AIEONME FOUNDATION

AI Ethics, Oversight & Norms for Multilateral Engagement

in partnership with AIEONME FRONTIER (Global South)

STRATEGIC PAPER — JUNE 2026

Post #3 of the Glasswing Trilogy · The Constructive Capstone

The Andean Protocol

A Blueprint for a Latin American Frontier-AI Safety and Evaluation Institute, Financed by Lithium-for-Compute

From diagnosis to architecture: how the region that mines AI's future can build a seat at the table that governs it

---

THE GLASSWING TRILOGY

#1 Claude Mythos — the two-tier problem · #2 From Spud to Scarcity — the pattern + lithium

#3 The Andean Protocol — the solution

June 2026

AIEONME Foundation | AIEONME Frontier | @aieonme

Author: Arq. Gustavo E. Cardozo | Founder & Executive Director

Ciudad Perico, Jujuy, Argentina (Lithium Triangle, NOA)

AIEONME-SP-2026-007

The Andean Protocol

A Blueprint for a Latin American Frontier-AI Safety and Evaluation Institute, Financed by Lithium-for-Compute

---

Publication Details

Publisher: AIEONME Foundation, Ciudad Perico, Jujuy, Argentina (NOA, Lithium Triangle).

Series: AIEONME Strategic Paper Series.

Document ID: AIEONME-SP-2026-007.

ISSN: [pending — application filed with CAICYT-CONICET, Centro Nacional Argentino del ISSN].

DOI: [reserved via Zenodo — https://doi.org/10.5281/zenodo.XXXXXXX].

Version: Version 1.0, June 2026.

Author: Arq. Gustavo E. Cardozo (ORCID: [0000-0000-0000-0000]).

Suggested Citation

```
Cardozo, Gustavo E. 2026. The Andean Protocol: A Blueprint for a Latin American
Frontier-AI Safety and Evaluation Institute, Financed by Lithium-for-Compute. AIEONME
Strategic Paper Series, AIEONME-SP-2026-007. Jujuy: AIEONME Foundation.
https://doi.org/10.5281/zenodo.XXXXXXX.
```

License

Contact: @aieonme · aieonme.ai · Ciudad Perico, Jujuy, Argentina.

About the AIEONME Foundation

The AIEONME Foundation (AI Ethics, Oversight & Norms for Multilateral Engagement) is an independent think tank founded on 1 February 2026 and based in Jujuy, Argentina, in the heart of the Lithium Triangle. It works to give the Global South — and Latin America in particular — an authoritative, evidence-based voice in the governance of frontier artificial intelligence. Its Spanish-language programme, AIEONME Frontier, addresses sovereign AI and frontier AI safety from a Global South perspective.

About the Author

Arq. Gustavo E. Cardozo is the Founder and Executive Director of the AIEONME Foundation. He writes from Ciudad Perico, Jujuy, Argentina.

Funding and Independence

This paper received no external, governmental, corporate, or AI-laboratory funding. It was produced with the AIEONME Foundation's own resources. The Foundation retains full editorial independence: no funder, partner, or third party had any right of review, comment, or approval prior to publication. The analysis and recommendations are those of the author alone.

Conflict of Interest

The author declares no financial or personal conflicts of interest in relation to the institutions, companies, or governments discussed in this paper.

Methodology

This is a desk-based strategic analysis. It synthesises primary sources (institutional budgets and mandates from the UK AISI, US CAISI, the EU AI Office, CENIA, CEPAL/ECLAC, CAF, and corporate and governmental announcements), official statistical series (USGS Mineral Commodity Summaries, ILIA 2025, provincial mining ministries), peer and policy literature on compute governance and green extractivism, and current reporting through June 2026. Where a quantity is contested across sources, the paper reports a range and names the disagreement rather than asserting a single figure. Forward-looking cost and revenue figures are explicitly framed as estimates. The cut-off date for all figures is June 2026; the policy environment is fast-moving and figures should be reverified before any formal citation.

Limitations

This analysis is constrained in three ways the reader should weigh. First, several frontier-AI access facts — notably the membership and country list of Project Glasswing — rest on single-source press reporting and could not be confirmed against a primary institutional document; they are flagged as such in the text. Second, the financing model depends on lithium-revenue and state-capture projections (the Chilean Codelco-SQM joint venture; Argentine provincial royalties) that are procyclical and politically contingent; a sustained fall in lithium prices or a change of government would materially alter the arithmetic. Third, the willingness of frontier laboratories to grant evaluation ("Trusted Access") to a Latin American institute is unknown and treated as a variable, not an assumption. None of these limitations is disqualifying, but each defines a boundary of confidence around the blueprint.

Contents

Bottom Line Up Front

This is the third and final paper of the Glasswing Trilogy. Post #1 (AIEONME-SP-2026-004) diagnosed Claude Mythos and Project Glasswing as the birth of a two-tier AI world. Post #2 (AIEONME-SP-2026-006, “From Spud to Scarcity”) confirmed the pattern with OpenAI’s GPT-5.5 and named the lithium-AI extraction asymmetry: Latin America mines the lithium that powers AI datacenters but is excluded from the frontier AI those datacenters run. Post #3 moves from diagnosis to architecture.

The thesis is simple and, on the evidence, fundable: Latin America can build a credible frontier-AI evaluation institute — a LATAM-AISI — for roughly US$40–60 million per year, because evaluation requires two-to-four orders of magnitude less compute than training. A minimum-viable cluster of 128–256 H100-class GPUs is enough to run UK AISI–grade evaluations on Claude Opus and GPT-class models. The institute does not need a frontier training cluster; it needs a frontier evaluation cluster, and that is affordable.

The financing core — “lithium-for-compute” — ties a sliver of Andean lithium revenue to sovereign evaluation capability. It is feasible, but only if designed around hard constraints: Argentina’s 3% provincial royalty cap and Milei-era RIGI stability clauses, Chile’s far more aggressive Codelco-SQM state-capture model, and Bolivia’s stalled YLB. The Andean Protocol is therefore Chile-anchored, Argentine-provincially co-financed, Brazil-and-Colombia politically sponsored, and CAF-financed — not Bolivia-led.

The window is closing. With EU GPAI enforcement live on August 2, 2026, Project Glasswing expanded to ~150 organizations across 15+ countries (India the only non-Western member), and the US voluntary-review executive order of June 2, 2026, the two-tier pattern is hardening into law and treaty practice. Latin America has roughly eighteen months to convert mineral leverage into a governance seat, or accept permanent client status.

1. The Cost of Governance Is Far Below the Cost of Capability

The decisive insight that makes this blueprint realistic is the asymmetry between training compute and evaluation compute. Meta’s Llama 3 405B consumed roughly 30.8 million H100-hours across about 16,000 GPUs. By contrast, METR’s frontier R&D evaluations run on eight or fewer H100 GPUs per environment over an eight-hour budget — a full pass is on the order of dozens to a few hundred H100-hours. OpenAI’s MLE-Bench consumes roughly 1,800 lower-tier GPU-hours per run. Evaluation is two-to-four orders of magnitude cheaper than training.

This is why a resource-constrained region can govern frontier models even though it cannot build them. The UK AI Security Institute — the global reference — runs on about £66 million per year (~US$85M) with roughly 250 staff and privileged access to over £1.5 billion of national compute. But its evaluation framework, Inspect, released open-source under MIT license in May 2024, ships 200-plus pre-built evaluations and runs primarily as API calls plus sandboxed containers. Most of the workload is inference and orchestration, not training. A LATAM-AISI joining the international network would inherit that stack rather than build it.

The comparative budgets confirm affordability. Singapore's AISI runs on about S$10 million per year. Canada pledged CA$50 million over five years (~US$7.3M/year). Japan's institute operates with roughly twenty-three staff. A 2024 CSIS audit concluded that the annual budgets of network members hover around US$10 million, with some exceptions. A LATAM-AISI at US$40–60 million per year would be the largest in the Global South and credible enough to sit at the network table beside the UK and US.

Table 1 — What an AI Safety / Evaluation Institute Costs

Institute	Annual Budget	Staff	Compute Model
UK AISI	~£66M (US$85M)	~250	Isambard-AI + Dawn (>£1.5B)
US CAISI	~US$15M (proposals to scale)	NIST-housed	Shared federal
Singapore AISI	~S$10M/yr (S$50M grant)	NTU-housed	Project Moonshot toolkit
Canada AISI	~US$7.3M/yr (CA$50M/5yr)	ISED-housed	CIFAR-partnered
LATAM-AISI (proposed)	US$40–60M/yr (mid-tier)	60–80 to 150	128–512 GPU eval cluster

2. The Proof of Concept Already Exists: Latam-GPT

Latin America has already demonstrated that regional cooperation can produce a frontier-relevant artefact cheaply. CENIA-Chile delivered Latam-GPT — a 70-billion-parameter model on the Llama 3.1 architecture, trained on more than 8 terabytes of curated regional data and 2.6 million documents from 20 countries, built by 200-plus collaborators across 15 countries — launched February 10, 2026 in Santiago in the presence of President Boric. It was seeded by US$550,000 from CENIA plus CAF co-financing, with training on a US$5–10 million supercomputer at Universidad de Tarapacá in Arica.

If a sovereign large language model is buildable for under US$15 million of capital expenditure, an evaluation institute — which needs even less compute — is unambiguously fundable. Latam-GPT is not the destination; it is the existence proof. It also gives LATAM-AISI its single most defensible unique contribution to the global network: Spanish, Portuguese and Indigenous-language evaluation of frontier models, a capability no Anglophone institute possesses and that CENIA's 2.6-million-document regional corpus makes immediately viable.

3. The Asymmetry That Justifies the Protocol

The Latin American Artificial Intelligence Index (ILIA 2025), produced by CENIA with ECLAC/CEPAL, quantifies the gap. The region generates 14% of global visits to AI tools and ranks third worldwide in generative-AI downloads. It accounts for 6.6% of global GDP. Yet it attracts only 1.12% of global AI investment. The bottleneck is not demand or talent; it is institutions. ILIA flags that national strategies lack budgets, implementation and evaluation, rendering them statements without action.

Brain drain is structural. Brazil, Mexico, Chile, Argentina and Colombia together produce 90% of regional AI research output, but only Costa Rica and Uruguay are net talent importers. The region trains researchers who then evaluate frontier models in London, Washington and Singapore — never from home. The Andean Protocol is, among other things, a talent-retention instrument: a credible institute gives the region's safety researchers a reason to stay or return.

4. The Financing Architecture: Lithium-for-Compute

The novel mechanism of the Andean Protocol is to tie a small, ring-fenced fraction of Andean lithium revenue to sovereign AI evaluation capability. The design must respect three very different national realities.

4.1 Argentina — Provincial, Capped, and Milei-Constrained

Argentine national lithium exports reached approximately US$905 million in 2025, a 40.3% year-on-year increase and the highest annual figure on record (this national lithium-only figure is distinct from the US$953 million in total provincial mining exports reported for Jujuy in Post #2 of this trilogy, which covers all minerals). But provincial royalties are capped at 3% of mine-mouth value under the Mining Investment Law, and the federal RIGI regime guarantees 30-year stability plus a 3-year export-duty exemption after registration — meaning federal lithium revenue will fall sharply from 2027. Jujuy province collected only an estimated US$11–16 million in lithium royalties in 2024; the three producer provinces (Jujuy, Salta, Catamarca) together captured roughly US$50 million in 2025 via royalties and JEMSE-style equity dividends.

The implication is decisive: Argentina's contribution to the Protocol must be sub-national. The Milei administration has structurally chosen deregulation and will not lead a multilateral governance institute. But Argentine provinces and scientific institutions (CONICET, UBA, Universidad Nacional del Litoral) can sign on independently — the "Andean Hub" model, anchored in Jujuy.

4.2 Chile — The Anchor

Chile's Codelco-SQM joint venture, Nova Andino Litio SpA, held its first board meeting on December 29, 2025 (Codelco majority, per the Comptroller's >50% requirement). It targets 280,000–300,000 tonnes per year of lithium carbonate equivalent through 2060 and is projected to deliver US$25–45 billion in lifetime state revenue, capturing about 70% of operating margin in 2025–2030 and 85% from 2031. Chile's state-capture model is an order of magnitude more powerful than Argentina's royalty cap. A ring-fence of roughly 0.2% of this JV's state revenue would alone fund the institute's mid-tier budget. Chile is therefore the financial anchor, with CENIA as the technical seat.

4.3 Bolivia — Political Signatory Only

Bolivia holds approximately 23 million tonnes of lithium resources — about a fifth of the global total — but YLB produces under 2,000 tonnes per year against 15,000-tonne nameplate capacity, and US$1.4 billion in Russian and Chinese contracts remain stalled in the Assembly. Under President

Rodrigo Paz, a pivot to private investment is signalled but years from scale. Bolivia can be a founding political signatory; it cannot be a financial anchor in the 2026–2030 window.

4.4 The Multilateral Treasury

CAF — the Development Bank of Latin America and the Caribbean, AA+/Aa3/AA- rated, 23 shareholders — financed the Latam-GPT seed and is the natural treasury and disbursement vehicle for LATAM-AISI. The IDB and its IDB Lab can co-fund fellowships and the data partnership. CEPAL provides the policy host and the ILIA evaluation index. This multilateral backstop is essential because lithium revenue is procyclical: the institute's budget must be denominated in absolute US dollars with development-bank guarantees, never as a fluctuating percentage of royalty flows.

Table 2 — Illustrative Funding Mix for the US$40–60M Mid-Tier

Source	Mechanism	Target / yr
Chile (anchor)	CENIA host + Codelco-SQM JV revenue ring-fence (~0.2%)	**US$20M**
Brazil	PBIA governance + infrastructure carve-out	**US$15M**
Argentina (provincial)	Jujuy+Salta+Catamarca 1% royalty pool + JEMSE-style equity	**US$5–8M**
Colombia/Mexico/Uruguay/ Costa Rica	Cartagena Declaration signatory contributions	**US$5–8M**
CAF + IDB	Multilateral co-financing	**US$10–15M**
EU-LAC + UK AISI Alignment	Pass-through grants	**US$5–10M**

5. The Technical Scope LATAM-AISI Must Own

Drawing on the UK AISI research agenda and the International Network's joint testing exercises, the institute's evaluation portfolio should comprise: cyber and autonomous-systems evaluations (the direct successors to the joint UK/US tests that found frontier model safeguards could be routinely circumvented); bio and chemistry PhD-equivalent knowledge probes; agentic and autonomy evaluations using AgentHarm-class benchmarks; safety cases and control evaluations; persuasion and human-influence studies; and societal-scale resilience research.

Above all, LATAM-AISI must own the capability no Anglophone institute has: sovereign-language evaluation. Spanish, Portuguese and Indigenous-language testing of frontier models for harms, persuasion and jailbreak robustness is the region's unique and non-substitutable contribution to the global network. CENIA's Ethics Team and the Latam-GPT corpus make this viable from day one.

6. Governance Architecture

The recommended model is a hybrid of CERN (a multinational scientific consortium) and the IAEA (a safety and inspection function), housed politically at CEPAL, operationally at CENIA, and financially at CAF.

- Political seat (Santiago): CEPAL/ECLAC as Secretariat; Cartagena Declaration signatories as founding members; an annual Ministerial Council.
- Technical seat (Santiago / Arica): CENIA as operational host; the Tarapacá supercomputer as the regional evaluation cluster.
- Financial seat: CAF as treasury, with IDB co-financing.
- Brazilian co-lead (Petrópolis): the LNCC AI Institute as the Portuguese-language and bio/medical evaluation node.
- Argentine co-lead (Jujuy): a provincial-led “Andean Hub” leveraging the JEMSE equity model, specialised in Indigenous-language AI harms — working around the federal deregulation stance.
- Network seat: pursue formal membership in the International Network of AI Safety Institutes by Q4 2026, joining Kenya as a Global South anchor.

Political feasibility is realistic precisely because it does not depend on any single government surviving an election. Brazil and Chile will sponsor; Colombia (the Cartagena host) and Mexico can be brought in; Argentina participates sub-nationally. Permanence beyond electoral cycles is built in via the CEPAL host and CAF treasury — in Year 1 the Protocol is a CEPAL Specialised Programme with CAF financing, far harder to dissolve than a treaty organisation, which it can become later.

7. The June 2026 Frontier Context — Why Now

The two-tier pattern diagnosed in Posts #1 and #2 is hardening from market practice into law and treaty. Project Glasswing has expanded to approximately 150 organizations across more than 15 countries, with India the only non-Western member. The US executive order of June 2, 2026 formalises a US-centric voluntary pre-release review pathway. The EU AI Office begins enforcing GPAI obligations on August 2, 2026, with fines up to €15 million or 3% of global turnover. The International Network of AI Safety Institutes is consolidating from informal forum into a de facto governance body — and Latin America is not in the room.

Each of these is a closing door. The GPAI enforcement date in particular creates an opening: LATAM-AISI could position itself as the third-party assessor of choice for Spanish and Portuguese-language compliance under EU AI Act Articles 53 and 55 — a revenue-generating function with structural demand, independent of frontier-lab goodwill.

8. Counterarguments and Honest Feasibility

This blueprint is presented with its failure modes visible, because a proposal that hides its risks is a manifesto, not a plan.

Argentine royalties are provincial, not federal. Build provincially (Jujuy, Salta, Catamarca), bypassing federal politics. This is the Andean Hub.

Milei privatization and Súper RIGI. Without a “compute commitment clause” in RIGI adhesion (e.g., 1% of investment value earmarked for sovereign evaluation infrastructure), Argentina free-rides. Propose the clause; do not assume cooperation.

Lithium price volatility. Battery-grade carbonate swung from ~US$13,400/t in December 2025 to ~US$26,300/t by late January 2026. Royalty revenue is procyclical; the budget must be set in absolute USD with CAF/IDB backstops.

Brain drain. Counter with UK AISI Alignment Project pass-through, CIFAR-equivalent fellowships, rotating placements at northern institutes, and salaries pegged to regional public-sector ceilings plus international top-ups.

US/China geopolitical pressure. Both will resist a Latin American institute with genuine evaluation authority. Mitigate with dual-signatory architecture and explicit non-alignment in the founding charter.

Will frontier labs grant Trusted Access? Uncertain — Posts #1 and #2 documented systematic exclusion. Begin with open-weight evaluation (Llama, Qwen, DeepSeek, Mistral) and EU-compliance assessment; earn trusted access over two to three years. This is a viable plan B with its own demand.

The Indonesia precedent failed at the WTO. The Andean Protocol does not propose an export ban; it applies conditionality to the investment-incentive regime (RIGI-plus / CEOL-plus), which is sovereign discretion and survives WTO scrutiny.

9. Implementation Annex — Phased Architecture and KPIs

Phase 0 — Founding (June 2026 – February 2027)

1. Amend the Cartagena Declaration to commit signatories to a regional evaluation institute, CEPAL-hosted, CENIA technical seat.
2. Founding compact: Chile, Brazil, Colombia, Uruguay, Costa Rica as national signatories; Jujuy, Salta, Catamarca, Santa Fe as Argentine sub-national signatories.
3. CAF board resolution ring-fencing US$15M seed from the 2027 envelope; observer-status application to the International Network.

KPI (by Feb 28, 2027): 7+ national + 4 Argentine provincial signatories; CAF/IDB MoU; CEPAL Secretariat assigned.

Phase 1 — Minimum Viable Institute (March 2027 – December 2028)

4. Operational launch at CENIA, ~30 staff (60% PhD-level technical); secure 128–256 H100/H200 GPUs via NVIDIA sovereign partnership + Claro hosting + Tarapacá shared access.
5. Deploy the Inspect stack adapted for Spanish/Portuguese/Indigenous-language evaluation; first evaluations in bilingual cyber CTF, AgentHarm-LATAM, and bio/chem co-run with Fiocruz.
6. Open Trusted Access negotiations with Anthropic, OpenAI and Google DeepMind for Claude Opus 5 / GPT-6 / Gemini 3-class models.

Budget: US$15–25M/yr.

KPIs (by Dec 2028): 50+ technical staff; open-source Spanish/Portuguese Inspect fork; two joint evaluations with UK AISI or EU AI Office; one published pre-deployment evaluation; formal Network membership.

Phase 2 — Credible Regional Institute (2029 – 2032)

7. Scale to US$40–60M/yr via the Chile-anchored revenue ring-fence, Brazilian PBIA carve-out, and Argentine provincial pooling.
8. Stand up the Jujuy/Salta Andean Hub for Indigenous-language AI harms, funded by a 1% lithium-royalty contribution and JEMSE-equivalent equity in new Súper RIGI compute projects.
9. Position LATAM-AISI as third-party assessor for Spanish/Portuguese compliance under EU AI Act Articles 53 and 55.

KPIs (by Dec 2032): US$50M+ budget secured for 5+ years; 150+ staff; 5+ joint evaluations published; regional Trusted Access framework signed by ≥2 labs; Indigenous-language methodology adopted by ≥1 other AISI; regional AI-safety brain-drain net flow turned positive.

Phase 3 — Treaty Organisation (2033+)

If credible and funded, convert from a CEPAL Specialised Programme to a treaty-based international organisation modelled on CERN, with formal voting governance, permanent funding, and direct frontier-lab access agreements.

Thresholds That Would Change This Plan

- If Chile unwinds the Codelco-SQM state-capture model then shift the financing core to CAF-anchored with heavier Brazilian and multilateral shares; reduce Phase 1 to the US$15M floor.
- If Argentine federal politics shifts toward regulation post-2027 midterms then fold the Andean Hub into a federally-recognised national AISI.
- If Bolivia unlocks YLB investment by 2028 then add Bolivia as a fourth lithium-revenue contributor.

- If frontier labs refuse Trusted Access then pivot to open-weight evaluation and EU-compliance assessment as a revenue service.

Conclusion: From Landlord to Architect

Post #2 ended with an image: the region as landlord, not participant — supplying lithium, bearing the water and environmental costs, excluded from the frontier AI its minerals help power. The Andean Protocol is the answer to that image. It does not ask the frontier labs for charity, and it does not wait for an invitation that Posts #1 and #2 proved will not come. It proposes that the region build the one thing that converts raw-material leverage into governance standing: the capability to evaluate frontier models in its own languages, on its own compute, under its own institutions.

The arithmetic is favorable. Governance costs a fraction of capability. The proof of concept — Latam-GPT — already exists. The treasury — CAF — already finances regional AI. The political instrument — the Cartagena Declaration — already binds seventeen countries. The leverage — the largest lithium concentration on Earth — is in the region's hands. What is missing is the decision to assemble these pieces before the August 2026 enforcement date, before GPT-6, before the next Glasswing expansion closes the door.

AIEONME Foundation and AIEONME Frontier write this from Jujuy, in the Lithium Triangle, because the argument is only fully visible from here: the same salars that feed the batteries of the AI buildout can fund the institute that gives the region a voice in governing it. The region that mines the future of AI can choose to architect its governance rather than be governed by it. That choice has a deadline, and the deadline is now.

Disclaimer

Distribution

This paper and its companion 2-page Executive Briefing are available open-access at aieonme.ai and deposited with a permanent DOI on Zenodo. For formal correspondence, technical seminars, or institutional collaboration, contact the AIEONME Foundation at @aieonme.

AIEONME Foundation · AIEONME Frontier

AI Ethics, Oversight & Norms for Multilateral Engagement

@aieonme | Founded February 1, 2026 | Jujuy, Argentina

Independent analysis. Capstone of the Glasswing Trilogy. Figures current as of June 2026 and subject to a fast-moving policy environment; where quantification is contested, ranges are stated.