

Held in kōrero, not collapsed to a number — plural values, living organisations, and AI

John G. Stroh — My Digital Sovereignty Ltd

27 May 2026

The hui that doesn't end in a vote

A hui meets in a wharehau or around a table. A kaupapa is raised — a question, a proposal, a tension that needs working through. Kōrero unfolds: people speak in turn according to the tikanga of the marae or the entity. Each contribution carries who said it, when, and from what standing. The kōrero takes the time it takes. At some point the sense of the hui becomes clear. A whakatau is reached — sometimes a consensus, sometimes a formal vote, sometimes a deferral for further kōrero elsewhere. The meneti records the decision, with reference back to the kōrero from which it emerged.

What is the record of this meeting?

The record is the kōrero AND the decision. The decision without the kōrero is unintelligible — you would not know what was held in tension, what was conceded, what was unresolved. The kōrero is not preamble; it is co-authoritative.

Western governance platforms model the record differently. A motion is moved, seconded, debated, put to a vote, and either passes or fails. The vote tally is the record. The debate is courtesy, not authority.

The argument of this article is structural: hui-based deliberation is not “softer governance” that needs translating into a vote tally to count. It is the structurally correct form for holding plural values — values that cannot be reduced to a single number — through a living, participatory practice. Late in the Western philosophical tradition, two thinkers in particular began to articulate what hui has held continuously: Isaiah Berlin on pluralism, Christopher Alexander on living structure. The AI governance debate, where the question of what computers should do for communities is now being decided, has not yet taken seriously what Berlin and Alexander began to name. This article is an attempt to articulate the bridge.

What kāhui practice holds in place

The platform underpinning the kāhui in this article is Village Kāhui Māori, a working sovereign-records platform for Māori governance bodies. The structural shape of decision-making it carries is straightforward and not novel — it is what kāhui have done for generations. The four steps are:

Kaupapa. A topic, question, or proposal is raised for the kāhui to consider. It has a

title, a background, and an indication of who raised it and why. The kaupapa is not the answer; it is the question that opens the kōrero.

Kōrero. Members of the kāhui contribute their perspectives. The platform captures these as threaded discussion with attribution, dated, preserved in the record. Te reo of the kōrero is preserved unless the speaker chooses otherwise.

Whakatau. The decision. This may be a consensus, a formal vote tally, a deferral, a referral to a specific person or committee. The platform does not force a vote. It records what the kāhui does.

Meneti. The minute entry that carries the decision, with references back to the kōrero from which it emerged. The meneti is not the record. The kōrero PLUS the meneti is the record.

A second structural detail. The platform names roles in te reo because the roles carry meaning that English flattens:

- **Rangatira** — ultimate authority for the kāhui, with accountability back to the hapū that appointed them. Not “owner”.
- **Kaitiaki** — stewardship of the workspace, acts on behalf of the present and absent. Not “moderator”.
- **Mema** — participates in the kāhui. Not “user”.

These are not stylistic preferences. The English alternatives carry the wrong shape. “Owner” implies property rights and unilateral authority. “Moderator” implies content-policing. “User” implies a customer relationship. Rangatira / kaitiaki / mema each carry different mana, different accountability, different ways of being-in-the-kāhui. The platform is the wrong shape for the work if it flattens these.

A third detail. Crown-facing compliance — Charities Services reports, Māori Land Court records, Incorporated Societies registers — is available, but generated *from* the same record. The kāhui does its work in its own way; the compliance outputs are derivative artifacts. This matters because most platforms have the inverse: Crown compliance is the centre of the record, with cultural framing as decoration on top. Kāhui Māori inverts the emphasis.

None of this is a feature the platform “delivers”. It is the structural fact the platform must not collapse. The work was already there; the platform’s task is to carry it without flattening it.

Berlin — pluralism, irreducibility, the violence of the single score

Isaiah Berlin spent his life arguing that the deepest human values are plural and incommensurable. Liberty, equality, mercy, justice, solidarity, fidelity, courage, dignity — each is an ultimate good in human life. None reduces to any other. The choice between them, when they conflict, is a real choice between things that genuinely matter, not a calculation against a common scale.

This is value pluralism. It is Berlin’s principal philosophical contribution, and it is the contribution most consistently misread.

The misread is to treat pluralism as a sophisticated form of relativism — “all values are equally valid, your truth is yours, mine is mine”. That is not Berlin’s position. The position is much stronger: there really are plural goods; they really do sometimes conflict; the conflict is not solvable in advance by a master principle that ranks them; the practical resolution requires judgment, attention, and the willingness to bear the loss when one good is sacrificed for another. Hannah Arendt named something close to this *plurality* — the human condition of being many, with each life its own perspective, none reducible to any other. Joseph Raz reframed it in terms of practical reason: the reasons for action that flow from ultimate values are themselves plural, and the agent must hold them in tension rather than collapse them into a utility.

Berlin’s *The Hedgehog and the Fox* is a parable about this. The hedgehog knows one big thing — he organises the world around a single principle and pursues it. The fox knows many things and does not insist on unity. Berlin’s argument is not that foxes are smarter; it is that the world is fox-shaped. The hedgehog’s project of reducing everything to one big thing is a project of doing violence to the actual plurality of human goods.

Two Concepts of Liberty applies this to a specific case. Negative liberty is the absence of interference. Positive liberty is the capacity for self-mastery, autonomy, full participation in a meaningful life. Both are real. Neither reduces to the other. Reducing negative liberty to positive liberty produces paternalism — “you don’t really want freedom from interference; you want self-mastery, which requires us to guide you”. Reducing positive liberty to negative liberty produces a thin freedom that abandons the conditions of its own meaningful exercise. The argument is not that one wins. The argument is that both are values, and a serious politics must hold both in tension.

The structural implication for AI is direct, and it has not been taken seriously enough.

Any system optimised against a scalar objective function is doing structural violence to value plurality. The scalar is the hedgehog form. The benchmark is the hedgehog form. The reward model that returns a single number is the hedgehog form. Each of these collapses what is irreducibly plural into one big thing, then optimises against the one big thing.

The optimisation gate — the point at which a candidate output is accepted or rejected, a candidate policy promoted or demoted, a candidate instruction-text retained or discarded — is the locus of the violence. Wherever the gate is a single number, the violence has been done before the output ever reaches the person it was meant to serve.

Kāhui practice refuses this collapse. The record carries the kōrero precisely because the kōrero is where the plural goods live — the things that were held in tension, the considerations that were weighed without being scored, the dissent that was preserved. Collapse the kōrero into the vote tally and the violence has been done. Hold the kōrero alongside the whakatau and it has not.

A brief gallery of voices in this tradition, fellow travellers worth knowing about. **Bernard Williams** on moral luck and the limits of practical reason — there are situations where any choice is a wrong; the task is not to avoid the wrong but to bear it. **John Gray** on agonistic liberalism — Berlin’s most direct intellectual heir, who

pushes the pluralism argument further into the claim that liberal regimes themselves are one set of arrangements among others, not a universal end-state. **Charles Taylor** on multicultural recognition — the irreducibility of frameworks within which goods are intelligible. **Iris Murdoch** on attention as the ethical primitive — the capacity to see the other person as they are, not as our purposes need them to be.

None of these voices appears in most AI governance discussions. The fact that they do not is part of the gap this article addresses.

Alexander — living structure, participation, the quality without a name

Christopher Alexander was an architect, and his books are about buildings. He spent fifty years asking why some buildings, streets, rooms, and gardens feel alive while others — including ones built with considerable expertise and budget — feel dead.

His answer is that there is a structural property he calls *life*, or wholeness, or — in his most haunting phrase — *the quality without a name*. It is recognisable. You can stand in two rooms, one alive and one dead, and know which is which. It is not measurable. There is no metric that captures it; the moment you try to reduce it to a number, the thing slips away.

The Nature of Order, his four-volume late work, sets out the structural account. Wholeness emerges from the interaction of *centres* — geometric, perceptual, structural entities that strengthen or weaken each other through their relationships. A living structure has centres that support each other; a dead structure has centres that compete or are arbitrary. The work of making a living thing is the work of strengthening centres so they reinforce each other, and this work is done patiently, in many small steps, through participation by the people who will live in the result.

A Pattern Language and *The Timeless Way of Building*, his earlier work, articulate this as a craft methodology. A pattern language is a set of generative rules that lets the people who will inhabit a place participate in its making. Not by drawing a blueprint and handing it down; by adjusting and re-adjusting, one pattern at a time, through use and observation and care.

The applied claim Alexander makes, across all of this work, is that *living things cannot be made by top-down imposition*. They emerge from participation. They require ongoing adjustment. They depend on what rationalist design ideologies call “irrational” vectors — love, care, attention, the long patience of attending to a place over years and noticing what is alive in it and what is not.

When this is applied to communities and organisations, the implication is striking. A community is not a feature set. A community is a living thing in Alexander’s sense — its life emerges from participation, ongoing adjustment, embodied vectors that include love, grief, attention, the long patience of holding the kaupapa over time. An organisation that has Alexander’s quality is one where its members can recognise their own work in its shape; the centres are reinforcing because the people who live in them are participating in their strengthening.

The implication for AI is also direct.

An AI that serves a living organisation cannot replace its participatory life. It can only attend to it. The closed-loop optimisation — a fixed objective, a benchmark that judges, an artifact that drifts toward higher scores without the people the artifact serves participating in shaping where it drifts to — does the opposite. It severs the participatory loop. The optimisation runs without the community in it. The community is the beneficiary or victim of where it drifts, never the participant in shaping its direction.

Kāhui practice carries Alexander’s pattern naturally. The kōrero is the participatory making. The whakatau is the centre that gets reinforced or weakened depending on whether the kōrero supports it. Mātauranga — the body of knowledge a community holds — is not a frozen rulebook; it is the ongoing adjustment of practice as the kāhui meets new conditions. The platform’s task is to carry this work without freezing it into a static feature set. Alexander would have recognised the structural shape immediately.

The bridge — why optimisation-against-utility fails living organisations

Set Berlin and Alexander side by side and the diagnosis sharpens.

Berlin tells us why scalarisation is structural violence to plural values. Alexander tells us why closed-loop optimisation severs the participatory life of communities. Together, they identify two distinct structural failures in the dominant AI-governance frame, both fatal:

1. The frame collapses what is irreducibly plural — the values the AI is supposed to serve — into a scalar objective. (Berlin.)
2. The frame severs the participatory loop — the community the AI is supposed to serve becomes a recipient, not a participant. (Alexander.)

Either failure alone would be serious. Both together are diagnostic of an architectural mistake, not a tuning problem.

This is not a complaint about specific AI systems behaving badly. Specific systems will continue to behave better and worse depending on the engineering. The argument here is upstream of engineering: the frame within which the engineering is done cannot, at any level of skill or scale, do what governance for plural communities requires. The frame is structurally wrong twice.

Kāhui practice is the working counter-example. It does not scale by collapsing kōrero into a vote tally. It scales — when it does — by holding plural goods through ongoing participation. The members of the kāhui are not the recipients of a decision arrived at by an optimisation process; they are constitutive participants in the deliberation that arrives at it. When a new kaupapa arises that the existing tikanga has not yet faced, the kāhui’s response is not to consult an oracle that has been pre-trained on a corpus and ranks possible decisions by score; the response is to kōrero, hold the new kaupapa against what is already alive in the tikanga, and find a way through that the kōrero supports.

The contemporary Māori governance framework named in this work is **Te Kāhui Raraunga’s Māori Data Governance Model** — Tuia te korowai o Hine-Raraunga, the

eight pou — and the broader Māori AI Governance Framework that organisation has authored. This is the current operative reference for kaupapa Māori governance of data and AI; **Karaitiana Taiuru's** 20 September 2025 critical analysis (taiuru.co.nz/critical-analysis-mana-raraunga/) documents why the earlier Te Mana Raraunga six principles (2016–2018) require supplementation when applied to AI specifically. Tangata-whenua determination of Māori data governance frameworks remains with Māori; this article cites these sources because they are the current operative work, not as commentary on it.

The polycentric architectural commitment that follows from the bridge is straightforward in structure even though it requires care in implementation. Each kāhui, iwi, hapū, marae trust, whānau rōpū holds its own steering authority over the platform that serves it. There is no global optimiser. There is no single score. Plural goods are held by plural authorities — exactly the form Berlin's pluralism requires, in exactly the participatory mode Alexander's living structure requires.

Two architectural primitives, both in production on the Village sovereign-records platform, make this concrete.

First, **steering packs**. Each governance authority signs its own pack defining what the AI serving its tenant may pursue, what it must refuse, what cultural boundaries it must respect. The pack is community-defined; the platform attends to it; the community holds the authority to revise it through hui and kōrero.

Second, **runtime cultural-boundary checking**. Outputs are gated against the steering pack at the moment they are produced, not only at training time. Sir Hirini Mead's Tikanga Test — the requirement that an action be judged against the tikanga of the people the action affects, by those people — is the structural pattern. The platform's role is to apply the test the community has authorised, not to substitute its own.

Neither primitive is a feature. Both are structural commitments that follow from taking Berlin and Alexander — and the kāhui practice they belatedly articulate — seriously.

Where the technical literature is heading — and the gap it leaves open

A line of frontier research is converging on a substrate close to the one Village uses for the parts of its system that adapt over time. The shared pattern: a frozen base model, an evolving text artifact that carries the agent's instructions, an automated validation gate that accepts or rejects candidate variants of the artifact.

Stanford's **TextGrad** introduced the idea of gradient-style updates against text artifacts via natural-language critiques. The 2025 **GEPA** work extended this with population-based exploration. The 2026 skill-evolution literature — **SkillOpt**, **Trace2Skill**, **EvoSkill** — develops methods for discovering and refining skill artifacts incrementally. (Paper identifiers in the closing block below.)

These methods are real progress. They move beyond the older fine-tuning-against-fixed-objective paradigm. The artifact under optimisation is now text, plural in its possible shapes, programmatically adjustable, learnable from feedback. This is more participatory in shape than what came before. It deserves attention.

But the acceptance gate, across all of this work, accepts the higher-scoring variant. A single number ranks the candidates, and the higher-scored candidate wins. The gate is the optimiser, and the optimiser is the gate.

The gap this leaves open is the gap this article has been articulating throughout.

No kaitiaki on the artifact. The community whose work the AI is supposed to serve has no role in deciding which variant of the artifact survives. The candidate that scored highest against the benchmark survives; the community’s view of what the artifact should pursue is, at best, an input upstream of the benchmark, not a constitutive authority over the gate.

No iwi mandate on the trajectory. The artifact drifts somewhere over time. The drift is determined by the benchmark, not by the kāhui that the artifact ostensibly serves.

No community-defined boundary tested at the moment outputs are produced. Cultural boundaries — if they exist at all in the system — are training-time considerations baked into the score function, not runtime gates checked against the community’s own tikanga.

This gap is structural, not feature-class. The optimiser-landscape research does not have a community-authority layer because the underlying frame does not have a notion of community as a constitutive locus of authority. The frame is: pursue a metric, gate by score. Adding “community authority” is not a parameter tweak. It requires a different theory of where authority sits — community-defined, plurally-held, participatorily-revised. That theory is what Berlin and Alexander point at and what kāhui practice has had continuously.

Village’s architectural commitment is to where the gating sits. The frontier optimisation methods, where they prove useful, can sit *under* steering-pack authority rather than *as* the authority. The benchmark is downstream of the kaitiaki, not upstream. The compressed form of this argument appears as §5.5 of the *Te Kāhui Māori June 2026 briefing*; this article is the expanded form.

What follows from taking the bridge seriously

If the argument lands, several specific things follow.

An AI that should serve plural communities cannot have its gate at a benchmark score. The gate must be at community-defined authority — steering authority sitting where the living happens, revised by the people whose lives the AI’s outputs touch, applied at the moment of output rather than only at the moment of training.

Polycentric steering is the architecture this requires. Not a single global authority over the AI’s behaviour; not a centralised review board; not the AI’s developers as the de facto authority. Each governance authority signs its own pack. Each pack reflects the tikanga of the people it governs. The AI attends to whichever pack is operative for the tenant it is serving. No tenant’s pack governs another tenant’s tenant. Plural goods, plurally held, plurally revised.

Runtime checking is the gate this requires. Not only training-time alignment to a global objective; runtime application of the community’s own boundary at the moment

outputs are produced. Mead’s Tikanga Test is the structural pattern: an action is judged against the tikanga of the people the action affects, by them. The platform’s task is to apply the test the community has authorised, not to substitute its own.

Berlin and Alexander did not resolve any of this. They named what governance requires. They identified — Berlin from political philosophy, Alexander from the craft of building — that plural goods cannot be reduced and that living things cannot be imposed. The AI governance debate has spent most of the last decade looking past their work. The cost of that has been the production of a frame that, even at arbitrary scale and engineering sophistication, cannot do what governance for plural communities requires.

Kāhui Māori has been doing this work without needing Berlin’s vocabulary or Alexander’s diagrams. The work is what it is, regardless of who articulates it. This article is one attempt to articulate, for readers who may not yet have encountered the underlying practice, what is structurally at stake — and what alternative architecture is already in use, working, and worth attending to.

Further reading

Berlin — *Two Concepts of Liberty* (1958); *The Hedgehog and the Fox* (1953); *Four Essays on Liberty* (1969).

Alexander — *The Nature of Order* (four volumes, 2002–2004); *A Pattern Language* (1977); *The Timeless Way of Building* (1979).

Te Kāhui Raraunga — *Te Mana o te Raraunga Māori Data Governance Model* (Tuia te korowai o Hine-Raraunga); *Māori AI Governance Framework* (with summary report and use cases resource). kahuiraraunga.io.

Karaitiana Taiuru — *Critical Analysis: Te Mana Raraunga* (20 September 2025). taiuru.co.nz/critical-analysis-mana-raraunga/.

Other voices in Berlin’s tradition worth reading:

- Joseph Raz — *The Morality of Freedom* (1986); *Engaging Reason* (1999).
- Bernard Williams — *Moral Luck* (1981); *Ethics and the Limits of Philosophy* (1985).
- John Gray — *Berlin* (1995); *Two Faces of Liberalism* (2000); *Isaiah Berlin: An Interpretation of His Thought* (2013).
- Charles Taylor — *Sources of the Self* (1989); *The Politics of Recognition* (1992).
- Hannah Arendt — *The Human Condition* (1958).
- Iris Murdoch — *The Sovereignty of Good* (1970); *Metaphysics as a Guide to Morals* (1992).

Where the technical literature points

A closing block, not part of the argument above. The methods exist; the gate-locus question remains open.

- **SkillOpt** — arXiv:2605.23904. Microsoft + Shanghai Jiao Tong / Tongji / Fudan, 2026. Skill artifact optimisation.

- **GEPA** — arXiv:2507.19457. 2025. Population-based exploration over text artifacts.
- **TextGrad** — arXiv:2406.07496. Stanford, 2024. Gradient-style updates on text via natural-language critique.
- **Trace2Skill** — arXiv:2605.21810. 2026. Skill induction from execution traces.
- **EvoSkill** — arXiv:2603.02766. 2026. Evolutionary skill development.

Citations are by paper identifier. Most are after this article’s author’s reading cutoff; the identifiers are the authoritative ground, not paraphrase. What the existence of the literature demonstrates is that the substrate the Village platform uses — a frozen base model, an evolving text artifact that carries the agent’s instructions, an automated validation gate — is being independently arrived at by multiple research lines.

Where Village’s framework differs structurally

The optimiser-landscape methods listed above and Village’s polycentric-governance framework share substrate but diverge on every structural commitment that matters for governance. The differences are not in tuning, hyperparameters, or benchmark choice — they are in where authority sits.

Dimension	Optimiser-landscape frame	Village framework
Locus of the acceptance gate	Benchmark score — a single number ranks candidate artifacts	Steering pack — community-signed authority over what the artifact may pursue
Authority over artifact evolution	The optimiser (the gate is the optimiser; the optimiser is the gate)	The community (rangatira / kaitiaki / mema, revising packs through hui and kōrero)
Treatment of plural goods	Scalarised into the objective; collapsed into one number before optimisation runs	Held by plural authorities; no global optimiser; no single score across tenants
Runtime boundary checking	Training-time alignment to a global objective; no community-defined boundary at output time	Runtime gate against the steering pack at the moment outputs are produced; Mead’s Tikanga Test as the structural pattern
Community role in the loop	Beneficiary or victim of drift — community sits upstream of the benchmark or downstream of the output, never constitutive of the gate	Constitutive participant — the community is the locus of authority that defines what the artifact may become

Dimension	Optimiser-landscape frame	Village framework
What “alignment” means	The artifact’s outputs score higher against a benchmark proxy for what the community wants	The artifact’s outputs sit within the boundary the community itself has authorised, revisable on the timescale the community itself sets
Failure mode when the artifact drifts	The community has no constitutive role to correct drift; only training-data and benchmark revision, both upstream of the optimiser rather than constitutive of it	The community revises its own pack; correction is intrinsic to the architecture, not bolted on
What federation means	Sharing weights, prompts, or skill libraries between systems that each have their own optimiser	Sharing or recognising steering authority between communities that each hold mana over their own tenants
What scale means	More users served by the same optimiser against the same benchmark	More authorities, each with their own pack; the platform attends to whichever pack is operative for the tenant it is serving

The structural commitment is consistent across the table. In the optimiser-landscape frame, the community is **served by** an optimisation process that decides where the artifact goes. In the Village framework, the community **constitutes the authority over where the artifact goes**. This is the difference between a service relationship and a polity.

Two practical consequences follow.

First, the substrate of the optimiser-landscape methods is reusable under Village’s frame. If one of these methods proves useful for a narrow technical task — refining a domain-specific skill, exploring variants of an instruction text, inducing a skill from execution traces — Village can adopt the method with the steering pack as the gate rather than the benchmark. The optimiser sits *under* community authority, not as it. This is not hypothetical: it is the architectural commitment that lets Village take frontier methods seriously without losing what makes governance governance. The methods become tools the community wields, not processes the community is subject to.

Second, the optimiser-landscape methods cannot be retrofitted with community authority without changing their underlying frame. Adding a “community input” channel upstream of the benchmark is not the same as making the community the authority

over the artifact. The benchmark still runs; the optimiser still gates; the community is still a recipient. To put the community at the gate requires what the optimiser frame does not have: a notion of community as a constitutive locus of authority. The presence or absence of that notion is not a feature flag. It is the choice of what kind of system is being built.

This article has argued, throughout, where that locus belongs. The technical literature is pointing at the substrate Village uses. The structurally open question — where the gate sits — is the question the philosophical bridge above is meant to answer. Berlin's pluralism and Alexander's living structure say the gate cannot be a scalar and cannot be closed. Kāhui practice shows what it looks like when the gate is community-defined, plurally held, participatorily revised. Village is the architectural commitment to building the platform that way from the start.