Federate, Don’t Align

The safe path through the AI sovereignty contest — why federated communities, and federated inference, are the lowest-risk option for nations that will never win the capacity race

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

The binary is a risk decision, not a choice of allegiance

- Small nations are told there are two doors: the American stack or the Chinese stack
- The debate is run as allegiance — whose AI do you build your institutions on?
- Reframe it as risk: the three options don’t differ in virtue, they differ in how they fail
- One of them fails in a way you can recover from
- Take the capacity-vs-authority diagnosis as given (see the legitimacy paper); ask what authority you can safely hold

Option one — align with the American stack

- Capability now, in exchange for standing exposures
- Public data flows through infrastructure under another country’s law, including extraterritorial reach
- Dependency deepens with every integration — the cost of leaving rises over time
- Continuity hostage to commercial and policy weather you don’t control
- Risk: not catastrophic on any given day, but not reversible on your timetable

Option two — align with the Chinese stack

- China’s 2026 Implementation Guidelines for Intelligent Agents: a coherent, well-executed architecture
- A national intelligent-internet with a central agent-registration platform — every agent queryable
- Categorised and tiered governance: the state decides which classes of application get scrutiny
- Rule-embedding and behavioural fencing, verified centrally
- Sovereignty as capacity, done well — but built on a single substrate, with no exit

Option three — federate, and align with neither

- Hold your own data and models; reach others’ through bilateral, consent-bound, revocable channels
- Two real costs: a capacity ceiling (smaller models, borrowed substrate) and coordination overhead
- No central point whose compromise reaches you; no dependency you can’t withdraw from
- Every link is one you can cut — records and provenance intact on the way out
- The safe option is not the powerful one; it’s the only one whose worst case you can walk back

The mechanism: the signed envelope

- Two installations exchange information only through a signed envelope
- Consent-bound, scope-limited, addressed to a named recipient, carrying each record’s provenance
- No shared platform a third party pivots through; no authority that can revoke standing
- Withdrawal is immediate; exit is without penalty
- Contrast the central registry: efficient because it’s one thing — which is why its capture reaches everyone

Federated inference — the same channel carries the AI

- The envelope carries inference, not just records
- A community’s situated model reaches a peer’s corpus through the envelope, scoped to a fixed query shape
- It answers against that corpus without the data leaving the peer’s control or losing provenance
- Capability assembled by composition across consenting holders — not concentrated in one model
- Rightful authority made mechanical: a property of the envelope and the signed record, enforced by refusal

The authority layer, not the capacity race

- Federation does not close the substrate gap — weights, accelerators, compute remain foreign-sourced
- What it buys: the separability of governance sovereignty from substrate sovereignty
- Governance authority — over data, provenance, steering, behaviour — can be held in full today
- Capability at this layer is gained by composition, not by out-building
- The honesty about the substrate ceiling is the credibility

A non-aligned layer for AI

- In the Cold War, nations refusing the US/Soviet sorting called themselves non-aligned
- Not neutral, not passive — organised around the refusal, and stronger together for it
- A non-aligned layer for AI: the same posture rendered in architecture
- Small nations and indigenous polities gain capability by federating, not by aligning above themselves
- A recognised category of strategic behaviour with a track record — available now, to actors the capacity race has written off

A federated Aotearoa

- ~5 million people, no hyperscaler — on the capacity axis, it never registers
- On the authority axis: a constitutional settlement that already assumes plural, co-equal centres
- Picture the mesh: iwi, hauora, kura/wānanga, research groups — each its own Village, model, and records
- Each federates bilaterally → national-scale capability with no national registry and no foreign substrate
- Polycentric (Ostrom) + plural values (Berlin/Alexander ↔️ kōrero); kaitiakitanga as a signed field on every record

What this is, and what it is not

- Federation does not win the capacity race — it refuses to enter it
- It leaves substrate dependence where it is, and says so
- It is more work than buying a platform, and asks communities to hold authority rather than delegate it
- In return: no catastrophic, unrecoverable bet — lose a partner, a model, a provider, and recover
- The capacity contest was never one most of the world could win; the authority layer was always open — and it’s already running

Federate, Don’t Align

The safe path through the AI sovereignty contest — federated communities and federated inference give small nations rightful authority over their AI without winning a capacity race they were never going to win.