Pacific Compute

For funders
San Francisco · Singapore

Sovereign AI Safety Compute.

Pacific Compute operates sovereign compute for AI safety. Our Singapore cluster helps safeguard the open frontier with launch evaluations, research and ongoing technical support.

AI safety organizations need control of their own compute.

Capability and safety

The same GPUs that make models more capable are also required for safety work. When safety and capability research compete for compute access, capability research often takes priority.

Dependence on labs

Sovereign AI safety compute means that safety organizations don't need to rely on labs for the compute needed to safeguard the frontier. Pacific works to secure affordable, accessible and private compute for a safer AI future. We also support and partner with labs, so that critical AI safety compute is available for work internally.

Rental prices

Most AI safety organizations access compute at ever-higher prices; CoreWeave reported an approximately 25% price increase in July alone. By aggregating the demand for AI safety compute, Pacific helps AI safety orgs find affordable, sovereign compute for research, evaluations and incident response.

The organizations safeguarding the frontier need compute capacity to pursue their mission.

AI safety organizations need their own sovereign compute resources. We work with third-party evaluators, safety hubs, academic institutions and the labs themselves to make sure that dedicated AI safety compute is always available and affordable.

OUR FIRST PROJECT

Singapore, for the open model frontier.

We provide open model labs, evaluators and academics with dedicated compute for safety research and evaluation. We believe that the open model frontier presents both the greatest need and the greatest risk, so we're beginning there.

What it costs.

$200k

Per evaluation, all-in

A suite of safety evaluations costs $200,000 all-in, including researcher time, compute resources and the follow-up that turns a signal into useful actions.

> We scope campaigns to a three-week target and a four-week maximum from agreed kickoff, including findings delivery and the scoped retest.

WORK	PRICE TO THE USER	TERMS
Launch evaluation of a frontier open model	Free to the lab	$200k all-in per evaluation, funded by philanthropy
Compute for third-party evaluations of frontier open models	$1.55 per H200 GPU-hour	Set below published H200 rates
Academic, safety hub and independent safety research	At or below cost	Allocation depends on the proposed work and available capacity
Incident investigations	By arrangement	Supported with GPUs when capacity allows

Compute is billed per H200 GPU-hour. Access, scope and capacity are agreed before work begins, and pre-release launch evaluations take priority during release windows. 37% below the median on-demand H200 rate of $2.47 per GPU-hour on September 28, 2026 (GPU Finder, 31 single-GPU listings). The cheapest H200 in stock that day was $2.00 at PrimeIntellect; Runpod Secure Cloud listed $4.59.

6 / 19



6 of 19 coded records have qualifying independent evaluation evidence. Public reports, not Pacific Compute's campaign history. Another 21 families await review.

Evidence & methods · 2026-09-01 ↗

Risks and Cautions

Evaluation can also help a lab improve its models. We accept that trade-off where we expect the safety benefit to outweigh the capability benefit. Our case is that these models are likely to be released anyway, and that additional scrutiny before release can expose risks while there is still time to respond.

Funders receive no access to lab findings or operational records.

Safety scope & confidentiality ↗

Discuss programme funding.

info@pacificcompute.org

Email or WeChat contact accepted through the request form.