



The Sovereignty Assessment Instrument

Pacific Ocean off Amberley, North Canterbury · © John Stroh

Version 0.3 · revised 21 August 2026 · [this version as a PDF](#)

Drafted with AI assistance, then checked and revised by the author. The judgements and the errors are the author's.

Status of these claims

What this publication does not claim, and what is outstanding against it in the register.

A question this rests on is parked: What do we do with an application that refuses to run inside the boundary?

We do not claim that a member's existing applications can be made to run inside the boundary. We claim only that the boundary reveals which ones cannot.

What this is

A diagnostic, not an argument. Six tests, each with a pass condition, a stated evidence requirement, and — most usefully — a description of how the test can appear to pass when it has not.

It is deliberately vendor-neutral and party-neutral, and it works at any scale. Anyone may apply it: an accountancy firm assessing its practice software, a school board assessing a learning platform, an official assessing a procurement, a journalist assessing a policy announce-

ment, a party assessing its own platform. The questions are the same for a four-person business and a government department; only the consequences differ. It is written so that a competent person with no technical background can apply it and defend the result.

How to use it. Score each test PASS, PARTIAL or FAIL against the evidence requirement, not against an assurance statement. A test scores PASS only where the evidence exists and can be inspected. “The vendor states that...” is not evidence; it is the thing evidence is needed for.

How to read the findings below. Each carries a marker.

- **EVIDENCED** — sources are cited and can be checked.
- **NOT EVIDENCED** — we searched, describe what we searched, and did not find. Absence of a finding is not proof of absence, and it is reported as a gap rather than dressed as a conclusion.

A test can fail for either reason, and the difference matters. An organisation that cannot answer a question about itself is in trouble regardless of what the answer would have been — but saying so accurately is the difference between analysis and assertion.

The overall finding is the lowest score, not the average. Sovereignty is not a portfolio. An organisation that satisfies five tests and fails the sixth is captive by way of the sixth.

FIG-01 THE SIX TESTS • AOTEAROA NEW ZEALAND • AUGUST 2026

00	Discovery	FAIL
Can the organisation find out whether it is using AI at all?		
01	Jurisdiction	FAIL
Data located here AND operator beyond foreign compulsion		
02	Compute	PARTIAL
Affordable compute for public-purpose work, outside the research system		
03	Non-extraction	PARTIAL
Protection that is verifiable, and survives a change of terms		
04	Portability	FAIL
Complete operational state moves, with a measured time		
05	Refusal	FAIL
Decline a unilateral change of terms and still discharge statutory functions		
Overall		FAIL
The finding is the lowest score, not the average. Sovereignty is not a portfolio.		

FIG-01 Six tests for whether a country has AI sovereignty or only AI adoption, beginning with whether it can tell. New Zealand's current position, with each verdict against its own row. The finding is the lowest score, not the average.

Test 0 — Discovery

FIG-09 NOBODY CHOSE THIS



FIG-09 You did not buy AI. You bought a membership app, an accounting package, a course platform — and each one quietly sends your work to a model somewhere else. This is how you find out which ones.

Question. Can the organisation find out whether it is using AI at all?

Pass condition. The organisation can determine, for every system it operates, whether that system performs inference; which model, operated by whom, in which jurisdiction; under what data terms; and will be notified when the answer changes.

Evidence required. A register covering every product in the data path, with dated entries. Where a vendor will not disclose, that refusal recorded as the entry.

Why this is Test 0 and not Test 6. The remaining five tests assume the organisation knows it is procuring AI and is choosing a provider. That assumption held in 2023. It does not hold now. AI arrives as a property of software already bought — the membership app, the accounting suite, the case management system, the online course platform. Nobody procured it, nobody was shown terms, and it changes by subprocessor update notice.

Without Test 0, the other five are not failed. They are **unanswerable**, which is a different and more serious condition. An organisation reporting PASS on Tests 1 through 5 while unable to answer Test 0 has assessed the part of its estate it happens to know about.

Common false pass — assessing the application layer. A SaaS product hosted in New Zealand may route inference offshore. Data location assessed at the application layer is silent about where inference happens, and we have not found the question on a procurement form in current use.

Second false pass — the no-training term one hop away. The term exists between the model vendor and the software vendor, not between the software vendor and the customer. The customer's supplier may not know what their subprocessor does, and is under no obligation to find out.

New Zealand, August 2026: FAIL — EVIDENCED.

We searched the Government Procurement Rules (5th edition), the digital government *Procurement and GenAI* guidance, and Office of the Privacy Commissioner AI guidance. **No requirement was found for a supplier to declare AI embedded in software supplied to an agency.** What exists is adjacent: guidance on how to deliberately buy an AI system, and an expectation that buyers integrate privacy and algorithmic impact assessments into procurement. Neither addresses AI arriving inside software bought for something else. Industry commentary on “AI disclosure in procurement” concerns suppliers declaring whether AI helped write their tender — a different question.

For contrast, draft United States GSA acquisition terms would require contractors to disclose all AI systems used in performing a contract within thirty days of award, not only those sold to the government. That is what the requirement looks like where it exists.

And the consequence is documented here. Reporting during 2026, based on documents obtained under the Official Information Act, describes Microsoft Copilot becoming the default AI tool across New Zealand public sector agencies by being added to existing Microsoft 365 agreements rather than procured separately — under an umbrella agreement running from October 2024 to September 2027. One ministry IT manager is quoted saying Copilot was presented as a feature add-on covered by the existing cloud risk assessment, and that nobody asked whether AI processing changed the data sovereignty profile. Two subsequent episodes — a Department of Corrections rebuke over Copilot use on Extended Supervision Order casework, and a Ministry for Regulation report produced with Copilot that carried classification errors — followed.

This strand is reported in technology press rather than drawn from primary documents. The underlying OIA material, the Corrections statement and the Ministry for Regulation report are all obtainable and we have not obtained them. Treat the pattern as reported and the specifics as requiring verification.

Test 1 — Jurisdiction

FIG-02 LOCATION IS NOT JURISDICTION

	DATA LOCATED IN NZ	OPERATOR BEYOND COMPULSION
Hyperscaler region in Auckland Foreign parent; foreign assets, foreign legal process	✓	✗
Offshore SaaS with NZ contract Neither condition. Widely procured.	✗	✗
NZ-incorporated operator; NZ infrastructure No foreign parent, no foreign assets, no disclosure obligation	✓	✓

Only the third arrangement is sovereign. The first is sold as though it were — region-washing.

FIG-02 The two-part sovereignty test applied to three common arrangements. Only the third qualifies; the first is widely sold as though it did.

Question. Can sensitive data be stored and governed under an appropriate jurisdiction, with enforcement?

Pass condition. Data is physically located in the jurisdiction, **and** the entity operating the infrastructure is beyond foreign legal compulsion — no foreign parent, no foreign controlling shareholder, no foreign assets exposed to seizure, no supplier obligation to disclose on foreign demand. Both conditions. Either alone is insufficient.

Evidence required. Facility locations and replication targets. Constitution, share register, directors' interests. Audited statement of asset location. Supplier contract schedule with disclosure clauses extracted.

Common false pass — region-washing. A hyperscaler region in-country satisfies location and fails operator. The data is physically here and legally reachable, because compulsion follows the operator regardless of where the disk spins. This is the most widespread false pass in the market and it is usually made in good faith.

Second false pass — attestation as evidence. A contractual assurance that the operator is not subject to foreign compulsion is a claim, not proof. The evidence is the corporate structure, not the statement about it.

New Zealand, August 2026: FAIL — EVIDENCED as to the instruments; NOT EVIDENCED as to practice. Information Privacy Principle 12 governs disclosure to an overseas recipient; the usual route is a judgement that the recipient is subject to comparable safeguards, which in practice rests on the vendor's own attestation. There is no residency requirement. IPP 12 makes no provision for Māori data governance. Cloud First provides that RESTRICTED information should move over time to New Zealand-based data centres *where a suitable onshore service exists* — a term we have not found authoritatively defined, and which on a location-only reading a hyperscaler region satisfies.

What we did not establish: how often agencies rely on the comparable-safeguards route in practice, or whether the Commissioner's transfer prohibition power has ever been exercised. The second is the more significant gap.

Test 2 — Compute

Question. Can an organisation access affordable compute for legitimate public-purpose work?

Pass condition. A public body, iwi organisation, school or community organisation can obtain sufficient compute for its own purposes at a cost it can meet, without a research grant and without a commercial contract that fails Test 1.

Evidence required. A named access route, published pricing or allocation criteria, and at least one organisation of that size and type that has used it.

Common false pass — the national facility. A large national compute investment satisfies researchers and does nothing for a marae or a two-partner practice. Ask who can actually get an allocation, on what criteria, and what happens to an applicant with no research affiliation.

Second false pass — assuming scale is required. The premise that sovereign compute means a data centre is now out of date. Efficiency-driven model development has produced capable models that run on hardware available at retail. The question is not whether the country owns a building; it is whether a small organisation can run capable models on infrastructure it controls.

New Zealand, August 2026: PARTIAL — EVIDENCED. The Institute for Advanced Technology's AI Research Platform, up to \$70 million over seven years, was decided in May 2026 with funding commencing July. Its stated purpose is research capability and commercialisation. There is no published allocation route for a council, marae or community organisation. Separately, Cloud First restricts public service agencies and some Crown entities from investing in on-premise ICT infrastructure without meeting specified criteria or obtaining approval from the Government Chief Digital Officer — a direct barrier to agency-controlled compute.

Test 3 — Non-extraction

Question. Can an organisation build or deploy a model, or use an application, without surrendering data for vendor training?

Pass condition. The protection exists, is verifiable by the customer, and survives a change in the vendor's commercial position or a foreign legal process.

Evidence required. The contractual term. A verification mechanism the customer can exercise. Evidence of at least one audit performed by someone other than the vendor.

Common false pass — the term is available, therefore the test passes. It is not wishful thinking to obtain a no-training term; major vendors state that enterprise and API data is not used for training by default, and offer zero-retention arrangements. The term is real. The gap is that it is written by the vendor, variable by the vendor, unverifiable by the customer, and subordinate to legal process in a foreign jurisdiction.

How that gap behaves under pressure. When a United States court ordered a major vendor to preserve output logs in 2025, enterprise and zero-retention customers were excluded. The exclusion held. It held because a US magistrate judge decided it should — not because a New Zealand customer had a right, was a party, was consulted, or would have been notified.

New Zealand, August 2026: PARTIAL — terms EVIDENCED, verification NOT EVIDENCED.

The terms are purchasable and published: major vendors state that enterprise and API data is not used for training by default, and Microsoft's commitment that enterprise Copilot prompts and responses do not train foundation models is reported in the same coverage as the adoption pattern above.

We searched for any New Zealand or Australian agency audit of such a claim and did not find one. That is a gap in our search, not proof that none has occurred. What we can say positively is that no statutory protection exists for public-sector data entering vendor training, and no notification obligation attaches to a change of terms.

Test 4 — Portability

Question. Can data, applications and workflows move to another provider at will?

Pass condition. The organisation's complete operational state — records, adapted models, configurations, integrations and audit trail — can be exported in a usable form and stood up elsewhere, with a measured time to do so.

Evidence required. A documented export. A measured migration time. Ideally, evidence of a migration actually performed.

Common false pass — open weights. Open model weights give model portability. They do not give organisational portability. The state that accumulates inside a hosted service over three years — embeddings, fine-tunes, agent configurations, integration surface, audit history — is the thing that cannot be moved, and it is not addressed by the licence of the underlying model.

Second false pass — multi-cloud as a policy statement. A requirement to avoid over-dependence on any single supplier is not portability unless somebody has measured the exit.

New Zealand, August 2026: FAIL — policy EVIDENCED, enforcement NOT

EVIDENCED. Cloud First requires agencies to implement a multi-cloud approach that avoids over-dependence on any specific supplier. We found no enforcement mechanism attached to it and no published measured exit times; we did not establish whether any agency holds such measurements internally. The open-weights ecosystem available as a fallback is substantially maintained outside the countries usually named in this debate, which is a dependency in its own right.

Test 5 — Refusal

Question. Can an organisation refuse a vendor's terms without losing the ability to govern?

Pass condition. A council, school, health service or agency can decline a unilateral change of terms and continue to discharge its statutory functions.

Evidence required. An instance. This test is passed by precedent, not by architecture — somebody has to have done it.

Common false pass — nobody has tried. Absence of a refusal is read as absence of a problem. It is more often absence of an option.

Why this is the load-bearing test. The other five describe conditions. This one describes the consequence. Once an organisation cannot refuse, the vendor has become part of how the country is governed, and no one voted on it.

New Zealand, August 2026: FAIL — NOT EVIDENCED.

We searched for any instance of a New Zealand public body declining a unilateral terms change from a major platform vendor and retaining function. **We did not find one.** That is not proof that none occurred; refusals are rarely announced.

What the search did surface points the other way. In the Copilot pattern described under Test 0, the reported reason for adoption was that the capability was already inside an agreement signed for something else. Refusal was not weighed and rejected. It never arose as a ques-

tion — which is the condition this test exists to detect.

Scorecard — New Zealand, August 2026

TEST	FINDING	BASIS
0 · Discovery	FAIL	Gap evidenced; adoption reported second-hand
1 · Jurisdiction	FAIL	Instruments evidenced; practice not established
2 · Compute	PARTIAL	Evidenced
3 · Non-extraction	PARTIAL	Terms evidenced; verification not established
4 · Portability	FAIL	Policy evidenced; enforcement not established
5 · Refusal	FAIL	Not evidenced — searched, not found
Overall	FAIL	

Four of six findings rest partly on things we searched for and did not find. They are marked, and they are the places where a reader with better access than ours could change the result. We would rather publish a corrected instrument than a confident one.

New Zealand has AI adoption. On this instrument it does not have AI sovereignty. That is a statement about instruments and enforcement, not about the competence or intent of anyone currently working on it.

What would change each score

Test 0 — require an inference-path declaration in the software supply chain: a bill of materials extended from security to AI. Without this, every other remedy applies only to what an organisation already knows about.

Test 1 — define “suitable onshore service exists” by the two-part condition above, and make jurisdictional risk assessment binding rather than advisory.

Test 2 — publish an allocation route for public-purpose compute outside the research system, and exempt sovereignty-qualifying infrastructure from the on-premise investment restriction.

Test 3 — a statutory prohibition on public-sector data entering vendor training, with an audit right and a notification obligation on change of terms. Because training is irreversible, disclosure regimes cannot remedy it after the fact.

Test 4 — require measured, published exit times as a procurement condition. Not a portability policy: a number, and a date it was last tested.

Test 5 — Test 5 passes when Tests 0 to 4 pass, and not before.

Sources and verification

Findings above draw on: the New Zealand government Cloud First policy as published at digital.govt.nz (refreshed April 2023); the Cloud Jurisdictional Risk guidance (updated 2024); the All-of-Government Cloud Sourcing Strategy; MBIE's AI Strategy (July 2025) and the July 2024 Cabinet paper on the approach to work on artificial intelligence; MBIE and Beehive material on the Institute for Advanced Technology AI Research Platform; and the published enterprise privacy terms of major model vendors. Retrieved 20 August 2026.

Verify before citing. Government policy pages change without version notice, and vendor terms change without notice at all.

What we could not establish, listed so it is not mistaken for settled:

1. Whether the Privacy Commissioner's transfer prohibition power has ever been exercised. This is the strongest cross-border provision in the Act and its use or non-use would materially affect Test 1.
2. Whether any New Zealand or Australian agency has audited a vendor's no-training claim rather than accepting the attestation. Test 3.
3. Whether any agency holds measured exit times internally, unpublished. Test 4.
4. Whether any public body has refused a unilateral terms change and retained function. Test 5.
5. The primary documents behind the Copilot procurement reporting — the OIA material, the Corrections statement, the Ministry for Regulation report. Test 0's supporting evidence is currently second-hand.

Where a finding here is wrong, the instrument is more useful if you say so than if you work around it.

Use and licence

Creative Commons Attribution 4.0 International (CC BY 4.0).

Copy it, adapt it, fold it into a procurement standard, publish a corrected version. Attribution is the only condition: *Sovereignty Assessment Instrument*, *agenticgovernance.digital*. No share-alike, deliberately — a share-alike condition would give an agency or a standards body a reason to hesitate before incorporating this into their own material, and wide adoption matters more here than downstream control.

No endorsement. Applying this instrument does not imply that its authors endorse the result, and nobody may represent a self-assessment as approved, certified or verified by us. The instrument tests whether an organisation can answer the questions. It confers nothing.

Corrections are more useful than agreement. Where a finding here is wrong, saying so improves the instrument. It is published to be argued with, and a version corrected by someone else is a better outcome than a version defended by us.

Cite this version

The Sovereignty Assessment Instrument, version 0.3 (August 2026).
`agenticgovernance.digital/downloads/sovereignty-test-v0-3.pdf`

Licensed CC BY 4.0. Reuse is free, including commercially, provided the author is credited and the reuse does not suggest that the author endorses it. This link is frozen at version 0.3; [/sovereignty-test](#) always shows the current one.

PREVIOUS

**A · What happened to
your software**

Series contents

NEXT

**C · Shared AI
infrastructure**

Alongside: [questions and answers](#) · [glossary](#) · [sources and provenance](#) · [slides](#)