IKP Crypto Agility & PQC Migration Session

Mitchell, privacymage — IKP Co-Chair Reconstruction

Session summary

This session represents the operational foundation for what's become BGIN's most externally visible 2026 workstream. The lead presentation (originally delivered at NIST's Crypto Agility Workshop) laid out the structural problem we've been circling: PQC migration in blockchain isn't a cryptographic problem — it's a governance trilemma between decentralization, security, and operational overhead, with misaligned stakeholder incentives at every vertex.

What landed

The session established that BGIN's role in the NEDO-funded PQC competition is as a *neutral evaluation forum* — not selecting algorithms, not setting chain roadmaps, but designing the criteria and multi-stakeholder process through which migration approaches are assessed. The three-lane framing (NEDO as sponsor, NIST as standards authority, BGIN as evaluation process) that was later formalized in the steering committee submission was seeded here. The 10-node global testbed proposal (3 Asia, 5 Europe, 5 North America) with Japanese government initial funding is a concrete deliverable.

The presentation covered the blockchain PQC transition landscape systematically: ECDSA key sizes (33 byte verification key, 77 byte signature) vs. Falcon (897 byte key, 666 byte signature) — illustrating the throughput-decentralization tradeoff. Migration plans were categorized into: (1) legacy secret with ZKP, (2) hybrid signature with account abstraction, (3) pure PQC migration, and (4) migration by centralization. Each makes different tradeoffs across the security-decentralization-overhead triangle.

The stakeholder incentive analysis identified six actor classes (users, developers/ideologists, miners, investors, node operators, collective interest) with different — and often conflicting — priorities during migration. The dormant asset dilemma was flagged: mandatory migration freezes assets (sacrificing decentralization), backward compatibility leaves them quantum-vulnerable (sacrificing security).

Gap analysis: no agreed definition of "PQC-secure blockchain," no comprehensive migration requirements, no incentive-aligned migration strategy, and migration scope varies wildly between chains (Bitcoin focuses on transaction signatures, Ethereum on consensus layer BLS signatures, neither adequately covers scaling solutions or smart contracts).

Overlap with active IKP work

1. Taxonomy of Harms (IKP-FASE)

The Google/EF/Stanford reference paper's on-spend / at-rest / on-setup attack classification now gives us a quantum-specific harm taxonomy layer to integrate. Their Ethereum vulnerability table (Account, Admin, Code, Consensus, Data Availability) is essentially a sectoral harm taxonomy that we should absorb into our framework. The

dormant asset challenge — 1.7M BTC in P2PK scripts with no migration path — is a governance harm category we haven't yet named but need to: *cryptographic obsolescence harm*.

2. Proof of Personhood & Agent Duality

The stakeholder incentive analysis highlights that migration friction is highest for end users. This connects to our PoP work: if identity credentials or agent delegation chains rely on vulnerable curves (secp256k1, BLS12-381), then personhood verification breaks under quantum attack. The reference paper's note that BLS12-381 needs a "somewhat larger CRQC" but remains accessible to first-generation machines means our PoP framework must treat PQC migration as a prerequisite, not an afterthought.

3. ZKP / SR 0011

The session's recognition that migration must cover "other primitives" beyond transaction signing — polynomial commitments, scaling solutions, smart contracts — validates the scope of our ZKP study. The reference paper's analysis of KZG commitment vulnerability and the on-setup attack vector against Ethereum's Data Availability Sampling is exactly the kind of infrastructure risk our ZKP work should contextualize.

4. Threat Intelligence / ISAC

The Google/EF/Stanford paper's "responsible disclosure" approach (ZK-proof-substantiated resource estimates without publishing attack circuits) is a disclosure methodology our threat intelligence sharing framework should study and potentially standardize. This is a novel contribution to the forensics-analytics distinction: verified threat claims without operational details.

Pending actions triggered

- **IKP PQC scope statement** (from steering committee submission, Section 3.4): Draft due late April–May 2026. This session's content provides the gap analysis. Should incorporate the identity/key management angle — specifically, how PQC migration intersects with key hierarchy security (HD wallets, extended public keys), agent delegation credentials, and personhood verification infrastructure.
- **9 April IKP call:** Align talking points for the 16 April NIST workshop. Core message: BGIN complements NIST, doesn't substitute. No FUD, no single-calendar predictions.
- **Taxonomy of Harms update:** Add "cryptographic obsolescence" as a harm category; integrate the on-spend / at-rest / on-setup attack type classification as a temporal dimension; map Ethereum's five vulnerability vectors as a case study.
- **ICBC side event (June, Brisbane):** This session's focus on L2/ZK/proving system PQC exposure makes the side event even more timely. Should ensure the event explicitly covers the ZK-rollup vulnerability chain described in the reference papers.

Compression seal: 🔑⚛️⚖️ — *Keys face quantum, governance decides the migration path.*