Earning the Right to Propagate

How a human Village governs its own AI — what we built, why a new constitution for AI evolution validates it, and the one gate that constitution is missing

Précis. In June 2026 Vincent Boucher, president of QUEBEC.AI and MONTREAL.AI, published a constitution for governing how institutional AI is allowed to evolve: *GoalOS — The Proof-of-Evolution Constitution*. Its rule is austere and right — *do not put intelligence on-chain; put proof of intelligence on-chain* — and most of its primitives are ones we had already built, for human governance, and shipped: a public-proof / private-data boundary, signed proof-carrying records, maker-checker decisions, audit pages that are evidence rooms rather than marketing pages. That convergence is worth saying plainly, because a serious independent thinker arriving at the same architecture is the strongest external validation a design can get. This essay does three things. It explains what the Village already implements and why. It adopts the part of Boucher's discipline we were missing — a named, evidence-gated way to decide when our *own* AI earns the right to replace what is in service. And it parts company on two points: we refuse the blockchain his standard is built on, on sovereignty grounds his own neutrality clause permits; and we add the one gate his constitution does not have. His subject is AI capability. Ours is people. A model can answer, an agent can act, an institution must prove — and a community must remain the author.

A standard arrived this June that is worth a small nation's attention, not because it is finished but because of who arrived at it and where. Vincent Boucher — president of QUEBEC.AI and MONTREAL.AI — published *GoalOS: The Proof-of-Evolution Constitution*, subtitled *AEP-001: The Blockchain-Native Standard for Proof-Carrying Intelligence Organizations*. It is, by its author's own careful framing, a doctrine and a falsifiable research program, not a shipped product: its core self-references are manuscript drafts, and it claims a protocol architecture rather than any achieved capability. That honesty is the first thing we recognised in it, because it is the discipline we hold ourselves to.

The reason to read it closely is convergence. We have spent this programme

building a sovereignty platform for human communities — whānau, parishes, hauora providers, kāhui — in which AI is an instrument the community holds, never the authority it answers to. Boucher set out to govern how an institution's AI is allowed to improve. We started from opposite ends — he from the machine, we from the people — and met at the same small set of primitives. When that happens, it is a signal. This essay is an account of the agreement, the one place we refuse to follow, and the one place we go further.

What Boucher built

The constitution's spine is a single sentence: *do not put intelligence on-chain; put proof of intelligence on-chain.* The machine's reasoning, the private prompts, the customer data, the raw traces — these stay private. What becomes public and verifiable is the *proof* that a piece of work was done, evaluated, and allowed to influence what comes next. Boucher's public doctrine names four steps — **Aim → Act → Prove → Evolve** — and turns each into a signed, versioned object: an institution states its aim; a bounded agent acts; the run emits an attestable proof packet; and a candidate improvement may *evolve* — propagate into service — only by passing a selection gate.

That gate is the heart of it, and it is the part worth stealing. Boucher calls it the **Proof Gradient**, and he is emphatic that it is "a selection law, not a vibe, not a popularity contest, and not a marketing score." The score is advisory; the gates are mandatory. A candidate upgrade is promoted only if its proof is valid, it beats the incumbent baseline on registered evaluations under equal budget, its measured risk is under threshold, it carries a rollback target that has been drilled, it has a scoped canary, its promotion is authorised to a specific scope rather than released globally, and a challenge window has passed. No proof, no evolution. No eval, no propagation. No rollback, no release.

Two further pieces matter. The **Evidence Docket** is his name for the public proof page, and his line for it is one we wish we had written: *"a proof page is not a marketing page. It is a claim-bound evidence room."* The reader of such a page can see what was tested, what passed, what failed, which baselines were compared, which gates were enforced, and how to re-run the claim. And the **threat model** is adversarial by design — false proofs, evaluator collusion, reward hacking, privacy leakage, governance capture, rollback failure, overclaiming — because, as he puts it, a system that hides failure cannot be a proof system.

What we had already built

Read that list as a person who has spent two years building governance for human communities, and almost every line is familiar — because we had already shipped it, for a different subject.

The public-proof / private-data boundary is our architecture. The Village pub-

lishes did:web identities and append-only proof chains; the tenant's content and the reasons behind a decision are encrypted and never reach the public surface. Boucher's Proof-Carrying Artifact — signed, versioned, immutable, with a rollback target — is our signed sovereign record, with its provenance hash, content hash, and tamper-evident proof chain, exportable as a dossier. His SelectionCertificate, the signed object that admits or rejects a change under maker-checker separation, is our DirectorEngagement decision record: signed member positions, proposer not equal to approver, the whole thing minuted. His Evidence Docket — the claim-bound evidence room, not a marketing page — is the discipline we have been enforcing on ourselves all year: verify before claiming, no defensive-hedge sections, Decision Records that read as tamper-evident minutes rather than press releases. And his claim discipline — no AGI or ASI claims, empirical only through evidence — is our standing rule against fabrication and false legal equivalence.

We are saying this carefully, because claim discipline cuts both ways. What is **shipped** is the human-governance machinery above: the identities, the proof chains, the signed records, the maker-checker, the Decision Records. What we did **not** have, until now, was the same rigour applied to our own AI's evolution. That is the gap Boucher's work let us close, and we will not pretend it was already closed.

The gap it closes for us

The Village runs small, sovereign AI for each kind of community — modest per-community models and the lighter layers around them. Deciding when a new version earns the right to replace the one in service has been empirical but ad-hoc, and the empirical record is humbling. A retrain we evaluated was deliberately *not* deployed. A run of weight experiments degraded performance. We learned, the hard way, not to train aspirationally — that the consistently reliable improvements have come from lighter-weight adaptation, not heroic retraining. We had the evidence habit. We did not have a named law that said, in advance, what a change must prove before it is allowed to propagate.

So we have adopted one, as an internal governance standard. It takes Boucher's Proof Gradient almost verbatim and makes it the rule for our own AI evolution: a candidate model is `draft`, then `candidate`, then `canary`, then `active`, and `active` is always revocable to `rolled_back`. It is promoted only through the same mandatory gates — proof valid, eval beats baseline, risk under threshold, rollback drilled, canary scoped, scope authorised, challenge window cleared. The advisory score weighs verified value and quality against cost, risk, and rollback debt, exactly as he specifies, and never admits a candidate that fails a gate. This is the discipline we were missing, and we are clear that it is a discipline newly adopted, not a system already running in production.

The one thing we refuse: the chain

Here we part company, and the parting is not cosmetic. AEP-001 is, by its subtitle, *blockchain-native*. Its Evolution Ledger is an Ethereum-based evidence spine. Its contract suite is Ethereum end to end — a reward vault, a slashing court, evaluator staking, account-abstraction and token-bound-account standards, an attestation service. Capability is governed, in his design, partly by crypto-economic incentive: evaluators stake, bad actors are slashed, good work is rewarded, and propagation rights settle on-chain.

We adopt none of that. A public chain, US-anchored crypto infrastructure, and crypto-economic settlement are a direct violation of the Village's reason for existing: sovereign, self-hosted infrastructure under EU and New Zealand law, no public chain, no foreign substrate, no community's continuity hostage to a token economy. For a whānau or a parish, the idea that an AI change is gated by staking and slashing is not merely wrong infrastructure; it is a category error about what the institution is.

The good news is that we are entitled to take the doctrine without the substrate by Boucher's own terms. His design principles include a commercial-neutrality clause: *"any model provider, runtime, storage layer, evaluator, chain, enterprise, or sovereign institution can implement AEP if it respects the proof interfaces."* We respect the proof interfaces. We simply implement them on sovereign ground: an append-only proof chain on self-hosted storage instead of an Ethereum ledger; did:web identities and signed sovereign records instead of on-chain attestations; a maker-checker decision record instead of a selection-certificate contract. Where his design has a reward vault and a slashing court, ours has named human accountability and nothing else. The Village is a working existence proof that you can have the whole of AEP-001's proof-discipline with no blockchain at all.

The gate the constitution is missing

This is the part that matters most, and it is where we go beyond Boucher rather than merely diverging from him.

His constitution's subject is AI capability. He says so plainly in his conclusion: the next frontier is "institutions whose agents may act, but whose improvements must prove themselves before they propagate." The institution he imagines is, in effect, a network of agents, and the question his gate answers is *which AI improvement has earned the right to spread*. It is a good question, rigorously answered. But it is not the whole question for a human community, and his own threat model shows the seam: among the threats he lists is **governance capture — "token votes or insiders bypass evidence"** — which he mitigates with governance separation. The deeper point is that a capability gate, however strict, can promote a technically superior upgrade that quietly moves authorship away from the people whose community it is, and pass it through every gate it has.

So we add a gate his constitution does not contain. Call it **ServeAuthorized**, and it asks two things no capability metric can answer. First: does the change keep the community the *author* of its own decisions? A change is refused if it makes the people less the authors — even if it scores higher, costs less, and clears every other gate. Proof-of-evolution is necessary; it is not sufficient. Authorship is not a metric. Second: where a change touches a community's data or framing — most sharply, te ao Māori data governance — does it have the consent of those it affects? Boucher's design scopes a promotion by tenant and risk class; we extend scope to *consent and cultural fit*. A change with no proof does not propagate; neither does a change with no consent.

Notice what this also does to his governance-capture threat. He has to *mitigate* token-vote capture because his substrate has token votes. We *remove the substrate*: there is no stake to weight and no token to capture, because promotion authority sits with named human guardians under maker-checker, accountable as persons and not as wallets. The capture vector he guards against does not exist in our design. That is not a coincidence; it is what happens when the subject of the institution is people rather than agents.

Keeping AI the servant

There is a standing risk in writing any of this down, and it is worth naming because it does not go away. The language of "institutional AI-evolution" is seductive. Adopt enough of it and the Village starts to drift — from a sovereignty platform for human communities toward an AI-operations product with communities attached. Boucher's frame, brilliant for its purpose, has gravity in that direction, because for him the AI *is* the subject. Every time we invoke this discipline we have to hold the opposite line: the AI is the servant, not the subject. The proof gradient exists so that a community can trust the instrument it holds — not so that the instrument can become the thing the community is for. If applying the discipline ever starts to make the AI the point, the discipline is being misused, and the right response is to stop.

That is also why we are glad the convergence happened. It tells us the proof-carrying, public-proof / private-data, claim-bound shape we have been building is the right shape — confirmed independently by a serious researcher coming from the machine's side of the problem. We take his rigour gratefully. We take the lineage behind it gratefully too: the idea that an artifact should carry the proof of its own correctness is older than either of us, going back to Necula's *Proof-Carrying Code* in 1997. And we keep the one thing his constitution leaves out, because it is the thing the Village is for.

A model can answer. An agent can act. An institution must prove. And a community must remain the author.

Postscript — from constitution to product

Since this essay was written, Boucher has published the product layer of the same programme: *GoalOS Mission OS: The Proof OS for Autonomous AI Work* (June 2026). It turns the constitution into something an institution can commission a single run of — "set the objective; GoalOS runs until proof is done" — and its flagship deliverable is not a report but a *governed decision state*: in its own words, "the deliverable is not a document. The deliverable is a governed decision state." Three things in it bear directly on the argument above, and it is a stronger paper for being read closely.

First, it quietly settles the point we were most insistent on. The Mission OS product is largely substrate-neutral: the Ethereum machinery so prominent in AEP-001 — the reward vault, the slashing court, the staking — is gone from the product layer, surviving only as a background caution that the workflow "must not broadcast mainnet transactions." Boucher's own commercialisation travels without the chain. The doctrine and the substrate were always separable; his product is the demonstration, and a more persuasive one than our argument alone.

Second, it names a discipline the Village already runs. Mission OS states an "autonomous publication law": the public site "is not hand-edited [but] generated from proof-aligned source, checked by automation, reviewed by a human, and then published," and the pipeline "must not auto-merge … publish unsupported claims, or remove claim boundaries." That is, almost to the step, how these pages are made. It also gives a clean name to the hazard our own claim-discipline guards against — *proof debt*, "the risk that unverified AI output becomes institutional default."

Third, and most importantly, it sharpens rather than softens the one place we part company — and we say so with respect for how good the product is. Mission OS is, by its title, an operating system for *autonomous AI work*: it runs "until DONE," and it publishes. That is exactly the gravity this essay warned about — the subject sliding from the people to the machine's autonomy. Mission OS keeps one human gate, a review before publication; the Village keeps the human as the *author*, not the reviewer of last resort. We take the rigour gratefully and hold the line all the more firmly: AI is the servant, not the subject.

The Village is a running system, not a brochure — see it at mysovereignty.digital and weigh it against the claims made here. The Village holds the discipline described here as an internal governance standard — the Proof Gradient, adapted. Sources: Vincent Boucher, GoalOS: The Proof-of-Evolution Constitution — AEP-001*, v12.1, June 2026 (QUEBEC.AI & MONTREAL.AI); and* GoalOS Mission OS: The Proof OS for Autonomous AI Work*, v23, June 2026. Quotations are verbatim from those documents. Intellectual lineage: G. C. Necula, "Proof-Carrying Code," POPL 1997. — John G. Stroh, My Digital Sovereignty*

Ltd., June 2026.