MISSIONOPSAI

A POSITION PAPER · FRAMEWORK EDITION

The State of Sovereign AI

Why sovereignty is now a test, not a slogan — a framework, a proof, and a challenge to self-reflect.

PUBLISHED BY

MissionOpsAI Ltd
Sovereign AI governance & assurance

EDITION

Framework — v0.1
The findings follow · June 2026

THE CLAIM

"Sovereign" has become the most-claimed and least-verified word in AI

Across UK defence, critical infrastructure, healthcare and finance, organisations are told their AI is "sovereign." Most of the time the word is doing more work than the architecture behind it. In an era when decision-making, public services and commerce all run on data and AI, control of those two things is the capacity to control everything that depends on them. Sovereignty, therefore, is no longer a procurement nicety — it is a question of whether you, or someone else, holds the switch.

This paper does three things. It states a test any organisation can apply to itself. It publishes that test in full, before we score anyone, so the method can be checked rather than trusted. And it issues a challenge — not to name and shame, but to invite organisations to look honestly at their own exposure. The aggregate findings, once the assessment is complete, will follow in a later edition. This edition is the framework and the invitation.

WHY NOW — THE PROOF, NOT THE PREDICTION

The switch has already been thrown once

On 12 June 2026, the US Commerce Department placed two of the most capable AI models in the world — Anthropic's Claude Fable 5 and Mythos 5 — under export controls on national-security grounds. To comply, the company had to disable them for every user on earth, not only the foreign nationals the order named. A non-US organisation that depended on those models lost them overnight, through no fault of its own.

We draw no motive from this and predict nothing. The point is narrower and firmer: it demonstrates, on the record, that a government can lawfully switch off even a domestic provider's flagship capability, with global effect, in hours. That is not a forecast of what will happen — it is evidence of what can.

The structural reason. This is not an aberration. The leading AI developers and cloud platforms — through Pentagon frontier-AI contracts, classified-environment supply and enterprise cloud — are embedded with one nation's national-security apparatus. For any organisation outside that nation, the primary force behind its critical AI and data substrate is therefore a foreign national-security interest, present whether or not it is ever exercised.

Authorities engaged: the CLOUD Act and FISA §702 (compelled disclosure); IEEPA, OFAC sanctions and the Export Administration Regulations (compelled restriction or withdrawal, now reaching AI). Sources: contemporaneous reporting (TechCrunch, NBC News, TIME, CNN, Axios); US statute.

THE CORRECTION

Residency is not sovereignty — and sovereignty is a dial, not a gate

Most "sovereign" claims describe *residency* — data stored in the UK. But the defining question is not *where* the data sits; it is *who controls it*, and whether a foreign power can compel, withdraw or direct it. UK-stored data on a foreign-controlled platform is residency, not sovereignty.

Absolute sovereignty today is rare and demanding — but it is not the entry price. The achievable standard is **Sovereign Capable**: run governed, auditable AI now; use frontier models for non-sensitive work through a control layer that keeps sensitive work and data on sovereign substrate; and retain the ability to cut to UK or EU-jurisdiction or air-gapped capability the moment the requirement demands. Sovereignty becomes a dial you can turn, not a gate you cannot afford.

THE FRAMEWORK

The Compulsion Test — three questions any board can ask

Strip away the abstraction and sovereignty reduces to one scenario. *Assume a foreign power compels its AI, cloud and data companies — which it lawfully can — to withdraw their services, hand over the data they hold, and act through the companies they own.* Three questions follow.

Q1 · CONTINUITY

Could you still operate?

Scored layer by layer — productivity, identity, model, compute, storage, comms. The chain breaks at its weakest link.

Q2 · EXPOSURE

What do they already hold?

The mosaic: individually trivial exchanges that, aggregated in foreign-controlled hands, form an intelligence picture before any compulsion occurs.

Q3 · CONTROL

Can they make you?

Ownership and funding: a sovereign stack owned by a foreign-controlled entity can still be directed or acquired. Control rights matter more than headline percentage.

The test is deliberately operational. It does not ask whether a system *feels* sovereign; it asks whether it survives a switch-off, a data demand and an instruction routed through ownership. The full method — evidence tests, hosting tiers, control thresholds — is published openly so anyone can apply or check it.

THE TWO QUESTIONS FOR THE BOARDROOM

What the test tells you about yourself

> **Are you a commercial-continuity risk to your own clients?** If a foreign switch-off would halt you, then everyone who depends on you depends on a decision made in another jurisdiction. The absence of a sovereign strategy is itself the liability.

> **Are you, by architecture, a national-security exposure?** If your work matters to the nation and all of it passes through foreign-controlled platforms, the exposure is real — not by intent, but by design. This question applies where data sensitivity and foreign dependence coincide; it is not asked of everyone.

These are not accusations. They are questions every responsible board should be able to answer — and most cannot, yet.

THE DESTINATION

The North Star is reachable — it is simply rarely crossed

A sovereign substitute exists today for every critical layer: sovereign inference, UK or EU operator-controlled hosting, self-hosted identity and storage, and sovereign productivity suites in place of the foreign default. Absolute sovereignty is therefore achievable now by composition. The binding constraint is not the absence of components — it is the cost and will to swap the deepest lock-in, the productivity substrate, and to prove the result under the test. Absolute sovereignty is not unreachable; it is reachable and rarely crossed, and the Compulsion Test is the proof of crossing.

The sovereignty spectrum — Exposed → Residency → Sovereign Capable → Absolute Sovereignty. The control layer routes work by sensitivity; the substrate is swappable, the governance constant.

THE CHALLENGE — AND OUR OWN COMMITMENT

Test yourself. We have tested ourselves first.

THE CHALLENGE WE PUT TO EVERY ORGANISATION CLAIMING SOVEREIGN STATUS

Can you operate Sovereign Capable — and could you swap out the risk substrate if you had to?

We hold ourselves to this before anyone else. MissionOpsAI publishes its own scorecard against this method first; the rubric is open and reproducible; the assessment function is walled from our certification business under a published independence policy; and no organisation is named in any aggregate finding. Every organisation assessed is offered its own confidential result and a right of reply. A finding is never a verdict that you cannot be sovereign — it is an invitation to become so, with a defined route to verification.

What follows. Once a representative sample has been assessed against the published method, the aggregate *State of Sovereign AI* findings — how many UK organisations claim sovereign status, how many can substantiate it, and where the gaps concentrate — will be published as the next edition. This paper deliberately reports no such figures, because we have not yet earned them.

THE STANDARD WE HOLD

The exposure is real, and we will make it tangible — never frightening. Every claim rests on a named authority or a dated fact, calibrated to its true weight, and paired with a way to fix it. Control of data and AI is the power to control what depends on them. Sovereignty is simply the answer to one question: do you hold the switch, or does someone else?

Disclaimer. These are MissionOpsAI's own standards — the test we hold ourselves to first, and the basis on which we offer assessment to others. They are not regulatory or official guidance and do not represent the position of any government, regulator or standards body. The landscape and its guidance are developing; we will update these standards as formal guidance emerges. Nothing here is legal advice; organisations should take their own. A MissionOpsAI position paper, framework edition v0.1, June 2026; method published separately and in full; no audit statistics — those follow the assessment.