


The Alignment Project
Clarification Questions – V2.0



Purpose of this document

This document addresses **clarification questions** to support applicants. Most of the key information regarding the Alignment Project is provided in the Application Pack. Applicants are strongly encouraged to review the guidance documents before referring to this document.

We will continue to update this clarifications document as new queries arise, if you have a question that is not covered here, please contact us at aisialignmentproject@dsit.gov.uk. The deadline for submitting clarification questions is 13th August 2025.

Version Control

Version	Published Date	Updates Made
1.0	30 th July 2025	Initial questions included
2.0	13 th August	Webinar questions included (final section in document)

Contents

Purpose of this document.....	1
Version Control.....	1
Programme Details	5
1. How many projects are you planning to fund?.....	5
2. Is the four-to-six-month project duration a strict limit?	5
Application Details.....	6
Due Diligence and Eligibility Criteria	5
3. Who is eligible to apply for the Alignment Project?.....	5
4. What qualifies as an academic institution?	5
5. What qualifies as a non-profit organisation?	5
6. Can international organisations apply?	5
7. Can consortia or joint applications apply?	5
8. Are university spin-offs eligible to apply?	5
9. Can for-profit companies apply?	6
10. What is the £315,000 grant funding threshold?	6
11. How do I determine if the £315,000 threshold applies to my organisation?	6
12. What does “economic activity” mean under subsidy control?.....	6
13. What happens if my organisation is over the £315,000 MFA threshold?	6
14. What due diligence checks are required?	6
15. What happens if an applicant fails the due diligence check?.....	6
16. I am an international applicant; thus the Charity Commission / Companies House Number is not applicable to me, what should I indicate?	6
17. Is there any restriction on applicants’ nationality?.....	6
18. Can unsuccessful applicants from the Systemic Safety Phase 1 programme or the AISI Challenge Fund use their previous application for the Challenge Fund?	7
19. How many applications can a host organisation submit?	7
20. How do we define "individuals affiliated with eligible organisations"?.....	7
21. I am affiliated with a university, what should I choose as the application type?.....	7
22. Are PhD students considered "affiliated to academic institutions" in the context of the Alignment Project? Can they be lead applicants?	8
Expression of Interest.....	7
23. What are the stages of the application process?.....	7
24. What is the purpose of Expression of Interest (EOI)?.....	7
25. Can I submit multiple EOIs?.....	7
26. Will early submission increase my chances of successful application?	7
27. Can I exceed the word-count limit in my EOI?.....	8

28.	I am an international applicant, what should I include as a legal registration number?	8
29.	What type of Research Publications should I include? Do I need to embed links to all of them?	8
30.	Will I be required to provide a detailed budget in the EOI form?	8
31.	How do we submit the EOI? Is there a form for it?	8
Proposal Iteration and Full Application Form.....		9
32.	How will I be assigned with an AISI Research Sponsor?	9
33.	What are applicants required to develop during proposal iteration?.....	9
34.	How long will the proposal iteration take place?	9
35.	What type of support can we expect from the Research Sponsors?	9
36.	How often should I meet with the Research Sponsor?	9
37.	I have a concern about my assigned Research Sponsor. What should I do?	9
Assessment Details		9
38.	How will my application be reviewed?	9
39.	What are the selection criteria? How will my application be assessed against these areas?	10
40.	How long does the review process take?	10
41.	Will I receive individual feedback?	10
42.	Can I reapply in future rounds if my application is unsuccessful?.....	10
43.	Will my Research Sponsor review my full application?	10
44.	Should I check with my Research Sponsor on the outcome of my application?	10
45.	How will the selection board balance projects with immediate impact vs. those with longer-term transformative effects?	10
46.	Will requesting a higher budget affect chance of being selected for funding?	10
Research Priorities.....		11
47.	Can multi-disciplinary projects address multiple priorities outlined in the RFT?	11
48.	Do applications need to be technical, or are international relations, governance, and policy studies accepted?	11
49.	Can an application focus on multiple alignment problems, like sycophancy and deception?.....	11
50.	Do you prefer theoretical or empirical studies?	11
51.	Is software/hardware security in AI within the scope of the call?	11
52.	Are AI data privacy issues relevant to the call?.....	11
Funding Agreement and Project Initiation.....		12
53.	What happens if the project takes longer than the agreed timeframes?	12
54.	What happens if the project cannot be completed?	12
55.	Who retains ownership of the intellectual property (IP) after project delivery?.....	12
56.	If I change university during the project delivery period, can the funds be transferred to the new university?.....	12
57.	Are the Funding Agreement T&C's negotiable?	12
Performance Monitoring and Reporting		13

58.	How often will the applicant and the Research Sponsor engage during project duration?	13
59.	Will the AISI Research Sponsor conduct 'lab visits'?	13
Final Evaluation Report and Submission		14
60.	What do I need to provide to evidence that the project has been completed?	14
61.	What types of deliverables are acceptable?	14
Payment.....		15
62.	What is the funding range for projects?.....	15
63.	When will the grant be paid out?	15
64.	Can I request milestone-based payments for long projects?	15
65.	Can the payment be provided in advance of activity?	15
66.	What types of costs are eligible?	15
67.	What costs are NOT eligible for funding?	16
68.	Is co-funding allowed?	16
69.	What happens if the project requires additional funding?	16
70.	Will the funding cover all estimated project costs?	16
71.	Can funds be used for salary costs, such as hiring a researcher on a fixed-term basis?	16
72.	Which organisation is ultimately funding the Alignment Project?	16
73.	Do collaborators need to be affiliated with the host organisation and how will the payment to the collaborators be managed / distributed?.....	16
74.	For Stage 1 – EOI Form – Budget Narrative, how accurate should we estimate our required funding? Do we get an opportunity to update it in Stage 2?	17

Programme Details

How many projects are you planning to fund?

The Alignment Project has a total of £15 million funds available for distribution to awarded research proposals. The exact number of projects to be awarded depends on the scope and required funding of each application (£50,000 to £1 million per project). Please note the coalition may consider funding projects that exceed £1 million.

Is the four-to-six-month project duration a strict limit?

No. Projects funded directly by AISI are recommended to last for 4 to 6 months. Projects funded by The Alignment Project in general are recommended to last for 6 to 24 months.

Due Diligence and Eligibility Criteria

Who is eligible to apply for the Alignment Project?

Applicants must be affiliated with academic institutions, research bodies, relevant non-profit organisations, or relevant for-profit entities as defined in the Eligibility Statement (see Application Pack, Eligibility Statement).

What qualifies as an academic institution?

An academic institution is an organisation that is eligible to receive funding from its country's primary research funding agency. Examples include:

- UK: Research Council-funded institutions;
- USA: Institutions eligible for NSF/NIH funding;
- EU: Institutions eligible for ERC funding.

What qualifies as a non-profit organisation?

A non-profit is a legally registered entity that operates for collective, public, or social benefit. It must be registered as such in its country of operation. Examples include:

- UK: Charities registered with the Charities Commission;
- USA: 501(c)(3) organisations;
- Other jurisdictions: Equivalent legal entities under local laws.

Can international organisations apply?

Yes. International applicants can apply but must confirm the Eligibility Statement and pass the Alignment Project's due diligence checks (see Application Pack, Eligibility Statement).

Can consortia or joint applications apply?

Yes. Consortia are allowed, but the host organisation to which the lead applicant is affiliated with must be eligible (see Application Pack).

Are university spin-offs eligible to apply?

Yes.

Can for-profit companies apply?

Yes. The Alignment Project is made up of a coalition of funders. Certain funders will accept applications from for-profit companies. These applicants may be subject to additional due diligence checks.

What is the £315,000 grant funding threshold?

This is relevant to organisations which are classified as enterprises under §7 of the Subsidy Control Act 2022. Applicants must not have received more than £315,000 in grant funding under the UK Subsidy Control Act 2022 within the last three fiscal years if they are considered an enterprise.

How do I determine if the £315,000 threshold applies to my organisation?

The MFA (Minimal Financial Assistance) threshold applies only if the affiliated organisation is classified as an enterprise under §7 of the Subsidy Control Act 2022.

- If your organisation is an academic institution or a non-profit, MFA conditions generally do not apply.
- If your organisation is engaged in economic activity, offering goods or services on a market, MFA may apply.

What does “economic activity” mean under subsidy control?

If your organisation sells goods or services in a competitive market, it may be considered an enterprise under UK subsidy control.

- Example: A research body selling AI models commercially may be classified as an enterprise.
- Non-example: A university conducting fundamental research without commercial sales is not an enterprise.

What happens if my organisation is over the £315,000 MFA threshold?

- If you are considered an enterprise, you may be ineligible if you have received more than £315,000 in MFA grants in the past three fiscal years.
- If you are not an enterprise, the MFA cap does not apply.

What due diligence checks are required?

The Alignment Project team will conduct mandatory due diligence checks on all applicants. This includes Spotlight checks, legal registration verification, and financial viability assessments. Applicants may be subject to further due diligence checks.

What happens if an applicant fails the due diligence check?

The application will not proceed if due diligence checks indicate financial instability, fraud risk, or failure to meet eligibility criteria.

I am an international applicant; thus the Charity Commission / Companies House Number is not applicable to me, what should I indicate?

Charity Commission / Companies House number is a unique registration number issued by regulatory authorities in the UK. You should ensure that you include a legal registration number in the country where the organisation is based (see Application Pack, Eligibility Statement).

Is there any restriction on applicants' nationality?

No, the fund is open to international applicants.

Can unsuccessful applicants from the Systemic Safety Phase 1 programme or the AISI Challenge Fund use their previous application for The Alignment Project?

Yes, unsuccessful applicants from the Systemic Safety Phase 1 and the AISI Challenge Fund are encouraged to seek opportunities within The Alignment Project to pursue their research projects.

How many applications can a host organisation submit?

There is no limitation on the number of applications from a single host organisation, so long as these applications are distinct.

How do we define "individuals affiliated with eligible organisations"?

Affiliation means the individual has an employment contract with the host organisation over the duration of the funding.

I am affiliated with a university, what should I choose as the application type?

This will depend on what the university is registered as. In this case it will either be limited or non-limited company.

Expression of Interest

What are the stages of the application process?

The application will include a two-step process:

Stage 1: Eligibility Statement and Expression of Interest (EOI)* Submission: This stage allows us to assess preliminary eligibility and determine whether applicants' EOIs are suitable to proceed to the next stage of development and assessment based on pre-defined criteria (see Application Pack, Assessment Criteria).

Stage 2: Shortlisting and Invitation to Submit a Full Application: If successful, shortlisted applicants will be notified and they will have the opportunity to work closely with a Research Sponsor to refine their proposals and help strengthen their submissions before final assessment.

What is the purpose of Expression of Interest (EOI)?

The EOI allows The Alignment Project team to evaluate the application against the Research Priorities. It is an essential step in shortlisting applicants.

Can I submit multiple EOIs?

Yes. Multiple EOIs are allowed if each proposal is unique and distinct. However, if multiple proposals are successful, there should be a justification of appropriate capacity and resources to run concurrent projects.

Will early submission increase my chances of successful application?

Early submission does not guarantee progression to the next stage as all applications will go through quality assessment. However, we will not accept applications after the deadline (10th September 2025), thus, you are encouraged to apply as soon as you are ready.

Can I exceed the word-count limit in my EOI?

No, a word-count limit is embedded in the application portal. Thus, you cannot submit your application if this is exceeded.

I am an international applicant, what should I include as a legal registration number?

For applicants based outside the UK, please provide the equivalent registration number issued by the relevant regulatory authority in your country. If known, a DUNS number is encouraged to be provided for a more streamlined due diligence process.

What type of Research Publications should I include? Do I need to embed links to all of them?

Preferably, you should include publications relevant to the area of your research. It is also recommended to include links to online publications for further checking during assessment.

Will I be required to provide a detailed budget in the EOI form?

In the Budget Narrative field, you will only be required to provide a short justification of your requested funding. This may include a breakdown of the forecasted expenses directly related to research projects and estimated costs. If shortlisted, you will then be required to provide a detailed budget.

How do we submit the EOI? Is there a form for it?

The Expression of Interest form is part of the Stage 1 Application form which can be accessed via the DSIT Funding Portal. The required fields and information for the EOI can be reviewed in the Application Pack before submission of application.

Proposal Iteration and Full Application Form

How will I be assigned with an AISI Research Sponsor?

The Alignment Project team will assign a Research Sponsor based on the remit of your proposal. If shortlisted, applicants will receive a notification from the Alignment Project team and should expect their assigned Research Sponsor to reach out to them to arrange collaborations.

What are applicants required to develop during proposal iteration?

Shortlisted applicants will be invited to complete a full application form that consists of an updated EOI, project plan, and required resources and costs. The Alignment Project team will provide shortlisted applicants with further guidance on this.

How long will the proposal iteration take place?

The duration of proposal iteration will depend on the nature and scope of the project, but it typically takes 3 weeks. The deadline for submitting the Stage 2 application is 15th October.

What type of support can we expect from the Research Sponsors?

At proposal iteration stage, Research Sponsors will work closely with applicants to refine their proposal before submission of the full application. They will provide their expert advice on improving the proposal which may include clarifying research scope and objectives, providing feedback on research methodology, potential risks, and feasibility, aligning the budget with research needs, conducting a final check to ensure the application meets The Alignment Project's standards before the applicant submits the full proposal.

How often should I meet with the Research Sponsor?

The frequency and nature of the collaboration will be agreed between the Research Sponsor and the applicants.

I have a concern about my assigned Research Sponsor. What should I do?

For any concerns, please reach out to us through our mailbox at aisialignmentproject@dsit.gov.uk

Assessment Details

How will my application be reviewed?

At Stage 1, we will conduct preliminary reviews to determine applicant eligibility. Eligible applications will then be forwarded to the Research Sponsors for EOI quality review based on pre-defined selection criteria (see Application Pack, Assessment Criteria). Conflicts of interest will be considered from this stage onwards to ensure fair assessment.

At Stage 2, a formal assessment panel, independent of the Research Sponsors, will review the full applications for final recommendations. Funding decisions will then be determined by the Selection Board Panel.

What are the selection criteria? How will my application be assessed against these areas?

Quality of applications will be assessed based on five pre-defined criteria: relevance, feasibility, innovation, actionability, and capacity. Applications must pass all criteria to be successful (see Application Pack, Assessment Criteria section). Certain funders within the Alignment Project may apply additional scoring methodologies to your application.

How long does the review process take?

For Stage 1, review outcomes will be provided by 18th September. For Stage 2, assessment results are to be expected within 3-5 weeks of submission. (see Application Pack, Timeline).

Will I receive individual feedback?

Unsuccessful applications will not receive individual feedback. If shortlisted, however, applicants will have a chance to receive feedback and work closely with an assigned Research Sponsor to iterate and improve a full application.

Can I reapply in future rounds if my application is unsuccessful?

Yes. Applicants can revise their proposal and resubmit in future funding cycles. However, we strongly advise against the submission of duplicate (exact copy of) proposals as these will be rejected during initial checks.

Will my Research Sponsor review my full application?

The assigned Research Sponsor will only review the full application before submission. They will not form part of the final assessment panel to avoid advocacy bias and maintain a transparent and fair assessment process.

Should I check with my Research Sponsor on the outcome of my application?

The assigned Research Sponsor is not permitted to directly inform applicants of the outcomes. Results of assessments will be officially communicated within 3-5 weeks of submission.

How will the selection board balance projects with immediate impact vs. those with longer-term transformative effects?

Both types of outcomes will be considered. There is no inclination on one over the other.

Will requesting a higher budget affect the chance of being selected for funding?

The Alignment Project provides funding between £50,000 to £1 million per research project, but we will consider projects that exceed £1 million. During application, applicants will be asked to indicate the required funding for their research project, with an opportunity to justify this in the Stage 1's EOI (Budget Narrative Section), and to further detail out in the Stage 2's Required Costs and Resources. Hence, the amount of requested funding will not affect the chance of approval as long as it is fully justified and detailed out during application stage.

Research Priorities

Can multi-disciplinary projects address multiple priorities outlined in the research priorities?

Yes. Cross-cutting proposals that address multiple priorities are encouraged.

Do applications need to be technical, or are international relations, governance, and policy studies accepted?

Non-technical applications are welcome, as long as they align with the Alignment Project priorities and pass the selection criteria.

Can an application focus on multiple alignment problems, like sycophancy and deception?

Yes, as long as the scope remains clear and focused.

Do you prefer theoretical or empirical studies?

No preference; both are viable.

Is software/hardware security in AI within the scope of the call?

If linked to AI safety and security risks, then yes.

Are AI data privacy issues relevant to the call?

If linked to AI safety and security risks, then yes.

Funding Agreement and Project Initiation

What happens if the project takes longer than the agreed timeframes?

All activities are required to be completed within the timeframes specified in the Funding Agreement. The recipient should notify The Alignment Project if there is a risk of delays in activities.

What happens if the project cannot be completed?

If the project cannot be completed, the recipient must notify The Alignment Project team as soon as possible. A formal review will determine next steps, which may include the funder(s) terminating funding or requesting a partial refund unless valid reasons are provided.

Who retains ownership of the intellectual property (IP) after project delivery?

The recipient owns any intellectual property developed using the fund, but the Alignment Project is granted a non-exclusive, irrecoverable and royalty-free, sub-licensable, worldwide license to use the intellectual property.

If I change university during the project delivery period, can the funds be transferred to the new university?

The Funding Agreement (GFA) is between the funder(s) and the affiliated university. If the applicant moves to another university during the project delivery period, the GFA will be terminated.

Are the Funding Agreement T&C's negotiable?

The Funding Agreement (GFA) is typically not negotiable.

Performance Monitoring and Reporting

How often will the applicant and the Research Sponsor engage during project duration?

The cadence of post-award performance monitoring will be confirmed as part of the proposal iteration stage, based on the level of input and support required from the Research Sponsor.

Will the AISI Research Sponsor conduct 'lab visits'?

There are no mandatory lab visits, but the Research Sponsor may arrange in-person or virtual meetings to review progress and discuss findings.

Final Evaluation Report and Submission

What do I need to provide as evidence that the project has been completed?

A final report detailing project outcomes, key findings, and deliverables must be submitted, along with any required financial and impact assessments. A final evaluation report template will be provided for consistency.

What types of deliverables are acceptable?

Acceptable deliverables include research papers, code, simulations, experiments, services, prototypes, or a combination.

Payment

What is the funding range for projects?

Grants range from £50,000 to £1 million per project but we may consider applications that exceed £1 million. Applicants are required to justify their requested budget through a Budget Narrative in the EOI form. If shortlisted, you will then need to complete a budget outlining the required costs and resources for the project.

When will the grant be paid out?

Payments will be made monthly in arrears, subject to submission of milestone-based reporting. In exceptional circumstances, depending on need, we may consider one payment in advance of activity. This will be decided on a case-by-case basis and reflected in the Funding Agreement. Recipients are encouraged to engage with their institution's finance team to ensure alignment with these terms.

Can I request milestone-based payments for long projects?

Payments will be made monthly in arrears, subject to submission of milestone-based reporting.

Can the payment be made in advance of activity?

Typically, payment is provided monthly in arrears based on milestones. Depending on need, we may consider one payment in advance of activity. This will be decided on a case-by-case basis and reflected in the Funding Agreement.

What types of costs are eligible?

Eligible direct costs typically include:

- **Personnel costs** - The payroll costs of all staff, full or part-time, who directly work on the project, and whose time can be supported by a full audit trail.
- **Materials and minor equipment costs** - The costs of consumables, materials, and minor equipment directly used for project delivery.
- **Licensing and access costs** - The costs related to acquiring licenses for software, datasets, or computing resources needed for the project.
- **Travel and subsistence costs** - The costs for travel, accommodation, and meals necessary for project activities.

Eligible indirect overhead elements must be restricted to those that will **directly impact and support the project** but cannot be identified as a direct cost. This may include office rent and utilities, administrative support (HR, Finance, Legal), and shared IT infrastructure (email, servers).

You may refer to the above guidance when calculating the high-level estimate of your required funding as part of Stage 1 application (EOI form – Budget Narrative).

We have published a more detailed Cost Guidance which will be more relevant to shortlisted applicants as part of the proposal iteration (Stage 2). In this stage, applicants may update their estimated budget and detail out each element of cost through the budget.

What costs are NOT eligible for funding?

As a general rule, any costs that cannot be categorised as eligible cost (see previous Q&A) must be considered ineligible for the purpose of funding.

The examples of ineligible costs include but not limited to:

- Using funds to fund lobbying (via an external firm or in-house staff)
- Capital expenditure for long term use primarily beyond the duration of the project
- Training and development; as per the nature and scope of projects, teams are expected to have existing capabilities and skills to undertake project activities
- Any other costs of activities not related to the aims of the project or not sufficiently justified.

As per previous Q&A, you may refer to above guidance when calculating the high-level estimate of your required funding as part of Stage 1 application (EOI form – Budget Narrative).

We have published a more detailed Cost Guidance which will be more relevant to shortlisted applicants as part of the proposal iteration (Stage 2). In this stage, applicants may update their estimated budget and detail out each element of cost through the budget.

Is co-funding allowed?

Yes. Applicants may secure co-funding from other sources, but it must be declared.

What happens if the project requires additional funding?

Additional funding may be available however; this is not guaranteed. Requests must be submitted with a clear rationale, budget details, and an impact assessment for consideration.

Will the funding cover all estimated project costs?

The funding will cover 100% of eligible costs. Please see Q#s 5 and 6 to understand what we consider eligible and ineligible costs.

Can funds be used for salary costs, such as hiring a researcher on a fixed-term basis?

Yes, salary costs are eligible costs.

Which organisation is ultimately funding the Alignment Project?

The grant programme is funded by a coalition of funding partners. It is led by AISI UK which is funded by the Department for Science, Innovation and Technology (DSIT).

Do collaborators need to be affiliated with the host organisation and how will the payment to the collaborators be managed / distributed?

Collaborators don't necessarily need to be affiliated with the host organisations. They may also be individual consultants.

It's important to note, however, that the Funding Agreement (GFA) will be between the funder(s) and the host organisation only. In addition, funds will be directly paid to the host organisation, who will then manage further distribution to and/or contractual agreements with their collaborators. Host organisations should ultimately bear the risks and responsibilities around their chosen collaborators.

For Stage 1 – EOI Form – Budget Narrative, how accurate should we estimate our required funding? Do we get an opportunity to update it in Stage 2?

The Budget Narrative should provide us with a rough estimate of your required funding and a justification as to why your organisation will need financial assistance from the Alignment Project to undertake the research activities. Hence, we do not expect the required funding at this stage to be accurate, rather a high-level estimate only based on high-level guidance provided in Q5 & 6.

In Stage 2, shortlisted applicants will be given an opportunity to refine and update their cost estimate and detail each element of their costs through a Cost Guidance and a Budget section in the DSIT Funding Portal.

Webinar Q&A

The informational webinar was held on 11th August 2025 and the recording is available on The Alignment Project website. Below are answers to questions that had not already been addressed in V1.0 of the Clarification Document.

Is the funding an unrestricted gift or contract?

The Alignment Project funding operates as grant funding, which is typically distinct from both an unrestricted gift and a contract. Grants are issued subject to a Funding Agreement (or Grant Funding Agreement), which outlines conditions such as milestone-based reporting, project deliverables, monitoring, and payment schedules.

Does EoI require organisational approval or is it just for final application?

We encourage applicants to engage early with relevant Research Office / Finance Departments and other parties required to approve the project and receive the funds. However, you can submit an EOI in advance of gaining all required approvals.

Is it possible to start the project in March 2026, or does it have to start within 2 weeks of signing the funding agreement (so still in 2025)?

We prefer projects to start as soon as possible after the funding agreement is signed. However, if a later start date—such as March 2026—is essential (e.g., to align with a future financial year), applicants should state this in their application. Such requests will be reviewed and accommodated where possible. It is a preference, not an absolute requirement.

How should clarification questions be submitted post this webinar?

Clarification questions should be emailed to the Alignment Project mailbox. Individual responses will not be provided if they could give an applicant an unfair advantage. Instead, the team will remove any project-specific details and publish a generalised answer in the publicly available clarification documents. The deadline for clarification questions is 13 August.

For Stage 1, we are expected to provide high-level budget requirements, which can be refined in Stage 2. Still, would we have to run those through our Universities Research Offices to account for overheads etc?

At the Expression of Interest (EOI) stage, applicants should provide an estimated budget, including a budget narrative, with as much certainty and detail as possible. It is advisable to involve your university's research office early, particularly to account for overheads, though methods for calculating these may vary between institutions. The budget can be refined at Stage 2, so changes after the EOI stage are acceptable.

Given the short expected length of the projects, and the fact that they are supposed to start quite soon after decision, is it advisable to cost PDRAs on those projects (which typically need time to recruit) or is it more expected to cost PI time and other related costs?

Funding can cover PhD students, postdoctoral researchers (PDRAs), and related staffing costs, provided they are working on the funded project. If additional time is needed to recruit researchers, this should be factored into the proposed project timeline. We have no explicit preference between costing PI time or PDRAs, but the budget should align with the project's needs and delivery plan.

Can we request funding for a PhD scholarship even if this will extend beyond project duration?

Funding can support PhD students working on the project; however, it is not intended to cover a full four-year PhD. The funded PhD work should fall within the project's duration and contribute directly to its objectives.

Are projects investigating new threats admissible, or are you looking only on project focusing on "defenses"?

Projects addressing new threats are admissible, provided they are relevant to AI safety and can demonstrate potential impact or usefulness. This includes work in different settings (e.g., vision models as well as text models) if the contribution is innovative and addresses identified gaps or blind spots.

Can the start of funding be delayed? (for new faculty members who are still hiring)

A delayed project start is possible if necessary, including for new faculty still hiring. Any required lead time—such as starting in a future financial year—should be clearly stated in the application so it can be considered during review.

The application guidance says "as a recipient you will receive.. compute access and venture capital" (p3). Please can you explain how these two will work?

Compute: The programme has a substantial compute budget with AWS and other providers. Applicants may request dedicated hardware if needed for security or safety reasons, provided the request is proportionate to the project's scope and needs.

Venture capital: VC support will follow a standard venture capital model, with investment provided in exchange for an equity stake. Operational details—such as AWS credit allocation processes or the timing of VC offers—will be shared in due course.

Are guidelines provided for suggested salaries for budgets?

Yes. Cost guidance, including eligible and ineligible spend categories, is provided in a PDF on the programme website.

When would the grant start?

There is no fixed start date. Following the EOI deadline (10 September) and full application deadline (15 October), successful applicants will be notified between 5 and 19 November. Once the funding agreement is signed, projects are typically expected to begin within around two weeks.

Do you have a specific definition of AI? For example is computer vision fair game? Or are you more interested in text-based?

Yes. Projects involving non-text models, such as vision models, are eligible if they can demonstrate relevance and potential impact for AI safety.

You are asking for a project start date within 2 weeks of funding decision around Nov.19th. That does not allow time for recruitment of any researchers. Any delay in recruitments then eats into the 4-18 month project duration. How can we address this?

Recruitment lead time can be built into the proposed project timeline. Proposals will not be rejected solely due to recruitment needs if the overall agenda is strong.

What kind of research credentials should applicants have? Can career changers without (much) prior research experience apply?

Applications are assessed case-by-case. Relevant expertise and evidence of successfully delivering work in the field are important; without such evidence, approval is less likely.

How would it work if we need ethics approval for the project? Would the funding be delayed until ethics is approved?

Ethics approval timing will be handled on a case-by-case basis, and applicants should note any required ethics approvals in their application so the funders can provide guidance on whether funding can start conditionally or must wait until approval is granted.

How does the initiative plan to facilitate cross-border collaboration and data sharing.

Researchers may be based in any country, subject to due diligence and anti-money-laundering checks. Mechanisms for cross-border data sharing will be agreed at the project iteration stage.

How will the initiative ensure that research outputs systems are both practical for both regulators and developers and accessible to the public.

One assessment criterion is “actionability,” meaning outputs should be clear, tangible, and relevant to real-world needs. No specific operational measures to ensure accessibility were given.

Can 1-2 people teams apply?

Yes. Independent researchers with a good idea can apply. Researchers may be matched with a host organisation for grant administration purposes but this will be determined on a case by case basis.

Is there a preference for a larger team vs single independent researcher?

No preference. Selection is based on the quality of the research proposal and how well it meets the assessment criteria.

Are guidelines provided for suggested salaries for budgets?

The project provides a cost guidance PDF on the website covering eligible spend and ineligible costs. Since this is the first round of the Alignment Project, there is no historical salary data specific to past grants, but the guidance PDF should be followed for budgeting.

For overheads, you provide 3 options. For Option 3 where we have to provide calculation of overheads, can you advise what breakdown you are looking for?

A: Overhead calculation methods vary by institution. Information is provided in the Cost Guidance and further details will be addressed during Stage 2.

What are the reporting requirements?

Reporting and post-award monitoring will be proportionate to the duration and value of the project and may be tied to project milestones and payment points. This typically includes monthly calls with the research sponsor and agreed financial and non-financial metrics. Requirements will vary by project.

Could you please give an estimate of the expected success rates of applications in phase 1, ie how many applications you expect to pass to phase 2?

Since this is the first round of the Alignment Project, there is no historical data to draw on. Stage 1 EOIs will be shortlisted for Stage 2 based on the likelihood of funding success, but the final decision rests with the review committee after Stage 2. No specific percentages or quotas are provided.

Can you advise if desirable research solutions/outputs should be domain specific or can be focused on priority sectors like national security and defence, health, CNI, etc?

Both are acceptable. Novel technical methods are welcome, especially with concrete use cases, and cross-sector solutions are encouraged.

Of the five assessment criteria, which tend to be the most decisive for shortlisting EOIs?

EOIs are evaluated holistically, and proposals that clearly demonstrate relevance, feasibility, and practical impact—while being technically sound and novel—are most likely to be shortlisted.

Are research outputs publishable and/or commercialisable or would they need to be pre-approved by the funders for such?

This will likely be determined on a case-by-case basis. Different funding sources (government grants vs. venture capital) may have different expectations around publication and commercialisation. Applicants should clarify any requirements with the funder before disseminating results.

The two questions posed...

- 1. How can we prevent AI systems from carrying out actions that pose risks to our collective security, even when they may attempt to carry out such actions?**
 - 2. How can we design AI systems which do not attempt to carry out such actions in the first place?**
- ... are quite broad compared to the listed priority research areas. Are grants likely to be granted only in the priority areas listed?**

The fund seeks alignment with stated research priorities but is open to innovative topics relevant to AI safety.

Are there any restrictions on publication (e.g. do publications need any kind of approval)? I partly ask this given that it sounds like there are multiple funders in the background.

Any publication restrictions will be determined on a case-by-case basis, depending on the funding source and project type. Standard grant reporting rules will generally apply, while VC-backed projects may have additional considerations around commercial or sensitive information. Applicants should check with the Research Sponsors if they anticipate needing to publish results.

Will there be a recording available after the meeting?

Yes. The webinar is recorded and will be available to rewatch.

As the for profit and not-for-profit requests are being treated separately do we need to specify this somewhere in the application? I've not seen anywhere to specify yet.

Yes. For-profits may receive venture capital investment and not-for-profits receive grant funding. As part of the initial section on the EOI form you will be asked to select the entity type e.g. Limited Company.

Is the form of funding different for nonprofit vs for-profit receivers?

Yes. The initiative provides funding through different mechanisms depending on the applicant type:

Nonprofit or academic entities typically receive grant funding, which does not require giving up equity.

For-profit entities may receive venture capital-style funding, where the fund invests in exchange for an equity stake in the company. Applicants should ensure that their chosen funding route aligns with their organisational type, as the source and conditions of funding will differ.

What would you recommend as the best path for unaffiliated researchers willing to leave for-profit jobs to work on research?

Independent researchers are welcome to apply and will usually be paired with a sponsoring organisation to receive and manage the grant.