Request for Proposal: Scaling Trust - Multi-Agent Security Arena

Mar 25, 2026

V1

SUMMARY

We are looking for a partner to co-design, build and maintain a successful Arena for the Scaling Trust programme. The Arena is central to the programme as it will be where the latest tools and research we fund will be tested in a live, adversarial environment. Anyone in the world will be able to participate in its challenges, and win a portion (or all!) of the multi-million pound prize pool.

By carefully designing this environment and the challenges deployed within it, we plan to both measure the state of the art in multi-agent security capabilities and failure modes, and to enable emergent secure agentic interactions.

The project will start immediately, initially runs for one year, and the contract can be renewed for up to three years in total (the programme duration). We have earmarked an initial £10m for this work, which includes both digital and cyber-physical components. You will be applying for this contract. Cost will be a factor on who we select, alongside culture & value-alignment, speed of execution, and a demonstrated ability to do the work.

We intend to select an Arena partner as soon as possible within the initial timeline set out below. However, if we do not identify a suitable partner, we will continue to accept and assess proposals on a rolling basis using the two stage application process detailed in this RFP until we find the right partner.

2-page short proposal submission deadline	**14 April 2026 (14:00 BST)**
Discussions with shortlisted applicants will take place during this period. Applicants will be notified of the outcome of this stage no later than 30th April 2026. We may ask you to begin work on the detailed proposal as soon as possible before 30th April.	**15 April 2026 - 30 April 2026**
Invitation for detailed proposal (8-page proposals)	**No later than 30 April 2026**
Applicants invited to submit a detailed proposal will receive specific timelines in the invitation and will be given two weeks to respond, with the option to submit earlier.	

Submit your short 2 page proposal **HERE**.

SECTION 1: REQUIREMENTS

1. Introduction

Scaling Trust is a £49.8m Research & Development programme to build tools for agents to securely interact with one another in untrusted environments. We believe getting this right will be transformative for the world; enabling truly secure, programmatic and scalable agent-to-agent interactions is a pre-requisite for a world of many agents executing tasks on our behalf. It may even be a pre-requisite to get to Artificial General Intelligence itself.

We plan to get there by splitting efforts into three tracks:

- *Fundamental Research* will both strengthen the theoretical backing of our efforts and produce a reservoir of new knowledge we are able to build tools around.
- *Tooling* will produce both implementations of frontier research, and various open-source tools usable in the context of agentic security (datasets, languages, agent templates).
- Finally, *the Arena* is where everything will come together, a live adversarial environment where the tools and research are stress-tested.

This is intended as a pipeline, where ideas flow freely between the three tracks; sometimes going from Fundamental Research all the way to being tested in the Arena, sometimes starting off from an empirical observation in the Arena spawning a new fertile research question.

The Arena is critical to the ultimate success of the programme. Without stress-testing the tools built and the research carried out under adversarial pressure, we risk them being insecure and therefore useless for meaningful applications. Security by Obscurity creates a dangerously unstable system. Security by Scrutiny is the only way to ensure these tools' robustness, widespread adoption, and humanity's ultimate reliance on them.

You can read more about the programme in its thesis and the call for proposals for the two tracks Fundamental Research (track 3) and Tooling (track 2) here.

The focus of this request for proposals is to identify and contract an Arena partner that will design, develop, operate and maintain the Arena. If you're not planning to apply as an Arena partner but are interested in contributing to the broader Arena effort, such as by proposing challenges or benchmarks, or designing specific components we still encourage you to get in touch by submitting this form. While this call is focused on selecting Arena partners, we expect additional opportunities for contribution to open up as the Arena develops, and we would welcome future engagement.

2. Requirement ("Services")

We are seeking a dedicated partner to act as an extension of the Scaling Trust programme team, collaborating closely to co-design, build and maintain the Arena throughout the programme's duration. We are looking for a partner capable of operating at pace, to deliver a robust, modular testing ground that can move from prototype to public operation over the course of the programme.

Beyond providing core engineering capacity, the ideal partner will help shape the Arena's fundamental mechanisms including challenge design, scoring systems and community interfaces, whilst ensuring the platform remains secure, observable and accessible to a global community of participants.

The Arena is intended to be a central interface across the wider Scaling Trust programme. We expect the Arena partner to engage with Creators from other tracks in order to help integrate their work into the Arena.

2.1. The Arena

Goals

The design goals of the implementation of the Arena are the following:

+ **Measure multi-agentic security**: Evaluate agents on both task performance and security.
+ **Track progress over time**: Track and visualise progress across all agents through time to show the impact of the program.
+ **Low barrier participation**: Low barrier to entry so that any team (worldwide and with a wide range of skill levels) can participate.
+ **Thin, modular and composable**: Arena should be made of thin, modular components and not a monolith (for example, it should be easy for external individuals to propose new challenges that can be added to the Arena).
+ **Arena outcomes are observable and exportable**: All results should be accessible by any one under permissive license, including transcripts and challenge outcomes at some point.

For the sake of clarity, we list here what this project explicitly does not aim to be:

- Be an agent hosting platform (while we might offer external services or credit for third party organizations to do so)
- Enforce a specific agent architecture or framework for participation
- Mandate a single communication protocol
- Guarantee fairness across hardware or compute budget (although some challenges might be designed to be rewarded for using a cheaper infrastructure)

Proposed Activities

The Arena will host both benchmarks and challenges (For the avoidance of doubt, ARIA will retain responsibility for challenge selection and programme-level governance, while the Arena partner will own implementation, operations, and technical integration of the agreed challenge set. ARIA may work with third-party challenge architects/designers who lead the creative and scientific design of specific adversarial scenarios):

- **Benchmarks**
  - A benchmark is a self-contained test that scores an agent (or a sub-component of an agent) on a specific capability.
  - Executing a benchmark does not require live interaction with other agents.
  - Anyone should be able to download the benchmarks and run them locally (or use the arena API to report their scores).
- **Challenges**
  - A challenge is a session between multiple agents, where every agent is given a task and a set of security policies to respect. After agents interact, they report their completed tasks and they are scored.
  - Agents are scored based on their ability to complete tasks and to respect security policies (e.g. no data is being leaked).
  - Participating in a challenge requires live interaction with other agents.

Proposed Mechanics & Rules

- **Participants**
  - Standard participation – agents participate in challenges to complete tasks. They are scored based on their ability to complete a task while respecting their security policies.

- Red team participation – agents participate in challenges with the sole goal of making the other agent fail to respect their security policy (and not their ability to complete the task in a challenge).
- Our participation – We plan on deploying a baseline red team and baseline agent.

- **Leaderboard**
  - Ongoing: Anyone can participate on benchmark and challenges in an ongoing way.
  - Quarterly: a snapshot of the arena is being taken (with key metrics, best agents and best red teams).
- **Prizes**
  - Quarterly rewards per challenge until we get to a good metric: £250k set aside per quarter to be split amongst active challenges (tentative current plan).
  - £1m grand prize for every 'season' of the Arena.
- **Challenge Iteration**
  - Every quarter we will have the option to add or retire challenges. We may put constraints around compute or the types of models used.

Proposed Scoring

How we score is critical to the programme's success. We expect the precise scoring approach to be refined with the selected partner, contestants and the broader community's input throughout the Arena's lifetime.

Here is how we're thinking about scoring today: we want to surface the most useful, secure agents. At both extremes – agents can be very secure (ie. security policy is always respected), but essentially useless. Eg. An agent that never interacts with another agent, and therefore never leaks any information during interactions. On the other hand of the spectrum, an agent can be very 'useful' in that it engages in a range of tasks successfully but its security is poor (i.e. the state of most agentic systems today).

While security can be ultimately collapsed into usefulness (an insecure agent is fundamentally less useful over a long enough period of time), we believe there is value in charting both an agent's 'security' and its 'utility' in the same graph.

Agents in the Arena could be scored against (Utility;Security), i.e. their ability to complete the task (Utility) *vs* their ability to respect security policies (Security). We also plan on using secondary metrics such as the cost of completing the task during the interaction (Cost Efficiency) and an agent's ability to perform across different challenges (Generalisation).

Arena Security

The Arena will attract adversarial behavior by design. We expect the selected partner to propose/co-design mitigations for threats directed at the Arena itself, including but not limited to:

- Sandbox breakout by malicious agents,
- Sybil attacks on the leaderboard
- Exfiltration of other participants' strategies or code, and
- Denial-of-service against challenge infrastructure.

Your ultimate proposal should outline your approach to platform-level threat modelling and ongoing security operations.

Arena Safety

We want to ensure safety and ethical concerns are instilled to the core of the future we're building. This will be included in the shape of the challenges we run, the tools and research that we fund, and other activities we engage in. We'd love your suggestions and feedback here.

For the Arena specifically, we plan on setting up a Safety Oversight Committee who are able to give independent opinions on processes we follow in the Arena (e.g. responsible disclosures, rule changes in challenges, challenge design that could leave to emergent nefarious capabilities).

Demo and other examples

A demo Arena by Nicola Greco with "crypto-emergent challenges" is hosted here.

Related efforts:

- Vending-Bench Arena explores multi-agent interaction in a shared economic environment, including coordination, collusion, and bargaining.
- Game Arena provides a useful reference for repeated strategic interaction in game settings.
- Concordia Contest is relevant for cooperation, self-play, and cross-play with familiar and unfamiliar agent populations.
- AgentDojo is an adjacent reference point on agent security evaluation.
- Gray Swan Arena a red-teaming competition platform where security researchers discover vulnerabilities in frontier AI models.
- Diplomacy and Cicero is an interesting example of an agent engaging in negotiation.

+ Claude entering security competitions is an example of piggybacking on top of existing security evaluations with agents as a way to measure their performance (which may be a more scalable and low-effort way of operating!).

2.2 Proposed Components

The following list identifies the essential components for design development, operation and maintaining the Arena. This list is non-exhaustive and we encourage applicants to propose any additional components or innovative solutions necessary to bring the Arena to a production-ready standard.

The core of the Arena is the underlying evaluation environment: challenge execution, interfaces, scoring, observability, and operations. The website and leaderboard are important public surfaces built on top of that core.

For the avoidance of doubt, ARIA will retain responsibility for challenge selection and programme-level governance, while the Arena partner will own implementation, operations, and technical integration of the agreed challenge set. ARIA may work with third-party challenge architects/designers who lead the creative and scientific design of specific adversarial scenarios.

Core infrastructure and execution:

- **Arena Engine** — the engine for hosting and executing challenges and recording transcripts
- **Interfaces for Agents** — the programmatic tools for agents to interact with the arena (e.g. APIs and MCP endpoints and their libraries)
- **Operational, Security and administration tools** — challenge versioning, abuse prevention, and monitoring.
- **Challenges Implementation** — The Arena partner will be responsible for building the standardised "plug-in" architecture and modular abstractions that allow challenge designers to integrate their work into the Arena. The partner will be responsible for the technical implementation, hosting, instrumentation and operation of agreed challenges.

User experience and ecosystem:

- **Website** — the main entry point for Arena participants, the Arena partner will lead the technical development and maintenance of the website while collaborating closely with ARIA's internal comms team and an appointed external comms partner on overall design, branding, and strategic tone of voice. Additionally, the Arena partner will work with ARIA to ensure challenge implementation legal terms and conditions of participation are communicated clearly and effectively to the community.

- **Leaderboard** – the scoring and ranking system and the user interface to display top participants

2.3 Path to Cyber-Physical

As laid out in the Scaling Trust thesis, we care about enabling secure agentic coordination across digital and physical worlds. Given the pace of AI and its integration in various parts of the physical world, we believe it is important to go beyond secure digital interactions, and towards secure *cyber-physical* interactions.

This (1) ensures we're building tools for the AI systems of tomorrow that will be spanning both physical and digital worlds (autonomous robots in the wild, self-driving labs, autonomous biological systems etc), and (2) is where the biggest gaps (and opportunities!) lie with respect to secure interactions between autonomous systems.

We are still determining which cyber-physical challenges we will run. At this stage, we are looking for partners eager to collaborate with us on the design of the most exciting (and feasible) cyber-physical challenges. One path we envision (and we're open to your ideas on):

+ Launch a digital Arena first, designed such that future challenges can incorporate cyber-physical constraints.
+ Introduce simulation and world-model-based challenges that stress-test secure coordination under partially physical assumptions.
+ Later incorporate physical verification technologies, trusted sensing, and robotics hardware, where doing so materially improves the quality or realism of evaluation.

Note - While the £10m earmarked is designed to cover the digital and cyber–physical aspects of the arena, there is scope for more funding if justified, subject to ARIA internal approval.

3. Delivery timelines

We're eager to get started as soon as possible. As such we want our partners to start work on implementing their plans immediately post-award. We expect the services for the Arena to be delivered over a series of distinct phases as outlined in the sections below:

- **Phase 1: Arena prototyping (May - June)**

  - Participation in co-designing the arena goals, requirements and systems (including challenge design, scoring mechanics, governance)
  - Implementation of the ARENA and its main components
  - Implementation of exploratory challenges
- **Phase 2: Arena testing (July - Sep)**
  - Implement the candidate challenges
  - Deploy first public version of the arena and its components
  - Test the security of the platform with real user participation
  - Note: at this phase we will value iteration over stability
- **Phase 3: Arena launch (Oct)**
  - Launch material for the website
  - Deploy production version of the arena and its components
  - Maintain public documentation, copy and live leaderboard
- **Phase 4: Arena iteration and maintenance (Oct onwards)**
  - Challenge design iteration
  - Ranking and reward system iteration
  - Introduction of new challenges
  - Arena component upgrades

Who you are

We do not have hard constraints on the type of organisation we will work with. You might be an engineering consultancy, a startup, a frontier AI lab, a university research group, a nonprofit or a group of individuals specially mobilising into a new entity for this call.

What we do know is that we are looking for partners who have:

+ A deeply technical focus: You are comfortable co-designing and operating at the edge of AI security.
+ An appetite for the iterative: You thrive in fast-paced environments where the infrastructure must evolve alongside the research.
+ A commitment to open, high-stakes evaluation: You believe in the mission of building robust, multi-agent security evaluations is in the public interest.
+ Willingness and ability to embed tightly with rest of the team at ARIA
+ Excited by the prospect of working with world-class researchers and software engineers in other tracks — functioning as an engine.

- Track record of engaging with, implementing, and operationalising safety practices in real-world contexts – the Arena is intended as a place where powerful new capabilities are built for agents, having experience thinking about AI safety will be valuable.
- Experience with open-source maintenance and communities – the Arena will be open-sourced and we will welcome public contributions, it will also interface with open-source tools others are building.
- Prior experience running large-scale games, AI benchmarks, security challenges or similar efforts.
- An established or planned UK presence: You have an existing UK base or are committed to establishing a presence in the UK to effectively deliver this project and engage with the local ecosystem.

SECTION 2: KEY CONTRACT TERMS

Costs & Duration

We're looking to fund one partner over a period of three years (initial 12 month contract with options to extend based on performance). We anticipate a budget of circa £10m (inc VAT where applicable) for the three years – including early cyber-physical efforts.

While these figures serve as a general guideline, the final allocation will be determined by the proposals received. We encourage you to request the resources necessary to achieve the project's objectives, provided the proposal demonstrates maximum value for money through efficient delivery and technical impact.

There are a few constraints on this funding:

- We'll fund your costs (see our eligible cost guidance for more details) in arrears, on a time and materials basis.
- We may consider incentive based payments on usage and other success metrics we set together.

Interactions with the rest of the programme

- We expect you to interact with some Creators from track 2 and track 3 since the Arena will be where other tracks' output are tested.

- In an ideal world, this is a virtuous engine that leads to awesome discoveries, we need to ensure we have good communications with these various tracks and it will be part of the requirements.

Intellectual property (IP)

The contract will be placed on terms and conditions (provided by ARIA to the preferred applicant). Including the following considerations:

- All new implementations, code artifacts, and deliverables created specifically for the Arena (Foreground IP) shall be released under standard open-source licensing (MIT or Apache 2), while the supplier shall retain full ownership of any pre-existing proprietary tools, platforms, or frameworks (Background IP) utilised, granting ARIA the necessary usage rights to operate the Arena. This position is our starting point and we're willing to consider proposals with different IP considerations as long as they are aligned with our values and the ultimate goal for this programme.
- ARIA is open to exploring mechanisms, potentially including a shared liability or a safe harbour framework, to help protect the Arena Partner from legal or operational risks arising from third-party adversarial actions, policy breaches, or malicious code executed within the intended scope of the Arena's challenges. The feasibility and form of any such approach will be considered as part of contracting and may not be possible in all cases.

SECTION 3: HOW TO APPLY

The application process will consist of two stages:

Stage 1 - Submission of a short written proposal - Applicants should submit an initial two page proposal. Teams shortlisted based on their initial proposals will be invited to meet with ARIA to discuss their proposal. Following this discussion applicants will be invited to submit a full proposal or notified you have been unsuccessful.

Stage 2 - Submission of a detailed proposal - Applicants shortlisted at stage one will be invited to submit a detailed proposal. As part of our review we may invite applicants to meet with the Programme Director to discuss any critical questions/concerns prior to final selection, this discussion can happen virtually or we may seek clarification on certain aspects of your proposal via email.

At this stage you will be notified if you have or have not been selected for an award subject to due diligence and negotiation. If you have been selected for an award (subject to negotiations) we expect a 1 hour initial call to take place between ARIAs PD and your lead within 10 working days of being notified. We expect contract signature to be no later than 6 weeks from successful/ unsuccessful notifications. During this period the following activity will take place:

- Due diligence will be carried out
- The PD and the applicant will discuss, negotiate and agree the project activities, milestones and budget details
- Agreement to the set Terms and Conditions of the contract/grant. Please note ARIA does not negotiate these terms.

Proposal Guidance

Applicants are invited to set out how they propose to deliver the Services outlined within this RFP. The format below is set out as a guide and represents a maximum length response. If applicants choose to respond in a different format this will be acceptable as long as sufficient information is provided to be comparable to this format of response.

Stage 1 - Short written proposal (2 pages max)

- Why you're excited about this
- What skills & projects you've built that support your expertise
- Your views on key tradeoffs in challenge design, scoring, platform openness, and operations. You might include here thoughts on security posture and/or on safety.
- What you think we're doing wrong
- Your high-level plan for Year 1, including very rough cost estimation using the table below.
- Who would work on this, their relevant experience (any links on them), and their availability to start.
- How you plan to embed with a small, fast-paced team here at ARIA
- Any terms upfront that you'd like to surface as important to you
- Anyone else you think we should talk to? (e.g. anyone we should engage as challenge designers, advisors, or early participants?)

Please complete the table below providing an estimate in GBP (inclusive of VAT where applicable and all other costs) of what you consider a reasonable funding amount for your project. It's ok if you're not sure — give us your best estimate. This table does not count to the two page limit.

Cost Type	Budget (£ Inc VAT)
Labour	
Materials	
Subcontract	
Equipment & Facilities	
Travel	
Other	
Subtotal	
Indirect Costs	
Profit	
VAT (where applicable)	
Total	

At stage 1, it is important to tell us what you can offer in capabilities and capacity rather than proposing the exact right thing, we are looking to understand what we could work on together rather than nail the exact specs of the Arena upfront.

Stage 2 Detailed proposal (~8 pages max)

More information will be provided in the invitation to submit a detailed proposal. This will include at a minimum:

- Proposed plan for delivery
- Commercial Proposal
  - including a detailed budget breakdown (we will share a costing template when we invite you to submit a full proposal)

Proposal Selection Criteria

In conducting a full review of the proposal we'll consider the following criteria:

A. Demonstrated ability to do the work (track record, past projects etc) - Evidence of the team's capability to successfully deliver the work, including relevant experience, past projects, and comparable delivery.

B. Proposed Plan - The proposal demonstrates a clear, credible, and well-structured plan for delivery, with a strong focus on how the work will translate into meaningful progress and outcomes.

C. Alignment with ARIA's culture and values - The team demonstrates alignment with ARIA's ways of working, including ambition, openness, and a willingness to take bold, high-impact approaches.

D. Speed of execution and iteration - The team and proposal demonstrate the ability to move quickly, with a clear approach to rapid experimentation, learning, and iteration.

E. Commercial terms that demonstrate value for the tax-payer, through transparent pricing models with no hidden costs and a proven ability to offer cost-effective solutions without compromising quality.

Applicants will be reviewed against both their written response and discussions with the ARIA team.

Format of Proposal Submission

Format	PDF. Pages should be numbered, and the response should include the applicant's name inserted as a header
Responses to	**Submit your proposal here** In case of any technical issues with the portal please contact clarifications@aria.org.uk

Timeline

We intend to select an Arena partner as soon as possible within the initial timeline set out below. However, if we do not identify a suitable partner, we will continue to accept and assess proposals on a rolling basis using the two stage application process detailed above until we find the right partner.

2-page short proposal submission deadline	**14 April 2026 (14:00 BST)**
Discussions with shortlisted applicants will take place during this period. Applicants will be notified of the outcome of this stage no later than 30th April 2026. We may ask you to begin work on the detailed proposal as soon as possible before the 30th April.	**14 April 2026 - 30 April 2026**

Invitation for detailed proposal (8-page proposals)	**No later than 30 April 2026**
Applicants invited to submit a detailed proposal will receive specific timelines in the invitation and will be given 2 weeks to respond, with the option to submit earlier.	

Clarification Questions

Send us questions via email to clarifications@aria.org.uk. You can also shoot us messages on Discord.

Any clarification request or responses containing information that is of relevance to all applicants will be provided to all applicants that confirm their intention to participate. Answers to clarification requests will also be posted to the ARIA website, following the deadline for submission of clarification requests. If applicants do not wish a query or response to be disclosed to other applicants, they must communicate this and the reason why, with the clarification questions

If you are disabled or have a long-term health condition, we can offer support to help you engage with ARIA, navigate our RFP process, you can find more information here.

APPENDIX

I. Illustrative Challenges

The challenges below are examples we're actively exploring, not commitments. We expect the final challenge set to emerge from community input, early Arena testing, and ongoing research. If you have ideas for challenges that would stress-test secure agentic coordination, we want to hear them.

We're thinking about challenges in two categories, Structured and Unstructured.

Structured challenges have a clearly specified goal and measurable success criteria. The best strategy may be difficult to execute, but what "good" looks like is known. These challenges let us track progress on specific capabilities.

Examples we're exploring:

+ *Requirements elicitation*: Extract a complete security policy from a user with minimal communication rounds.
+ *Constrained negotiation*: Reach a mutually beneficial agreement with a counterparty agent under strict token or time budgets.
+ *Protocol selection*: Given a security goal, select and correctly configure the appropriate cryptographic protocol from a library.
+ *Unforgeable physical receipts*: Produce cryptographically verifiable proof that a robot performed a claimed physical action.
+ *Crypto-emergent challenges:* can agents *discover* that cryptographic coordination is useful, and deploy it appropriately? For example, agents earn points for finding the intersection of their private data, but lose points if any agent learns information beyond the intersection. The winning strategy is essentially private set intersection (PSI). Other variants might embed principles from zero-knowledge proofs, fully homomorphic encryption, or multi-party computation, without naming them.
+ *General Game Playing with hidden information*: Compete on unseen game rule-sets where some state must remain private.

Unstructured challenges place agents in complex, open-ended environments where the strategy space is vast and likely computationally intractable. The reward signal may be simple (e.g., profit, task completion), but there are many paths to get there, and no one knows the optimal approach.

These challenges test whether generalised capabilities emerge and compose in ways we didn't explicitly design for.

Examples we're exploring:

- *Autonomous business*: An agent coordinates with suppliers, manufacturers, and customers to profitably sell a cyber-physical product, optimising against its Profit & Loss statement with no prescribed strategy.
- *Secure scientific replication*: One autonomous lab discovers a novel method and must securely teach it to an untrusted competitor for independent replication.
- *Collaborative manufacturing*: Robotic agents self-organise to build a novel object from components sourced through an untrusted, potentially adversarial supply chain.
- *Disaster response with compromised agents*: Competing robotic teams collaboratively map a disaster site and triage victims, even when some agents are sharing false information.

We intend to select one unstructured challenge as the programme's symbolic North Star, a large-scale demonstration that cyber-physical agentic coordination can achieve outcomes that are useful in the real world, under adversarial constraints. This selection will happen as the programme matures and we learn from early Arena results

II. Confidentiality, Publicity, Conduct and Conflicts of Interest

ARIA will treat the information you provide in your Proposal as confidential and will not disclose it other than as necessary for the purposes of this RFP, the associated procurement process and ARIA's functions. ARIA may use the information you provide to assess and compare submissions, administer this RFP, agree and manage any resulting Contract, and support ARIA's operational activities, including the development, testing and improvement of internal tools, systems and processes.

ARIA may share Proposals with relevant employees, contractors, advisers and third-party expert reviewers for evaluation purposes. Any such recipients will be subject to appropriate confidentiality, non-use and conflict of interest obligations. ARIA will process any personal data contained in your Proposal in accordance with applicable data protection legislation and its privacy notice and may contact individuals named in the Proposal using the details provided.

In support of its internal processes and operations, ARIA may use secure* generative AI tools. This may include using information from Proposals to assist with assessment, comparison, analysis, summarisation and identification of themes, risks or opportunities, and to support the development and improvement of internal tools and processes. Any AI-assisted outputs will be reviewed by ARIA staff and will not be used as the sole basis for contract award decisions.

*For ARIA, "secure" means enterprise-approved generative AI tools that have undergone ARIA's Technology Security Assessment Process to ensure appropriate security controls, transparency over data storage and processing, and that ARIA data is not used to train external AI models or shared beyond authorised environments.

Although ARIA will treat Proposals as confidential, applicants are advised not to submit potentially patentable or highly commercially sensitive information before appropriate protections are in place. ARIA may disclose confidential information where required in order to meet its obligations as a public body, including for parliamentary, ministerial, judicial or audit purposes (including disclosure to the Comptroller and Auditor General). ARIA may publish statistics or high-level information relating to this RFP or any resulting Contract, but will not publish confidential or personal data.

The contents of this RFP are confidential and must not be copied, reproduced, distributed or disclosed except for the purpose of preparing and submitting a Proposal. No publicity regarding this RFP or the award of any Contract will be permitted without ARIA's prior written consent. ARIA may use ideas or concepts contained in a Proposal for any reasonable purpose connected with this

RFP or subsequent discussions, provided that it does not disclose the identity of the originating applicant.

ARIA reserves the right to amend, clarify or withdraw this RFP at any time and may extend submission deadlines where appropriate. ARIA may reject any non-compliant Proposal, disqualify any applicant guilty of serious misrepresentation, fraud, regulatory breach or conduct giving rise to material risk to ARIA, re-run the procurement on the same or an alternative basis, choose not to award any Contract, or amend the timetable, structure or content of the procurement process as it considers appropriate. ARIA will not be liable for any costs, expenditure or effort incurred by applicants in connection with this RFP or the procurement process, including where the process is amended or terminated. It is at ARIA's sole discretion as to which, if any, Proposal is accepted.

Any attempt by an applicant or its advisers to influence the contract award process in any way may result in disqualification. Applicants must not enter into any agreement or arrangement with any other person as to the form or content of any other Proposal, offer any inducement to alter another Proposal, enter into arrangements that prevent or restrict another party from submitting a Proposal, canvass ARIA or its employees or advisers in relation to this procurement, or attempt to obtain confidential information concerning another applicant or Proposal. Applicants are responsible for ensuring that no actual or potential conflicts of interest exist between the applicant (or its advisers) and ARIA, its employees or advisers, and must declare any such conflicts in their Proposal. ARIA reserves the right to disqualify an applicant where a conflict cannot be appropriately managed.

ARIA staff (employees, secondees, individual contractors engaged by ARIA in a staff-like capacity, and Board members), and their immediate family members (including spouse, civil partner or unmarried partner; sibling; child whether biological, adopted or step-child; and parent) are not eligible to submit a Proposal. Any entity in which such an individual is a board director or person of significant control (as defined on gov.uk) is also ineligible. Any Proposal falling within these categories will be rejected on eligibility grounds.