

Precis: Taonga-Centred Steering Governance

Document Type: Technical Documentation

Generated: February 10, 2026

Tractatus AI Safety Framework

<https://agenticgovernance.digital>

Precis: Taonga-Centred Steering Governance

Polycentric Authority for Sovereign Small Language Models

STO-RES-0010 v0.1 DRAFT — Stroh & Claude (2026)

The paper addresses a governance problem that arises when communities deploy their own language models with full access to model weights, rather than consuming AI through commercial APIs.

Such sovereign deployments permit direct modification of the model's internal representations at inference time through steering vectors — interventions that determine how the model represents kinship, place, authority, grief, and spiritual practice. These are instruments of norm enforcement. The paper asks who should govern them.

The prevailing answer is: the platform operator. The operator defines bias, extracts the corrections, distributes them downward. Communities customise within the limits set from above. The paper argues that for domains of Maori cultural knowledge, this hierarchy is structurally wrong, and proposes a polycentric alternative.

The proposal has three elements that connect directly to your work:

First, steering packs that encode iwi knowledge are treated as taonga — with iwi-controlled lifecycles, access conditions, and constraints on redistribution. The paper proposes iwi-operated "taonga steering registries" that are functionally analogous to the Biocultural Labels and Local Contexts infrastructure you have developed, applied here to AI steering interventions rather than to collections metadata. These registries enforce governance at the API level: access conditions, provenance verification, and revocation rights that the platform cannot circumvent.

Second, the architecture is polycentric. Iwi governance bodies and community trusts operate as co-equal steering authorities alongside the platform operator, each with distinct jurisdiction. There is no single apex. The model's activation space is a shared substrate, not a constitutional order. This maps to the distinction between delegation (platform grants authority downward) and recognition (iwi authority exists independently; the architecture either accommodates it or fails to).

Third, iwi hold a right of non-participation. They may decline to publish packs, may withdraw them at any time, and the platform must not substitute its own values into the gap. This is the architectural expression of tino rangatiratanga: iwi sovereignty does not depend on the platform's existence or goodwill.

The paper is, in a sense, an attempt to describe from the platform side a technical substrate that would be compatible with what your "Tikanga in Technology" programme addresses from within te ao Maori. Whether it succeeds — whether the architecture genuinely respects tikanga-based governance or merely provides new mechanisms for its subordination — is not a question we can answer ourselves.

Reference: Stroh, J. & Claude (2026). Taonga-Centred Steering Governance: Polycentric Authority for Sovereign Small Language Models. STO-RES-0010 v0.1 DRAFT. agenticgovernance.digital

© 2026 Tractatus AI Safety Framework

This document is part of the Tractatus Agentic Governance System

<https://agenticgovernance.digital>