

Post-Quantum Authentication in Enterprise and Mobile Networks

Tirumaleswar Reddy.K

Nokia

k.tirumaleswar_reddy@nokia.com

Hannes Tschofenig

University of the Bundeswehr Munich

hannes.tschofenig@unibw.de

Position paper submitted to the IAB Workshop on Accelerating the Deployment of Post-Quantum Authentication (PQWS), Prague, 11–12 October 2026

Introduction

Most work on post-quantum (PQ) authentication has focused on the Web. Enterprise and mobile networks have different deployment constraints, and those constraints change what determines the pace of migration. In these networks, migration has a long, largely sequential critical path. Cryptographic modules must first support the new algorithms, products must integrate those modules, and the resulting implementations need to be validated or certified before deployment. Only then can new credentials be issued and rolled out across manufacturer and operator certificate hierarchies. The migration therefore cannot complete faster than the slowest of these steps.

This paper makes five observations:

1. **Manufacturing and trust-anchor decisions create migration gates that protocols cannot remove.**
2. **Post-quantum readiness varies by device class.**
3. **The credential scheme affects long-term crypto-agility.**
4. **Large PQ credentials expose transport reliability problems, not merely latency problems.**
5. **The security benefit of PQ authentication arrives only when traditional authentication can be rejected.**

This paper does not propose a particular algorithm or migration approach. Its main observation is that, for enterprise and mobile networks, the dominant barriers to post-quantum authentication are largely outside the authentication protocols themselves: hardware capabilities, product cycles, certification and validation, PKI dependencies, and coordination across organisational boundaries. Many of the decisions that determine the eventual migration timeline are being made now, as equipment is designed and purchased, rather than later when post-quantum deployment begins.

1 These networks are not the Web

Migrating away from a cryptographic algorithm requires two things: the authority to stop accepting it, and the ability to update every system that still depends on it. The Web and enterprise or mobile networks differ mainly in which of these is harder.

Updating deployed devices in enterprise and mobile networks depends on the vendor producing and qualifying an update first, then the operator rolling it out incrementally.

2 Two gates that no protocol change can remove

Many devices in enterprise and mobile networks receive their initial identity from the manufacturer, often before the eventual operator is known and years before the device is commissioned. In IEEE 802.1AR [2], the manufacturer provisions an Initial Device Identifier (IDevID), while the operational network may later issue a Locally Significant Device Identifier (LDevID). In 3GPP deployments, see Section 9 of [3], an equipment vendor may provision a vendor certificate and the operator later issue an operator certificate through its own PKI. The bootstrapping mechanisms differ, but the dependency is the same: a manufacturer-provisioned identity is used to establish enough trust for the operator to provision a local identity.

This creates two migration gates that cannot be removed by changing the authentication protocol.

- **The manufacturing gate.** The algorithm used in the manufacturer-installed identity is determined when the device is produced. A device manufactured today with a traditional credential will continue to rely on that credential whenever it bootstraps, unless the credential itself can later be replaced. Upgrading the operator's PKI does not change this dependency because the initial credential was issued by another organisation, potentially many years earlier. For equipment with a service life of ten or fifteen years, decisions made on the production line today can determine which devices will be able to participate in a future post-quantum migration.
- **The trust-anchor gate.** Trust must be migrated in both directions. Before an operator can authenticate a device using a post-quantum manufacturer credential, its enrolment infrastructure must trust the corresponding manufacturer CA. Before the device can authenticate a post-quantum enrolment server, it must trust the corresponding operator CA. These changes have to be sequenced carefully: if one side presents a certificate chain that the other does not yet recognise, enrolment fails. Some constrained devices have no way to update their trust anchor at all. Others can be updated, but if the update happens in the wrong order, the device can no longer validate its peer, and remote equipment with no on-site operator may need a physical visit to recover.

The PKI structure adds further dependencies. Manufacturers and operators may maintain multiple issuing CAs, regional intermediates, product-specific hierarchies, or separate PKIs for different classes of equipment. Each affected CA certificate and trust relationship must migrate while the existing hierarchy remains usable. PKI migration is therefore a multi-stage deployment problem, not merely a two-party protocol upgrade.

3 The choice of scheme determines whether a second post-quantum family remains available

Three approaches are commonly considered for post-quantum authentication credentials. They are usually compared in terms of credential size, operational complexity, and compatibility with existing certificate-processing and validation code. Another important difference is the degree of algorithmic independence they provide if a second post-quantum signature family is needed.

- **Composite credentials** combine ML-DSA [4] with a traditional signature algorithm in a single credential. The credential remains secure as long as at least one component remains secure. It presents a single certificate chain and validation path, and the verifier processes the composite signature as a whole.
- **Dual certificates** keep the traditional and post-quantum credentials separate and rely on the protocol or application to coordinate their use. This increases operational complexity, but the credentials remain independent and their pairing is not fixed by a composite algorithm specification.
- **Post-quantum-only credentials** remove the traditional component entirely. This avoids the complexity of a hybrid construction, but the security of the credential then depends entirely on the selected post-quantum scheme.

Current Composite ML-DSA combinations pair ML-DSA with a traditional signature algorithm such as RSA, ECDSA, Ed25519, or Ed448. They therefore do not provide an independent fallback

Table 1: Deployment constraints in the Web PKI and enterprise and mobile networks.

	Web PKI	Enterprise and mobile networks
Relying parties	Browsers, operating systems, and servers that are typically updated within days or weeks	Base stations, gateways, network functions, CPE, industrial equipment that may remain deployed for 5 to 15 years
Trust anchors	Managed and distributed through browser and operating-system root (store) programs	Often installed at manufacture or commissioning. Migration must be carefully sequenced to avoid lockout
Certificate issuance	Automated, for example through ACME, with relatively short certificate lifetimes	Often based on enrolment protocols such as CMP [1] and EST, sometimes with manual operational steps. Credentials typically have longer lifetimes
Transport	Primarily TLS over TCP or QUIC, where packet loss is handled by the transport	Also includes IKEv2 and EAP, where fragmentation can turn packet loss into retransmission of an entire authentication exchange
Key storage	Often software-based, with HSMs commonly used at certification authorities and other high-value endpoints	Frequently based on TPMs, TEEs, or secure elements whose memory and capabilities were sized for RSA and ECC rather than PQ signing
Disabling old algorithms	Browser vendors can disable algorithms centrally, although deployment still has to reach enough of the ecosystem first	Requires all equipment that still depends on the old algorithm to be upgraded, reconfigured, or retired
Coordination	Browser and operating-system vendors, CAs, and the broader Web PKI ecosystem	Equipment vendors, operators, enterprises, industry bodies, certification laboratories, and, in some sectors, national regulators

to a second post-quantum family. The traditional component protects the credential if ML-DSA is found to be weak before traditional algorithms become quantum-vulnerable. It does not provide the same protection if ML-DSA is found to be weak afterwards, because the traditional component would then no longer provide an independent security assumption.

One point is worth clarifying, since the opposite is often assumed. A composite credential does not require both component private keys to reside in the same cryptographic module; the limitation is which algorithm combinations are defined and deployed.

Artefact signing already requires an independent choice

Firmware and software signing have requirements that differ from interactive authentication. The Composite ML-DSA specification [5] notes that Composite ML-DSA is not SUF-CMA secure against quantum adversaries and advises against its use where that property is required.

Standalone post-quantum schemes are therefore relevant for artefact signing. Hash-based signatures are attractive because they provide a security basis independent of ML-DSA and are well suited to long-lived artefacts. A handshake signature is relevant only to an authentication exchange, whereas a firmware signature may be verified throughout the lifetime of a device, potentially by immutable or difficult-to-update boot loader code.

Hash-based schemes involve their own trade-offs. Stateful schemes such as HSS/LMS [6] and XMSS [7] offer relatively fast verification and moderate signature sizes, but require careful state management; this is discussed in “Hash-based Signatures: State and Backup Management” [8]. Stateless SLH-DSA avoids state management at the cost of larger signatures and slower operations. NIST’s draft SP 800-230 [9] explores additional SLH-DSA parameter sets for software, firmware, and certificate signing.

Deployments that use a standalone post-quantum scheme for long-lived artefacts also need a corresponding signing hierarchy and trust anchor. Adopting composite credentials for authentication therefore does not eliminate the need for an independent post-quantum signing hierarchy elsewhere in the system.

4 Some constraints are firmware, and some are silicon

Hardware limitations affecting post-quantum authentication fall into two categories with different time horizons.

The first is **algorithm support**. TPM 2.0 Library Specification version 1.85 [10], for example, adds ML-KEM and ML-DSA [4] but not SLH-DSA [11]. TCG’s PC Client Platform TPM Profile version 1.07 [12] distinguishes between PQC-ready and PQC-upgradable devices, reflecting that cryptographic capabilities can sometimes be added after deployment. Adding support still de-

pends on the specification, product, validation, and deployment cycle.

The second is more fundamental: **whether the hardware has sufficient resources to perform the operation at all**. ML-DSA signing and verification have asymmetric resource requirements: signing requires substantially more working memory than verification, and the polynomial arithmetic involved also costs more CPU cycles per operation. Many secure elements and lightweight HSMs were designed around the memory and processing budgets of RSA and ECC and may not have sufficient resources for post-quantum signing. The RAM and flash budgets available in such devices are discussed in “Adapting Constrained Devices for Post-Quantum Cryptography” [13]. Firmware can add an algorithm, but it cannot add the memory or processing power the hardware does not have.

Readiness assessments should therefore distinguish by module, not by network type. A base station or gateway may have ample general-purpose compute, but its private key is held in a TPM, TEE, or secure element, and that module’s own memory budget governs whether it can sign. For those modules, the answer was fixed when the silicon was designed.

The required operations also matter. In public mobile networks [14], constrained devices typically hold SIM credentials and authenticate using symmetric-key methods, such as EAP-AKA. These devices still need to verify signed firmware, but signature generation may not be required.

5 Large credentials can cause failures, not just latency

On the Web, a larger certificate chain mostly costs latency, since TCP or QUIC can recover an individual lost packet on its own. In enterprise and mobile networks, the same chain costs latency too, but it can also cause the whole exchange to fail, because IKEv2 and EAP cannot recover a single lost fragment.

IKEv2 [15] commonly runs over UDP and cannot rely on the transport layer to recover lost packets. RFC 7383 [16] provides IKEv2 fragmentation, but if any fragment is lost, the receiver cannot reassemble the original message. IKEv2 does not selectively retransmit only the missing fragment; recovery requires retransmission of the fragmented message after the retransmission timer expires. As authentication messages grow, the number of fragments increases and so does the probability that packet loss delays or prevents completion of the exchange.

TLS-based EAP methods face a related problem. RFC 9191 [17] documents difficulties caused by large certificates and long certificate chains, including practical limits on the amount of certificate data that existing EAP implementations can carry successfully. Post-quantum certificates make these limits more

relevant because both signatures and public keys are substantially larger than their traditional counterparts.

Work to improve the handling of large authentication exchanges is ongoing in the IPSECME [18] and EMU [19] working groups.

6 Certificate formats that reduce credential size are being explored

Merkle Tree Certificates (MTCs) [20] reduce certificate size by replacing an individual signature on each certificate with a compact inclusion proof relative to a signed Merkle tree. For post-quantum authentication, this can substantially reduce the amount of certificate data transmitted compared with a conventional chain.

The design originated in the Web PKI and is being standardised in the PLANTS working group. Work is also exploring its use in private PKI deployments, described in “Merkle Tree Certificates Deployment Use Cases” [21]. The benefit depends on the use case: reducing certificate size matters directly for protocols such as TLS, IKEv2, or EAP, where credentials are transmitted during authentication, but much less for firmware or software signing, where a few additional kilobytes accompany an artefact that may already be many megabytes in size.

The largest size reductions require relying parties to have access to sufficiently recent Merkle tree information. In the Web PKI, such information can be distributed independently of the handshake. Whether the same model works well for enterprise, mobile, and constrained-device deployments depends on how reliably those devices can obtain and maintain the required tree state.

MTCs therefore merit evaluation for these environments, particularly where certificate size drives the fragmentation problems described in Section 5.

7 The security benefit only arrives when the traditional path is retired

Deploying post-quantum credentials does not by itself provide effective protection against an active downgrade attack while the verifier still accepts a traditional authentication path. An attacker able to break the traditional scheme can force or present that path instead. Recent analysis of the Web’s post-quantum transition reaches the same conclusion [22]: meaningful post-quantum authentication requires the verifier to stop accepting the traditional alternative.

Enterprise and mobile networks may appear to have an advantage because a single organisation can often define the security policy. The harder problem is deployment reach. A network operator cannot disable the traditional path until every device that still depends on it has been upgraded, reconfigured, or retired.

8 Suggestions for the workshop

The workshop should treat post-quantum authentication as a life-cycle and infrastructure migration, not only as a protocol transition. In particular, migration guidance should:

- distinguish protocol issues from hardware, PKI, certification, and deployment constraints,
- account explicitly for manufacturer-installed credentials and trust-anchor migration,
- assess signing and verification capabilities separately, recognising that some constrained devices use asymmetric cryptog-

raphy only to verify, for example checking a signed firmware artefact, and never need to sign,

- evaluate credential-size reduction alongside improvements to transport robustness, and
- define migration completion in terms of when the traditional authentication path can be retired.

References

- [1] Brockhaus, H., von Oheimb, D., Ounsworth, M., Gray, J., “Internet X.509 Public Key Infrastructure – Certificate Management Protocol (CMP)”, RFC 9810.
- [2] IEEE, “Secure Device Identity”, IEEE Std 802.1AR-2018, 2018.
- [3] 3GPP TS 33.310, “Network Domain Security (NDS); Authentication Framework (AF)”.
- [4] NIST, “Module-Lattice-Based Digital Signature Standard”, FIPS 204, August 2024.
- [5] Ounsworth, M., Gray, J., Pala, M., Klaussner, J., Fluhrer, S., “Composite ML-DSA for use in X.509 Public Key Infrastructure”, Internet-Draft, draft-ietf-lamps-pq-composite-sigs (work in progress).
- [6] McGrew, D., Curcio, M., Fluhrer, S., “Leighton-Micali Hash-Based Signatures”, RFC 8554.
- [7] Huelsing, A., Butin, D., Gazdag, S., Rijneveld, J., Mohaisen, A., “XMSS: eXtended Merkle Signature Scheme”, RFC 8391.
- [8] Wiggers, T., Bashiri, K., Koelbl, S., Goodman, J., Kousidis, S., “Hash-based Signatures: State and Backup Management”, Internet-Draft, draft-ietf-pquip-hbs-state (work in progress), August 2026.
- [9] NIST, “Additional SLH-DSA Parameter Sets for Limited-Signature Use Cases”, SP 800-230 (initial public draft), April 2026.
- [10] Trusted Computing Group, “TPM 2.0 Library Specification”, version 1.85, March 2026.
- [11] NIST, “Stateless Hash-Based Digital Signature Standard”, FIPS 205, August 2024.
- [12] Trusted Computing Group, “PC Client Platform TPM Profile Specification”, version 1.07.
- [13] Reddy, K. T., Wing, D., Salter, B., Kwiatkowski, K., “Adapting Constrained Devices for Post-Quantum Cryptography”, Internet-Draft, draft-ietf-pquip-pqc-hsm-constrained (work in progress), June 2026.
- [14] 3GPP TS 33.501, “Security architecture and procedures for 5G System”.
- [15] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., Kivinen, T., “Internet Key Exchange Protocol Version 2 (IKEv2)”, RFC 7296.
- [16] Smyslov, V., “Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation”, RFC 7383.
- [17] Sethi, M., Mattsson, J., Turner, S., “Handling Large Certificates and Long Certificate Chains in TLS-Based EAP Methods”, RFC 9191.
- [18] Reddy, K. T., Smyslov, V., “Separate Transports for IKE and ESP”, Internet-Draft, draft-ietf-ipsecme-ikev2-reliable-transport (work in progress), August 2026.
- [19] Reddy, K. T., “Post-Quantum Enhancements to TLS-Based EAP Methods”, Internet-Draft, draft-ietf-emu-pqc-eap-tls (work in progress), April, 2026.
- [20] Benjamin, D., O’Brien, D., Westerbaan, B., Valenta, L., Valsorda, F., “Merkle Tree Certificates”, Internet-Draft, draft-ietf-plants-merkle-tree-certs (work in progress), July 2026.
- [21] Gray, J., Klaussner, J., Tindell, L., Mallaya, G., “Merkle Tree Certificates Deployment Use Cases”, Internet-Draft, draft-gray-plants-mtc-deploy-use-cases (work in progress), September 2026.
- [22] Benjamin, D., “Post-Quantum Authentication: Beating the S-Curve”, Security Area Open Meeting, IETF 126, July 2026.