

# Earning the Right to Propagate

---

How a human Village governs its own AI

John G. Stroh · My Digital Sovereignty Ltd. · 12 slides with presenter notes

# A constitution arrived in June

---

- Vincent Boucher — president of QUEBEC.AI and MONTREAL.AI — published *GoalOS: The Proof-of-Evolution Constitution* (AEP-001, v12.1, June 2026)
- By his own framing: a doctrine and a falsifiable research program, not a shipped product — its self-references are manuscript drafts
- It governs how an institution's AI is allowed to improve
- We govern how human communities hold AI — opposite ends, same primitives. The convergence is the story.

# Boucher's rule

---

- "Do not put intelligence on-chain. Put proof of intelligence on-chain."
- Keep the machine's reasoning and private data private; make public only the *proof* that work was done, tested, and allowed to influence what comes next
- Public doctrine: **Aim → Act → Prove → Evolve**
- Each step becomes a signed, versioned object

# The Proof Gradient — a selection law, not a score

---

- "The score is advisory; the gates are mandatory" — a high score never buys past a failed gate
- Promote only if: proof valid · beats baseline on eval · risk under threshold · rollback drilled · canary scoped · scope authorised · challenge window cleared
- "No proof, no evolution. No eval, no propagation. No rollback, no release."
- A proof page "is not a marketing page; it is a claim-bound evidence room"

## We'd already built it — for people

---

- Public-proof / private-data boundary → did:web identities + append-only proof chains; tenant content encrypted
- Proof-Carrying Artifact → signed sovereign records (provenance hash, content hash, tamper-evident, exportable)
- SelectionCertificate → DirectorEngagement decision record, maker-checker, proposer ≠ approver
- Evidence Docket → our Decision Records: verify-before-claiming, no defensive-hedge, minutes you can't quietly edit

# The gap it closed for us

---

- Our own AI evolution was empirical but ad-hoc — modest per-community models and the lighter layers around them
- Humbling record: a retrain tested and **not** deployed; a run of weight experiments degraded; "don't train aspirationally"
- We had the evidence habit — not a named law for when a change earns the right to propagate
- Now adopted: candidate → canary → active (always revocable), through Boucher's gates

# The one thing we refuse: the chain

---

- AEP-001 is blockchain-native: an Ethereum Evolution Ledger, staking, slashing, a reward vault, a token economy
- We adopt none of it — a public chain and US-anchored crypto infrastructure violate the Village's reason to exist (sovereign EU/NZ, self-hosted, no foreign substrate)
- His own neutrality clause lets us out: any institution may implement AEP "if it respects the proof interfaces"
- We respect the interfaces — on sovereign ground. Named human accountability where his design has staking and slashing.

# The gate the constitution is missing

---

- His subject is AI capability: "institutions whose agents may act, but whose improvements must prove themselves before they propagate"
- A capability gate can promote a superior upgrade that quietly moves authorship away from the people
- We add **ServeAuthorized**: does the change keep the community the *author*? Refused if not — even if it scores higher
- And consent + cultural fit: touching te ao Māori data governance needs the consent of those it affects, or it does not ship

# Governance capture: removed, not mitigated

---

- His threat model lists "governance capture: token votes or insiders bypass evidence" — mitigated with governance separation
- But the token-vote substrate *is* the attack surface
- We remove it: promotion authority sits with named human guardians under maker-checker — accountable as persons, not as wallets
- There is no stake to weight and no token to capture

# Keep AI the servant, not the subject

---

- The language of "institutional AI-evolution" has gravity — it pulls toward an AI-ops product with communities attached
- Hold the opposite line every time the discipline is invoked
- The proof gradient exists so a community can trust the tool it holds — not so the tool becomes the point
- If applying it ever makes the AI the subject, it is being misused — stop

## Postscript: from constitution to product

---

- Boucher's product layer — *GoalOS Mission OS* (June 2026): "set the objective; GoalOS runs until proof is done"
- Largely substrate-neutral — the reward vault / slashing court / staking are gone; his own product travels without the chain
- Its "autonomous publication law" (generate → QA → human-reviewed PR → publish) names the discipline the Village already runs; "proof debt" names our hazard
- Built for *autonomous AI work* that runs "until DONE" — exactly why we keep the human as the author, not the reviewer

# A community must remain the author

---

- A model can answer. An agent can act. An institution must prove. **And a community must remain the author.**
- Take the rigour gratefully; refuse the substrate; keep the gate the constitution leaves out
- Read the full essay: [agenticgovernance.digital/papers/earning-the-right-to-propagate.html](https://agenticgovernance.digital/papers/earning-the-right-to-propagate.html)

---

Sources: Vincent Boucher, *GoalOS: The Proof-of-Evolution Constitution — AEP-001, v12.1*, and *GoalOS Mission OS: The Proof OS for Autonomous AI Work, v23*, June 2026 (MONTREAL.AI & QUEBEC.AI). Lineage: G. C. Necula, "Proof-Carrying Code," POPL 1997.