WHITE PAPER · SOVEREIGN AI · JULY 2026

BRIEF

Sovereign AI: *the state of play*

Companies are increasingly taking control over the AI they run: keeping data within their borders, models no one can switch off, weights they can audit for themselves. We mapped over a hundred sovereign-AI companies and programs worldwide: what they are for, who funds them and impact to date . What follows is the overview, alongside the detailed landscape explorer.

$500–600B

of AI spend influenced by sovereignty requirements by 2030: 30–40% of the total (McKinsey, Dec 2025)

50+

countries with active sovereign-AI programs (NVIDIA, Feb 2026): up from 20+ a year earlier

$30B+

NVIDIA's sovereign-AI revenue in FY2026: more than tripling year-on-year, ~14% of total revenue

~110

organisations & programs profiled

DEFINITION

The landscape

Independence labs

Private companies focusing on strategic autonomy from the US and Chinese AI stacks. Frontier or near-frontier capability under domestic ownership.

Mistral (FR) · Cohere–Aleph Alpha (CA/DE) · Domyn (IT) · Cosine (UK)

National programs

Publicly funded models, anchored in language and culture. Models that speak the nation's languages and reflect its law and values.

OpenEuroLLM (EU) · Korea's MSIT program · ALLaM (SA) · SEA-LION (SG) · PLLuM (PL)

Sovereign infrastructure

Compute and cloud under local jurisdiction and vendors selling "sovereignty" as a product line.

EuroHPC AI Factories · Isambard-AI (UK) · STACKIT (DE) · NVIDIA, Oracle, Microsoft, AWS

What is Sovereign AI?

Sovereign AI is a nation's ability to build, run and control AI on its own infrastructure, data and talent instead of renting from foreign providers. AI now matters enough strategically that countries no longer want their critical systems sitting inside a handful of large tech companies, usually American or Chinese, that they cannot govern.

> *"Nearly every sovereign AI program on earth runs on chips from one company in California."*
>
> NVIDIA SUPPLIES 52% OF TRACKED SOVEREIGN INFRASTRUCTURE PROJECTS (CNAS, APR 2026)

There are several layers to consider. Infrastructure sovereignty is owning or controlling the physical compute: the data centres and GPUs sitting inside a country's borders. Data sovereignty: keeping a nation's data under its own laws and on domestic servers rather than letting it flow abroad. Models/capability: whether a country can train or fine-tune its own foundation models. Under all three sits the talent and supply chain: the researchers, engineers and access to critical components.

Beyond capability, nations are motivated by four factors. **National security**: governments don't want critical systems running on foreign infrastructure that can be cut off or surveilled. **Economic competitiveness**: AI is treated as foundational to future productivity, so importing it means exporting the value. **Cultural and linguistic concerns**: for countries whose languages are not the focus of the big commercial models. **Regulation**: local control makes domestic privacy and governance rules easier to enforce.

In reality, sovereignty is a spectrum of managed dependencies. Nations can decide themselves how much control each layer needs.

THE ARGUMENT

The case *for* — and the case *against*

For

Resilience. A state that cannot run critical AI workloads without foreign permission is strategically exposed.

Bargaining power. Mistral's Arthur Mensch frames sovereignty as leverage: without production capacity, Europe has "nothing to put on the table" — and risks becoming a "vassal state".

Value capture. If AI spending heads toward hundreds of billions a year, importing it all adds straight to the trade deficit.

Language and culture. Models trained on Anglo-American inputs impose foreign defaults.

Rule-of-law alignment. Putting critical AI beyond foreign legal reach and being in control to prevent surveillance, subpoena. Defence and health are the strongest ground for sovereignty. As seen in the Schrems II legislation France's air-gapped ASGARD system.

Against

Subscale economics. US hyperscalers guided ~$725B of AI capex for 2026; the EU's public gigafactory facility is €20B, the UK's whole programme ~£2B. Critics call the gaps unbridgeable.

Execution gaps. IndiaAI released under 4% of its budget in two years; Krutrim, India's first GenAI unicorn, quit foundation models; OpenAI paused Stargate UK; AI21 collapsed into a niche.

The hardware problem. Nearly every program runs on NVIDIA chips, TSMC fabs and ASML lithography. Sovereignty purchased from the same chokepoints.

Provenance theatre. Many "sovereign" models are fine-tunes of foreign bases, such as China's Qwen. Frontier capability also depreciates in months: a national model funded through a budget cycle is superseded before procurement completes.

Opportunity cost. New Zealand explicitly declined to build models; diffusion scholars argue adoption, not duplication of pre-training, determines national gains.

The European picture

Mistral AI is the strongest non-US/China frontier lab, valued at **€11.7B** after ASML, Europe's chip champion, led its €1.7B Series C (Sep 2025) and took 11%. Since then: ARR moved past $400M, an $830M debt raise for its Essonne datacenter (Mar 2026), a reported ~€3B round in talks at ~€20B, and a defence framework that makes Mistral the backbone of French military AI on the air-gapped ASGARD supercomputer. Around it sits a full stack: Kyutai's open-science voice lab, Yann LeCun's new AMI (a record **$1.03B seed**, Mar 2026), LightOn's air-gapped deployment platform, and OVHcloud selling shelter from the US CLOUD Act.

The boldest structural move is the **Cohere–Aleph Alpha merger** (announced 24 Apr 2026; ~$20B expected combined value at close, H2 2026): a transatlantic champion financed by Germany's Schwarz Group ($600M Series E lead), anchored on its STACKIT sovereign cloud, the flagship of the Canada–Germany *Sovereign Technology Alliance*. Schwarz is meanwhile building a **€11B, 200MW, ~100,000-GPU** campus in Brandenburg.

Brussels has moved from regulation to industrial policy with mixed delivery. Delivered: 19 EuroHPC AI Factories, 13 antennas; JUPITER, Europe's first exascale machine, now #5 worldwide; an adopted AI Act simplification that delays high-risk obligations to 2027–28. Pending: the €20B AI Gigafactories are **still unselected**, with 77 bids across 16 member states waiting. The Commission's **Frontier AI Grand Challenge** already picked its winner (Jun 2026): the Domyn-led EUROPA consortium will train a fully open **400B+ model in all 24 EU languages** with up to 2.5% of EuroHPC's entire compute.

National-language models now exist across the bloc: Spain's ALIA, Poland's PLLuM (now inside the mObywatel citizen app), the Netherlands' lawfully-sourced GPT-NL, Bulgaria's BgGPT — with Switzerland's radically open **Apertus** as the reference "public AI," already in clinical testing at CHUV. OpenEuroLLM, the flagship open-model consortium, has shipped only reference models against a 31 July 2026 deadline: compute, not ambition, is its bottleneck.

The British picture

Britain focused not on a state-built model, but **"control and leverage".** A state compute plus state equity in private champions. **Isambard-AI** (Bristol, 5,448 GH200s, 21 AI-exaflops, world #13); the **£500M Sovereign AI Unit**, an explicit state VC, wrote its first cheques in June 2026 with up to 1M GPU-hours per company; the **AI Security Institute** is a world-leading sovereign *evaluation* capability; and a **£1.1B AI Hardware Plan** (Jun 2026) puts £400M into direct government purchases of UK-made AI chips becoming the first real counter to a decade of losing champions (Arm and Graphcore to SoftBank, Faculty to Accenture for >£600M in Jan 2026).

The flagship bet is **Cosine's Lumen Sovereign**: a ~1.35-trillion-parameter coding model trained from scratch on Isambard-AI, co-designed by BT, Lloyds, NatWest, LSEG, BAE Systems, Leonardo and Babcock, deployable air-gapped, targeted for end-2026 — Britain's one true frontier-training exception. On the other hand, **OpenAI paused Stargate UK in April 2026**, citing energy prices and the still-unresolved copyright regime, weeks after ministers celebrated £31B of US "Tech Prosperity" pledges. Nscale's much-touted Loughton "sovereign supercomputer" site was still a scaffolding yard in June. At ~£2B, UK public AI spend is roughly 1% of what a single US hyperscaler lays out in capex each year.

SYNTHESIS

Five findings

01 From-scratch vs fine-tune is the fault line.

Most “sovereign” LLMs sit on foreign open bases — increasingly China’s Qwen (~1 billion downloads; more than half of global open-model downloads). Korea’s national program disqualified Naver *and* NC AI in round one for exactly this; only independently trained models advance. Meta’s Llama still underpins 36% of open-based national models (CNAS).

02 Sovereignty still means data, language and jurisdiction — not silicon.

But the exceptions got real in 2026: China’s GLM-5 was trained on ~100,000 Huawei Ascend chips (self-reported) and served entirely on domestic silicon; Korea fields Rebellions NPUs; Japan’s PFN runs its own MN-Core; the UK committed £400M to buy British chips.

03 The label is unstable — and increasingly audited.

Australia’s Maincode dropped “sovereign” as divisive; New Zealand declined to build models at all; US hyperscalers sell “sovereign cloud” to everyone while Microsoft told the French Senate, under oath, that it cannot guarantee French data stays out of US hands. Sovereignty by architecture is not sovereignty by ownership.

04 Language is the Global South’s driver.

Spanish and Portuguese are ~2–3% of mainstream training data; hence LatamGPT (launched Feb 2026), Brazil’s Sabiá, Nigeria’s N-ATLAS, Türkiye’s T3AI, and the Arabic families from Falcon to Fanar. This is the axis where a few hundred thousand dollars demonstrably buys relevance.

05 The economics are hard — now with receipts.

IndiaAI had released under 4% of its five-year budget by Feb 2026; Krutrim quit models for GPU rental; AI21 shed 60% of staff after an acquisition that never was; OpenAI paused Stargate UK over energy and copyright. The winners so far sell picks and shovels: NVIDIA’s sovereign revenue tripled to $30B+.

LANDSCAPE

The landscape, *region by region*

A summary of organisations and programs including what it is, who funds it, what it has achieved, and the sovereignty overview. Figures as of 6 July 2026.

Entries in our landscape research, by region — July 2026. Europe includes the EU, EU-level programs, UK, Switzerland and Norway.

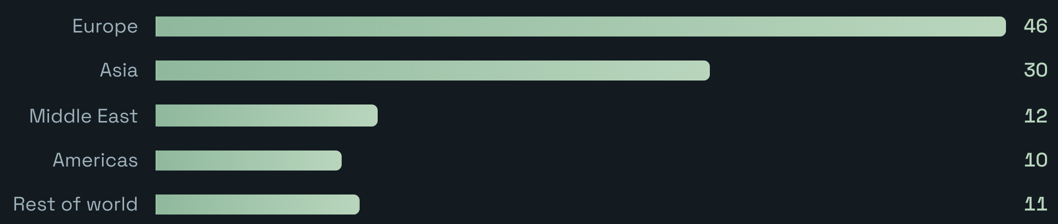

The capital extremes are striking. Saudi Arabia's **HUMAIN** is the most capital-intensive program anywhere with up to 600,000 NVIDIA GPUs over three years, a $10B AMD deal, $3B into xAI, and a 6.6GW datacenter roadmap. At the other end, **LatamGPT** launched in February 2026 across 15+ countries on roughly $550K. Korea bought 260,000+ GPUs nationally and tripled its AI budget to ₩10.1T; Japan runs GENIAC as a 70-project accelerator.

UNITED KINGDOM

Sovereign AI Unit / Fund ↗

WHAT IT IS & CAPABILITIES
DSIT state VC: direct equity (to ~£20M/company) + up to 1M GPU-hours each on Isambard-AI + £80M procurement pilots + £9M strategic-asset grants.

FUNDING
Up to £500M (Spending Review 2025); investing from Apr 2026

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
First cohort Jun 2026 (Callosum equity; Cosine, Prima Mente, Doubleword compute). Goal: UK-controlled capability in security-critical niches. Critics: subscale vs France's €109B.

Cosine / Lumen Sovereign ↗

WHAT IT IS & CAPABILITIES
AI-coding lab training Britain's first sovereign frontier model: ~1.35T-param MoE, from scratch on Isambard-AI, 30+ regulated workflows, air-gappable.

FUNDING
Sovereign AI Programme compute + private backing

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
Coalition: BT, Lloyds, NatWest, LSEG, PwC, BAE, Leonardo, Babcock. Target: deployment-ready end-2026. The UK's one true frontier-training exception.

Isambard-AI + AIRR ↗

WHAT IT IS & CAPABILITIES
UK flagship public AI supercomputer (Bristol): 5,448 GH200, 21 AI-exaflops; AIRR also spans Dawn (Cambridge; 6x expansion with AMD MI355X, 2026).

FUNDING
£225M machine; £300M AIRR phase 1; ~£1B compute pledge to 2030 (420 EF target)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
Live Jul 2025; TOP500 #11 (2025) → #13 (Jun 2026); >100 projects incl. BritLLM (UCL) and NHS-relevant screening models. Delivered core of UK sovereignty.

Nscale ↗

WHAT IT IS & CAPABILITIES
London-HQ AI "neocloud": ~200k GB300s contracted for Microsoft (up to ~$14B; Texas, Loughton, Sines, Narvik); Stargate Norway JV with Aker + OpenAI (100k GPUs, end-2026).

FUNDING
$1.1B Series B (Sep 2025) → $2B Series C @ $14.6B (Mar 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
Largest UK/EU Series B and C on record. Caveats verified: Loughton site still a scaffolding yard Jun 2026; timelines slipped to 2027; anchor customers are US.

Edinburgh supercomputer + AI Hardware Plan ↗

WHAT IT IS & CAPABILITIES
National supercomputer (online 2027; exascale spec dropped) + £400M direct government purchases of UK-made AI chips.

FUNDING
£750M (machine ~£300–330M) + £1.1B Hardware Plan (Jun 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
"Start here, scale here, stay here" — first explicit counter to the champion-acquisition drain (Arm, Graphcore→SoftBank; Faculty→Accenture >£600M, Jan 2026).

AI Security Institute ↗

WHAT IT IS & CAPABILITIES
State frontier-model evaluation (renamed from "Safety," Feb 2025); pre-deployment access to OpenAI/Anthropic/GDM models.

FUNDING
£240M; becoming arm's-length body

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ
World-leading sovereign evaluation capacity — sovereignty in judging models, not building them.

Growth Zones + Tech Prosperity Deal

WHAT IT IS & CAPABILITIES

Five AI Growth Zones (fast-track planning/grid); £31B US investment pledges (Microsoft $30B, NVIDIA 120k GPUs, Google $6.8B).

FUNDING

Private (mostly US) capital

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Compute located in the UK, not UK-controlled. OpenAI paused Stargate UK (Apr 2026) over energy costs + unresolved copyright — the strand's fragility, proven.

National Data Library ↗

WHAT IT IS & CAPABILITIES

Public data as strategic asset; + Health Data Research Service (£600M w/ Wellcome, Genomics England).

FUNDING

>£100M (within £1.9B data allocation)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Data-side complement to compute; early-stage, criticised as slow.

EU PROGRAMS

EuroHPC AI Factories ↗

WHAT IT IS & CAPABILITIES

19 AI factories + 13 antennas across three waves; shared sovereign compute for startups and science.

FUNDING

EuroHPC JU (EU + states); e.g. BRAIN++ €90M

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Operational and expanding (new call Apr 2026). JUPITER (Jülich): Europe's first exascale system — #5 on the Jun 2026 TOP500.

OpenEuroLLM ↗

WHAT IT IS & CAPABILITIES

~20 orgs (Hajič/Charles Univ. + Sarlin/AMD) building open models for all 24 EU languages.

FUNDING

€37.4M total; €20.6M EU grant; first STEP Seal

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Reference models only so far; ~8B flagship due 31 Jul 2026; compute is the admitted bottleneck. EuroLLM (Lisbon-led) shipped its open 22B in Dec 2025.

InvestAI / AI Gigafactories ↗

WHAT IT IS & CAPABILITIES

Mobilise ~€200B; €20B facility for up to five gigafactories (~100k chips each), EIB co-financed.

FUNDING

EU + EIB + private

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

No consortia selected yet: 77 bids, 16 states, >€230B proposed; formal call slipped to ~summer 2026. Known bids incl. France's AION (Iliad/Scaleway-led).

CADA + AI Act omnibus ↗

WHAT IT IS & CAPABILITIES

Cloud & AI Development Act proposal (Jun 2026): triple EU DC capacity, 4-level Cloud Sovereignty Framework for procurement. AI Act simplification adopted Jun 2026.

FUNDING

—

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

High-risk AI obligations deferred to Dec 2027 / Aug 2028. Regulation now openly subordinated to the sovereignty-through-competitiveness agenda.

Frontier AI Grand Challenge ↗

WHAT IT IS & CAPABILITIES

EC prize to build a fully open frontier model in all 24 EU languages.

FUNDING

Prize incl. up to 2.5% of EuroHPC compute for a year

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Won 19 Jun 2026 by Domyn-led EUROPA consortium (with Fraunhofer): 400B+ open model within ~a year — the EU's designated open-frontier bet.

FRANCE

Mistral AI ↗

WHAT IT IS & CAPABILITIES

Europe's frontier flagship. Open-weight Mistral Large 3 (Dec 2025: 675B/41B-active MoE, Apache 2.0), Magistral reasoning, Medium 3.5, Devstral 2; "Vibe" assistant (ex-Le Chat).

FUNDING

~$3.05B equity; €1.7B Series C @ €11.7B, ASML 11% (Sep 2025); $830M debt (Mar 2026); ~€3B @ ~€20B (reported, unclosed); ARR >$400M

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Armed-forces framework (AMIAD) on air-gapped ASGARD; Airbus/BMW/ASML/EDF deployments; Mistral Compute 40MW → Campus AI 1.4GW→3GW with MGX/Bpifrance/NVIDIA. The reference European alternative to US labs.

LightOn ↗

WHAT IT IS & CAPABILITIES

Paradigm: on-prem/air-gapped enterprise genAI platform for defence, gov, education.

FUNDING

First EU pure-play genAI IPO (Nov 2024)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Reality check: FY2025 revenue €1.7M, loss €7.4M, ~$36M market cap — pure sovereignty pitch, tiny commercially.

AMI ↗

WHAT IT IS & CAPABILITIES

Yann LeCun's post-Meta lab; world-models ("advanced machine intelligence"), not LLM-scale-chasing.

FUNDING

$1.03B seed @ $3.5B pre (Mar 2026) — largest EU AI seed ever; Toyota, NVIDIA, Samsung

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Instant frontier-adjacent research pole in Paris; sovereignty value = talent gravity, not products yet.

Linagora / LUCIE ↗

WHAT IT IS & CAPABILITIES

Maximal-transparency 7B: open weights + open data; 3T tokens on national Jean Zay compute.

FUNDING

Community (OpenLLM-France) + France 2030 context

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Pulled after ~3 days of ridicule (Jan 2025), relaunched instruct-tuned. Cautionary tale + open-commons pole.

Kyutai ↗

WHAT IT IS & CAPABILITIES

Open-science voice lab: Moshi full-duplex (~200ms), Hibiki, Pocket TTS (Jan 2026, CPU real-time).

FUNDING

~€300M from Niel, Saadé/CMA CGM, Schmidt

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Everything open source; member of AION gigafactory consortium. Europe's open research commons.

Pleias ↗

WHAT IT IS & CAPABILITIES

Copyright-clean small models; Common Corpus (2.27T fully open tokens; ICLR 2026 oral).

FUNDING

Small/undisclosed

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Sovereignty via legal data provenance — a distinct, exportable axis.

H Company / Harmattan ↗

WHAT IT IS & CAPABILITIES

H: computer-use agents on in-house Holo VLMs (Holo 3, 2026). Harmattan: defence drones-AI.

FUNDING

H: $220M seed (2024). Harmattan: $200M @ $1.4B, Dassault-led (Jan 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

French-funded agentic + defence-AI depth; sovereignty secondary for H, core for Harmattan.

French state layer

WHAT IT IS & CAPABILITIES

€109B AI Action Summit pledges (Feb 2025); Jean Zay 125.9PF; OVHcloud AI Endpoints (CLOUD-Act shelter pitch); S3NS (Thales+Google) SecNumCloud 3.2 (Dec 2025); Bleu live.

FUNDING

France 2030 / Bpifrance instruments

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The most complete national stack in Europe: champion lab + state compute + qualified trusted clouds + defence integration.

GERMANY

Cohere–Aleph Alpha ↗

WHAT IT IS & CAPABILITIES

Merger of Canada's enterprise-AI champion with Germany's PhariaAI sovereign "AI OS" (~200 staff), on STACKIT.

FUNDING

Schwarz Group $600M (€500M) Series E lead; ~$20B expected combined value at close

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Announced 24 Apr 2026; pending close H2 2026 (antitrust + AA shareholders). Flagship of the Feb 2026 Canada–Germany Sovereign Technology Alliance. Cohere ARR ~$240M.

Schwarz Digits / STACKIT ↗

WHAT IT IS & CAPABILITIES

Pure-German sovereign cloud (Lidl parent), expanding into a "German hyperscaler."

FUNDING

€11B Lübbenau campus — Schwarz's largest-ever single investment

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Up to 200MW / ~100,000 GPUs from ~2027; hosts PhariaAI; anchor cloud of the merged entity. Around it: SAP (>€20B sovereign-cloud program), Delos (Azure-based public-sector cloud; "OpenAI for Germany," 4,000 GPUs), T-Systems, IONOS.

OpenGPT-X / Teuken-7B ↗

WHAT IT IS & CAPABILITIES

Fraunhofer-led from-scratch 7B in all 24 EU languages, open weights.

FUNDING

BMWK ~€14M (ended Mar 2025)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Multilingual showcase, not frontier; sold by Deutsche Telekom as "sovereignty made in Germany."

Black Forest Labs ↗

WHAT IT IS & CAPABILITIES

FLUX — world-leading open image generation.

FUNDING

$300M Series B @ $3.25B (Dec 2025; Salesforce Ventures + AMP)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Germany's most valuable AI company; capital and customers largely US — ecosystem depth, weak sovereignty claim.

German state layer

WHAT IT IS & CAPABILITIES

National Data Center Strategy (Mar 2026): DCs as critical infrastructure, 4x AI capacity by 2030, ≥1 German gigafactory; "Made for Germany" €631B private pledge.

FUNDING

Federal + private

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Berlin's pivot from regulation-first to build-first, mirroring Brussels.

REST OF EU

ALIA / BSC (ES) ↗

WHAT IT IS & CAPABILITIES

Public open family (Salamandra 2B/7B/40B + ALIA-40B): 9.37T tokens, 35 languages, from scratch on MareNostrum5 (4,480 H100).

FUNDING

€90M compute upgrades + €150M integration (SEDIA)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Actively shipping (instruct releases Dec 2025/Jan 2026). "Europe's first public, open, multilingual AI infrastructure." Also ES: Latxa (Basque, beats GPT-3.5 on Basque); Multiverse Computing (€189M Series B; ~$580M Series C @ ~$1.75B reported, unclosed).

Domyn (IT) ↗

WHAT IT IS & CAPABILITIES

Colosseum-355B for regulated industries + Colosseum datacenter (S. Italy, ~6,000 Blackwells; claimed Europe's largest AI supercomputer).

FUNDING

Undisclosed; deep NVIDIA partnership

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Won the EU Frontier AI Grand Challenge (Jun 2026) — now the EU's designated open-frontier builder. Also IT: Minerva (first from-scratch Italian family, open); Almawave Velvet (expanded 2026).

GPT-NL (NL) ↗

WHAT IT IS & CAPABILITIES

Dutch LLM built exclusively on lawfully-sourced data (TNO+SURF+NFI).

FUNDING

Economy Ministry

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Beta with external orgs since Feb 2026; "Gem" municipal assistant serving ~2.8M residents; 2026 Dutch Privacy Award. The lawful-data sovereignty archetype.

PLLuM + Bielik (PL) ↗

WHAT IT IS & CAPABILITIES

State Polish family (8B–70B; Wrocław Tech-led consortium) + grassroots Bielik (SpeakLeash, 11B v3).

FUNDING

Ministry of Digital Affairs; volunteer + Cyfronet compute

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

PLLuM deployed in the mObywatel citizen app; Poland won two AI Factories. State + community two-track model.

INSAIT / BgGPT (BG) ↗

WHAT IT IS & CAPABILITIES

Sofia institute (ETH/EPFL partnership); BgGPT open Bulgarian models; BRAIN++ AI Factory (€90M) from 2026.

FUNDING

~€85M state/10yr + Google/AWS/DeepMind donations

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Flagship CEE success — an EU small state buying real research capacity.

Nordics

WHAT IT IS & CAPABILITIES

Silo AI (Poro/Viking Nordic families) — acquired by AMD for ~$665M (2024); AI Sweden's GPT-SW3 (dated first-mover); Denmark's Gefion (1,528 H100; Novo Nordisk Foundation 85% + EIFO); NORA.LLM (NO, academic).

FUNDING

Mixed state/foundation/corporate

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The linguistic-sovereignty blueprint — with the caveat that its best lab is now US-owned. Evroc (SE): pure-play EU hyperscaler challenger, platform live 3 Jul 2026.

EUROPE, NON-EU

Apertus / Swiss AI (CH) ↗

WHAT IT IS & CAPABILITIES

Fully open 8B/70B (weights + data + recipes), 15T tokens, 1,000+ languages incl. Swiss German, Romansh; trained on Alps (4,000+ GH200).

FUNDING

ETH Zurich + EPFL + CSCS

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The reference "public AI": v1.1 (Jun 2026), Meditron medical variant in clinical testing at CHUV, Ticino administration adoption. Radical openness on national compute.

Infomaniak EURIA (CH) ↗

WHAT IT IS & CAPABILITIES

Free multimodal assistant processed entirely in Swiss datacenters; waste heat → Geneva district heating.

FUNDING

Infomaniak (private, self-funded)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Full launch Dec 2025. No data leaves Switzerland; frequently mislabelled French — it is Swiss.

MIDDLE EAST

HUMAIN (SA) ↗

WHAT IT IS & CAPABILITIES

PIF's full-stack champion: DCs, cloud, models (ALLaM 34B / HUMAIN Chat), apps.

FUNDING

NVIDIA up to 600k GPUs/3yrs (KSA+US); AMD $10B/500MW; AWS $5B+; $3B into xAI (Feb 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Most capital-intensive program anywhere; roadmap 1.9GW (2030) → 6.6GW (~2034); first Riyadh/Dammam DCs due Q2 2026; $10B VC fund planned. MBS-chaired.

UAE complex

WHAT IT IS & CAPABILITIES

TII's open Falcon line (H1, Falcon Arabic — tops Open Arabic leaderboard; Perception, Apr 2026); G42 (Microsoft $1.5B; China divestment); Stargate UAE (1GW, ~$20B; first 200MW Q3 2026) in the 5GW US–UAE campus; MBZUAI's K2-V2 (from-scratch "360-open" reasoning, Jan 2026); Jais 2 Arabic models.

FUNDING

Abu Dhabi state + US partners; Microsoft $15.2B UAE commitment

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The China→US realignment vehicle. Note: K2 V1 (not V2) was the Qwen-based, ETH-criticised release; Jais trains on US (Cerebras) compute. Sovereign capacity, US-operated — the definitional caveat.

SDAIA / ALLaM (SA) ↗

WHAT IT IS & CAPABILITIES

National Arabic-English models (7B from scratch; 70B Llama-2-initialised); 160 govt agencies contributed data.

FUNDING

Saudi state

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Distributed via Azure and watsonx — a jurisdictional irony for a sovereign model. Vision 2030 instrument.

Qatar

WHAT IT IS & CAPABILITIES

Fanar (QCRI): Fanar 2.0 27B multimodal trained domestically on 256 H100s (Dec 2025); Qai (QIA) + Brookfield $20B JV; Ooredoo sovereign AI cloud (replicated in Kuwait).

FUNDING

State (QIA ~$524B AUM)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Fanar 3.0 due Dec 2026; no hyperscaler region in-country yet. Islamic-values alignment as explicit design goal.

Israel

WHAT IT IS & CAPABILITIES

National AI Program approved 16 Jun 2026 (5,000 GPUs/yr 2027–32; 1,000-B200 Nebius-built supercomputer, phase 1 live). DictaLM 3.0: best-in-class open Hebrew (on Mistral/Nemotron/Qwen bases).

FUNDING

Thin: directorate NIS 120M; "$30B" is a press estimate, not budget

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Approved but thinly funded; defence-adjacent. AI21: NVIDIA acquisition denied; >60% layoffs, Maestro-only pivot (May 2026). NVIDIA building a ~$1.5B Israel campus.

INDIA & SE ASIA

IndiaAI Mission ↗

WHAT IT IS & CAPABILITIES

Umbrella: subsidised compute (40,535 GPUs committed, 56% allocated; ₹65/GPU-hr), 12 funded model teams, datasets across 22 languages.

FUNDING

₹10,372 cr (~$1.25B)/5yrs — only ~₹401 cr released in 2 years

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The disbursement gap is the story. India AI Impact Summit (Feb 2026): 89-country declaration, $200B+ pledges claimed.

Sarvam AI ↗

WHAT IT IS & CAPABILITIES

Govt-anointed flagship: Sarvam 30B & 105B MoE (Feb 2026) — from scratch, Apache 2.0, 22 languages.

FUNDING

₹247 cr + 4,086 H100s (IndiaAI); $234M @ $1.5B, HCLTech-led (Jun 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Answered the "Mistral wrapper" embarrassment of Sarvam-M; capability gap remains (AA Index 18/12 vs frontier ~60s). Also: BharatGen (>₹2,200 cr, 1T-param goal, 22 languages); AI4Bharat open Indic commons; Krutrim's cautionary pivot out of models (May 2026).

SEA-LION (SG) ↗

WHAT IT IS & CAPABILITIES

Open SEA-language family (v4 multimodal on Gemma; Qwen-SEA-LION line) — base layer for others' national models.

FUNDING

NRF (NAIS 2.0 >S$1B/5yr)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The "refine, don't pretrain" archetype; #1 on Tamil/Filipino tasks. National AI Council formed Feb 2026, PM-chaired.

Indonesia / Malaysia

WHAT IT IS & CAPABILITIES

Sahabat-AI (Indosat+GoTo; 70B, local languages) on "GPU Merdeka" sovereign cloud; ILMU (YTL): tops MalayMMLU (87.2%), runs on GB200 Kulai campus.

FUNDING

Indosat ~$200M DC; YTL RM10B; MY govt RM5.9B AI budget; ID sovereign fund under Danantara planned

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

"Kedaulatan digital" / "100% Made in Malaysia" — data + GPUs on national soil, models on open bases.

Vietnam / Thailand / PH

WHAT IT IS & CAPABILITIES

VN: FPT $200M AI factory, Viettel's first B200 SuperPod (Feb 2026), from-scratch PhoGPT/KiLM. TH: Typhoon (SCB 10X) + new ThaiLLM national infrastructure. PH: DOST's NAICRI national AI center (Feb 2026).

FUNDING

Corporate + state mixes

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

NVIDIA's sovereign push runs through all three; the Philippines' "no national effort" gap closed in Feb 2026.

KOREA & JAPAN

Korea MSIT program ↗

WHAT IT IS & CAPABILITIES

National sovereign-model contest: independent training required, ≥95% of global leaders targeted. Round 1 (Jan 2026): LG #1; Naver AND NC AI eliminated (Naver for a Qwen-derived encoder); Motif added; Round 2 Aug 2026, final Dec 2026.

FUNDING

₩530B direct + ₩1.46T for 13,136 GPUs; 2026 AI budget ₩10.1T; NVIDIA alliance 260k GPUs (est. $8–10B); ₩1,000T mega-plan (Jun 2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The strictest independence bar anywhere — fine-tunes of foreign bases disqualified. LG EXAONE (K-EXAONE 236B, Dec 2025), Upstage Solar Pro 2/3 (first KR GenAI unicorn, Apr 2026), SKT A.X K1 (519B MoE on Rebellions NPUs — full-stack play).

Japan — GENIAC ↗

WHAT IT IS & CAPABILITIES

METI accelerator: 70 projects/4 cycles; now robotics foundation models + AI-ready data; AI Basic Plan (Dec 2025).

FUNDING

METI/NEDO; FY2026 ~¥1.2T incl. ¥387B physical-AI

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Outputs: Rakuten AI 3.0 (~700B open MoE, Mar 2026, "Japan's largest"); NTT tsuzumi 2 (from-scratch, single-GPU, Oct 2025); SB Intuitions' 1T-param plan on 4,000+ B200s; PFN's PLaMo on own MN-Core silicon (strongest homegrown stack); Sakana ($2.65B; US-heavy cap table). Govt "Gennai" platform serves 7 domestic models to 180k civil servants.

CHINA & TAIWAN

DeepSeek ↗

WHAT IT IS & CAPABILITIES

Open (MIT) frontier: V4 Preview (Apr 2026): 1.6T/49B-active + 284B Flash, 1M context.

FUNDING

High-Flyer; ~RMB 50B raised at >$50B valuation (2026)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The Jan 2025 sovereignty catalyst — cheap, self-hostable frontier under export controls; itself banned across Western government systems.

Zhipu / Z.ai ↗

WHAT IT IS & CAPABILITIES

GLM-5 (~744B MoE, MIT): trained on ~100,000 Huawei Ascend 910Bs (self-reported), served on 7 domestic chip platforms; GLM-5.2 tops open coding ranks.

FUNDING

HKEX IPO Jan 2026 (~$7.1B debut value; raised HK$4.35B)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The domestic-silicon sovereignty benchmark — no NVIDIA anywhere in the pipeline.

Qwen (Alibaba) + Moonshot ↗

WHAT IT IS & CAPABILITIES

Qwen: ~1B cumulative downloads, >50% of global open-model downloads, 180k+ derivatives — the base layer under other nations' "sovereign" models. Moonshot's Kimi K2.6 ties GPT-5.5 on SWE-Bench Pro; #2 on OpenRouter.

FUNDING

Alibaba internal; Moonshot ~$2B @ >$20B (May 2026, state-linked backers)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

China's open-weights soft-power vector. State layer: "AI+" initiative (70%/90% penetration targets), 15th Five-Year Plan self-reliance; Ascend 950PR mass production; US H200 sales resumed with a 25% surcharge — compute statecraft monetised on both sides.

Taiwan

WHAT IT IS & CAPABILITIES

TAIDE govt models (Llama bases); Foxconn's FoxBrain; AI Basic Act (Jan 2026); Nano4 supercomputer #33 TOP500 (Jun 2026); 10,000-Blackwell national system building.

FUNDING

NT$1.5T infrastructure framing; >NT$30B 2026 budget

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Explicit defence against PRC-origin models — language/values sovereignty on US bases.

AMERICAS

Canada

WHAT IT IS & CAPABILITIES

The named-sovereign national program: Sovereign AI Compute Strategy (CAD $2B) + "AI for All" (Jun 2026: ~$2B more; $925.6M for sovereign public AI infrastructure; TELUS deal; 44 compute awards).

FUNDING

Federal budgets 2024–26

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Cohere is the champion ($7B pre-merger; CAD $240M federal DC investment; 2026 IPO expected) — now merging with Aleph Alpha. Bell–Cohere: USD $220M, 2,304 GB200s at Merritt BC (Jun 2026) — "first fully domestic AI stack," live late 2026.

Brazil

WHAT IT IS & CAPABILITIES

PBIA (R$23B/4yr headline; ~55% is credit — actual public spend "considerably low" per academics); Santos Dumont upgraded to 18.85 PF; Maritaca's Sabiá-4 (PT-native, continual-pretrained, API-only); WideLabs Amazônia.

FUNDING

Federal + private

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

"IA soberana": state cloud with state-held keys by 2028; execution and composition caveats apply throughout.

LatamGPT (CENIA) ↗

WHAT IT IS & CAPABILITIES

Regional open model: Llama-3.1-70B continued-pretraining on 8TB+ regional data, 15+ countries.

FUNDING

~$550K total (CENIA + CAF)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Launched 10 Feb 2026 with President Boric; v2 due ~Sep 2026. The minimum-viable-sovereignty experiment — regional relevance at five figures per country.

United States

WHAT IT IS & CAPABILITIES

The stack everyone hedges against: AI Action Plan (Jul 2025) + full-stack AI export program; Anthropic Claude Gov + $1 federal offer; OpenAI Stargate (~7GW/$400B+ committed) + "OpenAI for Countries"; Palantir as classified-enclave layer.

FUNDING

Private frontier capital + federal procurement

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Washington's counter-move: reframe others' sovereignty as alliance-managed access to the American stack.

AFRICA · OCEANIA · RU/TR

Cassava (pan-Africa) ↗

WHAT IT IS & CAPABILITIES

Africa's first AI factory: ~3,000 NVIDIA GPUs live in South Africa (Jun 2025); ~12,000 planned (EG/KE/MA/NG).

FUNDING

Up to $720M + NVIDIA partnership

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

African-owned infrastructure keeping African data on-continent. Nigeria's N-ATLAS (UNGA80, Sep 2025): first African national LLM (Llama-3 base; Yoruba, Hausa, Igbo + ASR).

Russia

WHAT IT IS & CAPABILITIES

Sber's GigaChat 3 line open-sourced under MIT (Nov 2025); Yandex's Alice AI (strongest Russian model; mid-tier globally).

FUNDING

State-linked (Sberbank); sanctions-constrained

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Jun 2026: pitching "sovereign AI" to the Global South — "it will align with your values"; discussing Chinese chips. Sovereignty exported as a values wedge — the cautionary mirror.

Australia / NZ

WHAT IT IS & CAPABILITIES

Maincode's Matilda (from-scratch; dropped the "sovereign" label as divisive); Sovereign Australia AI (256 B200s; AU$10M copyright fund); Kangaroo LLM stalled. NZ: adoption-first strategy — explicitly no sovereign models.

FUNDING

Private (AU); policy-only (NZ)

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

The label-instability exhibit: one company de-badged, one nation opted out entirely.

Türkiye — T3AI ↗

WHAT IT IS & CAPABILITIES

First Turkish LLM (T3 Foundation + Baykar; 1,792 volunteers); beta Aug 2024.

FUNDING

Foundation + defence-champion backing

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

National LLM tied to the digital-sovereignty agenda.

VENDORS

NVIDIA ↗

WHAT IT IS & CAPABILITIES

The compute substrate of nearly every program above.

FUNDING

Sovereign revenue >$30B FY2026 (3x YoY, ~14% of total); 50+ countries; 52% of tracked projects

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Sovereignty that depends on one US chipmaker — and a vendor-defined category.

Trusted-JV middle path

WHAT IT IS & CAPABILITIES

S3NS (Thales+Google; SecNumCloud 3.2 qualified Dec 2025), Bleu (Capgemini+Orange), Delos (SAP) — US technology under local ownership/operation.

FUNDING

French/German public sectors

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Sovereignty via corporate structure; credible for compliance, untested against a determined US legal order.

Microsoft / AWS / Oracle

WHAT IT IS & CAPABILITIES

"Sovereign clouds": air-gapped model support (MSFT, Feb 2026); in-country Copilot in 15 nations; AWS European Sovereign Cloud GA (Jan 2026, €7.8B, separate partition); Oracle EU Sovereign Cloud (Frankfurt, Madrid) + Stargate builds.

FUNDING

Global

ACHIEVEMENTS, END GOAL & SOVEREIGNTY READ

Jurisdiction trumps geography: Microsoft's own Senate testimony concedes US legal reach. Residency ≠ sovereignty.

About Mirai Collective

Mirai Collective is an AI development studio. We build AI products end-to-end, train whole teams from zero to one, and guide organisations through AI adoption — practitioners, not theorists. The full landscape dataset behind this briefing (~110 profiled organisations with funding, capability and sourcing detail, plus the claim-level verification log) is available on request: hello@miraicollective.io

Sources. Company filings and press releases · EuroHPC Joint Undertaking · national ministries and parliamentary replies · TOP500.org · CNAS Sovereign AI Index · Stanford HAI · McKinsey · International Energy Agency · press of record (Bloomberg, CNBC, Fortune, Reuters, TechCrunch). Figures as of 6 July 2026.