

AI Safety - Australia & New Zealand: Year 2 Funding Proposal

Summary

This is an urgent grant proposal seeking funding to prevent AI Safety - Australia & New Zealand (AIS ANZ) from transitioning to a volunteer-only organization. Without funding secured in the next two months, we will be forced to discontinue our full-time operations despite our proven impact. We're requesting the Minimum Viable Scenario (\$121,070 USD) that would secure essential operations for 12 months.

Our mission is to grow and support a large, ambitious, and influential local community focused on preventing the most harmful impacts of AI. We do this by providing them with information, education, networks and advice. While we have specific metrics to track AIS ANZ's success, our long-term vision is to establish Australia as the Southern Hemisphere's hub for AI safety.

As Co-Founder and Director of AIS ANZ, my (Yanni Kyriacos) leadership has evolved through several phases. From 2022 to April 2024, Chris Leong and I ran the organization part-time (normally with less than 10 hours / week combined) while pursuing other activities. In May 2024, I transitioned to full-time leadership with funding from the Long-Term Future Fund (continuing until May 2025), allowing for significant growth over the past 10 months (Chris moved on to other interests).

The next proposed phase would build on our community growth and engagement success by maintaining essential operational capacity to serve our existing programs and support our growing community. The activities and results presented below focus primarily on achievements from the current full-time phase, demonstrating our ability to execute effectively and the value we provide to the AI safety ecosystem in Australia and New Zealand

Organisational Strategy, Activity Pillars & Tactics

Vision	A flourishing world where humans and AI coexist harmoniously
Mission	Grow and support a large, ambitious and influential local community focused on preventing the most harmful impacts of AI

Strategic Goals (the goals we need to meet to achieve our mission)	<ol style="list-style-type: none"> 1. Increase the quality and size of the AI Safety talent pool in Australia & New Zealand (ANZ) 2. Build community infrastructure & networks 3. Create AIS jobs, collaborations and organisations in ANZ that otherwise wouldn't have existed
--	--

Aligning our Key Activities to Community Engagement Stages

Community Engagement Stage	Audience Goal	What is the Key Activity?	Why are we doing it?
1. Inform	"Help me learn about locally and remotely available AI Safety opportunities"	Distribute content via our email database, LinkedIn company page and Facebook groups	Share local and remotely available job, educational, networking and funding/grant opportunities
			Make AI safety careers feel tangible and accessible to our ANZ audience through success stories
2. Connect	"Help me meet and build productive relationships with other AIS ANZ members"	Casual meetups in 6 major cities (Sydney, Melbourne, Canberra, Brisbane, Auckland and Wellington)	Acts as the soft-entry point for new community members that want an in-person introduction to the AI Safety community
		Monthly online networking events	Connect people in AIS across ANZ on a monthly basis
		Annual Careers Conference	Yearly tentpole event, designed to drive unique connections and concrete career plans

		Facebook chat groups	Our Facebook group has buzzing chat groups where people in ANZ go daily to discuss the latest news and career opportunities. A low effort way to bring people together digitally
		Sydney AI Safety Hub (SASH)	The main objective of SASH is to grow Australia's AI safety community by providing a dedicated space for AI safety-related networking, movement building activities, and productive work. It launched in March with bridge funding via Manifund
3. Grow	"Help remove the barriers stopping me from having a positive impact in AI Safety"	Career planning calls	Help community members develop concrete plans to transition into AI safety roles, matching their skills and interests with available opportunities
		Assistance with funding applications	Support promising projects and individuals in securing resources to pursue AI safety work that might otherwise go unfunded
		Triage community members / connect them with collaborators	Match people with complementary skills and interests to create new opportunities and increase the impact of existing initiatives

		Technical Alignment Research Accelerator (TARA)	Take motivated and proficient technical talent and upskilling them in safety relevant Machine Learning techniques
		Incubate projects	Provide guidance and support to help nascent AI safety initiatives get off the ground, turning ideas into actual work

How Our Key Activities Support Strategic Goals

Strategic Goal	How Our Key Activities Help Us Achieve It
Goal 1: Increase the quality and size of the AI Safety talent pool	"Inform" activities create awareness and interest in AI Safety
	Career planning calls develop talent
	TARA develops technical capabilities
	Careers conference enables learning from more experienced community members
Goal 2: Build community infrastructure & networks	Monthly city meetups create local hubs / schelling points
	SASH will provide physical space for informal networking and mentoring
	Online events connect dispersed community
	Facebook groups enable ongoing engagement
	Careers conference creates networking opportunities
Goal 3: Create AIS jobs, collaborations and organisations in ANZ that otherwise wouldn't have existed	Project incubation creates new initiatives
	Collaborator connections enable new projects
	Assistance traversing the funding ecosystem enables new projects and roles that otherwise wouldn't exist
	"Inform" activities let people know about funding and local / remote opportunities

Theory of Change

More talented and effective people working on AI safety increases the likelihood of achieving our vision (a flourishing world where humans and AI coexist harmoniously)

1. By identifying and supporting talented individuals in ANZ interested in AI safety, we:
 - Help them overcome barriers to entry
 - Provide paths to impact through technical, policy and movement building work
 - Connect them to opportunities locally and globally
 - Support their professional development and effectiveness
2. By creating a strong community:
 - Members support and learn from each other (e.g. through peer learning and informal mentorship)
 - People stay engaged in AI safety work
 - Resources and opportunities multiply through collaboration
 - Technical expertise informs policy positions
3. This leads to:
 - More people working effectively on technical, policy and movement building AI safety challenges
 - More new AI safety projects and organisations
 - A stronger collective voice in policy discussions, leading to increased influence on local and international AI governance

Organisational & Project KPIs

As Director, I will balance KPI tracking and measurement with program delivery. With limited capacity as the sole staff member, I will need to carefully prioritize between measurement and execution. The following KPIs are organized according to our three strategic goals, allowing us to directly measure our organization's progress:

Strategic Goal	Program	Measure	KPIs
1. Increase quality and size of AI Safety talent pool	TARA	Completion Rates	80%+
		Transition Rate into AI Safety Roles	>50% within 12 months
		Final Projects Submissions	70% of students submit a project
	Career Planning	Number of Calls	1/week
		Career Transition Rates	TBA: Measured at 6 & 12 months
	Marketing Channels	Newsletter Subscribers	966
		LinkedIn Followers	2000
		Facebook Group Members	789

2. Build community infrastructure & networks	Monthly Meetups	Attendees per City	≥10 per month
		Number of Active Cities	7
	Online Networking	Predicted Return Rate in Post-Event Surveys	≥80%
		Unique Attendees	≥80 over 12 months
	SASH	Unique Members	≥30 in year 1
		Unique Connections Created	≥100
		Desk Utilization	50-75%
		Member Satisfaction	>50 NPS
		Repeat Usage Rate	90%
	Careers Conference	Recommendation Score	90%
		Number of Connections Made	TBC
		Transition Rate into AI Safety Roles	>25% within 12 months, >50% within 24 months
3. Create AIS jobs, collaborations & organizations	Funding Support	Number of Successful Funding/Grant Applications	TBC
	Project Catalyzation	Number of Collaborative Projects Catalyzed	≥5
	Career Support	Number of Job Transitions Facilitated	TBC

Track Record

1. Phase 1 (2021-2023): Building foundations as Co-Founders, Chris Leong and I ran AIS ANZ part-time (<10 hours/week) while:
 - Testing 15+ event formats from casual meetups to unconferences
 - Facilitating 3+ AISF courses
 - Running 2 grassroots advocacy campaigns with Good Ancestors Policy
 - Launching local chapters in Brisbane and Melbourne
 - Growing our Facebook group to 294 members
 - Growing Newsletter subscribers to 118
 - Successfully beta-testing a farmstay retreat
 - Ran a successful city weekend retreat
2. Phase 2 (3 May 2024 - 27 March 2025): Yanni working full-time on AIS ANZ with LTFF funding, we achieved:
 - Growth in Outreach & Community Channels:
 - Newsletter: 118 → 454 subscribers
 - LinkedIn: 0 → 749 followers
 - Facebook: 478 → > 700 members
 - Launched region's first AI Safety co-working space: the Sydney AI Safety Hub ([SASH](#))
 - Event Success:
 - Inaugural AI Safety Careers Conference: Sold out 4.5 weeks early, 55+ attendees, 89% recommendation score
 - Possibly created 100+ unique connections if we conservatively estimate each person met ≥ 3 new people
 - Monthly meetups in 6 cities over 9 months: 350+ unique participants, creating >700 AI Safety connections at $\leq \$6$ per connection (\$3,600 total spend on marketing and F&B)
 - Launched Monthly Online Networking series with 76% satisfaction rate: 43 attendees across four sessions creating 110-215 unique connections, at \$0.93-\$1.82 per connection (\$200 total platform cost)
 - Launched Ask The Hiring Manager Q&A series (78% satisfaction rate) before discontinuing due to capacity constraints. The series connected job seekers directly with hiring managers, achieving:
 - 14 live attendees and 31 recording watches for LawAI (contributing to 65+ regional applications)
 - 5 live attendees and 58 recording watches for SPAR's hiring round, and
 - Helped progress one candidate to a work trial at Sage.
 - Series demonstrated clear value but required too much operational bandwidth to sustain with current staffing
 - Leadership & Coordination:

- Built and led team of 12 volunteers (Meetup Coordinators, Newsletter Editor, TARA Committee)
- Led ~75 advisory and prospective calls over 10 months: ~ 45 Career Advising, ~ 17 Community Engagement & Volunteering, ~ 13 Prospective Partnerships
- Organized AIS leaders dinner during EAGx
- Incubated multiple volunteer projects (Governance Newsletter, TARA, Grassroots Advocacy)
- Created, secured funding and launched [Technical Alignment Research Accelerator \(TARA\)](#), recruiting 21 participants and a teaching assistant for the upcoming course.
- Advocacy:
 - [Submitted formal recommendations to government AI safety inquiries](#)
 - Drove peak attendance (23 and 21 participants per online event) for Good Ancestors Project's public consultation events on AI regulation through targeted marketing campaigns.
 - As a test, I (Yanni Kyriacos) met with four Australian politicians after sending 21 emails (this demonstrates Australian politicians might be more accessible than in other countries).

Key Learning: Our growth metrics and satisfaction scores demonstrate our ability to execute effectively and scale operations. The transition from part-time to full-time leadership showed significant impact multiplication, suggesting further capacity (through additional staff) would enable even greater impact.

Minimum Viable Scenario (\$121,070)

This maintains current operations.

- Director Salary: \$96,000
- Marketing & Operations: \$25,070
 - Costs to market and run all the events and services I've described above, with 20% added on top to account for the unexpected. More details forecast [here](#)

Salary Justification for Director Position - AI Safety Australia & New Zealand

Proposed Salary: \$150,000 AUD (\$95,880 USD)

Personal circumstances / context:

- Sole income provider for family of three
- Supporting spouse with chronic health conditions
- Previous corporate salary was \$160,000 AUD (\$102,272 USD) (3 years ago)
- Market rate for equivalent corporate position today: \$150,000 - \$200,000 AUD (\$95,880 - \$127,840 USD)
- Below proposed salary (\$150,000 AUD / \$95,880 USD) means there's considerable pressure to go back into private sector

Market analysis:

Current roles in the AI Safety sector for comparison (ranging from junior to senior positions):

- Global Shield: Office Director (\$135,000-\$175,000 USD)
- Epoch AI: Data Product Lead (\$120,000-\$180,000 USD)
- Schmit Sciences: AI Institute Associate (\$120,000-\$160,000 USD)
- CEA Strategy Lead: University Groups (\$106,667–\$128,143 USD)
- FAR AI: Events Project Manager and Business Analyst (\$85,000-\$125,000 USD)
- Center for AI Safety: Operations Associate (\$70,000-\$90,000 USD)

The proposed salary is notably conservative compared to these market benchmarks, particularly considering the Director-level responsibilities of this role, e.g.:

- Organizational leadership and strategy
- Program development and oversight
- Community building across Australia/NZ
- Stakeholder management
- Fiscal responsibility

Summary:

The proposed salary of \$150,000 AUD (\$95,880 USD) represents a balanced approach between;

- Alignment with market rates (at lower end of corporate scale and below AI safety sector positions)
- Personal sustainability (minimum comfortable level for sole family income)
- Organizational responsibility (efficient use of resources while ensuring leadership stability)