

Opinion Article

Collective Privacy, Indigenous Data Rights, and the Computational Systems That Threaten Them

Karaitiana Taiuru*

Taiuru & Associates, New Zealand

*Corresponding author

Karaitiana Taiuru, Taiuru & Associates, New Zealand

Submitted: 21 March 2026

Accepted: 07 May 2026

Published: 08 May 2026

ISSN: 2333-7133

Copyright

© 2026 Taiuru K

OPEN ACCESS

Keywords

- Māori Data Sovereignty
- Indigenous Privacy
- Computational Systems
- Artificial Intelligence
- Treaty of Waitangi
- Collective Privacy
- Biometric Surveillance
- Genealogical Data
- Privacy Engineering
- Care Principles
- Postcolonial Technology Studies

Abstract

The dominant computational privacy paradigm rooted in individualist, Western liberal traditions is structurally inadequate for protecting Maori people, the Indigenous population of New Zealand. Maori privacy is fundamentally collective in nature: it covers genealogical records, traditional knowledge systems, culturally treasured resources, and the shared dignity of extended family groups (whanau), sub-tribes (hapu), and tribes (iwi). This article argues that contemporary computational systems including artificial intelligence (AI), biometric recognition platforms, cloud-based data storage, and algorithmic decision-making tools embed design assumptions that structurally exclude Maori privacy interests. Drawing on more than 30 years of practitioner experience and grounded in Maori customary law, cultural protocols, and the Treaty of Waitangi (New Zealand's founding constitutional document), this article identifies five critical failure modes in current computational architectures and proposes an Indigenous led framework for redesigning privacy into systems at source. It provides technical specifications for implementing collective privacy across three system classes: identity and access management, AI training pipelines, and cloud storage architectures. The article further engages with communitarian and postcolonial traditions within Western privacy theory, situates the analysis within science and technology studies (STS) and critical data studies scholarship, and grounds the five failure modes in documented case evidence. The findings have broad implications for any computational system operating in jurisdictions with Indigenous populations covered by treaty or constitutional obligations.

INTRODUCTION

The Māori concept of collective privacy extends substantially beyond the individualist conception embedded in most Western legislation and, by extension, in the computational systems that implement it. For Māori, privacy is bound to mana (personal authority and prestige), tapu (sacred restriction), and the collective dignity of whānau (family), hapū (sub-tribe), and iwi (tribe) connected across generations through shared genealogical lineage (whakapapa). When a computational system exposes personal data about a single Māori individual, it may simultaneously violate the privacy of that person's entire extended family network. This is not a metaphor: it is a structural reality that contemporary privacy-by-design frameworks have almost universally failed to address.

This article makes four core arguments. First, that mainstream computational privacy architectures are structurally Western and individualist, and that this is not a neutral design choice but an active exclusion of

Indigenous governance epistemologies. Second, that this exclusion produces measurable, documented harms for Māori communities across health systems, law enforcement, social services, and the commercial sector. Third, that remedying this requires more than a culturally sensitive policy layer atop existing systems: it requires computational redesign from the ground up, informed by Māori customary protocols and grounded in Treaty of Waitangi obligations. Fourth, a dimension added in direct response to editorial review that such redesign is technically feasible: this article provides concrete architectural proposals for implementing collective privacy in identity management systems, AI training pipelines, and cloud storage infrastructure.

The analysis draws on over three decades of practitioner experience at the intersection of Māori customary law, digital technology, and data governance, including published work on Māori data sovereignty, AI ethics, biometric systems, and the adequacy of existing legal instruments for protecting Māori privacy rights. It

also engages with an interdisciplinary body of scholarship spanning privacy engineering, science and technology studies (STS), critical data studies, postcolonial technology studies, and accountable algorithms research.

THE STRUCTURAL INCOMPATIBILITY OF WESTERN PRIVACY ARCHITECTURE

New Zealand’s Privacy Act 2020 is, by international standards, a modern and comprehensive instrument [1]. Yet the Act operationalises privacy as a right held by an individual in relation to information about that individual. It does not recognise that information about an individual Māori person may be simultaneously and inseparably information about that person’s whānau, hapū, and iwi.

Before presenting the Māori critique, it is important to acknowledge that Western privacy theory is not monolithic. Scholars including Amitai Etzioni [2], Charles Taylor [3], and more recently, Woodrow Hartzog [4], have articulated communitarian and relational dimensions of privacy that resist the purely individualist liberal model. Nissenbaum’s [5], theory of contextual integrity holds that privacy norms are constituted by the social contexts in which information flows, a framing with meaningful resonance for Māori customary protocols governing information sharing. Solove’s [6], taxonomy of privacy harms acknowledges aggregation harms that affect groups, not only individuals. These traditions demonstrate that the critique of computational individualism does not require rejecting Western theory wholesale: it requires insisting that its communitarian strands be taken seriously in system design, not merely in legal philosophy.

The deeper problem is not Western privacy theory in its full complexity but the highly selective version of that theory that has been operationalised in computational architectures. GDPR the dominant global reference standard for privacy-by-design encodes individual rights: access, rectification, erasure, portability, and objection. These rights are held by natural persons individually and exercised individually. The data models that implement GDPR compliance consent records linked to individual identifiers, subject access request workflows, erasure queues keyed on individual data subject identifiers have no structural analogue for collective rights-holding.

The CARE–GDPR Technical Incompatibility

The CARE Principles for Indigenous Data Governance — Collective Benefit, Authority to Control, Responsibility, and Ethics are structurally incompatible with GDPR-style architectures in ways that go beyond policy preference. The following table sets out the principal technical incompatibilities:

CARE Dimension	GDPR Architecture	Technical Incompatibility
Collective Benefit	Data minimisation: collect only what is necessary for the specified individual purpose	Collective benefit analysis requires understanding data in relational context across group members, which conflicts with the minimisation principle when applied atomistically to individual records.
Authority to Control	Individual consent, obtained from the data subject at point of collection	Collective consent requires group-level decision-making structures (e.g., iwi governance bodies) as legal principals in consent chains. Standard consent management platforms have no mechanism for collective principals.
Responsibility	Controller–processor obligations owed to individuals; no duty of care to non-data-subject communities	Responsibility to the broader Indigenous community affected by data use is architecturally invisible: GDPR’s accountability framework has no data model for community-level impact.
Ethics	Purpose limitation: data used only for specified, explicit, legitimate purposes	Indigenous ethical review (e.g., kaupapa Māori research protocols) involves ongoing relational obligations that cannot be reduced to a fixed purpose specification lodged at collection time.
Right to Erasure (GDPR Art. 17)	Individual can demand deletion of their personal data	For AI systems, erasure of specific training data after model training is technically infeasible. The trained weights encode statistical patterns across the corpus; no current method can surgically remove the influence of a specific data item post-training [7,8].

This analysis draws on access control theory to highlight a further incompatibility. GDPR’s access control model is essentially discretionary access control (DAC): the individual data subject holds the access right and may grant or revoke it. Māori data sovereignty, by contrast, requires a form of mandatory access control (MAC) in which access decisions are governed by policy set by the tribal authority, not merely by the preference of the individual whose record is at issue. No commercial identity and access management (IAM) platform currently implements tribal-authority MAC as a native data classification tier.

This structural analysis is consistent with the STS literature on infrastructure and classification. Bowker and Star’s [9], work on how classification systems encode and perpetuate social hierarchies is directly applicable: when GDPR’s individual-subject data model is instantiated in a database schema, it does not merely fail to represent collective rights, it makes collective rights structurally unrepresentable within that schema. Winner’s [10] foundational question whether artefacts have politics is answered affirmatively by any database schema that treats a genealogical lineage record identically to an email address.

FIVE CRITICAL FAILURE MODES IN CURRENT COMPUTATIONAL SYSTEMS

The following sections document five failure modes through case evidence, technical analysis, and engagement with the relevant empirical and scholarly literature.

Biometric and Facial Recognition Technology

Facial recognition technology (FRT) represents one of the most acute and documented threats to Māori privacy in New Zealand. In 2022–2023, commercial FRT deployments by Foodstuffs North Island attracted substantial public attention. Concerns centred on higher false-positive identification rates for Māori and Pacific Island people a bias well-documented in large-scale evaluations.

The National Institute of Standards and Technology's Face Recognition Vendor Testing (FRVT) programme has produced the most rigorous independent dataset on FRT demographic differentials. Grother et al. [11], found that across 189 algorithms tested, false positive rates for Black and Asian faces were 10 to 100 times higher than for white faces; no major FRVT study has specifically tested Māori demographic representation, itself a significant evidence gap. Buolamwini and Gebru's [12], Gender Shades audit demonstrated that commercial gender classification systems showed error rates of up to 34.7% for darker-skinned women compared to 0.8% for lighter-skinned men, a disparity attributed directly to training data composition. These findings are consistent with the Māori experience and apply with additional force to the specific case of tā moko (traditional Māori facial tattoo), a category of facial appearance that no major commercial FRT dataset has documented training coverage for.

In 2024, an independent trial of the Department of Internal Affairs' Identity Check system indicated that the trial report obscured risks to Māori, including heightened likelihood that people wearing tā moko would be misidentified or excluded from the system entirely. The Biometric Processing Privacy Code 2025, introduced by the Office of the Privacy Commissioner, represents a regulatory step forward, but it does not resolve the underlying architectural problem: training data composition.

From a privacy engineering standpoint, the Cavoukian [13], Privacy by Design framework's Principle 7 "respect for user privacy" cannot be operationalised in an FRT system without specific testing against the full demographic range of the deployment population. The computational implications are clear: any FRT system that undergoes accuracy validation without specific testing against the

full range of Māori facial diversity, including the variety of tā moko forms, should be considered architecturally inadequate for deployment in New Zealand.

Algorithmic Decision-Making in Public Services

Algorithmic systems are increasingly deployed across New Zealand's public sector to support decision making in social welfare eligibility, child protection risk scoring, and health service prioritisation. These systems are trained on historical administrative data that encodes decades, in some areas, more than a century of discriminatory policy application toward Māori.

Oranga Tamariki (New Zealand's child welfare agency) has deployed algorithmic risk scoring tools that have attracted sustained criticism from Māori communities and the Waitangi Tribunal (WAI 2022). Eubanks' [14], comparative analysis of algorithmic welfare systems in the United States documents how such systems consistently amplify existing social inequities: the Allegheny Family Screening Tool, the Indiana benefits eligibility system, and the Los Angeles homeless services algorithm each produced outcomes that disproportionately burdened racialised communities. The mechanisms are structurally the same as those operating in New Zealand's public sector algorithmic systems.

Privacy risks in this context extend beyond data security. Noble's [15], concept of "oppressive algorithms" applies directly: algorithmic systems that disproportionately flag Māori individuals for welfare review, police attention, or child services intervention expose those individuals and, by extension, their wider whānau networks to institutional surveillance without meaningful consent. Benjamin's [16], analysis of "New Jim Code" mechanisms further situates this pattern within a broader theory of racialised technological design. The New Zealand Government's Algorithm Charter [17], requires signatory agencies to identify and address bias, but adherence is voluntary, audit mechanisms are weak, and the Charter contains no specific requirements for a Māori cultural impact assessment prior to deployment.

Cloud Hosting and Data Jurisdiction

Research conducted across a 12-month period from 2024 to 2025 found significant noncompliance with Māori data residency principles, including among Māori organisations and advocates who publicly champion data sovereignty [18,19].

United States hyperscale cloud providers including Amazon Web Services, Microsoft Azure, and Google Cloud

Platform operate under the US CLOUD Act [20], which permits US law enforcement agencies to compel disclosure of data held by US companies regardless of where that data physically resides. Daskal's [21], legal analysis of the CLOUD Act demonstrates that the legislation creates a structural bypass of bilateral Mutual Legal Assistance Treaty (MLAT) frameworks, effectively subordinating the legal protections of the data's jurisdiction of residency to US law. For Māori data, particularly genomic data, health records, and genealogical databases stored with these providers, this creates a jurisdictional vulnerability that no policy statement about data sovereignty can resolve without technical enforcement measures at the storage and access-control layer.

Artificial Intelligence Training and Traditional Māori Knowledge

Large Language Models (LLMs) and other AI systems are trained on web-scale datasets that include substantial quantities of Māori cultural content such as te reo Māori (the Māori language), oral histories, traditional knowledge systems, and other culturally significant material without the consent of the Māori communities who created and hold kaitiakitanga (custodianship) over this knowledge.

Bender et al.'s [22] "Stochastic Parrots" paper and Denton et al.'s [23], work on dataset documentation each highlight that the ethical risks of AI systems begin at data collection and cannot be remediated at the model output level alone. Gebru et al.'s [24], Datasheets for Datasets framework and Mitchell et al.'s [25], Model Cards proposal provide partial infrastructure for documenting training data provenance, but neither has been implemented with Indigenous data sovereignty obligations as a first-order requirement.

The problem is compounded by the architecture of modern AI training pipelines, which make it technically infeasible to identify or remove specific items of training data after a model has been trained [8]. For Māori communities whose cultural knowledge includes material subject to tapu restrictions — ceremonial knowledge, sacred whakapapa records, or information shared only within specific cultural settings — current AI training practices represent an irreversible collective privacy violation with no adequate legal remedy under existing frameworks.

Genealogical Databases and Lineage Privacy

Whakapapa lineage records occupy a unique position in Māori society: they function simultaneously as personal

identity documents, collective historical records, legal instruments establishing land rights and tribal membership, and records encoding cosmological relationships between the living, their ancestors, and the natural world. Computational systems including commercial genealogy platforms, health research databases, and government population registries routinely hold these lineage records without any recognition of their collective nature or the customary restrictions (tikanga) governing their use.

The emergence of AI-powered genealogy tools capable of inferring lineage connections from fragmentary data presents a particularly acute risk. Humbert et al.'s [26], analysis of genomic privacy demonstrates mathematically that it is possible to infer relatives' genetic information from a single individual's genome at a level that constitutes a privacy violation for those relatives, even where those relatives have provided no sample. This inferential vulnerability scales directly to whakapapa inference: a system that can reconstruct family networks from publicly available records exposes tribal affiliations, land interests, and inter-familial relationships in ways that no individual can consent to, and that no current privacy framework is equipped to address.

TECHNICAL OPERATIONALISATION OF COLLECTIVE PRIVACY

The five failure modes identified in the previous section each points to a specific class of architectural deficiency for which a technical remedy exists. This section proposes concrete design patterns across three system classes: identity and access management (IAM), AI training pipelines, and cloud storage architectures, demonstrating that collective privacy is not merely a policy aspiration but an engineering objective that can be specified, designed, and validated.

Identity and Access Management: Extending to Collective Principals

Standard IAM architectures implement Role-Based Access Control (RBAC) or Attribute-Based Access Control (ABAC), in which access decisions are made on the basis of individual identity attributes. To support Māori collective privacy, a tribal-authority layer must be introduced as a first class principal in the access control model. The following architectural pattern achieves this without requiring wholesale replacement of existing IAM infrastructure:

Design Pattern 3.1-A: Collective Principal IAM Extension Extend ABAC policy engine to include an

IwiGroupAttribute alongside standard identity attributes. Define a CollectiveDataClass metadata tag for records carrying collective sensitivity. Enforce rule: access to CollectiveDataClass records requires not only the individual data subject's consent but a valid authorisation token issued by the designated iwi authority endpoint. Implement the tribal authority as a trusted identity provider (IdP) within the existing OpenID Connect or SAML federation, issuing short-lived authorisation assertions scoped to specific data classes. This preserves backwards compatibility with existing IAM infrastructure while inserting the tribal authority as a required policy decision point.

Verifiable credentials (VCs) under the W3C Decentralised Identifier (DID) specification provide a complementary mechanism. A tribal authority can act as a VC issuer, issuing credentials that attest to membership, consent, or authorisation at the collective level. These credentials can be presented to relying parties without disclosing the underlying identity data, preserving privacy while enabling collective consent enforcement [27]. Te Hiku Media's implementation of data sovereignty in te reo Māori speech technology using community controlled licensing rather than open dataset publication — provides an existing operational precedent for tribal-authority governance of AI data access [28].

AI Training Pipelines: Collective Consent and Data Provenance

The AI training pipeline offers several integration points for collective privacy enforcement. The following architecture describes a three-layer governance model:

Cloud Storage Architecture: Enforcing Data Sovereignty at the Infrastructure Layer

Māori data sovereignty requires that designated sensitive data such as genomic records, whakapapa databases, health data, and sacred knowledge be stored exclusively within New Zealand jurisdiction and accessible only under protocols established by the relevant Māori authority. This is achievable through a combination of encryption key management and geo-constrained storage configuration:

Design Pattern 3.3-A: Tribal Key Management Architecture Encrypt all MCH-class data at rest using AES-256. Store encryption keys in a Hardware Security Module (HSM) owned and operated by the iwi authority, not by the cloud provider. Configure the cloud storage service to call out to the tribal HSM for key retrieval on each data access request. This ensures that even under a US CLOUD Act compulsion order, the cloud provider cannot decrypt MCH-class data without the cooperation of the iwi authority, which is not subject to US jurisdiction. Store data in AWS NZ, Azure New Zealand North, or equivalent New Zealand-region infrastructure. Implement geo-fencing at the storage policy layer to prevent replication outside New Zealand without explicit iwi authority approval.

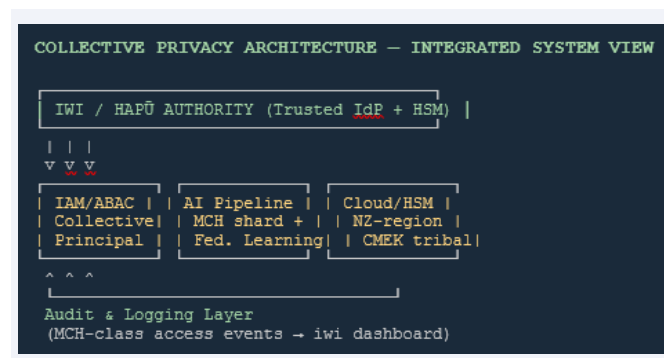
This architecture is technically analogous to customer-managed encryption key (CMEK) implementations already offered by major cloud providers for enterprise customers with regulatory requirements including HIPAA-regulated health data in the United States and GDPR-regulated personal data in the European Union. The Māori

Pipeline Stage	Collective Privacy Integration Point
Data collection and corpus assembly	Mandatory provenance tagging: each corpus item tagged with data class, source community, consent status, and custodian identifier. Items tagged as Māori Cultural Heritage (MCH) class require a collective consent assertion from the relevant iwi authority before inclusion. Implement as metadata schema extending Dublin Core with Indigenous data sovereignty fields (ENRICH extension; cf. Hollowell & Nicholas, 2009).
Training data governance	MCH-class items partitioned into a separately governed training shard. Training on this shard proceeds only under a time-limited data use agreement signed by the iwi authority. Implement using federated learning architecture [29]: tribal data remains at tribal infrastructure nodes; gradient updates (not raw data) are shared with the central model. This prevents raw cultural data from ever leaving tribal jurisdiction.
Model evaluation and model cards	Extend Gebru et al.'s [12] Datasheet and Mitchell et al.'s [25], Model Card frameworks to include mandatory Indigenous data impact fields: estimated proportion of training corpus derived from Indigenous sources; consent status of Indigenous-origin data; communities consulted in evaluation; identified cultural bias risks.
Deployment and ongoing governance	For LLMs deployed in New Zealand public sector contexts, a Māori cultural review gate is required before production deployment. This mirrors the Algorithm Charter's signatory obligations but operationalises them as a hard technical gate (deployment blocked without signed cultural impact assessment) rather than a voluntary policy commitment.
Machine unlearning (right to erasure)	Where MCH-class data is determined to have been included without valid collective consent, implement selective synaptic dampening using gradient ascent on identified training samples [7], or shard-based retraining [8]. Acknowledge that no current method guarantees complete erasure; this should be disclosed as a residual risk in the model card and addressed through upstream consent governance.

sovereignty use case requires an adaptation of the same pattern: the iwi authority acts as the key custodian rather than an enterprise IT department. The technical feasibility is established; the implementation gap is institutional and political, not engineering.

At the access control layer, a zero-trust network architecture (NIST SP 800-207) in which the iwi authority acts as the identity provider provides a further enforcement mechanism: no access to MCH-class data is permitted without a valid assertion from the tribal IdP, regardless of where the requesting party sits in the network. This approach has been prototyped in First Nations contexts in Canada [30], and provides a directly applicable model for Aotearoa New Zealand.

The following diagram illustrates the integrated architecture across all three system classes:



The audit and logging layer is not a peripheral addition: for collective privacy enforcement, the iwi authority requires real-time visibility into all access events affecting MCH-class data. Implementing this as a dedicated logging pipeline — with access logs streamed to an iwi-controlled dashboard in near-real-time — transforms data sovereignty from a policy statement into a live operational capability. This mirrors the approach taken by the First Nations Information Governance Centre’s OCAP® (Ownership, Control, Access, and Possession) principles, which require operational, not merely aspirational, implementation [31].

TOWARDS AN INDIGENOUS-LED FRAMEWORK FOR COMPUTATIONAL PRIVACY

Addressing these failure modes requires more than incremental reform. The following five principles, grounded in Māori customary law and adapted from existing Māori data sovereignty and Indigenous research frameworks, are proposed as design requirements for computational systems operating in contexts where Māori data rights are engaged. Each principle is stated at the level of policy and then operationalised as a technical design

requirement, responding to editorial feedback that the principles be demonstrated as technically feasible.

Principle 1: Collective Privacy by Design

Privacy impact assessments must be extended to evaluate collective privacy harms, not only individual ones. Systems processing data about Māori individuals must assess the implications of that processing for associated whānau, hapū, and iwi.

Technical operationalisation: Extend the standard privacy impact assessment (PIA) template to include a Collective Impact Assessment (CIA) section. The CIA requires system designers to identify: (a) all Māori collective entities whose interests are engaged by the data processing; (b) the nature of the collective interest (genealogical, cultural heritage, land rights, other); (c) the consent and governance mechanism through which that collective interest is represented. The CIA section should be completed with direct input from the relevant iwi authority, not merely reviewed against policy documentation.

Principle 2: Culturally Informed Consent Architecture

Consent mechanisms must accommodate collective decision-making protocols. Where data relates to shared cultural heritage, sacred knowledge, or whakapapa, consent must be obtainable at the collective level.

Technical operationalisation: Implement group consent workflows using the W3C Verifiable Credentials specification, with the iwi authority as the issuing entity. A group consent VC is issued by the iwi governance body following tikanga-compliant internal deliberation, scoped to a specific data use purpose and time period, and presented to the data processor as a cryptographically verifiable consent assertion. This approach is technically analogous to the B2B consent federation models used in regulated industries (e.g., financial data sharing under Open Banking) and does not require novel cryptographic infrastructure.

Principle 3: Mandatory Indigenous Representation in AI Training Governance

Any AI system trained on data including te reo Māori, traditional cultural knowledge, or data about Māori individuals must incorporate Māori governance of the training process.

Technical operationalisation: Implement using the federated learning architecture described in Section 3.2.

MCH-class training data remains at iwi infrastructure; the central model team receives gradient updates only. An iwi Data Governance Committee reviews and approves each training run against an agreed data use protocol before gradients are released. This is not merely consultation: the iwi authority holds a technical veto on data release.

Principle 4: Jurisdictional Sovereignty for Sensitive Indigenous Data

Computational systems must provide technical mechanisms for Māori designated sensitive data to be stored exclusively within New Zealand, under New Zealand legal jurisdiction, and accessible only under protocols established by the relevant Māori authority.

Technical operationalisation: Implement the tribal CMEK + geo-fencing architecture described in Section 3.3. All MCH-class data must be tagged at creation with a data residency policy label. Storage platform configuration enforces: (a) residency in New Zealand region; (b) encryption with keys held by the iwi HSM; (c) access blocked without a valid iwi IdP assertion. Cloud provider contractual terms must explicitly acknowledge and preserve these constraints. Regular automated audits verify that no MCH-class data has been replicated outside the approved residency boundary.

Principle 5: Genealogical Data Classification

Computational systems that process genealogical data must implement classification schemes that recognise Māori lineage records as a distinct data category carrying heightened collective sensitivity.

Technical operationalisation: Extend standard data classification schemas (e.g., Microsoft Information Protection labels, AWS Macie classification policies) to include a Whakapapa Sensitivity class at the highest protection tier, above standard 'Highly Confidential' classifications. Implement automated detection heuristics for whakapapa-pattern data (structured genealogical records, named lineage chains, land partition references) using named entity recognition models fine-tuned on Māori genealogical data structures. Data matching this classification is automatically tagged and subjected to the full MCH-class governance regime. Commercial genealogy platforms operating in New Zealand should be required to implement this classification as a condition of operating under the Biometric Processing Privacy Code 2025 and any successor instruments.

THE TREATY OF WAITANGI AS A BINDING COMPUTATIONAL STANDARD

The obligations of the Treaty of Waitangi, New Zealand's

founding constitutional document, signed between the British Crown and Māori tribal leaders in 1840 are not aspirational statements. The Waitangi Tribunal is a permanent commission of inquiry established under the Treaty of Waitangi Act 1975 to investigate alleged breaches of Treaty principles. Its 2023 report (WAI 2522), which specifically examined Māori data governance, found that the Crown's obligations extend to the digital domain and that failure to ensure Māori self-determination over Māori data constitutes a Treaty breach.

For computational scientists, system designers, and technology policy makers, this has direct implications. Any system deployed by or on behalf of the New Zealand Government including contractor-built systems, publicly funded research infrastructure, and AI tools embedded in public services must be assessed against Treaty obligations before deployment. This is a model for how computational systems should be assessed against treaty and constitutional obligations in any jurisdiction where such obligations exist, including Canada under Section 35 of the Constitution Act 1982.

The critical point for this journal's readership is that Treaty compliance is a technical requirement, not merely a legal or political one. A system that cannot implement culturally appropriate consent protocols, cannot enforce data residency for sensitive Indigenous data, and cannot be audited for disproportionate algorithmic impact on Māori has a design deficiency, one that the architectural patterns described in Section 3 are specifically designed to address.

This article's position is consistent with the postcolonial technology studies literature's critique of the politics of infrastructure. Couldry and Mejias' [32-44], analysis of data colonialism frames the extraction of data from marginalised communities by dominant technological actors as a structural continuation of colonial resource extraction. The architectures proposed here are, in part, a technical response to that structural critique: by placing encryption keys, consent governance, and training data control within the hands of iwi authorities, they shift the power geometry of data governance rather than merely inscribing sovereignty aspirations in policy documents.

CONCLUSION

The privacy challenges facing Māori communities in the age of computational systems are not peripheral concerns at the margins of the discipline. They represent a fundamental test of whether computational science is capable of designing systems that serve all people rather than encoding the assumptions of the majority.

The five failure modes identified in this article: biometric misidentification, algorithmic bias, inappropriate data jurisdiction, unconsented AI training, and inadequate genealogical data classification are the product of architectures designed without Indigenous communities at the table.

This article has demonstrated, in response to editorial review, that collective privacy is not merely an aspirational framework but an engineering objective. The collective principal IAM extension, the federated learning governance architecture, the tribal CMEK and geo-fencing model, and the Whakapapa Sensitivity classification scheme are concrete, technically feasible proposals that can be specified, implemented, and audited. Each draws on established engineering patterns already in production use in analogous regulated-data contexts; the innovation lies in their adaptation to the specific requirements of Māori data sovereignty, not in the novelty of the underlying technical mechanisms.

Remedying the structural failures documented here requires Indigenous leadership in computational system design, not as cultural advisors appended to projects already under development, but as equal partners holding technical veto rights in the earliest stages of design. It requires the Treaty of Waitangi to be treated as a binding technical standard, not a political courtesy. In addition, it requires the computational science community to recognise that privacy is not a solved problem when it is solved only for those whose value systems built the current frameworks.

Whether computational science becomes part of the solution or, by default, part of the problem depends on choices being made right now in system architecture, training data governance, and privacy-by-design standards. This article calls on the computational science community to make those choices consciously, and with full awareness of what is at stake.

REFERENCES

1. New Zealand Government. Privacy Act 2020. New Zealand Parliamentary Counsel Office. 2020.
2. Etzioni A. *The Limits of Privacy*. Basic Books. 1999.
3. Taylor C. *Sources of the Self: The Making of the Modern Identity*. Harvard University Press. 1989.
4. Hartzog W. *Privacy's Blueprint: The Battle to Control the Design of New Technologies*. Harvard University Press. 2018.
5. Nissenbaum H. Privacy as Contextual Integrity. *Washington Law Review*. 2004; 79: 119-157.
6. Solove DJ. *Understanding Privacy*. Harvard University Press. 2008.
7. Cao Y, Yang J. Towards Making Systems Forget with Machine Unlearning. *Proceedings of the IEEE Symposium on Security and Privacy*. 2015.
8. Bourtole L, Chandrasekaran V, Choquette-Choo CA, Jia H, Travers A, Zhang B, et al. *Machine Unlearning*. *Proceedings of the IEEE Symposium on Security and Privacy*. 2021.
9. Bowker GC, Star SL. *Sorting Things Out: Classification and Its Consequences*. MIT Press. 1999.
10. Winner L. Do Artifacts Have Politics?. *Daedalus*. 1980; 109: 121-136.
11. Grother P, Ngan M, Hanaoka K. *Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects*. NIST Interagency Report (NISTIR) 8280. 2019.
12. Buolamwini J, Gebru T. *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*. *Proceedings of the Conference on Fairness, Accountability and Transparency (FAT*)*. 2018.
13. Cavoukian A. *Privacy by Design: The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario. 2009.
14. Eubanks V. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press. 2018.
15. Noble SU. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press. 2018.
16. Benjamin R. *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity Press. 2019.
17. New Zealand Government. *Algorithm Charter for Aotearoa New Zealand*. Stats NZ. 2020.
18. Taiuru K. *State of the Nation: Māori Data Sovereignty and Data Governance in 2025*. Taiuru & Associates Ltd. 2025.
19. Taiuru K. *When Clinical Algorithms Don't See Us: Māori Data Sovereignty Approaches to Detecting and Mitigating Bias in Health AI*. Taiuru & Associates Ltd. 2025.
20. United States Congress. *Clarifying Lawful Overseas Use of Data (CLOUD) Act*. Pub. L. 2018. 115-141.
21. Daskal J. Microsoft Ireland, the CLOUD Act, and International Lawmaking 2.0. *Stanford Law Review Online*. 2018; 71: 9-16.
22. Bender EM, Gebru T, McMillan-Major A, Shmitchell S. *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Proceedings of FAccT 2021*. 2021.
23. Denton E, Hanna A, Amironesei R, Smart A, Nicole H. *On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet*. *Big Data & Society*. 2021; 8.
24. Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Daumé H, et al. *Datasheets for Datasets*. *Communications of the ACM*. 2021; 64: 86-92.
25. Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. *Model Cards for Model Reporting*. *Proceedings of FAccT 2019*. 2019.
26. Humbert M, Ayday E, Hubaux JP, Telenti A. *Addressing the Concerns of the Lacks Family: Quantification of Kin Genomic Privacy*. *Proceedings of ACM CCS 2013*. 2013.
27. Sporny M, Longley D, Chadwick D. *Verifiable Credentials Data Model v1.1*. W3C Recommendation. 2022.
28. Jones PL, Te Hiku Media. *Kaitiakitanga: A Model for Indigenous Data Sovereignty*. 2021.
29. McMahan HB, Moore E, Ramage D, Hampson S, Agüera y Arcas B. *Communication-Efficient Learning of Deep Networks from Decentralized Data*. *Proceedings of AISTATS 2017*. 2017.

30. First Nations Technology Council. Reclaiming Our Digital Futures: Indigenous Data Infrastructure in Canada. First Nations Technology Council. 2023.
31. First Nations Information Governance Centre (FNIGC). Ownership, Control, Access and Possession (OCAP®): The Path to First Nations Information Governance. FNIGC. 2014.
32. Couldry N, Mejias UA. The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism. Stanford University Press. 2019.
33. Anderson J, Christen K. "Decolonizing" Attribution: Traditions of Exclusion. *J Radical Librarianship*. 2019; 5: 113-152.
34. Office of the Privacy Commissioner. Biometric Processing Privacy Code 2025. Office of the Privacy Commissioner of New Zealand. 2025.
35. New Zealand Government. Digital Identity Services Trust Framework Act 2023. New Zealand Parliamentary Counsel Office. 2023.
36. Carroll SR, Garba I, Figueroa-Rodriguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE Principles for Indigenous Data Governance. *Data Science J*. 2020; 19: 43.
37. Hudson M, Farrar D, McLean L. Tribal Data Sovereignty: Whakatohea Rights and Interests. In T. Kukutai & J. Taylor (Eds.), *Indigenous Data Sovereignty: Toward an Agenda*. ANU Press. 2016.
38. Taiuru K. Compendium of Maori Data Sovereignty, Version 2. Taiuru & Associates Ltd. 2020.
39. Taiur, K. A Māori Framework for Describing Artificial Intelligence Agents. Taiuru & Associates Ltd. 2023.
40. Taiuru K. Critical Analysis of the Maori Data Sovereignty Network's Data Principles. Taiuru & Associates Ltd. 2024.
41. Taiuru K. An Indigenous Privacy Perspective from New Zealand. Presentation at the International Association of Privacy Professionals ANZ Summit 2024, Melbourne, Australia. 2024.
42. Te Mana Raraunga (Māori Data Sovereignty Network). Principles of Māori Data Sovereignty. 2018.
43. Waitangi Tribunal. Oranga Tamariki Urgent Inquiry (WAI 2522). Legislation Direct. 2023.
44. Walter M, Andersen C. *Indigenous Statistics: A Quantitative Research Methodology*. Left Coast Press. 2013.