Transformativ IP

TECHNOLOGY LEADERSHIP BRIEFING

Q-PKI Mesh FAQ

For CTOs & CIOs

A plain-language guide to the strategic, technical, and compliance questions leadership teams typically ask before greenlighting a machine identity and data protection platform.

Table of Contents

Section titles are hyperlinked in full; questions with longer, more detailed answers are hyperlinked below each section for quick reference.

1. Strategic Overview

Q: What is Q-PKI Mesh, in one sentence?

It's a decentralized, software-only replacement for traditional Certificate Authority (CA) infrastructure — one that handles machine identity *and* data encryption in a single lightweight layer, built from the ground up on NIST post-quantum algorithms.

Q: Why should this be on our roadmap now, and not in two years?

Two clocks are running out at the same time.

1. First, quantum computing is close enough that "harvest now, decrypt later" attacks — where adversaries capture your encrypted traffic today so they can crack it once quantum computers are available — are already a real risk to any data with a multi-year shelf life.
2. Second, certificate lifespans are being compressed industry-wide (down to 200, 100, and even 47 days, with domain validation every 10 days). Legacy, server-heavy PKI wasn't built to issue and rotate credentials at that pace without breaking something.

Q: What's the core architectural bet here?

That trust doesn't need a central authority to be verifiable. Instead of every device phoning home to a CA server, database, and revocation list to prove who it is, each node validates identity locally using math — peer-to-peer, no round-trip required.

2. How It's Different From What We Have Today

Q: How is this different from what we already run (Legacy & Vendor-A-style platforms)?

Those platforms are excellent at what they were built for: centralized certificate lifecycle management. They issue and track X.509 certificates from a hierarchical Root/Intermediate CA structure, usually backed by physical or cloud HSMs. **Q-PKI Mesh takes a different architectural approach — it removes the central CA server, the sync databases, and the HSM dependency entirely, and pushes trust validation out to the edge node itself.**

Q: Does that mean we'd be ripping out our existing PKI?

Not necessarily. A common pattern in the source material is a **hybrid split**: keep your existing platform as the "control plane" for standard enterprise TLS and machine identity, where it already works well, and deploy Q-PKI Mesh at the "data plane" — inside edge devices, IoT/medical firmware, or anywhere payload-level encryption and offline validation matter most.

Q: What actually breaks in a traditional PKI model that this is designed to fix?

Three things, specifically:

1. **Single points of failure** — if the central CA server, sync database, or revocation service goes down, identity verification stalls everywhere downstream.
2. **The payload blind spot** — traditional PKI secures the pipe (TLS in transit), not the file once it lands. If an attacker gets past the handshake, the data behind it is often unprotected.

3. **The rotation-velocity wall** — hierarchical PKI with manual or semi-automated certificate management — wasn't designed to survive 2026's 200-day, 2027's 100-day, 2028's 47-day, and domain 10-day compliance cycles without outages.

3. Quantum & Cryptography

Q: Is this "quantum encryption" or "post-quantum cryptography"? Those aren't the same thing.

Correct, and the distinction matters for due diligence. **Quantum key generation** is about entropy — using true physical randomness to generate an unpredictable key. **Post-quantum cryptography (PQC)** is about the math — using new algorithms (like lattice-based cryptography) that a quantum computer can't break, even though they run on ordinary classical hardware. **Q-PKI Mesh is a PQC implementation: it runs standard NIST-approved PQC algorithms on conventional processors, and not quantum hardware.**

Q: Which specific NIST algorithms does Q-PKI Mesh use?

Three, all federally standardized:

- **ML-KEM (FIPS 203)** — key encapsulation, i.e., securely establishing shared secrets across a network. Implemented at ML-KEM-768 or ML-KEM-1024.
- **ML-DSA (FIPS 204)** — digital signatures that replace RSA/ECC for machine identity. Implemented at ML-DSA-65 or ML-DSA-87.
- **AES-256-GCM (NIST SP 800-38D)** — symmetric encryption for the actual bulk data payload, paired with SHA-384 hashing.

Q: Why not just use PQC for everything, including the bulk file encryption?

Performance. Asymmetric PQC math (ML-KEM, ML-DSA) is computationally heavier than symmetric encryption. The architecture uses PQC only for the handshake and identity proof, then hands bulk data off to hardware-accelerated AES-256-GCM (via AES-NI instructions most modern CPUs already have). That's how it avoids the throughput hit that pure-PQC file encryption would cause.

Q: We're not using any proprietary or "black box" cryptography here, right?

That's the Q-PKI Mesh design principle — every primitive maps to a published NIST standard (FIPS 203, FIPS 204, SP 800-38D). That matters directly to your compliance story, that is covered in Section 4 below.

4. Compliance & Audit

Q: Can we get FIPS 140-3 validation without buying racks of Hardware Security Modules?

FIPS 140-3 has always had a software validation path (Level 1/Level 2), not just the hardware path (Level 3/4) that traditional PKI vendors lean on. Because this platform is software-only, it's positioned to qualify under **FIPS 140-3 Level 1 or 2 (Software Module)**, typically by running inside an already FIPS-validated OS cryptographic boundary (e.g., RHEL FIPS mode). That's a meaningfully cheaper and faster audit path than standing up HSM infrastructure — but confirm current lab validation status before you rely on this for a live compliance deadline.

Q: What does this mean for FDA 21 CFR Part 11 (electronic records/signatures)?

The design fuses the identity signature (ML-DSA) directly into the data payload (AES-256-GCM) at the moment of encryption, rather than treating "who signed it" and "what was signed" as separate systems. This produces a self-contained, tamper-evident record — if the file is altered, the signature breaks immediately, without depending on a centralized log that an insider could tamper with.

Q: Where does traditional PKI still have the edge, compliance-wise?

Enterprise-grade audit trails. Established platforms have mature, centralized reporting showing exactly which machine held which identity and when certificates rotated — auditors are very familiar with that pattern. If you go hybrid, the recommended approach is feeding the decentralized platform's logs into your existing SIEM or logging engine so you get one unified audit report instead of two disconnected ones.

Q: Any compliance caveat we should walk into this with eyes open on?

Yes — FIPS only validates algorithms it has approved. If any custom or proprietary math sneaks into the implementation, it forfeits FIPS eligibility entirely. Any vendor conversation should confirm the deployment runs in a strict "FIPS-approved mode" using only the standard algorithms listed above — and we will provide validation certificates.

5. Lifecycle Automation

Q: What's the "Rapid Rotation Engine," practically speaking?

It's automation that rotates keys on a tiered schedule — 200-day and 100-day cycles for core enterprise infrastructure, 47-day for higher-risk internet-facing endpoints, and edge/medical firmware — without manual scripting or server restarts. New identities are provisioned and swapped in asynchronously, overlapping the old and the new, so no packets are dropped mid-rotation.

Q: And the 10-day domain validation requirement — how's that handled?

Via an integrated ACME-style handler that answers DNS/HTTP validation challenges automatically at the local node or router level, every 10 days, without a human touching a DNS record. If a domain can't validate (e.g., it's been compromised or rerouted), the node is designed to fail closed — isolating itself from the rest of the mesh rather than staying trusted on stale validation.

Q: Will any of this cause downtime during rotation?

Q-PKI Mesh is designed with zero-downtime rotation via asynchronous, overlapping key swaps in the background. That's a claim you can validate directly against your own uptime SLAs during a trial, rather than taking it at face value.

6. Security Posture

Q: How does this specifically address AI-driven credential theft — a threat legacy PKI wasn't designed for?

Legacy PKI has a structural weakness here: it trusts any valid key regardless of *where* or *how* it's being used. An autonomous AI agent that scrapes a static private key from memory or a config file can replay it from anywhere and be treated as legitimate. This platform's signatures instead bind identity to real-time context — hardware fingerprint, clock synchronization, and geolocation — at the moment of validation. A stolen key replayed from a different machine or location fails the check instantly, even though the cryptographic key itself is technically valid.

Q: Does this replace our identity provider (Okta, Active Directory)?

No — it's designed to sit downstream of your IdP, not replace it. The intended flow is: your existing IAM authenticates the user or machine via standard OAuth/OIDC or SAML, and this platform converts those permission attributes into localized, quantum-safe cryptographic tokens and edge-level certificates.

7. Business Case & Footprint

Q: What's the resource footprint on constrained devices — IoT, medical hardware, embedded firmware?

The stated footprint is under 500 KB for the software module and under 1 KB of runtime memory on constrained firmware, positioned specifically for hardware too small to run traditional PKI agents or ACME clients.

Q: What's the cost/TCO argument?

The core savings lever is eliminating the infrastructure traditional PKI requires: dedicated CA servers, sync databases, and physical or cloud HSMs. Materials cite a directional ~70% infrastructure savings and deployment timelines shrinking from months to days — treat these as vendor-stated targets to validate against your own environment during a pilot, not guaranteed outcomes.

Q: What does a realistic first step look like?

A phased rollout: pilot on a single silo or device class, expand to hybrid cloud, then move to full mesh deployment — rather than a rip-and-replace of existing infrastructure on day one.

Bottom Line for Leadership

The core trade-off to evaluate: you're weighing the operational simplicity and edge resilience of a decentralized, serverless model against the audit maturity and ecosystem familiarity of established centralized PKI. For most enterprises, the pragmatic path validated in this material is **hybrid** — keep centralized PKI where it's already trusted and well-audited, and deploy this platform where legacy PKI structurally can't reach: constrained IoT/medical firmware, offline/air-gapped edge sites, and anywhere payload-level (not just transport-level) encryption is the actual requirement.

Transformativ IP