GDC 2026 — Meeting Report

Post-Quantum Cryptography Migration on DLTs:
Deployment, Agility & Governance

Blockchain Governance Initiative Network (BGIN)
IKP — IAM, Key Management and Privacy

3 September 2026 | 17:00–17:50 | Geneva

This record is written under the Chatham House Rule. Information may be used; other participants are not identified.

Agenda

- Opening: BGIN; PQC as deployment, agility, and governance for digital-asset stacks, not a new-algorithm contest.
- Scene-setting: quantum vulnerabilities in cryptocurrencies; ECDLP; error correction and interconnects; attack classes; migrate now.
- Competition: Japan-hosted, internationally open prize on PQC migration for blockchain networks; NIST schemes, not new primitives.
- Open discussion: slowing quantum; PQ verifiable credentials; signals of a step change; AI in cryptanalysis and resource estimates; industry roadmaps and Q-day as a probability.
- Close: BGIN Block 15 (Washington, D.C., October 2026); hybrid.

Session Description

Wallets and credentials rely on long-lived cryptographic trust. NIST-standardized post-quantum primitives exist. What still lacks an agreed playbook for public ledgers is how to stage a hybrid then a cutover when every node must agree on validity rules, how to keep verifying old signatures while funds remain live, how to rotate keys without stranding dormant accounts, and how to tell a genuine PQ-ready stack from a screenshot.

This 50-minute breakout held resource estimates and migration design in one room. It built on BGIN IKP crypto-agility work, a Google Quantum AI scene-setting on quantum vulnerabilities in cryptocurrencies, a Japan-hosted multi-year migration prize (BGIN coordinating evaluation; NIST remaining the primitive authority), and industry remarks on protocol roadmaps. Follow-on drafting continues at BGIN Block 15.

Speakers

The session was conducted under the Chatham House Rule. Floor interventions may be used; neither the identity nor the affiliation of other participants is recorded here.

On the programme: Mitchell Travers (BGIN), moderating; Adam (Google Quantum AI), scene-setting on quantum hardware progress, attack classes, and why ECDLP-based public-key cryptography is the exposed assumption; Shin'ichiro Matsuo (BGIN), announcing a Japan-hosted internationally open prize on PQC migration for blockchain networks; and Conor (Project 11), on protocol-community sentiment, hybrid then cutover, and DLT-specific knock-on costs.

Floor discussion included whether quantum development can be slowed, post-quantum recommendations for verifiable credentials, what signals a step change if qubit count is the wrong metric, AI-assisted factoring and cryptanalysis, public resource-estimate leaderboards, and how long protocol versus hardware roadmaps actually are.

Slides

Session slides were presented from a BGIN Beamer deck combining DLT-versus-TLS cutover frames, resource estimates versus migration/agility, the Japan prize lanes (METI/NEDO, BGIN, NIST), discussion questions, and a Block 15 call with session, event, and registration links.
Session slides (PDF): https://bgin-global.org/documents/meeting-reports/gdc26/GDC26_PQC_Migration_DLTs_Slides.pdf

Notes & Key Points

The assumption at the bottom of the stack is the failure point. Most deployed public-key cryptography, including on-chain and off-chain finance, rests on the elliptic-curve discrete logarithm problem being hard. A cryptographically relevant quantum computer makes that problem easy. The talk located that assumption under protocols, software, and markets: if it fails, the rest of the stack fails with it. Contagion paths already exist—stablecoins backed by short-term government debt; a planned NYSE-linked on-chain trading system; an EU pilot for cross-border educational credentials using on-chain identities.

Qubit count is the wrong progress metric. Google's Sycamore (53 qubits, 2019 supremacy demonstration) to Willow (105 qubits, 2024 error-correction demonstration) looks like a slow doubling if one only counts qubits. The five years were spent on capabilities that make larger devices functional. Quantum error correction was presented as a performance discontinuity: unprotected superconducting qubits lose information in about 100 microseconds; with error correction, lifetime scales exponentially in a distance parameter. A next discontinuity discussed was coherent interconnects between processors. Progress was described as escaping one scalability barrier then hitting the next, platform-specific and not Moore's-law linear. Compiling and resource reduction (illustrated as shrinking space-time volume for magic-state production) were treated as equally important and often less visible than hardware headlines.

Attack classes are not one race. In-flight attacks on a public mempool need a fast cryptographically relevant machine (minutes-scale window once a public key is revealed). At-rest

attacks can use slower machines against keys already on-chain. A third class, attributed to the recent *PRX Quantum* paper, separates exploit manufacture from exploit use: one party attacks public protocol parameters and can sell a classical exploit that others run from a laptop against many victims. The implied threat model is not only "who owns a quantum computer." Harm paths named included theft of on-chain assets, rewriting history on proof-of-work consensus, breaking confidentiality in systems such as Zcash, inflation attacks, collapsing stablecoin pegs, and forging cryptographic votes. The close of the scene-setting was that devices of that class do not exist yet; the way off the bleak path is to identify vulnerable protocols and replace them with post-quantum cryptography.

Do not wait for a grokable Q-day signal. From the floor: if progress is a step change, what should non-specialists watch? The reply was that papers and a few blogs help, but late entrants may not announce themselves, known labs may go quieter, and some state-run takeovers of private labs have already reduced public output. The recommended stance was to assume it will happen and start PQC work now, rather than wait for a visible "this is the moment."

Japan-hosted prize: migrate networks, do not invent a primitive. A multi-year, internationally open competition (Japanese government as host and funder; BGIN coordinating; path into a BGIN standard and then ISO, including TC 307 and potentially JTC 1) will score efficient methods to migrate blockchain networks to quantum-resistant signatures and advanced key management, using NIST-class schemes. It is not a new-PQC contest. Comparative targets named in the room were Bitcoin and Ethereum as two accounting models. Evaluation axes: theoretical security of the protocol, and implementation efficiency. Anyone worldwide may apply or join evaluation. This year: define the specification of a "PQC-ready" chain and completion conditions, including a one-day academic workshop alongside SSR (Security Standardization Research) in December. Next year: roughly nine to ten months to build submissions, application close around October, then about a year of global mathematical and engineering evaluation; winner targeted by March 2029. Block 15 is where an evaluation committee starts to be stood up.

Verifiable credentials already have a NIST-shaped answer. Asked for a PQ recommendation for VCs, the quantum speaker deferred; the chair noted ML-DSA as the relevant recommended scheme and that identity and asset stacks share the same cryptographic base, so agility lessons transfer.

AI accelerates both sides, and that is a forcing function for agility. There is no known campaign to slow quantum hardware; AI is being used to speed it. Inside a quantum lab, AI already writes high-performance classical control code, is hoped to help find error-correcting codes, circuits, and decoders, and has been used to merge proof techniques in algorithm papers. From the floor and from Project 11: AI cryptanalysis is already in the wild (including a recent hardware-wallet incident and a reported break of HAWK, a NIST PQC scheme). A one-shot jump from elliptic curves to ML-DSA was rejected: the next primitive can also break. Crypto agility—rotating from EC to ML-DSA to hash-based signatures to something not yet invented without a wholesale redesign—was presented as the actual design goal, because the industry should expect many such migrations over decades.

Public resource-estimate leaderboards need nuance. ECDSA.fail (and related "auto research") was described as an AI-led competition unpacking published resource estimates, with contributors who are often not quantum specialists. Asked whether a public curve toward zero is a policy-grade benchmark, the quantum speaker said much of the work is classical reversible compilation inside a quantum outer shell, accessible to computer scientists; whether better attacks use truly quantum primitives is an open question, and most published discrete-log / factoring estimates assume the classical-subroutine approach.

Industry is at "it is real; what now," not "it is fake." Project 11's 2024 pitch was often dismissed; a Google paper shifted sentiment. Ethereum has published a roadmap; other foundations are moving; Bitcoin coordination exists without a foundation model. The hard questions after "yes, migrate" are algorithm choice, throughput, block size, node sync, and social consensus. Minimum protocol-side timeline to a fully post-quantum chain was given as about two to three years on the aggressive end—not because the cryptography cannot be designed quickly, but because testing, AI cryptanalysis, governance, testnets, and then user-driven spends from address A to address B have undefined tails. Hybrid (classical and PQ side by side) then later removal of elliptic curves was the preferred path. Elliptic curves were stated as still fine to use today.

Q-day is a probability, not a date on a slide. The quantum speaker would not bet personal money on a cryptographically relevant machine by 2030, would bet around 2032 (on the order of 50–two-thirds), and would put 2029 at about 5% via a confluence of hardware and error-correction research that is not yet a matching architecture. For a cybersecurity owner, that 5% was treated as already demanding action. A same-day viral RSA-260 factorization was classical AI, not quantum; the market still asks "is this quantum?" every time. Sentiment risk for banks and institutions was named as a reason to assume a quantum computer soon and start the migration conversation, rather than wait for a clean announcement.

Session outcomes

- A shared problem statement: NIST primitives exist; the DLT gap is deployment, consensus-wide validity rules, long-lived verification, dormant funds, and comparable evidence of "PQ-ready," not a rival algorithm list.
- Scene-setting that qubit count, a single Q-day date, and "wait until they factor 21" are the wrong monitors; error correction, interconnects, compiler progress, and a 5% near-term tail already justify starting now.
- Attack distinctions the room used: in-flight versus at-rest versus parameter/exploit-manufacture; contagion into traditional finance via stablecoins and planned on-chain market infrastructure.
- Crypto agility as the design objective: hybrid then cutover; do not freeze a one-time EC-to-ML-DSA rewrite; expect repeated migrations, including under AI cryptanalysis.
- A Japan-hosted, internationally open migration prize (NIST schemes; Bitcoin/Ethereum as evaluation anchors; security plus efficiency; BGIN standard then ISO), with 2026 definition

work, an SSR-adjacent workshop in December, applications next year, and Block 15 to stand up evaluation.

- Industry inflection: protocol roadmaps are appearing; remaining work is architecture and social consensus, plus an unpredictable user-migration tail of two to three years at best on the protocol side.
- A named continuation: BGIN Block 15 (15–16 October 2026, Washington, D.C.; hybrid), including Keynote: PQC (15 October, 09:20–09:50) and IKP: PQC Migration (15 October, 09:50–11:20 and 11:30–13:00).

Event: https://bgin-global.org/events/20261015-block15
Registration: https://www.eventbrite.com/e/bgin-block15-tickets-1990274282957
Keynote: https://bgin-global.org/events/20261015-block15/sessions/1-2
Session: https://bgin-global.org/events/20261015-block15/sessions/1-4