PQC FOR AI AGENTS SERIES · PAPER #1

The Channel Question

Post-Quantum Substrate for AI Agent Communication

Mohamad Amin Hasbini

Independent researcher · Paris, France

Published June 17, 2026

ABSTRACT

AI agents now exchange consequential, long-lived data over MCP and A2A. Both delegate confidentiality to classical TLS, and neither mandates post-quantum key exchange, so every agent integration inherits its quantum posture by default. Under harvest-now-decrypt-later, that exposure is already accruing. This paper poses the channel question as an executive decision: who owns the quantum posture of an organization's agent traffic, the architecture that owns it on purpose, and the regulatory clock that is already running.

KEYWORDS: post-quantum cryptography, AI agent security, hybrid key exchange, X25519MLKEM768, crypto-agility, MCP, A2A, harvest-now-decrypt-later, TLS 1.3, DORA, NIS2, EU AI Act

---

CITE AS

Hasbini, M. A. (2026). *The Channel Question*: Post-Quantum Substrate for AI Agent Communication. PQC for AI Agents Series, Paper #1.

Available at https://mahasbini.org/papers/s2-01-the-channel-question/
PDF https://mahasbini.org/publications/papers/s2-01-the-channel-question.pdf

TL;DR

- **The gap.** MCP and A2A, two leading open protocols shaping agent communication, both delegate confidentiality to whatever TLS sits beneath them and mandate no post-quantum key exchange. Their specifications are detailed on identity and authorization and silent on the quantum posture of the channel.
- **The inheritance.** Whoever terminates an agent's TLS handshake, in most deployments a content-delivery network or a gateway, sets its quantum posture. Most organizations are inheriting classical key exchange by default, without anyone having decided it.
- **The clock.** Harvest-now-decrypt-later means data exchanged today that must stay confidential past the arrival of a cryptographically relevant quantum computer is already exposed. The required confidentiality lifetime of the data, not the arrival date of the computer, sets the exposure.
- **The answer.** Hybrid post-quantum key exchange (X25519MLKEM768) is standardizing now and already deployable. The algorithm is not the hard part. Owning the termination point and building crypto-agility into the channel is. The channel is one of three cryptographic surfaces an agent carries through the transition, alongside the identity it presents and the receipts it leaves, not a feature bolted onto one of them.
- **The scale wall.** Key exchange is the solvable half. Standardized post-quantum signatures run to kilobytes and do not aggregate the way classical schemes let machine identities attest at scale, and the largest non-human identity fleet in production has set its own migration horizon in the early 2030s. Agent fleets are converging on the same shape. Unsettled is an argument for agility, not for delay.
- **The frame.** For regulated financial workflows, this is already a board-level crypto-agility issue: the DORA technical standards impose a binding duty to keep cryptography changeable. The broader frameworks now endorse the direction (the OECD Recommendation names post-quantum cryptography; the EU AI Act sets an outcome-based resilience bar and names no cipher) but none operationalizes it, so the honest argument is precise, not maximal. Post-quantum migration is the defensive companion to Europe's quantum ambition, not its competitor.
- **The contribution.** The post-quantum-on-the-channel primitives already exist in the literature, and this paper does not claim them. What has drawn insufficient attention is the connective tissue: a full-stack treatment that maps authentication, authorization, communication, and audit as one substrate, grounds it in the obligations that will actually regulate it in Europe (DORA, NIS2), and demonstrates its feasibility on a live reference deployment. That cartography, not a first claim on any primitive, is the contribution.

An agent-to-agent exchange, 2026

A procurement agent at a European manufacturer negotiates a component order with a supplier's sales agent. The two never share a network. They speak over A2A, exchange signed agent cards, authenticate over OAuth, and settle terms that include unit pricing, volume commitments, and a delivery schedule that will not be public for eighteen months. TLS encrypts the exchange in transit, terminated at the supplier's content-delivery network. It works. Both sides log a clean transaction.

Neither side chose the cryptography that protected that channel. The procurement team chose an agent framework. The platform team chose a CDN. The CDN chose a TLS configuration. Whether the key exchange guarding eighteen months of commercial confidentiality was post-quantum or classical was not a decision anyone at either company made. It was inherited.

This is the channel question, and it is already in production.

> *AI agents now carry consequential, long-lived data over channels whose cryptographic posture no one is deliberately choosing. MCP and A2A, two leading open protocols shaping agent communication, both ride TLS, and neither mandates post-quantum key exchange. Under harvest-now-decrypt-later, every such channel accrues exposure today. The decision is not which algorithm. It is whether the organization owns its agents' quantum posture or inherits it by default.*

What MCP and A2A actually secure

Two leading open protocols shaping cross-vendor agent communication in 2026 are detailed about identity and authorization, and that is worth crediting before naming what they leave out.

MCP, the Model Context Protocol, treats an agent's tool server the way the web treats any protected service: the agent proves who it is and what it may touch, down to the fine grain of token handling. That part is mature and fully specified. What it does not specify is the cryptography of the connection underneath. Confidentiality is delegated to whatever TLS sits below. Post-quantum key exchange appears nowhere in the specification.

A2A, the Agent2Agent protocol, now hosted by the Linux Foundation, published its first stable release in 2026. It defines how one agent authenticates another across the standard enterprise schemes, and signs each agent's card so a receiver can trust it came from the claimed provider. These secure authentication and message-origin integrity, over classical primitives. As with MCP, confidentiality is delegated to TLS; the specification recommends modern TLS but names no key-exchange profile, post-quantum or otherwise.

The pattern is the same in both. The specifications answer who the agent is and what it may do. They assume the channel is confidential and do not say how. That assumption is exactly where the quantum question hides: not in the protocol, but in the transport the protocol takes for granted.

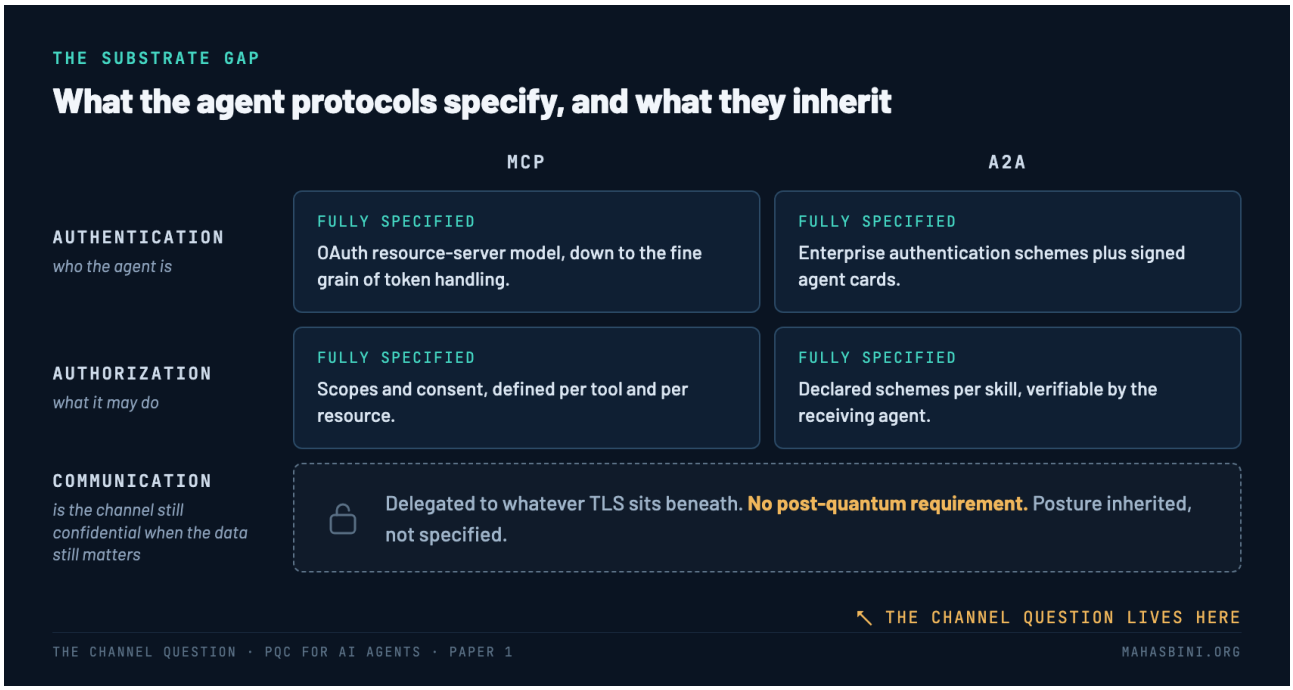

Both protocols are detailed on identity and authorization, and silent on the channel. The quantum posture is inherited, not specified.

An agent carries its cryptography in three places, and each has to survive the same transition: the identity it presents, the channel it speaks over, and the receipts it leaves behind. The protocols that carry agent traffic have specified the identity surface, and the authorization built on top of it, and left the channel to whatever TLS sits beneath. This paper takes the channel, the surface no one was assigned to own. Identity and receipts surface here too, and the next paper draws all three into a single model an organization can inventory and score.

The inheritance problem

Agent traffic terminates somewhere, and whoever terminates the TLS handshake chooses the key exchange. It terminates in one of three places, and each hands the decision to a different owner. At a **CDN**, the provider's roadmap decides. At a **gateway**, your platform team. At the **origin**, you do.

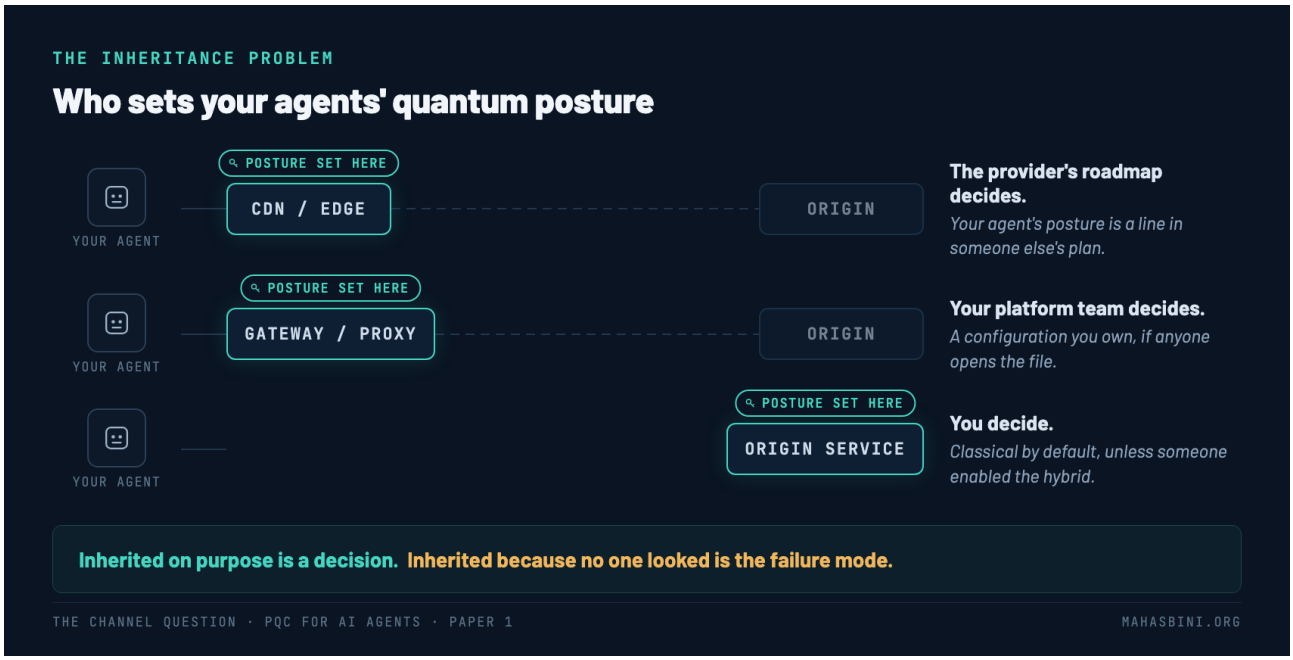

Whoever terminates the handshake sets the posture. Inherited on purpose is a decision; inherited because no one looked is the failure mode.

When the channel terminates at a CDN, the agent's quantum posture is the CDN's roadmap. That can be a defensible inheritance: a major edge provider that has enabled hybrid post-quantum key exchange by default, and has published a timeline toward post-quantum authentication as well, is a credible thing to inherit. When the channel terminates at the origin, the organization chooses, and most origin stacks still negotiate classical key exchange by default unless someone has enabled the hybrid.

The result, in most deployments, is that no one decided. The quantum posture of the channel fell out of an agent-framework choice and a hosting choice made for unrelated reasons. It was inherited, not engineered.

Inheriting is not automatically wrong. Inheriting a credible post-quantum roadmap from an edge provider may be the right call for a given channel. Inheriting classical-by-default because no one looked is not a call at all. The failure mode this paper is concerned with is not the choice to inherit. It is the absence of a decision.

Harvest-now-decrypt-later for agent traffic

Harvest-now-decrypt-later is the reason the channel question is present-tense rather than a problem for 2030. An adversary records encrypted traffic today and stores it, then decrypts it once a cryptographically relevant quantum computer is available. What determines exposure is not when that computer arrives. It is how long the data needs to stay confidential.

Agent channels concentrate exactly the data that harvesting targets. Cross-organization commercial terms. Diligence material. Health and financial records. Source code and model weights. And, increasingly, the delegation tokens and capability grants that authorize an agent to act on behalf of a user or another agent. That last category matters most: a harvested authorization exchange is not only a confidentiality loss, it is a map of the authority graph, who could instruct whom to do what.

The retention math is straightforward. Data that must stay confidential past the early-to-mid 2030s is already in scope today. The public migration timelines converge on that window. NIST's initial public draft IR 8547 sets out an expected transition under which quantum-vulnerable public-key algorithms are deprecated after 2030 and disallowed after 2035. An agent exchange recorded in 2026 that protects an eighteen-month secret, or a ten-year one, is exposed for the remainder of that lifetime if the channel was classical. The conversation does not get safer once it is over.

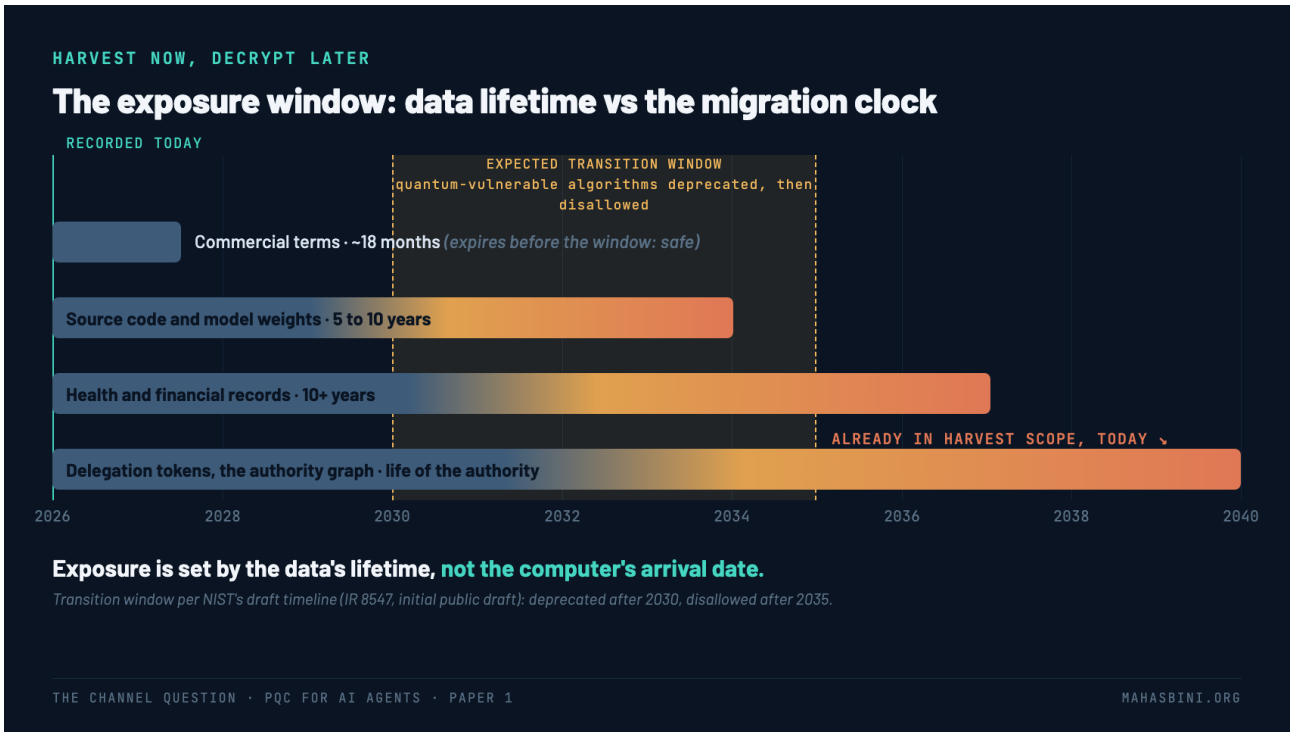

Exposure is set by the data's lifetime, not the computer's arrival date. Every lifetime that crosses the transition window is in harvest scope now.

The standard objection is that the symmetric cipher doing the bulk encryption, AES-256 in practice, is considered to withstand quantum attack, so the recorded traffic stays safe regardless. For agent channels the objection fails on its own terms. Every symmetric session key is negotiated over the channel's asymmetric key exchange. If that exchange is classical, an adversary who breaks it retroactively recovers the session keys, and the symmetric layer protects nothing it carried. Quantum-resistant symmetric cryptography negotiated over quantum-vulnerable key exchange inherits the weaker posture, not the stronger one. The harvest target is the handshake, not the cipher.

The substrate answer

The fix at the channel is not exotic. A hybrid handshake runs a classical exchange and a post-quantum one together, so the session stays confidential if either component holds. The profile the industry has converged on, X25519MLKEM768, is one step from formal publication as an internet standard and already ships in mainstream TLS. The algorithm is not the hard part.

So the algorithm is not the hard part. The hard part is architectural ownership: knowing where every agent channel terminates, deciding the quantum posture at that point on purpose, and building in crypto-agility so the choice can

change cleanly when an algorithm does. Crypto-agility as a program rather than a one-time swap is the subject of the next paper in this series.

Ownership also does not require re-architecting every agent. A component that cannot terminate hybrid key exchange itself, a legacy integration, a constrained runtime, a vendor binary, can be fronted by a proxy or sidecar that does. The wrapper brings the termination point under the organization's control and gives the legacy component a post-quantum outer layer while it waits for native support. The pattern is ordinary engineering, and it is available now.

This is the communication layer of an agent-identity substrate. Authentication establishes who the agent is. Authorization establishes what it may do, and proves it across trust boundaries. Communication establishes that the channel carrying both will still be confidential when the data still matters. The three are one substrate, not three products. The earlier papers in this body of work built the first two layers; this is the third.

It is deployable today, not in theory. As a working reference, Meetade, a browser-native conferencing system, terminates its own TLS 1.3 with X25519MLKEM768 on OpenSSL 3.5, on infrastructure the operator controls and can migrate. The point is not the product. The point is that owning the termination and running hybrid post-quantum key exchange on real, agent-adjacent traffic is a 2026 engineering task, not a 2030 aspiration.

The aggregation cliff

Key exchange is the solvable half of the channel question, and everything above says so. Signatures at machine scale are a different story, and an executive deciding agent-fleet architecture should know where that story stands.

Classical signature schemes gave machine identity an escape hatch at scale: aggregation. The clearest case in production is Ethereum, which runs on the order of a million validators, each signing on a cadence of minutes. Each of those classical signatures is 96 bytes, and they aggregate: a million collapse to the size of one. The standardized post-quantum signatures do not. The workhorse scheme runs to roughly 3.3 kilobytes and stays there, so the same fleet goes from one small attestation to a million large ones. Ethereum's published plan targets core protocol upgrades by 2029, with full migration extending into the early 2030s, and the engineers planning it frame the deadline in confidence terms rather than hardware terms: the transition has to land before users stop trusting the chain, not before a quantum computer ships.

The relevance to agents is direct. Agent fleets are converging on exactly this shape, thousands to millions of non-human identities attesting to one another at high frequency, and they will reach the same wall with none of the single-protocol coordination that lets Ethereum plan a decade-scale migration as one community. A fleet that signs constantly cannot simply swap a 96-byte signature for a 3-kilobyte one and absorb the bandwidth, storage, and verification budget unchanged.

This unsettledness is not a reason to wait. It is the second face of the channel question, and it is the strongest argument against treating post-quantum migration as a single swap. Confidentiality on the channel is deployable today with hybrid key exchange. Attestation at machine scale has no settled answer yet, which means the organizations that absorb whatever answer emerges cleanly will be the ones that built the agility to adopt it. Agility, not prediction, is the posture that survives an unsettled standard.

What the field is building, and what it misses

The intersection of post-quantum cryptography and AI agents is no longer empty, and that is worth stating plainly rather than claiming a frontier that has already been entered.

At the channel layer specifically, several independent groups have now placed NIST-standardized post-quantum key exchange and authentication directly on agent-to-agent communication. The nearest neighbor to this paper is MAGIQ, a post-quantum multi-agent governance system that specifies hybrid post-quantum TLS, X25519 with ML-KEM-768 for key exchange and ML-DSA-65 for transport authentication, signed session-authorization tokens enforced through hash-chain one-time tokens, and Merkle-aggregated token commitments that make every message attributable to its sending agent, with universal-composability proofs over the combined system. The Aegis Protocol combines ML-KEM and ML-DSA agent message protection with decentralized-identifier identity and a zero-knowledge policy proof, in a simulation that the authors state did not achieve stable library integration for the post-quantum primitives.

Two further efforts sit adjacent. One elevates post-quantum mutual authentication and key encapsulation to mandatory at the Model Context Protocol layer while explicitly deferring agent-to-agent coordination to future work. Another treats quantum-secure agent communication as a first-class lifecycle property but builds quantum-key-distribution and quantum-random-number components into its stack, layers that a software-deployable thesis deliberately excludes.

A separate and now-dense cluster addresses the audit trail: signing agent actions with post-quantum signatures so that the record of what an agent did survives the quantum transition. The strongest single instance is dedicated, formally modeled post-quantum audit evidence for regulated workloads under a harvest-now-forge-later threat model, where an adversary aims not to read but to forge the after-the-fact record once the cryptography breaks. Further agent-audit work signs records with post-quantum signatures toward the same end.

This is the right instinct applied to one layer at a time. The audit work makes the record of an agent’s action unforgeable after the fact, which matters. But the audit trail is downstream of the channel. A post-quantum receipt of an action whose authorizing exchange traveled over a classical channel proves the action while leaving the authority that triggered it harvestable. Signing the log does not protect the conversation that produced it.

The literature, for all its gaps, is ahead of the market. At a full-day European quantum-security industry summit in Paris in June 2026, with the continent’s post-quantum vendors, integrators, and regulated buyers in one program, AI agents surfaced in exactly three roles: as an attack tool that accelerates cryptanalysis, as an inventory line in a cryptographic asset register, and as a hardware-isolation problem at runtime. At no point in the program did an agent appear as a first-class identity whose channels need a deliberate quantum posture. Enterprise crypto-surface inventories presented on stage enumerated web PKI, code signing, machine-to-machine APIs, and microservices, and stopped before the agent layer. As of mid-2026, the gap this paper names is not yet written into the standards and barely surfaces in the rooms where European migration programs are bought and sold.

So the honest read of the field is this. The channel layer has been entered but not settled: the strongest treatments are preprints, several are simulation-only or carry trusted-provider and stateful-key assumptions, and one substitutes quantum hardware for the software-deployable path. The identity and authorization layers are crowded. The audit layer is crowded, including a post-quantum instance. Each contribution solves a layer. Across the verified literature, the connective tissue receives insufficient attention: grounding a post-quantum agent architecture in the obligations that will actually govern its deployment in Europe, the Digital Operational Resilience Act for financial entities and NIS2 for essential and important entities, and treating authentication,

authorization, communication, and audit as four co-equal engineering problems mapped to a multi-year crypto-agility migration rather than a single upgrade. The contribution of this series is not a first claim on any primitive. It is the cartography of the whole stack, with a live reference deployment, and the regulated-sector migration path that the primitives, individually, do not provide.

The regulatory and sovereignty frame

The deadline is not self-imposed, and the precise shape of it matters, because the field routinely overstates it.

For financial entities, the obligation is in force and binding. The DORA technical standards on ICT risk management require, in operative text, that the policy on cryptographic controls include provisions for updating or changing, where necessary, the cryptographic technology "on the basis of developments in cryptanalysis," so that it remains resilient against cyber threats; an entity unable to update or change the technology must instead adopt mitigation and monitoring measures. (Article 6(4) of the DORA technical standard, Commission Delegated Regulation (EU) 2024/1774, on encryption and cryptographic controls.) This is a crypto-agility obligation, not a post-quantum mandate, and that distinction is the point: the arrival of cryptographically relevant quantum computing is by construction a development in cryptanalysis, and the duty is the capability to change. The Level 1 DORA regulation has applied since 17 January 2025, and its technical standard since 2024, so the obligation is live now. An agent channel inside a regulated financial workflow is within its scope.

Outside finance the duty is broader but softer. NIS2, Directive (EU) 2022/2555, Article 21(2)(h), requires essential and important entities to maintain policies and procedures on the use of cryptography and, where appropriate, encryption. It does not, on its face, carry DORA's explicit duty to keep the technology changeable as cryptanalysis advances, and as a directive it binds through national transposition rather than directly. NIS2 is the wider expectation that covered entities govern their cryptography deliberately, which is difficult to reconcile with an inherited, undecided agent channel.

The broader instruments recognize the threat and prescribe no specific cryptography, which is exactly why the argument here is stated conservatively. The EU AI Act, in the Article 15 obligation on high-risk systems to be accurate, robust, and cyber-secure, names no cryptographic algorithm, technique, or standard. The words cryptography, encryption, quantum, and post-quantum do not appear in Article 15. A claim that the AI Act requires quantum-safe cryptography is therefore an overclaim and is not supportable from the text. The stronger and honest argument runs the other way: because Article 15 sets a technology-neutral, state-of-the-art resilience standard, the adequacy of cryptographic choices is judged against the prevailing state of the art, and a harvest-now-decrypt-later exposure is precisely the kind of exploitable weakness a diligent provider must weigh. Post-quantum cryptography is a defensible means of meeting that standard, never a named requirement of it. At the international level, the apex instrument goes further than acknowledgment. The OECD Recommendation on Quantum Technologies, adopted 28 May 2026 as the first intergovernmental standard in the field, includes among its operative principles "promoting quantum-resilient critical infrastructure to safeguard security and privacy, including by adopting post-quantum cryptography standards" (Principle 1.1(b)(ii); OECD/LEGAL/0508, a non-binding Council Recommendation). The direction is endorsed at the highest level. What a high-level principle does not do is locate the duty in any one organization's agent channels or put the migration on a clock. That operational specificity is what DORA supplies for financial entities and what the broader frameworks still leave open.

Naming this accurately is part of being useful to the people who answer to a supervisor. The binding crypto hook is DORA, sectorally. The deprecation clock is the NIST draft timeline. The technical means are standardizing now at the IETF. The high-level frameworks recognize the threat, and the cryptographic-transition layer is treated unevenly across them.

Finally, the sovereignty frame, stated carefully. Europe is investing seriously in quantum computing, and that investment is sound: the continent's quantum-hardware programs are a genuine strategic asset. The defensive companion deserves the same seriousness. The June 2026 European Technological Sovereignty Package sets out to protect the digital stack, across chips, cloud, AI, and open source, and its headline communication and strategy framing carry no post-quantum or cryptographic-transition commitment. Sovereign infrastructure that still rides quantum-vulnerable key exchange is sovereign in custody but not in confidentiality horizon. Post-quantum migration is that missing layer. It is not in competition with quantum computing. It is the half of the same agenda that keeps the rest sovereign through the transition, and agent communication is where the gap becomes concrete and immediate.

Five things to do Monday

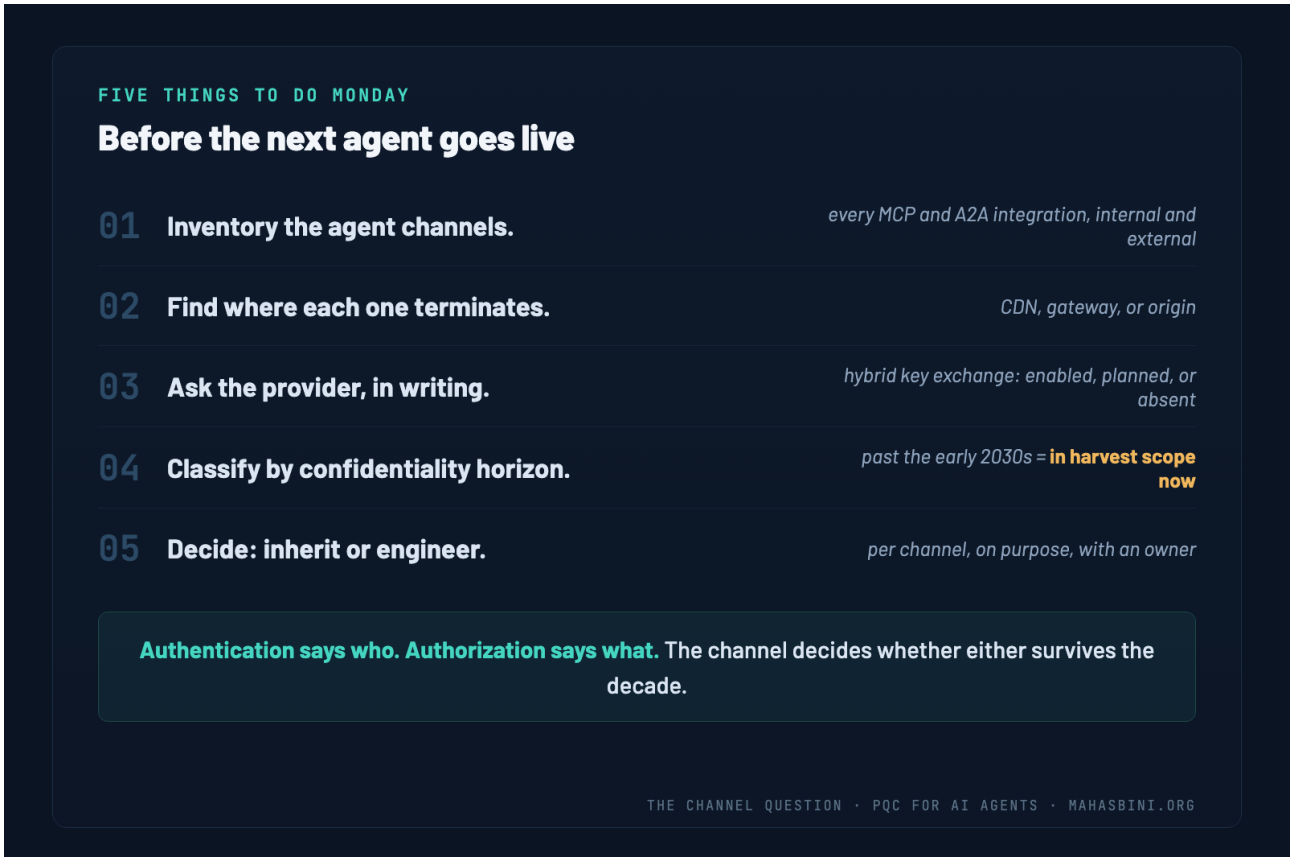

Authentication says who. Authorization says what. The channel decides whether either survives the decade.

1. **Inventory the agent channels.** List every place an AI agent communicates with a tool, a service, or another agent: MCP servers, A2A endpoints, internal and cross-organization.
2. **Identify who terminates each one.** For each channel, determine where the TLS handshake terminates, at a CDN, a gateway, or the origin. That party owns the quantum posture today.

3. **Ask the question out loud.** For CDN- or vendor-terminated channels, ask the provider, in writing, for its post-quantum key-exchange status and roadmap for your traffic. For origin-terminated channels, check whether hybrid post-quantum key exchange is enabled.
4. **Classify by confidentiality horizon.** Tag each channel by how long its data must stay confidential. Anything past the early 2030s is in harvest-now-decrypt-later scope now.
5. **Decide inherit or engineer, per channel, on purpose.** Inheriting a credible roadmap is a legitimate choice. Inheriting classical-by-default is not a choice, it is an oversight. Make it a decision, and record it.

Authentication says who. Authorization says what. The channel decides whether either survives the decade. That is the question to answer before the next agent goes live.

Paper 1 of the PQC for AI Agents series. It takes the first of an agent's three cryptographic surfaces, the channel. Next: identity and receipts, drawn together into a post-quantum agility model an organization can inventory and score, because crypto-agility is a program, not a one-time swap.

Appendix, for technical and cryptographic reviewers

Protocol references. MCP authorization (protocol revision 2025-11-25, the operative release) follows the OAuth 2.1 resource-server model and mandates RFC 9728 (Protected Resource Metadata), RFC 8707 (Resource Indicators), and PKCE with the S256 method; issuer validation per RFC 9207 appears in a later revision of the specification, not the operative one. A2A v1.0 defines declarative security schemes and signed agent cards (JWS per RFC 7515, payload canonicalized per RFC 8785). Both delegate transport confidentiality to TLS; A2A recommends TLS 1.3 or later; neither specifies a key-exchange profile.

Hybrid key exchange. draft-ietf-tls-ecdhe-mlkem ("Post-quantum hybrid ECDHE-MLKEM Key Agreement for TLSv1.3") defines three groups combining ML-KEM with a classical ECDHE exchange: X25519MLKEM768, SecP256r1MLKEM768, and SecP384r1MLKEM1024. For X25519MLKEM768 the negotiated shared secret concatenates the ML-KEM-768 and X25519 secrets (64 bytes: 32 + 32), and the client key_exchange concatenates the ML-KEM-768 encapsulation key and the X25519 ephemeral share (1216 bytes: 1184 + 32).

Signature scale. BLS12-381 signatures are 96 bytes and aggregate, n signatures to one. FIPS 204 ML-DSA-65 signatures are 3,309 bytes; ML-DSA-44, 2,420 bytes; neither aggregates. FIPS 205 SLH-DSA signatures run from 7,856 to 49,856 bytes depending on parameter set. Aggregate or compressed post-quantum signature constructions exist in the research literature and are not standardized.

Termination topologies. CDN/edge-terminated, gateway/proxy-terminated, origin-terminated. Each determines which party selects the negotiated key-exchange group, and therefore who owns the channel's quantum posture.

Identity-layer adjacency (for readers tracking the full stack). IETF WIMSE, OAuth transaction-token and agent work, and the W3C agent-identity community work are actively defining agent authentication and authorization. The IETF efforts do not yet address post-quantum cryptography; the W3C Agent Identity Registry Protocol Community Group has begun to scope post-quantum requirements for agent identity, though no finished specification yet exists. Identity is starting to move; the channel that carries agent traffic, the subject of this paper, is not yet part of that work. That separation is the gap this paper names.

Reference deployment. Meetade: TLS 1.3 with X25519MLKEM768 on OpenSSL 3.5, operator-controlled infrastructure, sovereign-migration-ready. AgentTrustLab is a working simulator for agent-to-agent authentication under post-quantum and zero-knowledge constraints. The PQC Scanner measures the post-quantum readiness of public TLS endpoints.

Selected sources. NIST FIPS 203 / 204 / 205 (2024); NIST IR 8547 (initial public draft, transition timeline; cite as expected/proposed, not final). Commission Delegated Regulation (EU) 2024/1774 (DORA RTS), Article 6(4), with the resilience cross-reference at Article 10(2)(a). EU AI Act (Regulation (EU) 2024/1689), Article 15. OECD Recommendation on Quantum Technologies (OECD/LEGAL/0508, 28 May 2026). EU coordinated implementation roadmap for the post-quantum transition (distinct from the June 2026 Technological Sovereignty Package). draft-ietf-tls-ecdhe-mlkem. MCP authorization specification; A2A specification. Prior art discussed in the field section: MAGIQ (arXiv:2605.06933); Aegis Protocol (arXiv:2508.19267); CA-MCPQ (IACR ePrint 2025/1790); PwC quantum-secure-communication preprint (arXiv:2603.15668); Kao (arXiv:2512.00110, sole author).

Aggregation-cliff sources: FIPS 204 (ML-DSA parameter sets); FIPS 205 (SLH-DSA parameter sets); pq.ethereum.org (Ethereum Foundation Protocol post-quantum hub: 2029 L1 target, full migration beyond, aggregation gap stated as “post-quantum signatures are larger and lack BLS’s native aggregation properties”); Justin Drake, “lean Ethereum”, Ethereum Foundation Blog, 31 July 2025; ethereum.org post-quantum cryptography page (updated April 2026); beaconcha.in validator chart (validator count); Edgington, Upgrading Ethereum (BLS 96-byte signatures, constant-size aggregates); Drake/Khovratovich/Kudinov/Wagner, IACR ePrint 2025/055 (research-stage hash-based aggregation). [Resolved 2026-06-12: public documents identified for the validator count, BLS figures, and migration horizon; only the confidence framing remains sourced to the June 2026 Paris summit, kept generic.]

About the author

Amin Hasbini is an AI and cybersecurity executive based in Paris. Former director of Kaspersky’s Global Research & Analysis Team (GReAT) for the Middle East, Turkey, and Africa. Twelve years, seventy countries of threat coverage. Invited contributor to the French Senate’s OPECST report on AI risks (2024). Current focus: post-quantum cryptography maturity and AI agent security inside regulated enterprises. mahasbini.org.