Self-Authenticating Addresses for Post-Quantum Delay-Tolerant Networking

Verification in the window before a bulletin arrives

Martin Tanchev

STINT 2026, Leuven

Goce Delcev University, Stip

Key management is the acknowledged gap

- BPSec protects a bundle once keys are in place. Key management is out of scope.
- DTKA fills it: a trusted collective issues signed, self-contained bulletins, valid however late they arrive.
- Bulletins travel at the speed of the contact graph.

LunaNet LNIS V4 supports BPSec and defers key management and cross-provider trust to a security specification still under development.

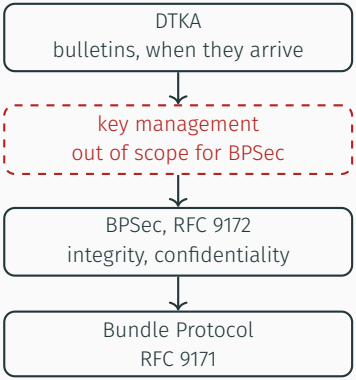

Post-quantum migration makes every bulletin heavier

- CNSA 2.0 mandates ML-KEM-1024 and ML-DSA-87.
- One node's bulletin entry: 88 B classical, 4184 B post-quantum. 48×.
- Heavier bulletins move more slowly through the same contacts.

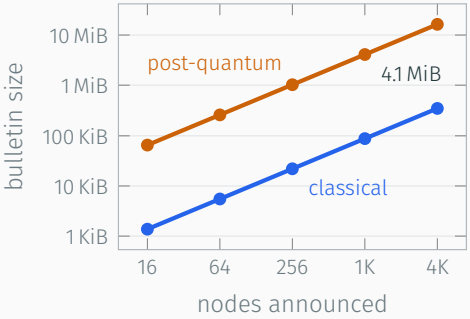

Entry = signing key + KEM key + metadata. Classical: Ed25519 + X25519. PQ: ML-DSA-87 + ML-KEM-1024. With ML-DSA-65: 3544 B, 40×.

The interval before a bulletin arrives gets wider under a migration that is already mandated.

Two rovers, two agencies, one encounter

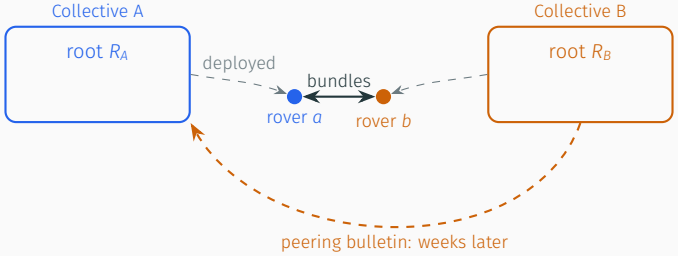

Their collectives agreed to federate *after* both rovers were deployed. The bulletin that would let either verify the other has not arrived. Discard, or accept unauthenticated? On a multi-week round trip, discarding costs a round trip, not a bundle.

The same window opens for an asset enrolled after the last bulletin, and at every key rotation. Three cases, the same band each time.

When is the window actually open?

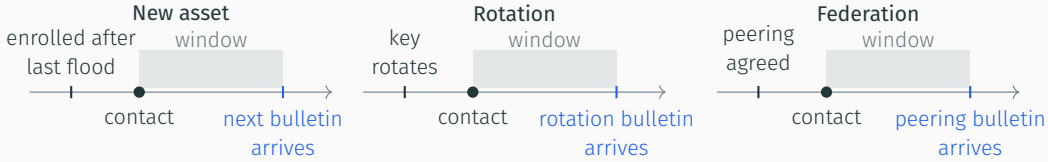

1. **New assets**, enrolled after the last bulletin propagation.
2. **Rotation**, which reopens the window for every peer.
3. **Federation** agreed after deployment. LunaNet's multi-provider case.

Not every contact. These three cases, and the same band each time. Within one well-run collective with a static population, pre-distribution covers almost everything.

An identifier the key material can prove

An identity begins with one public **inception record**: its first signing key K_0 and a hash commitment to its successor, $H(K_1)$. The **identifier**, call it Q, is the hash of that record: 32 bytes, fixed for life, provable from the record alone.

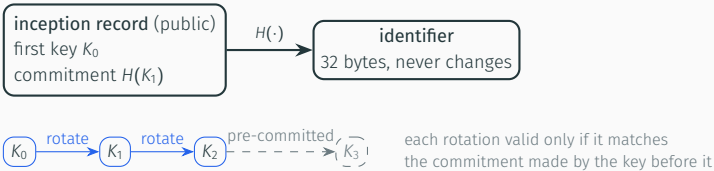

A thief with the current key cannot rotate the identity: the successor is pre-committed, generated offline, never exposed.

Hash-derived identifiers: SFS (1999), CGA (RFC 3972), HIP. Pre-rotation: KERI. The composition is mine.

The receiver's procedure, cheapest check first

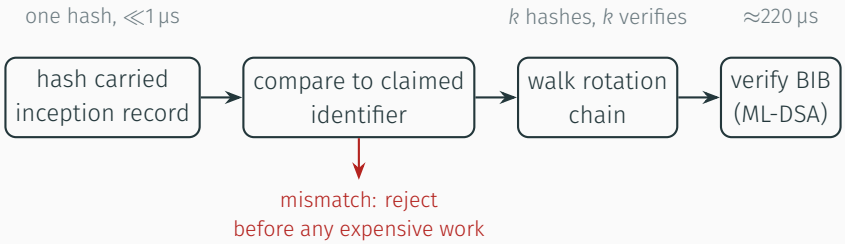

Everything checked is carried *in the bundle*: no lookup, no handshake, no reachable authority. The cheap first step is also the flood defense.

Timings measured on x86-64. Rad-hard flight processors run 10^2–$10^3\times$ slower, still far below bundle rates. Full compute table in Backup.

Provisional authentication: binding now, standing later

Binding: the bundle itself proves it came from the holder of Q. **Authorization:** a bulletin says Q is someone's asset.

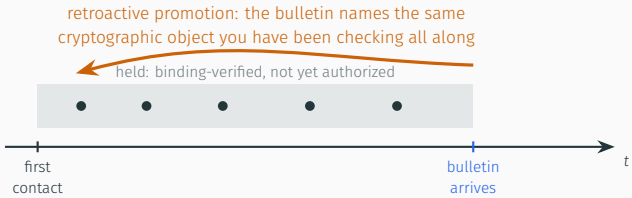

Inside the window the receiver has binding: every bundle under Q came from **one keyholder**, and a thief of the current key **cannot rotate Q away**. It does not have authorization. So bundles are *held*, not acted on. When the bulletin names Q, the whole held history promotes at once. No round trip.

Nothing offline can say Q is rover a. The alternative was discard, or accept blind.

Where it sits: additive, beneath DTKA

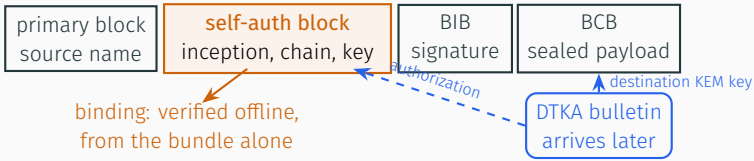

- DTKA stays the trust root. Its only new job: asserting “identifier *Q* is one of ours.”
- New on the wire: one extension block, one security context. Names and existing blocks untouched.

Binding from the bundle, now. Authorization from the bulletin, later.

Rules that keep authorization sound

- **Promote by receipt time.** A stolen key can backdate a creation timestamp. It cannot backdate my clock.
- **Admission is per link.** Identities are free to mint by design. Budgets attach to physical contacts, which are not.
- **No transitive trust.** Peering vouches one collective, depth one from my own root. Not inherited, and withdrawable.
- **Records are owned.** A member record belongs to the collective that vouched it. No peer can overwrite or revoke it.

Free minting is not a flaw. Gate identity creation behind an authority and binding is no longer verifiable offline, which closes the window this whole talk is about. For the same reason, a history-derived confidence score would be cheap to farm: the resource to price is the physical contact, not the observed history.

The contact plan behind the measurement

link	role	dense (lunar)	sparse (Mars)
ground A ↔ ground B	terrestrial	always	always
ground ↔ own orbiter	uplink	6 h every day	1 h every 2 d
orbiter ↔ own surface asset	relay	8 min every 6 h	8 min every 12 h
orbiter ↔ other agency's asset	cross-support	8 min every 12 h	8 min every 2 d
rover a_2 ↔ rover b_1	encounter	30 min every 7 d	30 min every 10 d

- Ground A, orbiter A, rovers a_1, a_2. Ground B, orbiter B, rover b_1, lander b enrolled on day 30. Bulletins move only on these contacts.
- One bundle per surface asset every 6 h (dense) or 12 h (sparse), lifetime 14 d / 30 d. Delivery measured at the first node of the other agency that accepts.
- Enrolment case: the lander's bulletin beat its first cross-agency contact in both plans. With a connected ground segment that case is small, federation after deployment is the one that costs.

The window, measured on a contact plan

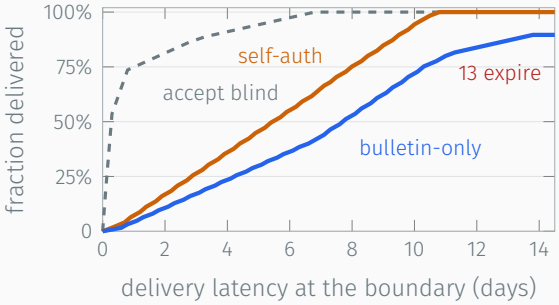

- Scheduled passes only, 2 agencies, 8 nodes. Deployed day 0, peering day 10. Flood itself: 1 to 18 h.
- 125 bundles crossed inside the window. Bulletin-only: mean 6.8 d, 13 expired, 1124 sends. Self-auth: 5.5 d, 0 expired, 125 sends.
- Bought: the wait for the next contact after the bulletin (median 11 h, max 6 d), and every resend.

The window is the ten days the agencies took to agree, not the flood. No mechanism shortens that. Sparse plan: 7.4 to 6.5 d, 172 to 46 sends. Rules and full table in Backup.

Finding: forward secrecy without a handshake

Epoch: a fixed interval, a month say, with its own KEM keypair, announced in advance in a DTKA bulletin and **erased** when the interval closes.

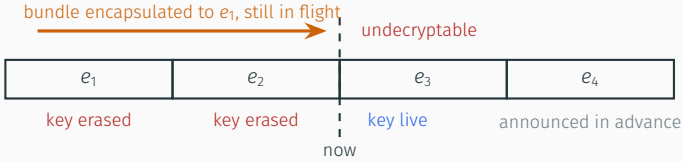

Erase on schedule and recorded traffic stays secret even after a later compromise. No interaction, ever.

The secrecy window can never be shorter than the bundle lifetime you tolerate.

Erasure must also reach ground-held state. Tested: a bundle outliving its epoch is undecryptable.

What it costs

	ML-DSA-65	ML-DSA-87
per bundle	3309 B	4627 B
first contact, total	8618 B	11 894 B
per rotation carried	5325 B	7283 B
confidential adds	1596 B	1596 B

Measured by the prototype's wire accounting. At 1 kbps one signature is ≈26 s of link. CNSA 2.0 approves only ML-DSA-87.

Monthly epochs cost ≈19 KiB per node per year of bulletins. Secrecy granularity is a bandwidth dial, limited by bundle lifetime.

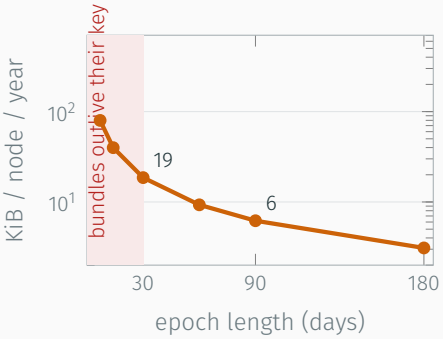

What I cannot yet answer

- **Duplicity.** One identity can show different chains to different receivers. Detection waits for a bulletin to reconcile.
- **Opportunistic terrestrial DTN.** Per-link admission leans on scheduled, scarce contacts. Where anyone can transmit, flooding reopens.
- **Scale.** The contact plan is synthetic: 8 nodes, periodic passes, one traffic pattern. A real plan (an ION contact plan file) and a custody-based baseline are next.
- **Rotation cutoffs vs delay.** Refusing superseded-key bundles cuts off a thief and drops legitimately delayed traffic. A dial, not a fix. Delay already broke sequence-based anti-replay once, key validity windows break it again.

Three things to take away

1. The window before a bulletin arrives is real, and the mandated post-quantum migration makes it wider.
2. Provisional authentication is worth having: binding verified from the bundle, standing settled by the bulletin, held history promoted without a round trip. And it is additive: one extension block, one security context, DTKA unchanged.
3. Forward secrecy in a delay-tolerant network needs no handshake, but it can never be finer-grained than the delay you tolerate.

martin.210215@student.ugd.edu.mk

Backup: rules of the measurement

- **Bulletin-only.** Reject until authorizable, the sender re-offers at every later cross-agency contact. Optimistic for the baseline: it assumes the sender learns of the rejection at once.
- **Self-auth.** One send. Binding verified, held at the boundary, delivered when the receiver can authorize. Never later than the baseline in either plan.
- **Blind.** Accept at first contact. The insecure floor.

in-window bundles	mean latency	expired	sends	cross-link bytes
dense, bulletin-only	6.8 d	13 of 125	1124	4.7 MiB
dense, self-auth	5.5 d	0	125	0.5 MiB
sparse, bulletin-only	7.4 d	0 of 46	172	0.7 MiB
sparse, self-auth	6.5 d	0	46	0.2 MiB

Per send: 1 KiB payload plus ML-DSA-65 overhead, 3309 B steady, 8618 B at first contact.

Backup: the window on an epidemic graph (earlier run)

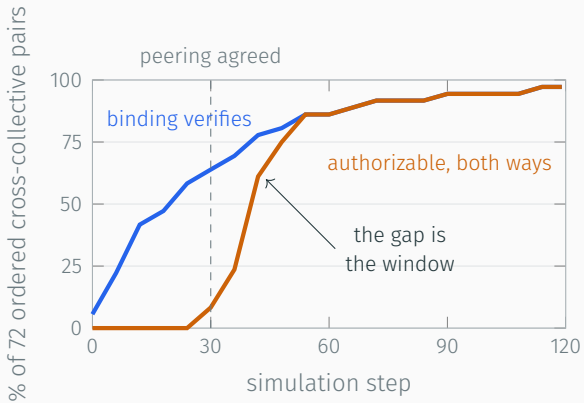

- 284 of 284 encounters: binding verified offline, every one *before* federation included.
- 70 of 72 ordered pairs met. 49 were authorized only retroactively. Mean window 26 steps.
- 115 bundles promoted that bulletin-only discards.

12 nodes, 2 collectives, random pairwise contacts, federation agreed mid-run. Shows the mechanism, the contact-plan run in the main line is the measurement

Backup: what this would mean in practice

layer	today	with self-auth
Bundle Protocol	RFC 9171	unchanged
BPSec	RFC 9172, its contexts	+ 1 context, + 1 extension block
DTKA	membership, peering, revocation	+ epoch keys, erasure, withdrawal
LunaNet security spec	unwritten	the federation case

- Nothing deployed has to move. Every addition is bulletin-native or a registered BPSec context.
- The one place this is not additive: erasure becomes a requirement on ground-segment state management.

Prototype: Rust, 64 tests including the full adversarial path, mock and FIPS-204/203 backends.

Backup: what an attacker can do inside the window

attacker	inside the window
no keys, forges as a known Q	rejected: BIB fails
stolen current key, rotates Q	rejected: successor pre-committed offline
joins an established exchange	rejected: continuity breaks
replays a bundle	deduplicated on bundle ID, now trustworthy
stolen current key, signs as Q	accepted. Residual risk, bounded by rotation and by revocation when the bulletin lands
mints fresh Q', claims to be a peer	accepted as *a* sender, never as an authorized one. Costs link budget. Never promotes

The two accepted rows are the honest residue. One is the rotation cutoff dial. The other is what per-link admission exists to price.

Backup: why not identity-based cryptography?

- IBE encryption and verification are offline too, and carry no key material at all. On that axis IBC is *better*.
- The cost is escrow: the PKG can forge any node's signature. Disqualifying when no agency may impersonate another's assets.
- Threshold PKGs remove escrow but need an online quorum to issue keys: the wrong trade where connectivity is the scarce resource.
- Boneh–Franklin expiry (identity ‖ date) requires periodic *confidential* delivery of private keys to disconnected nodes. Public assertions are strictly easier to distribute.

Backup: why not FN-DSA (Falcon)?

- Falcon-512: 666 B signatures, 897 B keys. About 5× smaller than ML-DSA-65. The bandwidth argument is definitely real.
- FIPS 206 is not final and CNSA 2.0 does not include it.
- Falcon signing needs constant-time floating-point Gaussian sampling, a hard ask on radiation-hardened flight processors. Verification is cheap. *Signing happens on the spacecraft.*

Position: revisit for bandwidth-critical links when FIPS 206 lands. Parameters live in the security context, so the swap is local.

Backup: compute cost, measured

	ML-DSA-65	ML-KEM-1024
keygen	293 µs	112 µs
sign / encaps	787 µs	124 µs
verify / decaps	222 µs	145 µs

Per bundle: one sign at the source, one verify per checking hop, one encapsulate/decapsulate pair when confidential. All sub-millisecond on x86-64 (200-iteration mean, FIPS-204/203 implementations).

The price of this design is bytes on the link, not cycles. A 10^2–$10^3\times$ flight-processor penalty still lands at $\sim$0.1–1 s per signature, far below DTN bundle rates.

Backup: replay and deduplication

- A replayed bundle is a valid bundle. Dedup on the bundle ID (source, creation timestamp, sequence), which verified binding makes trustworthy for the first time.
- Acceptance bounded by bundle lifetime, and ID memory the same way.
- Stated assumption: roughly synchronized clocks. Reasonable for spacecraft, weaker for improvised terrestrial nodes.

Backup: spec hygiene

- **Canonical encoding.** Identifier = hash of bytes, so: deterministic CBOR, fixed field order, version field, domain separation tag in the digest input.
- **Hash choice.** The registrant grinds his own record, so size for collision resistance: 256-bit digest, 128-bit strength, from the CNSA-approved SHA-2 family.
- **Chain growth.** Cold receivers need the full chain. At one rotation a year over a 15-year mission that is 15 short records. Deltas for warm peers, full chain otherwise.