

Fund evaluations before models ship.

Fund evaluations that help labs find and understand dangerous capabilities before a model is released. Support the benchmarks, replication and follow-up that make a finding useful.

2026 · INITIAL FUNDRAISING ASK

\$4M

Establish the team, evaluation tools and compute capacity, and run evals on a suite of open frontier model releases in Q4.

Programme funding keeps researchers available, evaluation tools maintained and compute ready for short release windows.

2027 · PROGRAMME BUDGET


\$12M

Sustain evaluation delivery and supporting open model safety research on both sides of the Pacific.

Safety research competes with capability research.

The same GPUs can make a model more capable or help researchers understand its risks. When those jobs share a budget, capability research can take priority. We give safety research compute of its own.

Why the GPUs matter



Endpoint-based evaluation · The first H200 node is rented and live.

What the funding delivers.

\$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.

Inside the findings package

Question

Model version, access mode and the safety question investigated.

Findings

Observations tied to supporting evidence, including negative or inconclusive results.

Reproduction & retest

Reproduction records and scoped retest results, or why a retest was not possible.

Test coverage

Tests completed, trial counts, conditions and work left untested.

Limitations

Uncertainty, scoring checks and what the evidence cannot establish.

Resources

Compute and endpoint use, staff time, storage/transfer and reconciled costs.

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