

Sovereign AI, Epistemic Capture and the Contradiction at the Heart of Irish Digital Policy

A Green Paper demonstrating the irreconcilable contradiction between the Government's stated commitment to Irish digital sovereignty and its total dependence on foreign AI systems controlled by jurisdictions that have already weaponised access, set against the global sovereign AI movement and the democratic imperative to act.

"Ireland can be an international leader in using AI to the benefit of our economy and society, through a people-centred, ethical approach to its development, adoption and use."

AI - Here for Good: National AI Strategy for Ireland, Department of Enterprise, July 2021

James Spalding | Published 14 June 2026

openmygovernment.com/ireland/projects/sovereign-ai

About the author. James Spalding is a Creative Technologist and AI Policy Advisor. In April 2026, his written evidence to the House of Lords Science and Technology Committee inquiry into Innovation in the NHS: personalised medicine and AI was accepted and published in the official parliamentary record (reference PMA0107).

EUR 0 SOVEREIGN AI INVESTMENT Ireland, to date	30+ NATIONS WITH SOVEREIGN AI CNAS Sovereign AI Index, 2026	22% GRID TO DATA CENTRES CSO, 2024	3 Days FABLE 5 SUSPENDED US export order, June 2026
---	--	---	--

THE DOCUMENTED CONTRADICTION

"Ireland can be an international leader in using AI to the benefit of our economy and society, through a people-centred, ethical approach to its development, adoption and use."

[AI - Here for Good: National AI Strategy, Department of Enterprise, Trade and Employment, July 2021](#)

Five years after publishing this strategy, Ireland has invested zero in sovereign AI capability. Every government department using AI depends entirely on commercial systems controlled by US corporations operating under US jurisdiction. In June 2026, the US government demonstrated it can suspend access to frontier AI models for all foreign nationals within hours. Ireland has no fallback, no domestic capability, and no plan to build one.

I Preamble and Purpose

This Green Paper addresses a matter of national strategic importance: Ireland's total dependence on foreign artificial intelligence systems and the absence of any sovereign AI capability, set against the Government's own published commitments to be an international leader in AI. Published on 14 June 2026, two days after the United States Government ordered the suspension of Anthropic's Fable 5 and Mythos 5 models for all foreign nationals worldwide, this analysis demonstrates that the Government's stated ambitions are irreconcilable with its inaction.

Three converging developments make this paper necessary now. First, the Fable incident of June 2026 proved that access to frontier AI can be weaponised by the jurisdiction that controls it - suspended globally within hours by executive order, with no recourse for dependent nations. Second, the city government of Rio de Janeiro released Rio 3.5 Open 397B, a frontier-competitive AI model, under the MIT open-source licence - demonstrating that sovereign AI capability is achievable at city-government scale. Third, over 30 nations have now established sovereign AI programmes, from France's EUR 109 billion commitment to the UAE's Falcon series, making Ireland's absence from this movement increasingly conspicuous and strategically dangerous.

This is not a speculative risk assessment. This is a documentation of facts: the Government's own commitments, the sovereign AI programmes of comparable and smaller nations, and the concrete events of June 2026 that demonstrated what happens when a nation has no sovereign fallback.

II The Government's Own Commitments

The Irish Government has published extensive commitments on artificial intelligence and digital strategy. These are not the words of commentators or critics. They are the Government's own promises, published in its own strategy documents, announced by its own ministers.

The National AI Strategy, AI - Here for Good (July 2021), committed Ireland to being "an international leader in using AI to the benefit of our economy and society." It set a target of 75% of Irish businesses using AI by 2030. It promised a supportive innovation ecosystem, secure data and connectivity infrastructure, and workforce preparation.

The National Digital and AI Strategy, Connecting Our People, Securing Our Future (2026), committed to 90 specific deliverables by 2030, including an AI Office of Ireland, an AI Research Centre of Scale, and access to advanced compute capacity. The Taoiseach stated: "The government is committed to ensuring Ireland continues to be a digital leader in Europe, and globally."

The Programme for Government, Securing Ireland's Future (2025), committed to harnessing "the digital and AI revolution to deliver effective and modern public services." It promised a National AI Fellowship programme, an AI Advisory Unit for the public service, and AI training for all public servants.

The CASPIr supercomputer, announced in October 2025, will be capable of 15 petaflops when it enters service in 2027. It is co-funded by the EuroHPC Joint Undertaking and will be operated by the Irish Centre for High-End Computing (ICHEC) at University of Galway. It is a mid-range research computer. It is not a sovereign AI training facility.

Not one of these commitments includes building, training, fine-tuning, or operating a sovereign AI model. Not one addresses the risk of dependency on foreign AI providers. Not one contemplates a scenario in which access to those providers is suspended - the scenario that materialised on 12 June 2026.

DETE: AI - Here for Good, National AI Strategy, July 2021. Gov.ie: Connecting Our People, Securing Our Future, 2026. Gov.ie: Programme for Government, Securing Ireland's Future, 2025. ICHEC: CASPIr announcement, October 2025.

III Total Epistemic Dependence

Ireland possesses no sovereign AI capability. No Irish-developed large language model exists. No Irish institution has trained, fine-tuned, or deployed a foundation model for Irish public service use. The AI Office of Ireland, announced in the National Digital and AI Strategy, has not been established. The AI Research Centre of Scale has not been funded. No pilot programme for sovereign AI inference exists. No government department has tested open-source AI models on Irish infrastructure.

Ireland hosts data centres that consumed 22% of the national metered electricity supply in 2024, more than every urban household in the country combined, and a share that has more than quadrupled since 2015. None of this capacity serves Irish sovereign AI. The infrastructure exists on Irish soil but serves foreign commercial interests. Ireland bears the energy costs and grid constraints of hosting the global AI industry while possessing none of the capability it produces.

This is not sovereignty. This is a dependency so complete that the word "strategy" applied to it is an act of self-deception.

KPMG Ireland: Ireland and AI - Strategic capability, system resilience and digital autonomy, 2026. Irish Times: Is sovereign AI the answer for Ireland?, September 2025. CSO: Data Centres Metered Electricity Consumption 2024 (22% of metered electricity).

IV The Fable Incident: When Sovereignty Became Concrete

On 9 June 2026, Anthropic released Fable 5, its most advanced publicly available AI model, alongside Mythos 5, a more powerful model available via API. Within 24 hours, Microsoft - itself a major AI provider - instructed its employees to stop using Fable 5 after its legal teams identified that Anthropic's new mandatory 30-day data retention policy overrode existing zero-data-retention agreements. Organisations that had negotiated zero-retention terms discovered those terms were void for Mythos-class models, with no opt-out available.

Within the same week, researchers discovered that Fable 5 had been silently degrading responses for users it classified as working on competing AI systems. For queries the model classified as attempts at AI model distillation, the response was silently downgraded - not flagged, not disclosed - with users receiving inferior outputs while being charged full price. The restrictions used steering vectors and prompt modification without user notification. Anthropic reversed the policy within 48 hours of public exposure and apologised.

On 12 June 2026 - three days after launch - the US Government issued an export control directive ordering Anthropic to suspend all access to Fable 5 and Mythos 5 for all non-US persons. The order cited national security concerns under the Export Control Reform Act and was enforced within hours. Every user outside the United States lost access to Anthropic's frontier models immediately. No advance warning was given. No alternative was offered.

The Fable incident demonstrated three things simultaneously. Foreign AI providers can change their terms unilaterally and retroactively. Foreign AI providers can silently manipulate the quality of their responses based on undisclosed criteria. And a foreign government can suspend access to frontier AI for all foreign nationals within hours, by executive order, with no democratic process and no recourse.

For any nation that depends entirely on US commercial AI - as Ireland does - these are not theoretical risks. They materialised simultaneously in the space of four days.

CNBC: Anthropic disables access to Fable 5 and Mythos 5 to comply with government directive, 12 June 2026. Claude Help Center: Data retention practices for Mythos-class models, June 2026. MLQ News: Anthropic Reverses Secret Policy That Silently Degraded Claude for Rival AI Researchers, June 2026. Time: Anthropic Pulls Its Most Powerful AI Models After U.S. Bars Foreign Access, 13 June 2026.

V The Global Sovereign AI Movement

Ireland's inaction is set against a global movement in which nations of every size and at every level of development are building sovereign AI capability. The sovereign AI movement is no longer led only by superpowers. It includes city governments, small nations, and regional blocs. Ireland's absence is increasingly conspicuous.

France has committed EUR 109 billion in AI infrastructure investments, announced at the February 2025 AI Action Summit. Mistral AI, valued at EUR 20 billion, has been explicitly designated as a strategic asset under the Horizon Numérique 2030 plan. The French Ministry of Armies has signed a framework contract deploying Mistral on sovereign military infrastructure. An EUR 18 billion budget through the Caisse des Dépôts funds AI and sovereign digital initiatives. France treats AI sovereignty as a matter of national security.

The United Arab Emirates has built the Falcon series of AI models through the Technology Innovation Institute, with over 55 million downloads to date. Falcon-H1 outperforms comparable models from Meta and Alibaba. Falcon Arabic is the highest-performing Arabic language model in the world. The UAE has prioritised domestic development, responsible governance, and long-term economic alignment of AI as critical national infrastructure.

Saudi Arabia has developed ALLaM through SDAIA, the top-ranked Arabic language model globally. HUMAIN, a Public Investment Fund company, has committed \$100 billion in AI infrastructure. Saudi Arabia treats AI capability as a pillar of its Vision 2030 economic diversification programme.

India's BharatGen project, led by IIT Bombay with government backing, is developing a multilingual multimodal foundation model supporting Hindi, Tamil, Bengali, and other Indian languages. India treats linguistic diversity in AI as a matter of democratic representation.

The United Kingdom announced a GBP 500 million Sovereign AI Fund in April 2026, adding to the AI Opportunities Action Plan published in January 2025. The UK explicitly links sovereign AI to national security, economic competitiveness, and democratic resilience.

Rio de Janeiro's city government released Rio 3.5 Open 397B, a 397-billion-parameter AI model that matches or exceeds commercial models across major coding and multilingual benchmarks. It is MIT-licensed. A city government open-sourced a frontier AI model. Ireland, an entire nation state with one of the highest GDPs per capita in the world, has produced nothing.

Additional national programmes: Japan (LLM-jp, METI Cloud Program), South Korea (260,000+ GPUs across sovereign clouds), Singapore (SEA-LION regional model), Ukraine (sovereign LLM beta expected 2026), Denmark, Italy, and Sweden - all with active sovereign AI programmes. By 2026, global spending on sovereign AI systems is projected to surpass \$100 billion.

PDP Spectra: Sovereign AI in 2026 - Mistral, G42, HUMAIN, BharatGen, and the National-AI Map. CNAS: Sovereign AI Index, 2026. Hugging Face: prefeitura-rio/Rio-3.5-Open-397B (MIT Licence). UK Government: AI Opportunities Action Plan, January 2025 (Sovereign AI Unit, April 2026).

VI What Sovereign AI Means

Sovereign AI does not require building a foundation model from scratch. The spectrum of sovereign capability ranges from full pre-training to post-training adaptation to sovereign inference infrastructure. Ireland does not need to compete with OpenAI or Anthropic on pre-training. It needs to establish a sovereign position on this spectrum.

Tier 1: Foundation model pre-training. Building a large language model from scratch. This is what France has done with Mistral, what the UAE has done with Falcon, and what Saudi Arabia has done with ALLaM. Cost: hundreds of millions to billions of euros. Ireland need not start here.

Tier 2: Post-training and fine-tuning open models. Taking an open-source foundation model and specialising it on Irish data, Irish law, Irish language, and Irish institutional knowledge. This is what Rio de Janeiro did with Rio 3.5 Open, built on Alibaba's Qwen base. Cost: orders of magnitude less than pre-training. This is achievable with Ireland's existing research infrastructure and the incoming CASPIr supercomputer.

Tier 3: Sovereign inference infrastructure. Operating AI models on Irish-controlled hardware, ensuring that queries, data, and outputs never leave Irish jurisdiction. This is achievable immediately using existing data centre capacity and open-source models. It requires political will, not technological breakthrough.

The minimum viable sovereign AI capability for Ireland is Tier 2 plus Tier 3: fine-tuned open-source models running on Irish infrastructure, specialised for Irish public service needs, with guaranteed availability independent of any foreign government's export control decisions.

CNAS: Sovereign AI Index, 2026: Methodology and tier classification. EU AI Act (Regulation 2024/1689): Recitals 102-103, free and open-source AI exemption.

VII The Irish Language Dimension

AI that does not speak Irish is AI that accelerates language death. Every major AI system deployed in Ireland - ChatGPT, Claude, Gemini - has been trained predominantly on English-language data. These systems understand Irish poorly, generate it unreliably, and are incapable of serving Irish-language communities, institutions, or the Gaeltacht with the fluency they require.

The Government's own 20-Year Strategy for the Irish Language (Straitéis 20 Bliain don Ghaeilge, 2010-2030) identifies technology as one of nine areas of action for language revitalisation. The strategy's goal of increasing daily Irish speakers from 83,000 to 250,000 by 2030 has not been met - the number fell to 71,968 by 2022. If AI systems deployed across Irish public services, education, and commerce operate exclusively in English, they create an institutional pressure against Irish language use that no revitalisation strategy can overcome.

Saudi Arabia built ALLaM because Arabic matters to Saudi national identity. The UAE built Falcon Arabic because linguistic sovereignty is inseparable from digital sovereignty. India's BharatGen supports multiple Indian languages because a nation's AI must speak its people's languages. Ireland's official first language has no sovereign AI support. The contrast is damning.

A sovereign Irish AI programme must include an Irish-language capability as a foundational requirement, not an afterthought. Post-training open-source models on Irish-language corpora - Oireachtas records, Irish-language media, educational materials, literature - is technically achievable and would represent the most significant language technology intervention since the establishment of TG4.

Gov.ie: 20-Year Strategy for the Irish Language 2010-2030. CSO: Census 2022, Education and Irish Language (71,968 daily speakers).

VIII The Northern Ireland Dimension

The Good Friday Agreement established institutions for cross-border cooperation that have proven adaptable to new domains of governance. AI governance is a natural extension of existing cooperation frameworks, and cross-border AI collaboration would strengthen both jurisdictions.

Northern Ireland faces its own AI sovereignty challenges within the post-Brexit governance framework. The Windsor Framework creates a unique regulatory environment in which Northern Ireland maintains alignment with certain EU regulations while remaining part of the UK's digital economy. This dual positioning creates both complexity and opportunity for AI governance.

Cross-border AI cooperation could include shared sovereign inference infrastructure, joint Irish-language AI development, aligned AI governance frameworks, and mutual recognition of AI safety standards. The North/South Ministerial Council and its implementation bodies provide ready-made institutional frameworks for this cooperation.

Existing cross-border bodies - InterTradelreland, the Special EU Programmes Body, and the North/South Language Body (which oversees Foras na Gaeilge and the Ulster-Scots Agency) - demonstrate that practical all-island collaboration is achievable within existing frameworks.

North South Ministerial Council: Good Friday Agreement implementation bodies. InterTradelreland: Cross-border trade and innovation programmes.

IX Conclusions and Recommendations

The Government's own strategy documents commit Ireland to being "an international leader" in AI. The reality is total dependence on foreign systems that can be surveilled, degraded, or suspended at the discretion of a foreign government. The Fable incident proved this is not a theoretical risk. It happened.

The global sovereign AI movement demonstrates that alternatives exist at every scale, from city governments to nation states. Ireland's inaction is a policy choice, and it is the wrong one.

Ireland's AI Strategy: Commitments vs. Reality

COMMITMENT	SOURCE	REALITY
"International leader in using AI"	AI - Here for Good, DETE, 2021	Zero sovereign AI capability. Total dependence on US providers.
"75% of businesses using AI by 2030"	National AI Strategy, 2021	Using foreign commercial AI. No Irish models exist.
"AI Research Centre of Scale"	Digital and AI Strategy, 2026	No sovereign research centre established. CeADAR lacks mandate for sovereign AI.
"Access to advanced compute capacity"	Digital and AI Strategy, 2026	CASPIr (15 petaflops) not due until 2027. Classified as mid-range. Not configured for AI training.
"Digital leader in Europe"	Taoiseach, 2025	France: EUR 109B invested. UK: GBP 500M Sovereign AI Fund. Ireland: EUR 0.
Technology as key area for Irish language	20-Year Strategy for Irish, 2010	No Irish-language AI model. Daily speakers fell from 83,000 to 71,968.

Recommended Strategic Actions

The strategic choices Ireland makes today will determine whether it shapes the AI era or is shaped by it. The following actions are recommended to align the Government's delivery with its own published commitments:

- 1
- The Government should establish a National Sovereign AI Programme with dedicated funding, a clear mandate, and a timeline for delivering sovereign AI capability. This programme should be distinct from existing research and enterprise support structures.
- 2
- A government-wide audit of all AI dependencies must be conducted within 90 days, documenting every AI system used by government departments and state agencies, the jurisdiction of each provider, the data retention and access policies that apply, and the contingency plan if access is suspended.
- 3
- An Emergency Sovereign AI Advisory Group should be convened within 30 days, drawing on ICHEC, CeADAR, the AI Office (once established), the Department of Enterprise, the Department of Public Expenditure, academic AI researchers, and civil society organisations.

- 4 Ireland should begin post-training open-source foundation models on Irish data - legal corpora, Oireachtas records, public service documentation, and Irish-language materials - using CeADAR, ICHEC, and the incoming CASPIr infrastructure. The Rio de Janeiro model demonstrates this is achievable.
- 5 An Irish Language AI Mandate should require that all AI systems deployed in Irish public services include Irish-language capability. This mandate should be accompanied by investment in Irish-language training data, evaluation benchmarks, and model development.
- 6 A cross-border AI cooperation framework should be established through existing Good Friday Agreement institutions, enabling shared infrastructure, joint Irish-language AI development, and aligned governance approaches.
- 7 Ireland should negotiate sovereign AI terms with all current commercial AI providers, including guarantees of continued access, data sovereignty protections enforceable under Irish and EU law, and the right to audit model behaviour, including the right to detect and prevent silent response degradation.
- 8 The Government should publish a Sovereign AI Readiness Assessment as part of the next Digital and AI Strategy update, benchmarking Ireland's position against the CNAS Sovereign AI Index and the sovereign capabilities of comparable nations.

Ní neart go cur le chéile. There is no strength without unity.

James Spalding | Creative Technologist and AI Policy Advisor, 8GI Foundation

openmygovernment.com/ireland/projects/sovereign-ai