schmiege, S., "Unbindable Kemmy Schmidt: ML-KEM is neither MAL-BIND-K-CT nor MAL-BIND-K-PK", April 2024, <<https://eprint.iacr.org/2024/523.pdf>>.
- [BosRS22] Bos, J., Renes, J., and A. Sprenkels, "Dilithium for Memory Constrained Devices", December 2022, <<https://eprint.iacr.org/2022/323.pdf>>.
- [FIPS203] "Module-lattice-based key-encapsulation mechanism standard", National Institute of Standards and Technology (U.S.), DOI 10.6028/nist.fips.203, August 2024, <<https://doi.org/10.6028/nist.fips.203>>.
- [FIPS204] "Module-lattice-based digital signature standard", National Institute of Standards and Technology (U.S.), DOI 10.6028/nist.fips.204, August 2024, <<https://doi.org/10.6028/nist.fips.204>>.
- [FIPS205] "Stateless hash-based digital signature standard", National Institute of Standards and Technology (U.S.), DOI 10.6028/nist.fips.205, August 2024, <<https://doi.org/10.6028/nist.fips.205>>.
- [FN-DSA] Fouque, P.-A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., and Z. Zhang, "Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU", October 2020, <<https://falcon-sign.info/falcon.pdf>>.
- [HQC] Gaborit et al, "Hamming Quasi-Cyclic (HQC)", August 2025, <https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf>.
- [I-D.ietf-pquip-hbs-state]
- Wiggers, T., Bashiri, K., Kテカlbl, S., Goodman, J., and S. Kousidis, "Hash-based Signatures: State and Backup Management", Work in Progress, Internet-Draft, draft-ietf-pquip-hbs-state-02, 16 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pquip-hbs-state-02>>.
- [I-D.ietf-suit-mti]
- Moran, B., Rテクningstad, O., and A. Tsukamoto, "Cryptographic Algorithms for Internet of Things (IoT) Devices", Work in Progress, Internet-Draft, draft-ietf-suit-mti-23, 22 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-suit-mti-23>>.

- [IEEE-802.1AR] "IEEE Standard for Local and Metropolitan Area Networks - Secure Device Identity", IEEE, DOI 10.1109/ieeestd.2018.8423794, ISBN ["9781504450195"], July 2018, <<https://doi.org/10.1109/ieeestd.2018.8423794>>.
- [ISO19790] ISO, "Information security, cybersecurity, and privacy protection 信息安全 Security requirements for cryptographic modules", February 2025, <<https://www.iso.org/standard/82423.html>>.
- [Li32] Bai, S., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., and D. Stehlé, "CRYSTALS-Dilithium: Algorithm Specifications and Supporting Documentation (Version 3.1)", February 2021, <<https://pq-crystals.org/dilithium/data/dilithium-specification-round3-20210208.pdf>>.
- [Lyu09] Lyubashevsky, V., "Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures", Springer Berlin Heidelberg, Lecture Notes in Computer Science pp. 598-616, DOI 10.1007/978-3-642-10366-7_35, ISBN ["9783642103650", "9783642103667"], 2009, <https://doi.org/10.1007/978-3-642-10366-7_35>.
- [NISTSecurityLevels] NIST, "Post-Quantum Cryptography: Security (Evaluation Criteria)", January 2017, <[https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-\(evaluation-criteria\)](https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/evaluation-criteria/security-(evaluation-criteria))>.
- [REC-KEM] Alagic, G., Barker, E., Chen, L., Moody, D., Robinson, A., Silberg, H., and N. Waller, "Recommendations for key-encapsulation mechanisms", National Institute of Standards and Technology (U.S.), DOI 10.6028/nist.sp.800-227, September 2025, <<https://doi.org/10.6028/nist.sp.800-227>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/rfc/rfc5280>>.
- [RFC8391] Huelensing, 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/rfc/rfc8391>>.

- [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/rfc/rfc8554>>.
- [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/rfc/rfc9881>>.
- [Stream-SPHINCS] Niederhagen, R., Roth, J., and J. W̄lde, "Streaming SPHINCS+ for Embedded Devices using the Example of TPMs", August 2021, <<https://eprint.iacr.org/2021/1072.pdf>>.

Authors' Addresses

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

Dan Wing
Citrix
United States of America
Email: danwing@gmail.com

Ben Salter
UK National Cyber Security Centre
Email: ben.s3@ncsc.gov.uk

Kris Kwiatkowski
PQShield
Email: kris@amongbytes.com