

Why sovereign AI

Better public institutions, on your own terms

AI can make government more responsive, better informed and easier to take part in. Sovereign AI means a country captures that upside on its own terms. It builds the skills, keeps control of its data and decisions, and runs the models its public institutions rely on. A country that rents all of it gets the tool and gives away the capability.



THREE REASONS

What a country gains, with the evidence behind each

01

Better government

More responsive services and sharper policy intelligence

AI can draft, translate, summarise and answer questions around the clock, in the country's own languages. It can read every consultation response instead of a sample. Running on infrastructure the state controls, it can also work with the records that matter most, which a rented service should never see.

26 min

a day saved on average, self-reported, by 20,000 UK civil servants given an AI assistant for three months. Worries about sensitive data held some of them back.

UK Government Digital Service, June 2025 [↗](#)

In-house

The Canton of Ticino translates sensitive government documents on Switzerland's open model Apertus, without relying on commercial providers.

Swiss National Supercomputing Centre, 24 July 2026 [↗](#)

02

National capability

Local talent, local languages, local industry

A country that runs its own AI trains people by doing it, and gives them work worth staying for. It can adapt models to its own languages, law and data. Its suppliers and startups grow with it, and the skills stay when a contract ends.

105B

parameters in Sarvam's open model, trained from scratch entirely in India on compute the government's IndiaAI Mission provides.

Sarvam AI, 6 March 2026 [↗](#)

13,500

AI fellowships India is funding: 500 PhDs, 5,000 postgraduates and 8,000 undergraduates. More than 500 have been awarded so far.

Government of India, answer to the Lok Sabha, 1 April 2026 [↗](#)

03

Strategic autonomy

Control over sensitive information and the choices it shapes

Data in a foreign AI service sits under foreign law. US providers can be compelled to hand it over wherever they store it. Models also carry their makers' values, and people who write with a slanted assistant drift towards its view without noticing. Sovereign AI keeps the data, the audit trail and the choice of model at home.

"No"

said Microsoft France's director of public and legal affairs under oath, asked to guarantee French data would never reach US authorities without French consent. He added it had never happened yet.

French Senate inquiry, 10 June 2025 [↗](#)

8 of 8

leading AI models advised more escalation for the US, UK and France than for China or Russia, in the same crisis scenarios.

CSIS Futures Lab, Critical Foreign Policy Decisions benchmark, 2025 [↗](#)

INSTITUTIONAL INNOVATION

What sovereign AI lets public institutions do differently

Act more *intelligently*

Read every submission, model the trade-offs before a decision, and check afterwards what a policy did. The UK government's own Consult tool sorted consultation responses about as accurately as civil servants on most measures, and over 120 times faster.

UK algorithmic transparency record, DSIT: Consult, May 2025 [↗](#)

Act more *democratically*

Broad listening lets a ministry hear thousands of citizens in their own words, then show them what it heard, with every finding linked to a quote. Taiwan's vTaiwan consultations, which shaped its first rules for Uber, have drawn more than 200,000 people.

Computational Democracy Project, vTaiwan case study [↗](#)

Act more *independently*

Own models, own logs, own tests. A ministry can change supplier without losing its tools or its memory, and can show anyone how an AI-assisted decision was made.

IN PRACTICE

Broad Listening & Simocracy

PL R&D's near-term bet. Broad Listening turns thousands of voices on a contested question into a shared picture anyone can check. Simocracy takes the next step: at Funding the Commons in San Francisco, structured deliberation drove a real budget allocation, with people making the final decision.

THE HONEST LIMIT

Sovereign still needs testing

The best open models still trail the strongest closed ones on the hardest tasks, and a model built at home can carry its own bias. A country still chooses, tests and answers for the models it relies on, whoever built them.

WHERE TO START

One problem, one pilot

Pick a public problem with a minister who owns it, run an open model on infrastructure the country controls, and measure it against the current process.

PL R&D partners with governments on exactly this: see Sovereign AI, government pilot partnership.

Also drawn on: US CLOUD Act, 18 U.S.C. § 2713 · Williams-Ceci et al., Science Advances, March 2026 · Pan and Xu, PNAS Nexus, February 2026 · Artificial Analysis Intelligence Index, October 2026