

Precis: Taonga-Centred Steering Governance

Document Type: Technical Documentation

Generated: February 11, 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, the paper proposes that steering packs encoding iwi knowledge be treated as taonga — with iwi-controlled lifecycles, access conditions, and constraints on redistribution that the platform cannot override. These packs are governed cultural objects, not plugins. The governance architecture that protects them must be structural, not policy-based — the platform cannot circumvent access conditions or substitute its own values when an iwi declines to participate.

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. This maps directly to the distinction between delegation and recognition that the paper develops: in the delegation model, the platform grants communities the ability to customise within limits it defines; in the recognition model, community authority exists independently and 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.

You have argued publicly that tikanga should shape how AI operates for and with Maori — what you have called "Tikanga AI." This paper is, in a sense, an attempt to describe from the platform side a technical substrate that would be compatible with that vision. It proposes architecture where tikanga-based governance has co-equal authority over model behaviour, rather than operating as a cultural overlay on top of a platform-sovereign system.

Whether it succeeds is not a question we can answer ourselves. The paper is written by a non-Maori author in collaboration with an AI assistant. The concepts from te ao Maori that it draws upon carry authority and obligation beyond what we can represent.

What makes your assessment particularly valuable is that you have done what few others have: built sovereign Maori digital infrastructure at scale. Digital Natives Academy, Digital Basecamp, the launch of Google Maori — and more recently your whanau's establishment of Native Tech as an NZQA-registered PTE — demonstrate that Maori-led digital infrastructure is not a theoretical proposition but an operational reality. Your collaboration with the Alan Turing Institute on AI and data justice research shows this extends beyond education into precisely the governance questions the paper addresses. The scholars whose review we are also seeking — Kukutai, Hudson, Carroll — have developed the governance frameworks the paper draws on. You know where those frameworks meet the operational realities of actually running Maori-owned digital systems. It is precisely this that makes your judgment about whether the paper's proposals are viable, misconceived, or somewhere between the two, of distinct value.

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>