

Safety compute for the open frontier.

Pacific Compute is a safety-only compute cluster in Singapore for open-weight labs on both sides of the Pacific. We use this compute to run evaluations on new model releases, and to work with the labs to troubleshoot safety issues as they arise.

WHY

Labs are torn between capability and safety work, and more often than not, they choose capability. Post-training is yielding such incredible returns, it's very difficult to commit to robust safety work as each RL run contributes directly to the success of these labs.

Compute incentivizes research direction, and everything is downstream from compute allocation. The compute can be spent only on safety. Inside that, no one tells a researcher which risks matter or which methods to use. We restrict the cluster, not the research.

HOW A CAMPAIGN RUNS

Before release, evaluations run against an endpoint the lab hosts; after release, on the open weights. When something fails, the lab hears first and decides what happens next.

We do not publish evaluation results; we provide the foundation for people to run them. Should an issue arise during an evaluation, we make resources available to the lab to close the problem and solve it for good.

TRUST, FROM BOTH SIDES

FOR THE LABS

No weights sit on our servers, and the lab sees every query. No government money; private philanthropy is the only funding.

FOR THE FUNDERS

Gradient updates are blocked by default; every exception is approved on its own and logged on the ledger. No owned hardware, so nothing is stranded if the project stops.

AT A GLANCE

WHAT	A safety-only compute cluster for open-weight labs, third-party evaluators and academic partners.
WHERE	Singapore. The first H200 node is rented and live.
WEIGHTS	Open weights only. No weights sit on our servers.
TRAINING	Gradient updates blocked by default. Exceptions approved by our technical team, one job at a time, and logged.
RECORD	Every job on a public ledger: who ran it, which model, how much compute, what category.
RESULTS	We do not publish evaluation results. The lab hears first.
POLITICS	Political work is not permitted. Any safety testing is fine, except work that revolves around politics.
FUNDING	No government money. Private philanthropy is the only funding.
COST	Launch evaluations free to the lab; research projects at or below cost. About \$200k per launch, an estimate.
SPONSOR	An established US 501(c)(3) has offered fiscal sponsorship; the agreement is being drafted.
CONTACT	info@pacificcompute.org

THE RECORD · AS OF 2026-09-01

Full table, with a source on every row: pacificcompute.org/ledger

6 of 19

■ independent evaluation ■ model card only

reviewed releases have an independent safety evaluation. None came with a dangerous-capability evaluation from its own developer. 20 more are queued.

42

widely adopted open-weight releases in the last year (1,000 or more Hugging Face likes).

113

open-weight releases in all in the last year (250 or more Hugging Face likes).