# We pre-train physical-AI models for nations, inside their borders.

The nation keeps the data, the compute and the model. We do the pre-training and hand it off.

# The models that will run a nation's ports, factories and defense are owned by someone else.

# 18 days

June 12 to June 30, 2026. A US export-control directive suspended every foreign national's access to a frontier AI model, worldwide. The switch is real, and someone else holds it.

If Washington or Beijing owns the models, then in a conflict or a tension they can shut them off. So nations are going to need to own their own.

**184 sovereign-AI projects in 67 countries.**

Governments already know this. Physical AI, the models that will run cranes, grids, plants and vehicles, is where owning the model stops being abstract.

Sources: Al Jazeera, Jun 14, 2026; Euronews, Jul 1, 2026; CNAS Sovereign AI Index, Aug 20, 2026.

# Nations bought the compute first. The models to run on it are the gap.

## 59%

of the 184 sovereign-AI projects are infrastructure. 9% are data. More compute projects were announced in Q1 2026 than in all of 2024.

## 3x

NVIDIA's sovereign-AI business more than tripled to over $30B in FY2026, then grew another 35% last quarter.

## Open

Physical-AI base models are open: Cosmos Predict 2.5, Cosmos Reason 2 and GR00T N1.6 shipped on Hugging Face on Jan 5, 2026. Continuing one on a nation's own data is now a 90-day project.

Sources: CNAS, Aug 2026; NVIDIA earnings calls, Feb 25 and Aug 26, 2026; NVIDIA, Jan 5, 2026.

# Orion pre-trains the model inside the nation's boundary, then hands it off.

### 1 Ingest

The nation's port, grid and factory data, with a provenance record for every file. Faces and plates blurred.

### 2 Curate

Deduplicate, mine the events that matter, fix the mixture. Versioned, with a reason for every dropped record.

### 3 Pre-train

Continue an open physical-AI base model on that data, on the nation's compute, at a fixed and logged budget.

### 4 Prove

Held-out cameras and months, matched compute, pre-registered threshold. No threshold, no claim.

### 5 Hand off

Checkpoints, model card, runbooks. The nation's lab re-runs it without us. That is the acceptance test.

Data never leaves the country. We do the pre-training, the half-art, half-science part. The nation does the post-training and truly owns the model.

Not for sale: deployment, robots, applications. The nation's existing vendors stay on side.

# One terminal, one task, one metric, ninety days. $250,000 fixed.

| | |
|---|---|
| **Scope** | One environment, one task, one metric, agreed in writing before training starts. |
| **Data** | Stays in-country. Orion works inside the customer's permissioned boundary. |
| **Compute** | The nation's own, or rented in-region under the nation's account. |
| **Deliverables** | Dataset manifest, pre-trained checkpoints with model card, evaluation report, runbooks, handoff. |
| **Rights** | The nation owns the model. Orion keeps the method. |
| **Then** | A national program: more environments, more tasks, in-country training capability, refreshed every cycle. |

**Pilot economics, expected case**

Price $250,000. Delivery cost $103,900. Contribution $146,100, 58%. Fully loaded margin 40%.

Compute is the customer's, so margin never depends on renting GPUs.

What compounds: the evaluation standard nations judge vendors by, recipes that travel between nations when data cannot, and lineage that brings every refresh back to us.

Planning figures from the Orion Operating Workbook, Oct 21, 2026. Not results.

# Our unfair insight: a port has already labelled its own data.

**The labels are already written.**

A terminal records what moved, when and where, in its operating system and its crane controllers. Align the camera archive with those logs and you get labels nobody has to draw.

**Curation is where the compute saving is.**

Port video is overwhelmingly redundant: fixed cameras, idle hours. DatologyAI reports 7.7x faster training from curation on text. We expect the same lever on port video, and we will measure it.

**Domain data goes in at pre-training.**

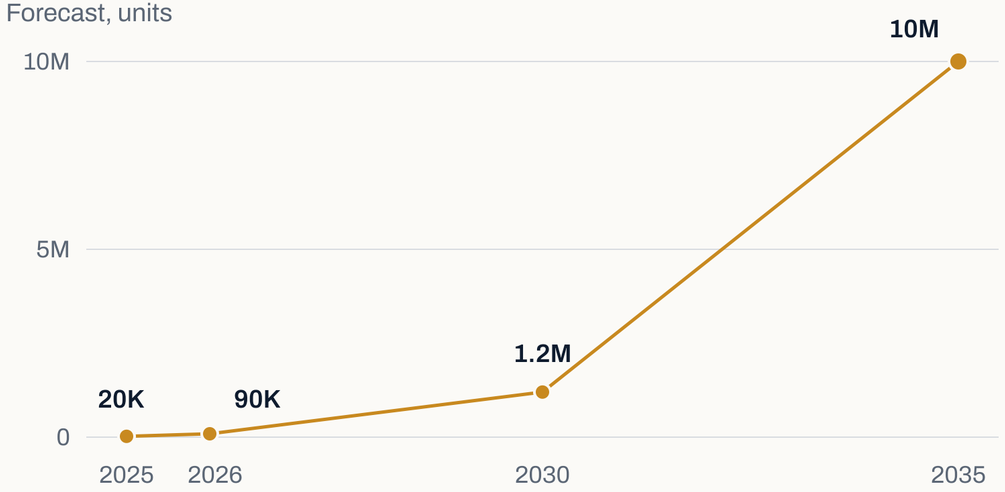
"We look for a lot of domain data to put into pre-training," Alexandr Wang, Meta, Oct 12, 2026. Nations have the domain data. We bring the pre-training.

Sources: DatologyAI, Aug 18, 2025; Cheeky Pint #35, Oct 12, 2026. First base model: an open self-supervised video encoder.

# Every nation funding sovereign AI will need physical-AI models it owns.

### Humanoid robots shipped per year

Forecast, units




BofA Global Research, Mar 12, 2026. 86% a year compounded from 2025.

## $30B+

NVIDIA's sovereign-AI revenue in FY2026, more than tripled in a year.

## 3.3tn won

Korea's public funding for physical-AI ports by 2035, announced Jul 23, 2026. One country, one sector.

## $138B

Goldman Sachs' 2035 humanoid market, raised from $38B in September 2026.

**Bottom up.** 67 countries fund sovereign AI; a handful host a frontier lab. One $250K pilot opens each door. A national program then covers every port, grid and plant the nation owns, and refreshes every cycle.

Sources: NVIDIA Q4 FY2026 call, Feb 25, 2026; Asia Economy, Jul 23, 2026; Goldman Sachs via BigGo, Sept 1, 2026; CNAS, Aug 20, 2026.

# Everyone else sells a nation access or compute. We make the nation the owner.

| Who | What they sell a nation | Why Orion wins |
|---|---|---|
| **Frontier labs** | OpenAI for Countries, Google, Anthropic sell in-country data centers and model access. The weights stay theirs, and so does the switch. | Ours is handed over. The lab can re-run everything without us. |
| **NVIDIA, clouds, neoclouds** | Compute and reference stacks. | Our substrate, not our rival. We train on the GPUs the nation already bought. |
| **National labs** | Research and post-training. | Our hand-off partner. They take over the day pre-training ends. |
| **Systems integrators** | Deployment and integration. | They don't pre-train. They ship what we hand off. |
| **Data-curation startups** | DatologyAI curates for labs, text first. | Not physical AI, not in-country, not for nations. |

We only do pre-training, so nobody the nation already pays loses a contract. That is why the door opens for a small team.

Why we win: a narrow job, a handoff the buyer can verify, and recipes that travel between nations while data never does.

# A small, deeply technical San Francisco team that only does pre-training for nations.

**Spencer Brown, founder and CEO.**

Founder of LVL3.AI, AI agent infrastructure, Palo Alto. Harvard master's in finance, in progress. Runs Orion's lab with an always-on agent swarm: nightly builds, 57 tests passing, one pull request per run.

**Hiring now.**

A technical lead for continued pre-training and the evaluation harness. A data and evaluation owner for permissions, splits and the gold set.

Nations with world-class industry and no frontier lab need a partner who will go in, pre-train, hand off and leave the capability behind. We are built to be that partner, and nothing else.

First conversations: Brazil, Argentina, the UAE, Saudi Arabia, Indonesia, Thailand, Cambodia and Central Europe.

Advisors and a design partner are being sought; none is named until confirmed.

# Raising a pre-seed to deliver the first accepted pilot.

**What exists today, Oct 21, 2026.**

The pilot offer and contract terms. A lab scaffold that runs nightly: provenance checks, four-pass curation, matched-compute arms, bootstrap intervals, 57 tests. Thirty target accounts researched, with sources.

**What the money buys.**

December 2026: a controlled result on public maritime data, with intervals. Q1 2027: the first paid pilot, accepted when the nation's lab re-runs the pipeline without us. Three paid pilots within 90 days of the first. Two hires.

The round is a SAFE. Size and terms on request.

No customers, no model results and no revenue are claimed in this deck. Every number carries its source and date.

Talk to Spencer

# Sources

| | |
|---|---|
| **Export controls** | Al Jazeera, "US asks Anthropic to block global access to top AI models," Jun 14, 2026. Euronews, "US lifts export controls on powerful AI models, Anthropic says," Jul 1, 2026. Access was suspended on Jun 12 and the lifting announced Jun 30, 2026. |
| **Sovereign AI index** | CNAS, "Sovereign AI's Second Wave Is Coming Into View," Aug 20, 2026: 184 projects in 67 countries through H1 2026; CNAS Sovereign AI Index, updated Aug 18, 2026: 59% infrastructure, 32% models, 9% data; more infrastructure projects announced in Q1 2026 than in all of 2024. |
| **NVIDIA** | Q4 FY2026 earnings call, Feb 25, 2026: sovereign-AI business "more than tripled year over year to over $30 billion"; physical AI contributed "north of $6 billion" of FY2026 revenue. Q2 FY2027 call, Aug 26, 2026: sovereign AI up 35% sequentially and more than tripled year over year. Press release, Jan 5, 2026: Cosmos Predict 2.5, Cosmos Transfer 2.5, Cosmos Reason 2 and Isaac GR00T N1.6 released as open models on Hugging Face. |
| **Physical AI market** | Bank of America Institute, "Physical AI, part 2: Humanoid robots," Mar 12, 2026: 20,000 shipments in 2025, 90,000 in 2026, 1.2 million in 2030, 10 million in 2035, an 86% compound annual growth rate. Goldman Sachs humanoid outlook via BigGo Finance, Sept 1, 2026: about 6.5 million units and $138 billion by 2035, up from 1.4 million units and $38 billion. |
| **Korea** | Asia Economy, "Ports to Become Giant AI Robots," Jul 23, 2026: Ministry of Oceans and Fisheries Physical AI Port Strategy, about 3.3 trillion won of public funds by 2035, Busan, Gwangyang and Jinhae New Port. |
| **Pre-training and curation** | Alexandr Wang, Cheeky Pint #35, Oct 12, 2026. DatologyAI, curation results on text, Aug 18, 2025: 7.7x faster training. |
| **Orion planning figures** | Orion AI Command Plan and Operating Workbook, Oct 21, 2026: pilot price, cost cases, milestones and the lab scaffold's test count. Planning figures, not results. |

Numbers are reported as published by the source named, with its date. Nothing here is an Orion result.