

Scaling Trust

Call for proposals

Rolling applications

Date: 4 September 2026

v2.0

Scaling Trust Track 2 and 3 Rolling Applications

ARIA is reopening applications for Tracks 2 and 3 of the Scaling Trust programme on a rolling basis.

This summary explains how applications will now be managed while the call is open on a rolling basis.

What has changed

Applications may now be submitted at any time while the call is open. This approach allows the programme to respond to emerging opportunities, evolving technical priorities and gaps in the developing portfolio, while maintaining a consistent review process.

Note that ARIA will not review proposals continuously as they arrive. Instead, proposals will be reviewed in quarterly batches.

Scope and programme focus

The scope of this rolling call is unchanged. Applicants should read the [Scaling Trust call for proposals](#) in full before applying. Applicants should continue to refer to it for the programme objectives, scope, eligibility requirements, proposal requirements, assessment criteria and funding terms.

The Scaling Trust team will also maintain [a community website](#) alongside the call. This website may share emerging areas of interest, lessons from the developing portfolio, links to funded projects and areas where further proposals would be especially valuable.

The website provides additional context to help applicants decide whether and when to apply. It does not replace, narrow or expand the original call scope.

Please note - the community website captures our most up to date thinking. We encourage all applicants to prioritise reading the information here to inform your proposal.

Rolling application process

Each quarterly cycle will follow the same rhythm, based on the timing of the first cycle set out below, rather than on calendar quarters.

- + Months 1 and 2: applications open
Applicants may submit proposals at any point.

- + End of month 2: proposal cut-off
Proposals submitted by the end of month 2 will enter the review process for that quarter. Proposals submitted after the cut-off will be considered in the following quarterly cycle.
- + Month 3: review
ARIA will carry out the review process set out in the original call for proposals. Further information on ARIA's project review and selection process can be found [here](#). ARIA may ask applicants for clarifications or invite discussion on critical questions before making final funding decisions.
- + End of quarter: selection and approval
Funding recommendations will be considered through ARIA's project selection and approval process. Applicants will then be notified of the outcome, subject to due diligence, negotiation and agreement of funding terms and conditions.

Timeline of the first quarter (the same pattern repeats every 3 months)

Applications open: **4 September 2026**

Proposal cut-off: **31 October 2026**

Review period: **1–30 November 2026**

Applicant notification: **30 November 2026**

ARIA expects to fund a small number of projects through this route and may choose not to make awards in a given review cycle.

ARIA may update cut-off dates, pause the call or close it early if needed, including in response to proposal volume or programme portfolio needs.

[APPLY HERE](#)

Please read the [portal instructions](#) here.

If you have any questions, please use the chat function on the funding call page for the quickest response. It can guide you to the right information or connect you with the ARIA team if needed.

If you are disabled or have a long-term health condition, we can offer support to help you engage with ARIA, navigate our funding application process, or carry out your project. You can find more information [here](#).

Summary

Thank you for your interest in applying to Scaling Trust's first call for proposals. This call for proposals is derived from the programme thesis [Scaling Trust](#), in turn derived from the ARIA Opportunity Space: [Trust Everything, Everywhere](#). **We strongly recommend reading both these documents alongside the Solicitation below.**

As you read through the document, if you have any questions, please use the chat function on the [funding call page](#) for the quickest response. It can guide you to the right information or connect you with the ARIA team if needed. You can also sign up to our [Discord community](#) or see [here](#) for more information.

Description	We are looking to fund projects across Track 2 - Tooling and Track 3 - Fundamental Research.
Track 2 Tooling	Open-source agents and reusable components that enable secure requirement capture, negotiation, protocol generation, and verification in multi-agent settings.
Track 3 Fundamental Research	Foundational work that turns empirical security into provable guarantees, and unlocks new cyber-physical trust primitives for agents.
Grant size	£100k - £3m
Total number of teams	Track 2: 4-6 teams Track 3: 3 'Research Centres' + 4-12 teams
Project duration	3 months - 18 months (with potential for renewal)
Submission Deadline	24 March 2026 (14:00 GMT)

Scaling Trust Track 2 and 3 Rolling Applications.....	2
Summary.....	4
SECTION 1: Programme Summary and Objectives.....	6
SECTION 2: Programme Structure.....	8
SECTION 3: Tracks.....	12
SECTION 4: Programme Duration and Project Management.....	22
SECTION 5: Eligibility & Application process.....	24
SECTION 6: Evaluation Criteria.....	26
SECTION 7: How to apply.....	28
Proposal Guidelines.....	29

SECTION 1: Programme Summary and Objectives

Scaling Trust is a £49.8m R&D programme to build tools for agents to securely interact with one another in untrusted environments while respecting the preferences of their users.

We believe this will unlock a set of new capabilities for the world; it will democratise access to advanced security primitives and their applications, will open up new secure interactions previously impossible for humans or traditional software, and will enable networks of personal agents to meaningfully coordinate.

We plan to get there by splitting efforts into three tracks: Fundamental Research will produce the reservoir of new knowledge and theoretical backing of our efforts, Tooling is informed by / are implementations of the research, and the Arena is the live adversarial environment where the tools and research are tested.

By the end of the programme, we aim to evidence real-world demonstration of the tools, confidence in their trustworthiness and evidence of high-impact usage.

We assume the reader is familiar with the [Scaling Trust programme thesis](#) where we lay out our rationale for this effort.

What we want to achieve

The core output of the programme are tools for agents to securely interact with one another, while respecting the preferences of their users. We can break this down into sub-components in order to help us understand what capabilities need to be built:




Fig 1. High-level core components

Sub-components

At a high level, we identified the following sub components (more information can be found in the [Thesis](#)):

1. **Requirement gathering** – input fuzzy user requirements → outputs a security policy
2. **Negotiation** – input individual security policy → outputs shared collective policy
3. **Security reasoner** – input security policy → outputs the final protocol
4. **Report** – input execution trace → outputs succinct convincing statement

This is a useful abstraction rather than necessarily the right split of components, depending on how the tools are built they might be integrated rather than separate. There is a large spectrum of possibilities on how each component can be built. For example the Security reasoner could be made of multiple components:

- + **Protocol designer** – able to reason through the needs of the user and generate a protocol (assuming a trusted third party), input security policy → output is the ideal functionality spec (or basic implementation)
- + **Cryptography solver** – given a precise security goal or ideal protocol, proposes a cryptographic implementation, input ideal functionality → outputs a cryptography protocol spec proven to be secure for validity
- + **Protocol implementer** – given a spec implements the protocol securely, input protocol spec → outputs an implementation

There is a large space of solutions for how to construct each of the components. We're interested in different approaches used by agents, with different strategies including but not limited to:

- + **Rule-based approach**: agents, or subcomponents, with a white-listed set of strategies, protocols, implementations or libraries.
- + **Theorem prover-based approach**: agents, or subcomponents, with access to theorem provers and a vast set of well-specified protocols to reason through.
- + **Learning-based approach**: agents, or subcomponents, that have been trained via state-of-the-art (e.g. via reinforcement learning, fine-tuning).

Success at the end of the programme

- + **Real-world demonstration of the tools**
 - Demonstrate autonomy for major interactions – Can be used without needing sophistication or regular human intervention, important for their democratisation.

- Demonstration of AI advantage – Enables interactions previously too expensive or impossible with traditional software/humans.
- Demonstration of generality – Can generalise across fields and tasks.
- + **Confidence in the trustworthiness of the tools**
 - Empirical confidence – Empirical results that show that the tools built are trustworthy, adversarially robust, cheap, efficient via benchmarks and competitions.
 - Scientific results – Science that makes us comfortable to use the tools, affording us formal guarantees.
- + **Evidence of high impact usage**
 - Community – a large community of builders and users that improve these tools together.
 - Industry adoption – a few teams have nailed the first versions of this technology, it becomes implemented in AI systems and starts creating value.
 - Customer-centric development cycle started – the cycle of tools improving based on market demand has started, and is likely to continue on its own.

SECTION 2: Programme Structure

We plan on getting there by splitting our efforts into three tracks:

- + **Track 1: Arena (Not part of this call for proposals)**
Adversarial testing grounds designed to scalably test AI systems capabilities in multi-agent coordination across digital and physical worlds.
- + **Track 2: Tooling**
Open-source coordination infrastructure usable by all in the arena and beyond, to steer innovation toward the most meaningful axes of progress.
- + **Track 3: Fundamental Research**
Flexible funding to create new fields of research and build a reservoir of new knowledge that future iterations of Tooling and the Arena can draw upon.

Why?


We can build tools in a vacuum, but if they're not tested in live, adversarial environments they're unlikely to be secure—this is why we need the Arena. Likewise, we can build tools for the use-cases we have in mind and iterate on them in the Arena, but empirical iteration without theory is guesswork. One impossibility result can eliminate an entire design space; one new primitive can unlock capabilities we hadn't imagined.

By having the three tracks working together, we create a shared environment conducive to breakthroughs. Empiricism is supplemented with theory, and theory is guided by empirical

research, all culminating in a live, adversarial environment where ideas are tested and iterated on.

Indicative roadmap

Keep in mind that we will change and pivot as we learn new information.



We've structured this programme to be highly iterative, this is because:

- + The space we're operating in (AI/software-ish) is one that moves quickly, a quarterly cadence is suited to the pace of the technology and ensures what we do remains relevant.
- + It allows for dynamic adjustments through time by building in regular feedback loops in what we do, ensuring we can remain grounded in impact with our funding.
- + It allows us to create regular milestones for people to rally around (e.g. the latest high score in a leaderboard)
- + It helps translation by building customer-centric iteration cycles that can graduate into industry more easily (in a way, our iterations are a way to bootstrap the kind of market dynamics we anticipate will continue running after the programme ends).

As such, the roadmap is divided into four phases:

1. **Bootstrap** – getting the programme up and running
We're building a v0 of the Arena and any key tooling in-house to get our hands dirty. This is to help us further understand requirements for our first solicitation, and is a useful tool for community building. We are also mobilising key members of the community to brainstorm the best set of initial challenges in the Arena and find key partners to work with. Throughout this phase, we are collaborating with our [Pre-Programme Discovery funded teams](#)¹.
2. **Test** – a first test of the arena live and of a 'season' of games
Once the first wave of funding is out, we aim to quickly test whether the Arena and the first version of tools we've funded are working as intended. We aim to go through a full Arena cycle, from launching a set of challenges to awarding prizes to the best contestants and charting performance across a set of key technical metrics. This will allow us to quickly adjust our methodology, and to decide whether we move on to phase 3 or re-adjust.
3. **Improve** – new improved challenges and agents
Once we've proven the programme concept is working and robust, phase 3 is about building momentum by running four seasons of the Arena (quarterly), improving the challenges and tooling at each quarter and charting their progress for all to see. The iterations will also inform us on how much progress we can aim for the technologies within the programme (we will have several data points of progress per quarter), and where we could double down – informing our second solicitation and any additional programmes that should be run.
4. **Scale** – setting the ground up for adoption and a next phase of the programme
This phase will see the Arena continuing running on a quarterly basis, while kicking off translation and adoption activities. This is the transition phase of the programme,

¹ Scroll to the 'Pre-Programme discovery projects' section in the [Scaling Trust Programme Page](#).

capitalising on the technologies that would've been built and ensuring the cycle of iteration that we have started can continue beyond the programme's lifetime.

SECTION 3: Tracks

This call is specifically for Track 2 and 3. While Track 1: Arena will be funded via a separate call than the present one, we describe it in the section below in order to help applicants understand the deployment environment for Track 2 and 3. If you're interested in delivering the Arena, please fill in the form [here](#).

Track 1 – Arena (Not part of this call for proposals)

Description	Adversarial testing grounds designed to scalably test AI systems capabilities in multi-agent coordination across digital and physical worlds.
Goals	(1) to surface the state of the art in secure agent to agent coordination, (2) track progress of the effort of the programme through time, (3) create a competitive environment for teams around the world to measure their systems.
Sub-tracks	- <i>1.1 Challenge Design</i> – designing, assessing and updating Arena challenges - <i>1.2 Digital Arena</i> – standing up the digital infrastructure needed for the Arena (website, logging tools, API, governance/oversight etc) - <i>1.3 Cyber-Physical Arena</i> – standing up the cyber-physical infrastructure needed for the Arena (physical environment, robotics, world models, simulation tools)
Award size	£8m initial prize pool for contestants

The Arena's purpose is to surface and improve the state of the art in secure agent to agent coordination. It will host challenges and participation will be global and open to all. The best contestants will be awarded prizes on a quarterly basis out of an initial prize pool of £8m.

All sub-tracks contribute to bringing the Arena and its challenges to life. We anticipate working with several service providers to set up and maintain the Arena, and with researchers and practitioners to design and iterate on Arena challenges. As indicated in the roadmap, we expect to have a live Arena by Q4, and a demo Arena soon.

Activities

The Arena will host both benchmarks and challenges:

+ Benchmarks

- A benchmark is a self-contained test that scores an agent, or a subcomponent, 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
- ARIA (or its contractors) will provide basic tools to participate in the challenges and run the engine required for running challenges.

Mechanics & Rules

+ Participants

- Standard participation – agents participate in challenges to complete tasks and they are scored based on their ability to complete a task and to respect their security policy (we'll be sharing more detailed information about how to participate in the coming months, sign up [here](#) to receive updates).
- 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 being part of the competition as the baseline red team and the 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.
- £1m grand prize for every 'season' of the arena.

+ Challenges


- Every quarter we will have the option (but not the obligation) to add or retire challenges. We may put constraints around compute or the types of models used.
- We may run challenges multiple times to obtain statistical significance.

Scoring

Agents in the arena will be scored against (Utility;Security), i.e. their ability to complete the task (Utility) vs their ability to respect security policies (Security). We want to surface the most useful, secure agents. 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).

Demo

To help bring this to life, the ARIA programme team has developed a demo to support the programme's design. The test arena is still a work in progress and is intended for illustrative purposes only: <https://arena.nicolaos.org/>



Track 2 – Tooling

Description	Open-source tooling that will provide the baseline infrastructure usable by all in the Arena and beyond
Goals	- Create basic agents to build on top of to participate in the arena

	- Build specialised components that can be utilised by many agents - Explore a diverse set of agent design strategies that can lead to production-ready/grade implementations - start the adoption phase for these technologies.
Sub-tracks	- 2.1 Agents – agents that can be used as ‘participants’ in the Arena, composed of a set of components - 2.2 Components – a specific tool that any agent can use - 2.3 Adoption – production software, integration and pilot efforts. Note: In this call for proposals, we are only concerned with 2.1 Agent and 2.2 Components. 2.3 Adoption will come in a later solicitation (expected mid-2027). We provide the information below to you for indicative purposes only.
Project size	£200k to £2m, per project
Expected no. of teams	4-6 teams We’re looking for Creator teams of any size for this track, moving at a significant pace given the progress we’re aiming for.
Project length	3 months to 1 year
Continuity	In mid-2027, alongside the solicitation for Track 2.3, we plan to review whether increased support to the most promising teams in tracks 2.1 and 2.2 would help them reach adoption and production grade systems.

Definitions

Agents are a combination of a set of components with some orchestration logic.

Components are specific tools usable by agents.

2.1 – Agents

We plan on funding two types of agents over time. In this first call for proposals, we are focused on ‘Basic Agents’.

Basic Agents

Basic agents are open-source agents designed to participate in the Arena. They will provide a baseline template ‘player’ for all Arena participants to iterate from. We are looking for agents with capabilities laid out in ‘What we want to achieve’ above. We’d like to fund several approaches to get there to find out what works best.

Production agents

As the Arena progresses, we will identify potential agent construction that we believe can become useful in production. We will select those from the top performing agents in the Arena as well as from those that with additional funding have the potential to be top performing.

2.2 – Components

There are individuals who will have little interest in building an agent, but will be interested in building a specific component. As described in Programme Objectives, we are interested in:

- + **Requirement gathering** – input fuzzy user requirements → outputs a security policy/goal
 - **Policy capture** – given a set of user goals, extracts a formal security policies
 - **Security policy elicitation protocols** – interactive methods that extract missing details, resolve ambiguities and help user discover their goals, input fuzzy user interactions → outputs security policy

Might include: security policy extraction tools, efficient communication elicitation protocols, user experience for policy discovery, datasets for training security policy elicitation.

- + **Negotiation** – input individual security policy → outputs shared collective policy
 - **Negotiation engine** – engine that can reason to maximise the utility of the agent, while respecting the security policy that propose or verify others proposals
 - **Contracting languages** – agreements for verification, dispute resolution, logging
 - **Negotiation safety** – communication with external parties opens up a new attack surface (e.g. jailbreaking, persuasion) and requires useful guardrails to prevent agents negotiating away from their goals and security policy

Might include: formal bargaining engines, negotiation simulations, benchmark for negotiation.

- + **Security reasoner** – input security policy → outputs the final protocol implementation

- **Protocol designer** – able to reason through the needs of the user and generate an idealised protocol (assuming a trusted third party), input security policy → output is the ideal functionality spec (or basic implementation)
- **Cryptography solver** – given a precise security goal or ideal protocol, proposes a cryptographic implementation, input ideal functionality → outputs a cryptography protocol spec proven to be secure for validity
- **Protocol Implementer** – given a spec implements the protocol securely, input protocol spec → outputs an implementation

Might include: implementation of different specialised AI models listed above or an end-to-end security reasoner, benchmark and datasets for each sub problem, libraries for cryptography, AI-assisted theorem provers,

- + **Report** – *input execution trace → outputs succinct convincing statement*
 - **Security Auditor** – given a protocol specification (or an implementation) determines that it was correctly implemented, *input protocol spec + implementation → outputs an audit report*

We expect to focus mostly on the Negotiation and the Security Reasoner components.

We're interested in different approaches used by agents, with different strategies including but not limited to:

- + **Rule-based approach**: agents, or subcomponents, with a white-listed set of strategies, protocols, implementations or libraries.
- + **Theorem prover-based approach**: agents, or subcomponents, with access to theorem provers and a vast set of well-specified protocols to reason through.
- + **Learning-based approach**: agents, or subcomponents, that have been trained via state-of-the-art (e.g. via reinforcement learning, fine-tuning).

Might include: MCP tools for cryptography, datasets for protocols (cryptography protocols, network protocols, mechanism design protocols), tools for fine-tuning or reinforcement learning, trained security reasoner models (trained with strategies described above)

2.3 – Adoption (not part of this call for proposals)

A core goal of the programme is to ensure these technologies are adopted. This isn't something we plan on funding in this first call for proposals. However at a later stage, we intend on:

- + Adopting a forward-deployed approach where we may provide funding for engineers to be embedded directly in partner organisations to customise and implement the programme's tools to help solve specific, real-life problems.

- + Stimulate synergies with capital partners (co-operate with industry partners, accelerators and funds)
- + Funding activities allowing promising teams to scale-up (e.g. to navigate regulatory compliance etc.)

Track 2 Performance Evaluation

We will evaluate success in this track based on the ability of agents and components to drive measurable performance in the Arena and achieve wider ecosystem adoption. If Arena competitors find these tools to be inefficient or insecure, we'll expect Creator teams to course-correct based on this feedback.

Metrics:

1. **Competitiveness** – The performance of components and agents in the Arena. The following metrics are key indicators of competitiveness:
 - + **Rank** - in the Arena Leaderboard
 - + **Impact** - where the introduction of components correlate with a measurable improvement in an agent's "Utility vs. Security" score
 - + **Velocity** - Demonstration of quarter-over-quarter improvements in performance metrics (e.g. outpacing the baseline Arena average.)
2. **Generality** – Tools demonstrating broad applicability (not over-fitting to single challenges)
 - + **Task breadth** - Production agents reliably completing a majority of distinct challenge types with positive utility and negligible security failures.
3. **Efficiency** - The computational and interactional efficiency of tools
 - + **Cost** - Operation costs for production agents must remain below the program's target cost per challenge (excluding negotiation overhead).
 - + **Negotiation Speed** - The speed in which agents can successfully negotiate (e.g. number of 'round-trips')
4. **Adoption** - Usage by the wider community
 - + **Market penetration** - The number of external organisations adopting the tools by the end of the programme.

Track 3 – Fundamental Research

Description	Theory that moves us from empirical to theory-driven guarantees and helps us design new security primitives that can aid agentic coordination.
Goals	- Bring scientific confidence to the trustworthiness of developed agents

	- Provide a scientific framework for formal AI security and generative security - Design new security primitives for agents to securely interact with the real world
Sub-Tracks	- 3.1 Formal AI Security: Formalisation of agentic adversaries and new security settings. - 3.2 Cyber-Physical Primitives: Security primitives that can aid cyber-physical agentic coordination - 3.3 Foundations of Generative Security: Automated protocol generation and verification. - 3.4 Bluesky: Open-ended research.
Project Size	£100k to £3m, per project.
Expected Team size	3 research centres (defined below) + 4 - 12 smaller teams
Project Length	6 months – 18 months (though applicants are encouraged to consider plans which may reach success (or failure) on faster timelines).
Continuity	We will double down on research projects in mid-2027 based on outputs from the first year.

We're looking to fund two types of Creators teams:

Research Centres (~£2-3m each): Three large funding awards to seed long-term nodes of expertise in Formal AI Security, Cyber-Physical Primitives, and Generative Security. We want these centres to become stewards of their fields—not just producing research, but attracting talent and anchoring a community in the UK.

Expected Research Centres' profiles:

- + an alliance of cross-international people with a base in the UK
- + a new centre created at a university
- + a new institute
- + a new foundation
- + a new startup that aims at attracting talent to build an r&d
- + a new unit of a large research centre

Exploratory projects (~£100-300k each): Smaller, faster funding awards for ambitious projects we can double down on later. These are for lean, cross-disciplinary teams of 4-5 researchers each or less, chasing ideas that wouldn't find funding elsewhere.

We expect to fund one research centre in each of Tracks 3.1, 3.2, and 3.3, and may also fund smaller projects within these tracks. In Track 3.4, we expect to fund smaller awards for experimental projects.

3.1 – Formal AI Security

Formal security definitions allow researchers to prove whether a system is secure under explicit assumptions, reason about what is possible (via feasibility and impossibility results, hierarchy of assumptions and guarantees), and provide building blocks for more complex protocols.

Although [early work](#) is taking place, we believe AI security today stands where information security stood in the pre-cryptographic era: we lack foundational definitions for core concepts such as intelligence, alignment, and robust communication. Without these definitions, we cannot prove security; we can only empirically observe failure.

This track seeks to establish Formal AI Security as a new discipline that applies the rigour of theoretical computer science to intelligent systems. We aim to move beyond empirical “red teaming” toward provable guarantees.

We’re specifically interested in the following areas:

- + **Foundational frameworks** - Formalisation of agentic adversaries and new security settings
- + **AI communication security** - Secure jailbreak proof communication and AI-to-AI efficient languages
- + **AI advantage** - Designing new primitives and protocols that leverage AI advantage primitives.

See more details in the [thesis](#).

3.2 – Cyber-Physical Primitives (‘Nature Crypto’)

As agents interact with the physical world, digital security primitives aren’t enough. How does an agent verify a sensor reading is authentic? That a manufacturing process occurred as claimed? That a biological sample hasn’t been tampered with? This track funds a new field of security that uses properties of nature—physical and biological—as foundations for trust.

Some of the topics we’re excited about here include:

- + **Verifiable physical processes** - Authenticated readings of the physical world via a variety of sensors and physical properties.

- + **Physical Trust Anchors** - Using physical properties (e.g., Physically Unclonable Functions, DNA watermarking) as roots of trust for digital systems
- + **Trusted Hardware** - Adapting secure enclaves and tamper-resistant hardware to various cyber-physical use-cases (e.g. making tamper-resistant sensors)

3.3 – Foundations of Generative Security

Agents operate in dynamic, context-specific environments, where they must be able to generate, negotiate, and verify their own security protocols on demand. This track aims to identify, formalise, and address the root research problems required to allow agents to autonomously design and verify cryptographic protocols that are provably secure. This is a non-exhaustive list of what we're interested in:

- + Efficient security policy gathering
- + Theory of agent-to-agent security negotiation
- + Succinct and inspectable proofs for correct execution and delegation
- + Automated protocol generation
- + Automated security proof generation for cryptographic protocols

3.4 – Bluesky research

We anticipate encountering novel research problems that we cannot foresee today. This track serves as our strategic reserve to address these 'unknown unknowns'. We are looking for proposals that:

- + **Address emerging bottlenecks** - tackling theoretical problems that we uncover as the programme progresses.
- + **Explore 'wildcard' concepts** - We invite applicants to propose high-value, radical research questions that could fundamentally shift our approach to agentic coordination, even if they currently lack a unified theory. If you see a critical theoretical gap in our [thesis](#), use this track to pitch the solution.

Track 3 Performance Evaluation

We ascribe to the philosophy of John Naughton and Bob Taylor in [Zen and the Art of Research Management](#) or Donald Braben in [Scientific Freedom: The Elixir of Civilization](#); namely hire excellent people, get out of their way, and avoid imposing short-term metrics on them. Therefore, unlike track 2, rather than applying a set of metrics/definition of success, we'll assess projects individually. We'll work with you to shape milestones that are appropriate for your project (see Section 4: Programme Duration and Project Management).

That said there are some things that we do want to track:

- + **Bottleneck resolution** - The successful ruling out of major theoretical impediments or the establishment of critical impossibility results.

- + **Community Seeding** - The successful creation or acceleration of new scientific communities, specifically in the fields of:
 - **'Nature' Cryptography:** Using physical/biological processes for security.
 - **Formal AI Security:** Rigorous mathematical study of intelligent systems.
 - **Autonomous Protocol Design:** Self-generating security standards.

Translation & Continuity

Across all tracks, we see creator support as extending beyond this initial funding round, with a focus on what comes next. We anticipate introducing a few initial mechanisms to help facilitate continuity:

- + **Adoption support** – an entire sub-track (2.3) dedicated to adoption where we plan on funding industry pilots, forward deployers and business development. A separate call for proposals will be run to find these partners.
- + **Challenges anchored on reality** – the challenges will reflect real-world valuable situations so that capabilities built are immediately useful.
- + For teams exploring the formation of a new company, the ARIA team is keen to engage early to discuss plans as they develop. While we cannot provide formal legal advice or fund incorporation costs (see eligible costs), we are very happy to share lessons learned and act as a thought partner throughout the process.

SECTION 4: Programme Duration and Project Management

Programme & Project Management

Progress for each project will be guided by clearly defined, jointly agreed milestones. These milestones will be proposed by the applicant ahead of the project start, agreed with ARIA, and designed to make progress easy to understand. To support this, milestones should be specific, measurable, and represent meaningful steps towards the overall programme goals. Further guidance on setting ARIA milestones can be found [here](#).

The programme team will work alongside Creators throughout the project, maintaining regular and open dialogue to refine goals and adjust direction where helpful. This approach is supportive rather than punitive: research is inherently uncertain, and we expect plans and milestones to evolve as new insights emerge.

As a minimum, quarterly check-ins will take place between Creators and the programme team to talk through progress against milestones and any emerging details of the work. These sessions are intended as “thought-partner” conversations rather than formal reviews.

Quarterly conversations will typically cover progress against target milestones, the key technical risks and the most effective activities to address them, expected outcomes and learnings, and any current dependencies on earlier phases or parallel efforts. If a project no longer feels like the right thing to continue, any decision to bring it to a close would be a shared and well-communicated one, made thoughtfully and without surprise, after exploring the available options.

Creator and Community events

We want to create an environment where you can meaningfully connect with each other—not just through us. We'll run:

Build weeks: Quarterly or bi-yearly physical gatherings where Track 3 researchers explain their work to Track 2 engineers, and Track 2 teams onboard Arena competitors onto their tools. Optional hackathons included.

These events will be essential for teams to exchange updates, share data and tools, and work together to solve cross-cutting challenges. Active participation and a willingness to share progress, challenges, and "negative results" openly within this community are essential and applicants should include estimated travel costs in their budget proposals (events will be UK based).

We'll also run:

- + **Bug bounties:** throughout the programme across both Arena and Tooling. If an Arena competitor finds a vulnerability in a Track 2 tools, they will be rewarded accordingly.
- + **Community platforms:** Discord for day-to-day chat, GitHub for code, and regular touchpoints via events with the broader ecosystem—including VCs and potential customers.
- + **Community efforts and engagement:** regular community events, and several touchpoints, that will include VCs and potential customers.

IP Approach

You are building the foundational infrastructure for secure agentic interactions. If these tools and research are proprietary, the ecosystem will not adopt them, and the programme will fail to scale. Therefore, transparency here is a requirement.

To maximise adoption and interoperability, all software produced under Track 2 (Tooling) and Track 3 (Fundamental Research) must be released under a permissive open-source licence. We require a dual-licence approach under **MIT and Apache 2.0**. This ensures compatibility with the broadest range of downstream users, including commercial and

academic entities. In all instances you retain ownership of your pre-existing IP (Background IP). We only require open licensing for the specific deliverables (Foreground IP) funded by this programme. In exceptional circumstances, you can request some results not to be published. ARIA will consider these requests on a case-by-case basis.

Exception (for Track 1): While this call for proposals focuses on Tracks 2 and 3, it is important to understand the IP dynamics of the Arena environment you are building for:

- + **During Competition:** Arena competitors retain ownership of their agent weights and strategies to maintain a competitive advantage during the season.
- + **Post-Competition:** To claim prize money, competitors are **required to open-source** their winning agents. This ensures that your Track 2 tools can be eventually integrated into a public library of state-of-the-art agent architectures, creating a compounding cycle of innovation.

SECTION 5: Eligibility & Application process

Who are we looking for

We welcome applications from across the R&D ecosystem, including individuals, universities, research institutions, small, medium and large companies, charities and public sector research organisations. What is most important to us is whether you are excited about the programme and its thesis.

Discord Community

You are also welcome to discuss with other applicants and to join our [Discord](#) community here. The purpose of this forum is to provide a space for collaboration and discussion amongst ARIA funded projects as well as the wider community, where members can share insights, ideas, and build in public.

Please note that we will NOT be able to answer any questions regarding the call for proposals in this forum. If you have any questions, please use the chat function on the [funding call page](#) for the quickest response. It can guide you to the right information or connect you with the ARIA team if needed.

Finding potential collaborators and teaming

For those seeking specific expertise to support their proposal, we have created a teaming request form to facilitate finding potential team members who have registered their interest in this programme. After a quick registration you will gain access to a list of other

individuals seeking to find/share their expertise and a dedicated teaming channel on our community Discord.

[\[Sign up for teaming\]](#)

Application Process

The application process for Track 2 and 3 consists of the submission of a detailed proposal including:

- + **Project & Technical information** to help us gain a detailed understanding of your proposal
- + **Information about the team** to help us learn more about who will be doing the research, their expertise, and why you/the team are motivated to solve the problem
- + **Administrative questions** to help ensure we are responsibly funding R&D.
Questions relate to budgets, IP, potential COIs etc

You can find more detailed guidance on what to include in a proposal [here](#).

For more details on the evaluation criteria we'll use, click [here](#).

Non-UK funding

Our primary focus is on funding those who are based in the UK. However, funding will be awarded to organisations outside the UK if we believe it can boost the net impact of a programme in the UK. In these instances, you must outline your proposed plans or commitments that will contribute to the programme in the UK within the project's duration.

If you are successfully selected for an award subject to negotiations this proposal will form part of those negotiations and any resultant contract/grant.

More information on the evaluation criteria we will use to assess your answers can be found later in the document [here](#).

We have provided some additional guidance on non-UK funding in our [FAQs](#) including available visa options.

SECTION 6: Evaluation Criteria

Proposal Evaluation Principles

To build a programme at ARIA, each Programme Director directs the review, selection, and funding of a portfolio of projects, whose collective aim is to unlock breakthroughs that impact society. As such, we empower Programme Directors to make robust selection decisions in service of their programme's objectives ensuring they justify their selection recommendations internally for consistency of process and fairness prior to final selection.

We take a criteria-led approach to evaluation, as such all proposals are evaluated against the criteria outlined below. We expect proposals to spike against our criteria and have different strengths and weaknesses. Expert technical reviewers (both internal and external to ARIA) evaluate proposals to provide independent views, stimulate discussion and inform decision-making. Final selection will be based on an assessment of the programme portfolio as a whole, its alignment with the overall programme goals and objectives and the diversity of applicants across the programme.

Proposal evaluation process and criteria

Proposals will pass through an initial screening and compliance review to ensure proposals conform to the format guidance and they are within the scope of the call for proposals. At this stage we will also carry out some checks to verify your identity, review any national security risks and check for any conflicts of interest. Prior to review of applications Programme Directors and all other reviewers are required to recuse themselves from decision making related to any party that represents a real or perceived conflict.

Where it is clear that a proposal is not compliant, outside the scope and/or does not pass a quality assurance review, these proposals will be rejected prior to a full review on the basis they are not compliant or non-eligible.

Proposals that pass through the initial screening and compliance review will then proceed to full review by the Programme Director and expert technical reviewers (this may include the use of AI. Further information on ARIAs proposal review process can be found [here](#) and the use of AI in the conditions of the call available [here](#)).

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

1. Worth Shooting For:

- a. The proposed project uniquely contributes to the overall portfolio of approaches needed to advance the programme goals and objectives.
- b. It has the potential to be transformative and/or address critical challenges within and/or meaningfully contribute to the programme thesis, metrics or measures.

2. **Differentiated** – The proposed approach is innovative and differentiated from commercial or emerging technologies being funded or developed elsewhere.
3. **Well defined** – The proposed project clearly identifies what R&D will be done to advance the programme thesis, metrics or measures, is feasible and supported by data and/or strong scientific rationale. The composition and planned coordination and management of the team is clearly defined and reasonable. Task descriptions and associated technical elements provided are complete and in a logical sequence with all proposed stage-gates and deliverables clearly defined. The costs and timelines proposed are reasonable/realistic. The proposal demonstrates the team's ability to operate at pace, with organisational and administrative processes that enable rapid progress within the programme's timelines, including readiness to hire and start immediately.
4. **Responsible** – The proposal identifies major ethical, legal or regulatory risks and that planned mitigation efforts are clearly defined and feasible.
5. **Intrinsic motivation** – The individual or team proposed demonstrates deep problem knowledge, have advanced skills in the proposed area and shows intrinsic motivation to work on the project and key individuals are dedicating sufficient time to the project. The proposal brings together disciplines from diverse backgrounds and shows a collaborative mindset. Including a willingness to work across disciplines and Tracks, actively seeking input from other Creators incorporating feedback into the development process. An entrepreneurial approach is evident, characterised by initiative, ownership, comfort with uncertainty and a bias towards experimentation and delivery.
6. **Benefit to the UK** – There is a clear case for how the project will benefit the UK. Strong cases for benefit to the UK include proposals that:
 - a. are led by an applicant within the UK who will perform the majority (>50% of project costs spent in the UK) of the project within the UK
 - b. are led by an applicant outside the UK who seeks to establish operations inside the UK and perform a majority (>50% of project costs spent in the UK) of the project inside the UK and present a credible plan for achieving this within the programme duration.

For all other applicants we will evaluate the proposal based on its potential to boost the net impact of the programme in the UK. This could include:

- c. A commitment to providing a direct benefit to the UK economy, scientific innovation, invention, or quality of life, commensurate with the value of the award;
- d. The project's inclusion in the programme significantly boosts the probability of success and/or increases the net benefit of specific UK-based programme

elements, for example, the project represents a small but essential component of the programme for which there is no reasonable, comparably capable UK alternative.

When considering the benefit to the UK, the proposal will be considered on a portfolio basis and with regard to the next best alternative proposal from a UK organisation/individual.

SECTION 7: How to apply

Before submitting an application we strongly encourage you to read this call in full, as well as the [general ARIA funding FAQs](#).

If you have any questions, please use the chat function on the funding call page for the quickest response. It can guide you to the right information or connect you with the ARIA team if needed.

Any questions or responses containing information relevant to all applicants will be provided to everyone that has started a submission within the application portal. We'll also periodically publish questions and answers on our website which you can find [here](#).

Please read the portal instructions below and create your account before the application deadline.

If you are disabled or have a long-term health condition, we can offer support to help you engage with ARIA, navigate our funding application process, or carry out your project, you can find more information [here](#).

Application [Portal instructions](#)

APPLY [HERE](#)

Proposal Guidelines

How to Format your proposal

- + Page count: max 10 pages, (including diagrams, excluding references).
 - *It can be a little as 1-pager if you think it contains all the information you need (e.g. if it is a 3-month project).*
- + Format: single line spacing, standard character spacing (neither expanded nor condensed)
- + Font: Arial. Colour: black. Size: 11-point font or larger
- + Margins: At least 0.5" margins all around
- + File type: PDF only

Expected outline (this is purely indicative, feel free to adapt it to your proposal):

Section 0: Summary

- + Summary of your proposal in 500 words max

Section 1: Programme & Technical

The aim of this section is to gain in-depth, technical information about the project being proposed. This should include:

- + Which Track you seek to pursue (2.1, 2.2, 3.1, 3.2, 3.3, 3.4). Applicants wishing to apply across multiple tracks should submit separate proposals for each track, with cross-references to any related proposals.
- + A detailed explanation of the proposed idea/solution, how it supports the technical objectives of the chosen pathway.
 - This should be supported by visual aids, data and/or strong scientific rationale for why what you are proposing would work.
 - Please include any required technical information, as specified in sections 2 and 3 of the call for proposals document.
- + A comprehensive list of the known technical risks/unknowns standing in the way of achieving the stated goals.
- + How the proposed approach is differentiated, e.g. from commercial or emerging technologies being funded or developed elsewhere.
- + A description of the proposed activity of work, key metrics and milestones and any dependencies and assumptions.
- + Estimated timelines - applicants should provide a Project Plan for the lifecycle of the project, showing what you plan to achieve for each period of the project.

Section 2: The Team

This section includes information about the proposed individuals or teams that will conduct the research and management structures. This must include:

- + Details of the project team - we want to know who will be doing the work (not just the principal investigator or project lead) and what portion of their time will be dedicated to this project (we usually prefer any lead or key researchers to be spending at least 50%, ideally 80%, of their time on the project).
- + Whether they are already in place? If not, how long after project kickoff are they likely to start?
- + You could include short bios about each team member (we discourage you from submitting CVs).
- + If you intend to collaborate with or rely on any third parties, sub contractors/grantees, who are they and which elements of the project they will support/deliver.
- + How you intend to coordinate and manage the teams including any collaborations with third parties.
- + Any potential gaps in your core competency which would be required in order to achieve the overall goals.
- + We also want to know what motivates you or the team to want to do this project and why you are the right person/team to work on this project.

Section 3: Administrative Response

This section includes information about the budget, intellectual property that you intend to rely on, any perceived conflicts of interest and for non-UK applicants how the proposed project may benefit the UK.

In completing your application you must also provide answers to the following questions. Answers to these questions are not included in the 10 page cap. You should complete these questions in the application portal so there is no need to format these in a specific way.

Application	Guidance
How much funding do you need?	Please provide a cost breakdown by completing the spreadsheet here. In your proposal you may submit your budget using yearly, quarterly, or monthly phasing. Prior to completing this template you should review ARIA's Eligible cost guidance here.

	<i>If your proposal is successful, prior to contract signature when the scope of work has been agreed, you will be required to provide a monthly cost breakdown.</i>
Are you proposing to contribute funding?	<i>If you or your organisation are proposing to contribute funding to the project please let us know how much funding you plan to contribute, who is contributing the funding, is the funding already secured and any other relevant details.</i> <i>ARIA will fund 100% of project costs and contribution of funding is not essential however, we welcome proposals that contribute funding in cases when such funding will strengthen the potential success. In these cases, this funding contribution will be considered as part of the overall strength of the project proposal.</i>
Does your proposal depend on background IP (pre existing)?	<i>If Yes, give us an indication of: What background IP is required, Whether you currently have rights to that IP.</i>
Have you already secured funding for a similar project or are you currently in the process of seeking support from other funding sources for the same project?	<i>If yes, tell us more about the funding you already have or are applying for.</i>
Any other factors or restrictions that might impact your freedom to operate and deliver the project?	<i>Please provide a detailed description of any perceived conflicts of interest with the programme director, import/export or security restrictions that you are aware of</i>
How do you envision commercialisation of the proposed project?	<i>Please complete and upload a commercial hypothesis for your project using the guidelines here.</i>
Are you proposing to perform the majority of the proposed project outside of the UK?	<i>Our primary focus is on funding those who are based in the UK. For the vast majority of applicants, we therefore require the majority of the</i>

	<i>project work to be conducted in the UK (i.e. >50% of project costs and personnel time).</i> <i>However, we can award funding to applicants whose projects will primarily take place outside of the UK, if we believe it can boost the net impact of a programme.</i> <i>In these instances, you must outline any proposed plans or commitments in the UK that will contribute to the programme within the project's duration. Please provide a detailed description of any proposed plans (including a timeline) or commitments).</i>
Has a suitably authorised member of your Organisation approved the submission of this proposal?	<i>In the application portal, please select the option that best describes your situation and provide details where required.</i>
Have you read and understood our funding terms?	<i>Our goal is to ensure your research can get going quickly, so we want to ensure a fast negotiation and award process. We aim to have agreements signed within 6 weeks, which we recognise can be much faster than standard at some organisations. Before proceeding, please confirm that you have read and understand our funding terms. If you are unsure which terms apply to you, you can find more guidance here.</i>
Additional questions about you/your organisation that can be found in the application portal.	