

A Greek SME is looking for research, evaluation, pilot and ethics partners to develop artificial intelligence that plans, reasons and explains its actions (Horizon Europe research cooperation)

Summary

Profile type

Research & Development Request Greece

Company's country

POD reference

RDRGR20261002009

Profile status

PUBLISHED

Type of partnership

Research and development cooperation agreement

Targeted countries

• All countries

Contact Person

[Enrico FRANZIN](#)

Term of validity

2 Oct 2026
2 Oct 2027

Last update

8 Oct 2026

General Information

Short summary

A Greek SME developing open-source software for artificial intelligence agents is preparing, as technical partner, a European Innovation Council Pathfinder Challenge proposal under Horizon Europe. The system has separate knowledge, planning and verification layers, so actions can be explained and refused when evidence is missing. Research, evaluation, pilot and ethics partners from two other countries or more are sought.

Full description

A Greek small and medium-sized enterprise that develops open-source software for running artificial intelligence agents under controlled and logged rules is preparing a proposal for the European Innovation Council Pathfinder Challenge DeepRAP (Deep Reasoning, Abstraction and Planning towards trustworthy cognitive artificial intelligence systems), part of Horizon Europe. The enterprise acts as technical partner. A prime Greek research centre with substantial EU grant track record has agreed to take the official coordinator role. Partners from at least two other Member States or Horizon Europe associated countries are sought.

Partners sought: Research partners: universities or institutes in knowledge representation, planning, formal verification or trustworthy artificial intelligence.

Evaluation partner: experienced in designing tests for artificial intelligence. Pilot partner: in regulatory compliance, cybersecurity or public decision support, providing realistic workflows.

Ethics and legal partner: expert in the European Artificial Intelligence Act.

Optional infrastructure partner: computing for tests on European open-weight models and local devices.

Companies of any size and research bodies may apply; one organisation may cover several roles.

Current artificial intelligence agents can use software tools and give fluent answers. They often fail when a task needs reasoning across different fields, planning that adapts when conditions change, or a clear explanation of why an action is justified. They also rarely refuse to act when information or rules are insufficient. This limits their use in regulated sectors such as compliance and public administration.

The project will design a cognitive system that splits these abilities into three layers that can be tested separately. The first layer organises what the system can do as linked concepts, so that skills transfer between fields. The second layer builds step-by-step plans with alternatives and revises them when conditions change. The third layer checks each action before and after it is carried out, and refuses and explains when evidence or rules are insufficient. A coupling mechanism joins the layers, and a curated memory of past experience supports planning. Five research hypotheses will be tested.

The system builds on existing open-source software for controlled execution and logging, so the funding is used for new research. It is designed to work with different language models, including European open-weight (openly available) models and small models running on local devices. Results will be tested on an open benchmark and shown in a laboratory demonstrator for regulated decision support, with cybersecurity incident triage (sorting security alerts by urgency) as a second case. Open code, data and a reproducibility package will be published. The project will not create a new foundation model (a very large general-purpose model), fully verify arbitrary text or operate on live critical infrastructure.

The call uses single-stage submission and a lump-sum grant of up to 4,000,000 euro per project. It expects an integrated system demonstrated at technology readiness level 4 (technology validated in the laboratory) and joint activities with the other funded projects, such as a shared benchmark. Partners join as full beneficiaries under a research cooperation agreement and share work packages, results and publications.

Timescales: expressions of interest by 19 October 2026; call deadline 28 October 2026, 17:00 Brussels time; project duration about 42 months.

Advantages and innovations

Society: the system is designed to support decisions in regulated sectors by explaining its actions and refusing to act when evidence is lacking. This reduces the risk of wrong automated advice and supports the European Artificial Intelligence Act. Human oversight is part of the design.

Market: European organisations gain agent technology that is not tied to a single provider and can run on European or local models. The open benchmark and shared technical formats can be reused by other projects.

Environment: short, efficient execution is a design goal, with the aim of lowering the computing needed per task. This will be measured during the project.

Innovation: current agents rely on one model that reasons and acts without visible structure. Here, abstraction, planning and verification are separate layers, each tested against a stated hypothesis, and refusal is treated as correct behaviour. A curated memory of past experience is tested as a way to improve plans before they are carried out.

In numbers: up to 4,000,000 euro of funding, about 42 months, five research hypotheses, one open benchmark and one demonstrator at technology readiness level 4.

Technical specification or expertise sought

Knowledge representation: published work on knowledge graphs (networks of linked facts), ontologies or analogical reasoning, with experience linking them to machine learning.

Planning: hierarchical task network planning (breaking goals into ordered sub-tasks), planning under uncertainty, or combinations of language models and planners, tested in changing environments.

Formal verification: model checking (automatic proof that a system respects stated rules), policy or contract verification, and awareness of European regulation.

Evaluation: benchmark design, statistical evaluation and open-data practice.

Pilot: anonymised or realistic workflows, and willingness to define pilot success measures jointly.

Ethics and legal: European Artificial Intelligence Act, data protection and societal impact assessment. Optional: computing infrastructure for tests on European open-weight models and local devices.

Stage of development

Under development

IPR Status

No IPR applied

IPR Notes

Sustainable Development goals

- **Goal 9: Industry, Innovation and Infrastructure**
- **Goal 16: Peace and Justice Strong Institutions**

IPR Notes

Partner Sought

Expected role of the partner

Research partners lead the scientific work packages: concept graphs, planning and verification. The evaluation partner co-leads the open benchmark and its measures. The pilot partner hosts the laboratory demonstrator and helps define what counts as a correct refusal. The ethics and legal partner leads the ethics work package and co-writes the societal impact statement. An optional infrastructure partner runs tests on European models and local devices.

Type of partnership

Research and development cooperation agreement

Type and size of the partner

- **University**
- **Big company**
- **SME 50 - 249**
- **Other**
- **R&D Institution**
- **SME <=10**
- **SME 11-49**

Call Details

Framework program

Horizon Europe

Call title and identifier

DeepRAP: Deep Reasoning, Abstraction & Planning towards trustworthy Cognitive AI Systems (HORIZON-EIC-2026-PATHFINDERCHALLENGES-01-03)

Submission and evaluation scheme

Single-Stage

Anticipated project budget

Coordinator required

4.000.000

No

Deadline for EoI

15 Oct 2026

Deadline of the call

28 Oct 2026

Project duration in weeks

182

Web link to the call

<https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/HORIZON-EIC-2026-PATHFINDERCHALLENGES-01-03>

Project title and acronym

ACRA – Accountable Cognitive Reasoning Architecture (working title)

Dissemination

Technology keywords

- **01003001 - Advanced Systems Architecture**
- **01003003 - Artificial Intelligence (AI)**
- **01003006 - Computer Software**

Targeted countries

- **All countries**

Market keywords

- **02007016 - Artificial intelligence related software**

Sector groups involved

- **Digital**