


Sovereign AI

A summary for funders

PL R&D's Sovereign AI program helps governments run AI on their own terms. In each country we take one public problem to a working pilot on an open model the government controls, and train the local team that will run it. We aim for three to five countries over 24 months. Funding pays for the pilots, their independent evaluation and the lessons the next country can reuse, so a government sees AI work for its priorities before it commits its own budget.



WHY FUND IT

Three reasons, with the evidence behind each

01

Governments are asking

Capacity is the help they want first

Ministries want to choose, run and check AI themselves. Most offers they get are a product or a plan. We build the skills on a real public task.

1st

In more than 1,500 written submissions to the UN's first Global Dialogue on AI Governance, governments were the only group to rank capacity-building first. Most others put safety first.

UN press release, Geneva, 6 July 2026 [↗](#)

02

The tools are ready

Open models already do real public work

A government can now host and adapt an open model, such as Switzerland's Apertus, for work like translation and drafting. What most countries lack is the team to run one safely.

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 [↗](#)

03

Control is at stake

Rented AI answers to someone else

Whoever runs a government's AI shapes what its officials see. Renting all of it hands over the data, the audit trail and the choice of model. Capacity built now keeps them 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 [↗](#)

WHAT FUNDING PAYS FOR

In each country, around the government's own team

01

The pilot

An open model hosted where the government chooses, adapted to its languages and data, and engineers to build the service with its team.

02

Independent evaluation

Academic partners test models on the country's priorities and languages, then measure the pilot against the current process. We publish the results, good or bad.

03

A local team

Training and side-by-side work for the officials and engineers who will run and improve the system after the pilot ends.

04

Shared lessons

A program manager, gatherings where pilot teams compare notes, and an open playbook the next country can start from.

HOW THE PROGRAM RUNS

In each country, from a named problem to a decision

1

Diagnose

Months 0 to 3. One public service or decision to improve, its owner, and a measure of success.

2

Design

Months 3 to 6. Data rules, oversight and red lines first, then the model and where it runs.

3

Pilot

Months 6 to 12. A live service used by over 1,000 people for at least 30 days, tested against the current process.

4

Decide and hand over

Scale, change or stop, on the evidence. The code, the know-how and the team stay with the government.

What every funder gets: a measure of success agreed before work starts, results reported against it, good or bad, and open code and lessons that any country can reuse.

HOW FUNDERS CAN TAKE PART

Anchor the program. Fund the shared core: the first pilots, their evaluation and the program team.

Back one country. Fund a pilot in a country or region where you already work.

Give compute. Credits or hosting for open models, so each pilot runs where its government chooses.

Scale what works. Help a pilot that works become a national service, with the government's budget or a development bank.

WHY PL R&D

Open by default. Open models, open standards and open code that a government can inspect, keep and run without us.

From research to deployment. A decade of turning research into open infrastructure used around the world: IPFS, Filecoin and libp2p.

Governance experience. Broad Listening and Simocracy, AI-assisted deliberation that has driven real budget allocations, and AI for Public Goods, a research field PL R&D co-leads, co-funded with Gitcoin.

A network behind each pilot. Nearly 400 teams across Protocol Labs, and more than 22,000 researchers, builders and residents from over 80 nations in our network.

THE HONEST LIMIT

Not every pilot will work

Some pilots will end at the decision step, and we will publish why. One pilot does not make a country self-sufficient in AI. That takes people who stay, a budget line and years of use. The program pays for the first working system and the evidence a government needs to fund the rest.

START HERE

Back the first pilots

Tell us what you want your funding to change, and we will show you where it fits in the program. research@protocol.ai

David Dao and James Tunningley, co-leads, Economies & Governance.