EU-LAC Supercomputing Network for AI

Part of the Team Europe Initiative

“EU-LAC Digital Alliance”

Call for proposals

Call for LAC researchers for joint use cases:
EUROHPC-EULAC4UC-2027

Version 3.0
November 2026



EuroHPC Joint Undertaking

CALL FOR PROPOSALS

Contents

0. Introduction 4

1. Background 4

2. Use Case Objectives — Selection Criteria — Collaboration Expectations — Expected Outputs 6

EUROHPC-EULAC4UC-2027-1 — Topic 1 (use case 1) Unified methodology for HPC/AI tracing, energy-aware performance monitoring, and engagement of LAC HPC centres 6

Use Case Overview 6

Background 6

Selection Criteria 7

Expected Outputs 9

Collaboration Expectations 9

EUROHPC-EULAC4UC-2027-2 — Topic 2 (use case 2) Large Language Models adapted to Latin American language varieties and the healthcare domain 10

Use Case Overview 10

Background 10

Selection Criteria 11

Expected Outputs 13

Collaboration Expectations 14

EUROHPC-EULAC4UC-2027-3 — Topic 3 (use case 3) High-Resolution Simulation of Earthquakes and Tsunamis for Risk Assessment 14

Use Case Overview 15

Background 15

Selection Criteria 16

Expected Outputs 17

Collaboration Expectations 17

EUROHPC-EULAC4UC-2027-4 — Topic 4 (use case 4) Impact of Climate Change on Coffee Production 18

Use Case Overview 18

Background 18

Selection Criteria 19

Expected Outputs 21

Collaboration Expectations 21

3. Available budget 21

4. Timetable and deadlines 22

5. Admissibility and documents 22

6. Eligibility 23

Eligible participants 23

Consortium Composition 23

Eligible actions and activities....23
Geographic location (target countries)....24
Duration....24
Project budget....24
7. Evaluation and award procedure....24
8. Award criteria....25
9. Set-up of the Funding Grant Agreements....25
10. How to submit an application....26
11. Help....27

0. Introduction

This is a call for proposals under the **EU-LAC Supercomputing Network for AI - Part of the Team Europe Initiative "EU-LAC Digital Alliance"**.

The call covers the following **topics**:

- **EUROHPC-EULAC4UC-2027-1 — Topic 1 (use case 1) Unified methodology for HPC/AI tracing, energy-aware performance monitoring, and engagement of LAC HPC centres**
- **EUROHPC-EULAC4UC-2027-2 — Topic 2 (use case 2) Large Language Models adapted to Latin American language varieties and the healthcare domain**
- **EUROHPC-EULAC4UC-2027-3 — Topic 3 (use case 3) High-Resolution Simulation of Earthquakes and Tsunamis for Risk Assessment**
- **EUROHPC-EULAC4UC-2027-4 — Topic 4 (use case 4) Impact of Climate Change on Coffee Production**
- Each project application under the call must address only one of these topics. Applicants wishing to apply for more than one topic, must submit a separate proposal under each topic.

We invite you to read the **call documentation** carefully.

The document provide clarifications and answers to questions you may have when preparing your application:

- the Call document outlines the:
  - background, objectives, scope, activities that can be funded and the expected impact (sections 1 and 2)
  - timetable and available budget (sections 3 and 4)
  - admissibility and eligibility conditions (including mandatory documents; sections 5 and 6)
  - evaluation and award procedure (section 7)
  - award criteria (section 8)
  - set-up of the Funding Grant Agreements (section 9)
  - how to submit an application (section 10)

1. Background

High-performance computing (HPC) is a key driver of scientific and technological advancement and is increasingly recognized as essential for research and development (R&D) across most scientific disciplines. With the rapid advancement of artificial intelligence (AI) and its growing integration into virtually all sectors that rely on digital technologies, access to advanced computational capacity is becoming a strategic necessity, not only for scientific research and innovation but also for economic development, productivity, and competitiveness. Countries lacking sufficient computational capacity risk becoming increasingly dependent on technologies and

solutions developed elsewhere, limiting their ability to address local challenges, strengthen domestic innovation ecosystems, and fully benefit from the opportunities offered by the digital and AI-driven economy. Consequently, HPC is gaining prominence on national and international policy agendas. However, in Latin America and the Caribbean (LAC), investment in HPC remains highly uneven across countries, with several countries investing significantly less than the United States and Europe. The development and expansion of HPC capabilities in the region are shaped by several challenges and opportunities, including the following:

- **Limited access to HPC infrastructure**: High-performance computing systems in LAC are distributed across public research institutions, universities, government laboratories, and private companies. While some of the largest systems are operated by private actors, many of the region's most important HPC resources are publicly funded and primarily serve academic and research communities. Nonetheless, overall HPC capacity in the region remains limited, with fragmented access that often requires reliance on shared national facilities or international collaborations. This can slow down research and innovation, and limit scalability of services and business models. Strengthening HPC in LAC therefore requires sustained public investment, supported by public-private partnerships for acquisition, operation, and long-term maintenance.
- **Opportunities for coordinated investment and shared infrastructure:** The European experience, notably through the EuroHPC Joint Undertaking and shared supercomputing infrastructures such as LUMI, demonstrates how coordinated investments can help pool financial and technical resources, achieve economies of scale, and expand access to advanced computing capabilities beyond what individual countries might be able to achieve independently. In Latin America and the Caribbean, HPC investments remain largely driven by national initiatives, with varying levels of resources and coordination across countries. Building on existing national and regional efforts, greater cooperation and investment in shared HPC infrastructure could offer opportunities to strengthen research and innovation ecosystems, broaden access to computational resources, and support scientific, technological, and economic development across the region.
- **Human resources**: A prerequisite to the successful uptake of HPC in the public and private sectors are skills and competences. Investments in skilled technical researchers, engineers, and technicians to develop and use HPC applications are critical to achieve the intended return of investment in HPC infrastructure.

The establishment of the EU–LAC Supercomputing Network for AI will help address these challenges and capitalize on emerging opportunities by fostering regional HPC integration, enhancing interconnectivity through infrastructure such as the BELLA programme's EllaLink cable, and building on existing political and technical momentum under the EU-LAC Digital Alliance. To ensure maximum impact and technical feasibility during this phase, the project initially focuses on engagement with seven specific countries that already possess sufficient computing infrastructure and active research communities for immediate and effective participation in the actions. This focused strategy will speed up the digital transformation of the targeted region while creating new opportunities for scientific collaboration, industrial innovation, and economic growth across both continents.

The expected outcomes of the action are the development of joint use cases that showcase the practical applications and societal benefits of EU-LAC collaboration in High-Performance Computing (HPC) and Artificial Intelligence (AI). These use cases will address a wide range of challenges, including climate change, natural hazard prediction, large language models, and health-related applications leveraging HPC infrastructures, AI methods, and large-scale datasets.

The use of Copernicus publicly available data will be encouraged through collaboration with the Regional Copernicus Centres established in the region. This collaboration will leverage the use of earth observation data and services to support the development of innovative AI applications and solutions.

Each use case will generate open scientific outputs, including datasets, models, software tools, and technical guidelines. These results will be made available to the wider research community through open repositories and collaborative platforms.

To further broaden participation and maximise impact, the project is launching this competitive open call inviting entities from Latin America and the Caribbean (LAC) to join as collaborating sites. Selected organisations will contribute to the development and implementation of the joint use cases by engaging local researchers, engineers, and technical staff to work alongside European partners. They will provide regional expertise, local knowledge, and access to relevant datasets and infrastructures, thereby strengthening the bi-regional collaboration and enhancing the relevance and reach of the project's outcomes.

2. Use Case Objectives — Selection Criteria — Collaboration Expectations — Expected Outputs

EUROHPC-EULAC4UC-2027-1 — Topic 1 (use case 1) Unified methodology for HPC/AI tracing, energy-aware performance monitoring, and engagement of LAC HPC centres

Use Case Overview

Use Case ID	UC1
Title	Unified methodology for HPC/AI tracing, energy-aware performance monitoring, and engagement of LAC HPC centres
Lead Contact	INRIA
Short Description	UC1 develops a common, open methodology and toolkit to collect, store, anonymise and pre-process traces of running jobs on heterogeneous HPC and AI systems, covering energy and power telemetry, I/O behaviour and resource utilisation. The toolkit includes dashboards that show performance problems and energy consumption per job and per workflow, and it is used to instrument the executions of the other use cases. Selected LAC HPC centres join UC1 as collaborating sites: together with the EU partners they shape the trace format and the collection and anonymisation process, deploy the toolkit on their own systems, and contribute anonymised traces to a public repository. For the region, the outcome is local capacity in energy-aware operation of HPC and AI systems, a shared evidence base that includes LAC systems.

Background

Energy has become one of the main concerns in operating HPC systems. As machines grow more powerful, they also draw more power, and electricity has become one of the

largest costs of running them. At the same time, AI is pushing consumption further, with workloads that keep large numbers of accelerators busy for long periods. Reducing the energy consumed by applications, and by AI applications in particular, has therefore become necessary.

Any effort to reduce energy consumption starts with measuring it reliably. Energy traces on their own, however, do not tell the whole story: they show how much a job consumed, but not why. For that, they have to be combined with system-level traces, such as CPU, GPU and memory utilisation, and with application-level traces, such as the job's I/O and compute phases. Together, these traces make it possible to pinpoint how a job behaves and to work towards improving it.

This is the objective of UC1: to develop an open methodology and toolkit to collect, store and pre-process these traces, and to use them to understand and improve the behaviour of jobs.

This call aims to bring together LAC HPC centres interested in taking part in this effort. The plan is to develop the methodology together: the participating centres work with the EU partners on the trace format and on how traces are collected and anonymised, deploy the toolkit on their own systems, and contribute anonymised traces to a shared public repository. Their staff also work alongside the EU partners on the other joint use cases, whose executions the toolkit will instrument.

Because the toolkit runs on the centres' own systems, taking part requires the cooperation of their system administrators, since some measurements need privileged access. The traces also contain information about users and workloads, so raw traces stay at the centre: only anonymised data is published, under a governance framework aligned with GDPR and equivalent national regulations and defined together with the participating centres.

Selection Criteria

This call aims to select one consortium composed of one to five LAC HPC centres, which together form the network of collaborating sites for UC1. The EU partners work directly with the technical staff of every member, and each member is responsible for instrumenting its own systems. Participants may be located in the same country. If a proposal is selected for funding, only those participating entities receiving financial support will be required to sign a grant agreement and will receive funding directly. Since the funding is expected to support only a limited level of personnel effort, applicants should clearly justify the role, responsibilities and expected contribution of each participating centre, as well as the allocation of the requested budget among consortium members.

Technical Skills and Experience

Mandatory

- Operational responsibility for at least one HPC or AI system in production use by a research user community, managed through a resource manager or batch scheduler.
    - description of the system (architecture, size, scheduler, user community).
- Experience operating monitoring, telemetry or job-accounting infrastructure on that system.
    - short description of the tools and data currently in use.
- Named technical staff assigned to UC1 for the funded period, with the role, profile and effort of each person.
    - names and CVs.
- Ability to install and maintain third-party software on the target system,

including the UC1 toolkit.
    - description of the centre's software deployment practice (for example Spack, EasyBuild, Guix, containers or environment modules).

Desirable (Non-Mandatory)

- Experience with performance tracing or profiling tools (for example Darshan, Score-P, Extrae, TAU or PAPI).
- Experience collecting energy or power telemetry at node, rack or facility level (for example RAPL, IPMI, NVML, PDU or BMC readings).
- Experience with parallel file systems and I/O characterisation.
- Experience handling research data under data protection rules, including anonymisation or data governance.
- Local expertise in AI or data-intensive workloads running on the centre's systems.

Regional Relevance

Desirable (Non-Mandatory)

- The proposed system serves a national, regional or multi-institutional user community rather than a single research group.
- Established links with the national research and education network, or with regional bodies such as SCALAC and RedCLARA.
- Institutional interest in energy-aware operation, for example an existing efficiency or sustainability initiative at the centre.

Collaboration Capacity

Mandatory

- A letter of support signed by the direction of the member centre, confirming institutional commitment to the activities described in the application and to the assignment of the named staff.
- A written statement from the system administration team confirming that the access required for instrumentation can be granted, and describing any constraints that apply. Applicants should state limitations plainly rather than overstate what is feasible: adapting to different system realities (different types of available information, lack of privileges to run certain codes, etc.) is one of the challenges tackled by the project, thus these constraints may actually add diversity and enrich the work.
- A commitment to take part in the decision-making discussions for defining the trace format and the plan for collection and anonymisation together with the EU partners and the other participants.
- A commitment to contribute anonymised traces from the instrumented systems to the public repository, under the governance framework developed jointly within UC1. Within that framework, each centre decides the scope of what it publishes.
- Availability of the named staff for regular technical meetings with the EU partners and the other member centres.

Access to infrastructure and data

Mandatory

- Access, for the duration of the project, to at least one production HPC or AI

system on which the UC1 collectors can be installed and run.

Desirable

- Energy measurements beyond the node, at rack, PDU, facility or cooling level.
- The possibility of collecting traces from real production jobs, not only test runs, subject to the centre's own data protection rules.

Consortium-level criteria

Mandatory

- Between one and five eligible participant centres.
- Complementarity: together, the members' systems cover a range of architectures, schedulers and storage systems.
- A coherent joint plan: a clear division of roles and effort between participants, and a realistic schedule aligned with the UC1 milestones.

Expected Outputs

Output Type	Specific Requirements	Open Access Repository
Datasets	Anonymised traces of production jobs from each member's instrumented systems, following the UC1 trace schema and including a description of the system; released under an open licence agreed in the governance framework (for example CC BY 4.0).	UC1 public trace repository
Software	Deployment configurations for each member's systems, and any extensions or fixes to the UC1 toolkit developed during the project, released under the toolkit's open-source licence	UC1 code repository
Guidelines/Reports	One site report per member: what was instrumented, access constraints, anonymisation choices and lessons learned. These reports feed the methodology document.	UC1 public repository

Collaboration Expectations

Role of the European Partner

INRIA leads UC1, with contributions from BSC and CINECA. The European partners:

- lead the design of the methodology, trace schema and anonymisation process, jointly with the member centres;
- develop and maintain the UC1 toolkit and dashboards, and provide deployment packages and documentation;
- support deployment on the members' systems through onboarding sessions

and regular technical meetings;
- integrate the members' traces into the public repository and credit the contributing centres as authors of the datasets;
- connect member staff with the teams of UC2, UC3 and UC4.

Role of the LAC Applicant

Each participant centre:

- takes part in defining the trace format and the collection and anonymisation process;
- deploys and operates the UC1 toolkit on its systems, and secures the access this requires;
- collects, anonymises and contributes traces, and decides within the governance framework what it publishes;
- assigns the named staff, who attend regular technical meetings and report issues with the toolkit;
- writes its site report;
- where appropriate, hosts use-case components that benefit from regional proximity, such as local data staging or validation experiments.

EUROHPC-EULAC4UC-2027-2 — Topic 2 (use case 2) Large Language Models adapted to Latin American language varieties and the healthcare domain

Use Case Overview

Use Case ID	UC2
Title	Large Language Models adapted to Latin American language varieties and the healthcare domain
Lead Contact	BSC
Short Description	Adaptation of open-source language models to the language varieties of Portuguese and Spanish spoken in Latin America and specifically to an application scenario in the healthcare domain.

Background

This use case will address the matter of linguistic representation in the LAC region, an important structural dimension of digital transformation and technological sovereignty in AI systems. Spanish and Portuguese are among the most spoken languages globally, and most speakers are in Latin America and the Caribbean, with over 400 million speakers of several regional varieties of Spanish, and over 200 million speakers of Brazilian Portuguese. However, most large-scale foundation language models are trained predominantly on English-centric datasets, and the majority also include a heterogenous mix of unlabeled Spanish dialects, which leads them to generate a synthetic amalgam that blends features from multiple varieties into an artificial "average" Spanish that no real community speaks. As a result, AI systems deployed in

Latin America frequently exhibit reduced accuracy and fluency, contextual misinterpretation and limited domain relevance. In parallel, there are over 400 Indigenous languages spoken by millions of people in the region, and those remain largely absent from digital AI infrastructures, contributing to their linguistic exclusion in emerging AI-powered services.

This use case aims to develop a language model adapted to the regional varieties of Spanish and Portuguese spoken in Latin America and the Caribbean, leveraging locally sourced and region-specific data and building on an open reference model. Rather than training a foundational model from scratch, the work will leverage this existing open architecture and apply rigorous refinement and alignment techniques. The primary objective is to maximize the usability, safety, and contextual accuracy of these models for the diverse linguistic ecosystem of the LAC region.

To broaden the practical usability and the socio-economic impact of this use case, the model will also be adapted to a specific application scenario in the healthcare domain. This application will be defined alongside local institutions according to their perceived needs and thoughts on what would be most useful for the region; the only requirement is that the application should rely on a text-based, instruction-tuned language model adapted to the language varieties of Latin America that would then be further adapted to the restricted domain of the application.

The impact of the use case is thus twofold,

- by generating both a general-purpose, regionally adapted model (UC2.1);
- and a domain-specific (healthcare domain) model used as building block for an useful tool (UC2.2).

The availability of open, regionally adapted LLMs is expected to improve linguistic accuracy and contextual performance in AI applications deployed across Latin America. Enhanced representation of regional language varieties will support more reliable AI-assisted public administration, education, legal documentation, agricultural advisory systems and citizen-facing digital services. The version adapted to the healthcare domain will likewise be uniquely positioned to empower biomedical research and clinical applications based on local language varieties. By selecting LAC institutions to participate through this competitive open call, the project unlocks opportunities to address major clinical challenges through continuous collaboration between the BSC and local experts. Potential avenues of exploration include the adaptation of models to regional medical terminology, which varies significantly across different countries, as well as the translation of clinical protocols and health campaigns into Indigenous languages to combat linguistic exclusion in public health. Given the scale of digital public services and the rapid growth of AI adoption in the region, even incremental improvements in model accuracy and contextual alignment can have wide-reaching effects on usability, trust and service efficiency.

Selection Criteria

Two separate projects will be selected, and the selection process will consider their complementarity and compatibility, particularly in relation to the two use cases (please state UC2.1 or UC2.2 explicitly in your application):

- UC2.1: General-purpose adaptation of an open language model to Latin American varieties of Spanish and Portuguese, and potentially also to Indigenous languages of the region;
- UC2.2: Domain-specific adaptation of the general purpose model (from UC2.1) to serve an application scenario in the healthcare domain.

At least one of the selected projects must involve an application tailored to the healthcare domain. Therefore, the selection committee will particularly value applicants' relevant experience in healthcare and their association or collaboration with healthcare institutions.

The selection process will also consider the applicants' alignment with the principles of open science and the EU Artificial Intelligence Act, including their commitment to using open models, open-source frameworks, and openly licensed datasets, or resources made available by partners with appropriate permissions for the development of an open-source language model.

The feasibility of the proposed work and its compatibility with the technical objectives of the project will also be considered, including the suitability of available resources for adapting a base model to be jointly selected with the BSC team.

Technical Skills and Experience

Mandatory

Staff assigned to either UC2.1 or UC2.2 must have the following technical skills and experience:

- Intermediate to advanced programming skills in Python.
- Basic Linux skills.

Specific for UC2.1:

- Experience with end-to-end multilingual language modeling (e.g. pretraining, data collection, evaluation).
- Experience with methods and pipelines for language model instruction-tuning (Supervised Fine-Tuning, SFT) and alignment (Reinforcement Learning, RL).
- Experience with methods and algorithms for supervised fine-tuning (SFT) and reinforcement learning (RL) for language models.
- Experience in machine learning and deep learning techniques and Python libraries such as TensorFlow, NumPy, Pandas.
- Knowledge of version control tools (Git and GitHub/GitLab).

Specific for UC2.2:

- Experience with the usage of text-based language models for healthcare-related purposes.
- Thorough familiarity with the specific sub-domain that the project proposal falls in and understanding of the benefits and risks that such an application may pose.

Desirable (Non-Mandatory)

- Experience with post-training frameworks such as NeMo-RL, Slime, TRL, or similars.
- Experience working in HPC clusters with Slurm queues.
- Knowledge of GPU-based computing, including multi-GPU and multi-node parallelization techniques.

Specific for UC2.2:

- Experience with the anonymisation and sharing of healthcare data that may contain Personally Identifiable Information (PII).

Regional Relevance

Mandatory

Specific for UC2.2:

- Association with a research institution or group that has experience with AI applications in the healthcare domain.

Collaboration capacity

- Letters of support from any other institutions that would be involved in the project under the condition of not receiving any financial support (e.g. private companies, clinics or hospitals).

Access to Infrastructure/Datasets

Desirable

Specific for UC2.1:

- Access to local data in Brazilian Portuguese and/or Latin American variants of Spanish that could be used for general-purpose instruction-tuning or a proven strategy to obtain or generate such data.

Specific for UC2.2:

- Access to local data in the healthcare domain.
- If the data contains Personally Identifiable Information (PII), an established strategy for anonymisation would also be appreciated.
- Access to local compute resources for testing and executing smaller workflows.

Expected Outputs

Output Type	**Specific Requirements**	**Open Access Repository**
Model	Instructed and aligned language model for Portuguese and Spanish varieties spoken in Latin America.	HuggingFace
Technical report	Technical report detailing the general-purpose adaptation process and the adaptation to a domain-specific healthcare application.	ArXiv
Conference publication	Potential conference paper on any part of the adaptation process that might be worth a publication.	Repository corresponding to the chosen conference

Datasets	Partial training corpus (Apache 2.0 license).	HuggingFace
Benchmarks	Evaluation data.	HuggingFace
Training code	Framework(s) used for training.	GitHub
Evaluation code	Framework(s) used for evaluation.	GitHub

Collaboration Expectations

Role of the European Partner

This use case is led by the BSC, specifically by the COPLA team (Collaboration Projects with Latin America) in the AI Institute, and its execution is also under the supervision of INRIA.

The "selected partners" henceforth refer to the partners that will be selected through this call, and "the team" henceforth refers to the existing BSC team as well as the partners that will be selected through this call, i.e. all of the staff who will be working on the project together.

The BSC team:

- Coordinates the execution of the project.
- Reports on the progress of the development of the use case.
- Establishes regular meetings (weekly or bi-weekly) with the selected partners to coordinate on the development tasks and follow up on their progress.
- Regularly informs the selected partners of any project updates and events.

Role of the LAC Applicant

The LAC applicants are expected to propose a project that they would like to carry out within the scope described in the call, and cite all evidence that backs the claims made throughout the proposal, e.g. any recent publications in similar fields, letters of support, and CVs of the staff that would be assigned to work on the project.

The selected partners will be expected to:

- Aid in the definition of the work plan alongside the rest of the team.
- Collaborate with the team on the development of the models.
- Attend regular online meetings with the team to align on the tasks and provide updates on their progress.
- Identify and inform the BSC team of any potential issues that may hinder the execution of the project or the timely deliverable of the expected resources at the end of the project.

EUROHPC-EULAC4UC-2027-3 — Topic 3 (use case 3) High-Resolution Simulation of Earthquakes and Tsunamis for Risk Assessment

Use Case Overview

Use Case ID	UC3
Title	High-Resolution Simulation of Earthquakes and Tsunamis for Risk Assessment
Lead Contact	Barcelona Supercomputing Center (BSC), Wave Phenomena Group, CASE Department
Short Description	This use case will develop HPC-enabled, physics-based, multi-hazard simulation models using earthquake rupture, ground motion, and tsunami generation/inundation. These simulations will aid to generate surrogate models in the region. Its goal is to move beyond current hazard assessment toward high-resolution, multi-scale simulation that quantifies compound earthquake-tsunami risk to coastal communities- and infrastructure-along dense coastal municipalities. It directly benefits the LAC region by transferring BSC-developed HPC/AI workflows to the partner institution, building local capacity in high-performance computational seismology and tsunami modelling, and delivering decision-support tools for civil protection, land-use planning, and infrastructure retrofitting.

Background

Earthquake-prone regions, particularly those located along active subduction zones and crustal fault systems, are exposed to significant seismic and tsunami hazards. These tectonic environments have produced destructive megathrust and crustal earthquakes, including tsunami-generating events, with potentially severe consequences for coastal communities and critical infrastructure. Many coastal areas are characterized by complex geological and geotechnical conditions, including soft sedimentary basins, liquefaction-prone soils, and high groundwater tables, which can amplify seismic impacts. The concentration of population, tourism, ports, and essential infrastructure in these areas further increases their vulnerability to cascading hazards associated with strong ground shaking, landslides, coastal deformation, and tsunami inundation. Historical earthquakes across Latin America and the Caribbean (LAC) and other seismically active regions provide valuable opportunities for validating advanced simulation workflows under diverse tectonic and environmental conditions.

Recent advances in computational geophysics enable coupled, physics-based simulation of earthquake rupture, seismic wave propagation, seafloor deformation, and tsunami generation within unified HPC/GPU-accelerated frameworks, increasingly combined with AI/ML techniques to accelerate simulation, interpolate sparse geotechnical data, and quantify uncertainty. This use case will apply BSC's UCIS4EQ (HPC physics simulation coupled with AI/ML surrogate models, MLESmaps) and the Unified CyberShake Workflow (UnifiedCSWF, originally developed for the Southern California Earthquake Center and already adapted to Iceland, southeastern Spain, and northeastern Italy) to validate against real earthquakes and generate synthetic megathrust and crustal rupture scenarios, producing high-resolution multi-hazard maps and infrastructure vulnerability assessments for coastal risk management. Potential validation events include the 2012 Mw 7.6 Nicoya, 2009 Mw 6.1 Cinchona, and 1991 Mw 7.7 Limón earthquakes, as well as other well-documented events from LAC.

Selection Criteria

Technical Skills and Experience

Mandatory

- Experience in earthquake rupture modelling, seismic wave propagation simulations, and/or tsunami modelling and inundation.
- Experience developing or applying machine-learning/AI models to accelerate or approximate physics-based seismic hazard simulations.
- Programming skills in Python for processing geophysical simulation data, including large meshes and NetCDF/HDF5 datasets.
- Experience running scientific simulations on CPU/GPU-based HPC systems or clusters, including job scheduling and performance monitoring.

Desirable (Non-Mandatory)

- Familiarity with geological, geophysical, and seismic datasets from earthquake-prone regions, or demonstrated ability to acquire, process, and integrate relevant datasets for workflow validation.
- Experience with seismic simulation software, CyberShake-type frameworks, or HPC workflows for physics-based seismic hazard analysis.
- Knowledge of infrastructure vulnerability, fragility curves, and risk assessment for critical infrastructure, such as ports, hospitals, and transport networks.
- Experience using Earth observation and remote sensing data (e.g., Copernicus Sentinel products and Global Earthquake Model datasets) for exposure mapping, coastal change analysis, or ground-deformation monitoring.
- Experience in scientific visualization and reproducible workflows using tools such as Jupyter, containers, and workflow managers.

Regional Relevance

Mandatory

- Demonstrated institutional mandate or track record in seismic/tsunami hazard monitoring, disaster risk reduction, or coastal infrastructure planning

Desirable

- Direct professional or research experience in the Nicoya Peninsula, Cinchona, or other central Pacific/Caribbean coastal areas targeted by this use case.

Collaboration capacity

- Letter(s) of support or a data/collaboration agreement from the hosting institution and, where applicable, from the national disaster management agency confirming access to seismic/geotechnical datasets and willingness to co-develop rupture scenarios and validate model outputs.

Access to Infrastructure/Datasets

Mandatory

- Possibility to obtain institutional access to local geotechnical borehole data (or tomographies), fault databases, and digital elevation/bathymetric models for the selected coastal areas. The applicant will be expected to use BSC HPC resources (MareNostrum) under the project's computing allocation and, where

needed, prepare and submit applications for additional computing resources, e. g. through RES (Red Española de Supercomputación)..

Desirable

- Experience in using local/regional HPC or GPU computing resources for pre-processing and preliminary testing; access to Copernicus Earth observation products or in-country seismic/tsunami monitoring networks (e.g., OVSICORI, RSN).

Expected Outputs

Output Type	Specific Requirements	Open Access Repository
Datasets	Curated 3D velocity/site-effect models, fault geometry database, and validated rupture scenario catalogue for e.g. Nicoya and Cinchona.	Zenodo, EPOS/ORFEUS seismic data portal
Models/Software	Regional configuration of the UCIS4EQ/UnifiedCSWF coupled earthquake-tsunami workflow, including ML surrogate (MLESmaps) site-effect interpolation module	GitHub/GitLab (BSC organisation)
Guidelines/Reports	Technical report on compound earthquake-tsunami impacts and infrastructure	Project website, Zenodo
Other	High-resolution multi-hazard maps (ground shaking, liquefaction potential, tsunami inundation) and infrastructure damage/exposure assessment for selected coastal cities	Project website / national disaster-management agency GIS repository

Collaboration Expectations

Role of the European Partner

BSC (Wave Phenomena Group, CASE Department) will provide access to the workflows used for seismic hazard assessment, computing allocation on MareNostrum via applications for HPC resources, technical mentoring on physics-based rupture and tsunami simulation, and co-supervision of scenario design, validation, and uncertainty quantification. BSC will also support integration of AI/ML surrogate modelling (MLESmaps) and co-author resulting guidelines and publications.

Role of the LAC Applicant

The LAC applicant will provide, or prepare datasets of high resolution topo-bathymetry, local geology and seismotectonic data; 3D seismic velocity models, geotechnical/site effects, and co-develop, simulate and validate earthquake rupture and tsunami

scenarios against validation use cases such as the Nicoya and Cinchona validation events; contribute to interpreting simulation outputs. Co-develop simulations and surrogate models for hazard/risk assessment.

EUROHPC-EULAC4UC-2027-4 — Topic 4 (use case 4) Impact of Climate Change on Coffee Production

Use Case Overview

Use Case ID	UC4
Title	Impact of Climate Change on Coffee Production
Lead Contact	CINECA
Short Description	The UC will exploit HPC resources and AI techniques to generate high-resolution, regionally consistent climate information for specific coffee-growing areas, generating scenario-based spatial risk maps and a publicly accessible web application for stakeholders. The UC ambition is to develop a comprehensive coffee production model integrating additional non-climatic factors, such as soil characteristics, crop management practices, pests, logistics, and socio-economic variables, depending on data availability and accessibility. LAC researchers will provide regional expertise in coffee agronomy and climate-crop interactions.

Background

The Use Case focuses on assessing the impact of climate change on coffee production, addressing a challenge of significant societal and economic relevance in major coffee-producing regions worldwide, particularly in Latin America and the Caribbean (LAC).

Coffee production plays a strategic role in agricultural economies, contributing substantially to export revenues, rural employment, and the livelihoods of millions of people across farming, processing, logistics, and trade. Many coffee-producing regions are highly vulnerable to climate variability and long-term climate change, with potential consequences for agricultural productivity, crop quality, economic stability, and food-system resilience.

Coffee cultivation is particularly sensitive to changes in temperature, precipitation patterns, and the frequency and intensity of extreme weather events. Even relatively small increases in average temperature (e.g., 1–2°C) can significantly affect yields, crop quality, geographical suitability, and the prevalence of pests and diseases, depending on local growing conditions and coffee varieties. Improving the accuracy and spatial resolution of climate-risk assessments and yield predictions could therefore provide substantial economic benefits, supporting agricultural planning, adaptation strategies, and more resilient coffee value chains.

Despite the strategic importance of coffee production, significant limitations remain in the availability and integration of high-resolution, AI/HPC-enabled climate-risk assessment tools tailored to the specific characteristics of coffee-growing regions.

Existing approaches often lack the spatial detail, regional consistency, or integration of local agronomic observations required to reliably assess climate impacts on coffee productivity.

The gap to be addressed is twofold: (i) the limited availability of fine-scale, regionally consistent climate projections suitable for coffee-growing areas, and (ii) the insufficient integration of crop-impact models with observed agronomic, phenological, and yield data across different production environments.

This use case combining climate modelling, AI-based downscaling, and agricultural impact assessment will directly contribute to the objectives of strengthening EU-LAC cooperation in HPC- and AI-enabled applications.

Starting from reanalysis datasets (e.g. ERA5), AI-based downscaling models will be trained to produce fine-scale climate and weather (seasonal and sub-seasonal) indicators, enabling more accurate assessment of climate risks, yield variability, and adaptation strategies. These indicators will be translated into scenario-based spatial maps to support evidence-based decision-making for producers, cooperatives, and industrial stakeholders. The results will be made accessible through a dedicated web application, allowing stakeholders to explore climate scenarios and related impacts on coffee production.

While the initial focus of the use case is on climate impacts, the longer-term ambition is to evolve towards a more comprehensive coffee production model integrating additional factors, such as soil characteristics, crop management practices, pests, logistics, and socio-economic variables, depending on data availability and accessibility.

The use case will benefit from the collaboration of ILLYCAFFÉ, which will contribute domain expertise and connections with local producers and cooperatives. This collaboration will help maximize the practical relevance, validation, and potential impact of the results for the coffee sector.

Selection Criteria

Technical Skills and Experience

Mandatory

- Expertise in agronomy and/or agricultural science (crop physiology, yield modelling, pest/disease management, or related disciplines)
- Demonstrated capability in empirical yield–climate analysis
- Experience in geospatial data management (georeferenced datasets, spatial data formats)

Desirable (Non-Mandatory)

- Experience in working with global (e.g. ERA5) and regional climate or weather model outputs.
- Data quality assurance: ability to flag outliers, missing values, known data limitations, uncertainty level, etc.
- Access to, or established relationships with, coffee producers and cooperatives that can facilitate data collection and field validation activities.
- Experience in high-performance computing (HPC) environments
- Development of web-based interfaces for the exploitation and dissemination of results.

Regional Relevance

Desirable

- Direct professional or research experience in coffee-producing regions (e.g. Minas Gerais), including familiarity with local agricultural practices, environmental conditions, and climate-related challenges.

Collaboration capacity

Mandatory

- Track record of relevant national and international research collaborations

Desirable

- Collaborations with coffee producers/cooperatives.

Access to Data

Desirable

- Access to historical agronomic, production and environmental data:

  A · Meteorological & Climate Observations
  A1 Daily Temperature (min, max, ave)
  A2 Daily precipitation
  A3 Humidity
  A4 Mean sea level pressure
  B · Coffee specific data
  B1 · Coffee yield records
  Annual yield in kg/ha (cherry or green bean equivalent) by farm ID / cooperative block / municipality. ≥10-year series preferred.
  B2 · Coffee quality indicators
  Cup score (SCA protocol), physical defects (type & count), screen size, moisture content at harvest. Linked to farm ID and harvest year.
  B3 · Phenological observations
  Key BBCH-analogous stages: bud break, full flowering, grain-filling onset, physiological maturity / harvest date. Linked to farm, variety and year.
  C · Agronomic Practices
  C1 · Irrigation records
  Method (drip, sprinkler, pivot, furrow), annual volume applied (m^3/ha/yr), scheduling calendar, water source (river, groundwater, harvesting).
  C2 · Fertilisation records (N, P, K)
  Annual application rates in kg/ha/yr by nutrient; nitrogen source (urea, ammonium sulfate, etc.); micronutrient applications if available.
  C3 · Plant age, density & renovation cycles
  Planting year or mean plant age (yrs), planting density (plants/ha), pruning method (esqueletamento, recepa), renovation schedule.
  C4 · Coffee variety / cultivar
  Dominant cultivar per farm: Arabica (Catuaí, Mundo Novo, Bourbon, Catucaí…); Robusta (Vitória Incaper, Conillon Clone…). Required for phenological parameterisation.
  C5 · Pest & disease incidence

Incidence/severity records for coffee berry borer (broca, H. hampei), leaf rust (ferrugem, H. vastatrix), cercosporiosis, nematodes. Linked to farm and year.

D · Soil & Terrain Data

D1 · Soil physical properties

Texture (clay/silt/sand %), bulk density (g/cm³), georeferenced.

D2 · Soil chemical properties

pH (H_2O & $CaCl_2$), organic carbon (%),available P and K, Ca, Mg

For each data type, the following aspects should be reported in your application:

- Availability (free of charge, fee based, …)
- Temporal & Geographic Coverage (e.g. 2000–2026, state/municipality)
- Format & Sharing Modality (e.g. CSV, GeoJSON, shapefile)
- IP / Legal Limitations

Expected Outputs

Output Type	Specific Requirements	Open Access Repository
Guidelines/Reports	Technical report	GitHub, Zenodo
Other	High-resolution risk maps	GitHub, Zenodo

Collaboration Expectations

Role of the European Partner

CINECA team provides computational resources and scientific expertise in AI model training and data driven climate / weather forecasts downscaling.

Role of the LAC Applicant

LAC team provides expertise in coffee agronomy and climate-crop interactions.

3. Available budget

The available call budget is **EUR 250.000**.

Specific budget information per topic can be found in the table below:

Topic	Topic budget
1 — EUROHPC-EULAC4UC-2027-1	**EUR 50.000**
2 — EUROHPC-EULAC4UC-2027-2	**EUR 100.000**
3 — EUROHPC-EULAC4UC-2027-3	**EUR 50.000**
4 — EUROHPC-EULAC4UC-2027-4	**EUR 50.000**

We expect to fund 1 project in topic EUROHPC-EULAC4UC-2027-1, 2 projects in topic EUROHPC-EULAC4UC-2027-2, 1 project in topic EUROHPC-EULAC4UC-2027-3 and 1 project in topic EUROHPC-EULAC4UC-2027-4. We reserve the right not to award all available funds or to redistribute them between the call priorities, depending on the proposals received and the results of the evaluation.

4. Timetable and deadlines

Timetable and deadlines (indicative)	
Call opening:	23 November 2026
Deadline for submission:	25 January 2027 – 17:00:00 CET (Brussels)
Evaluation:	February 2027
Information on evaluation results:	March 2027
GA signature:	April 2027

5. Admissibility and documents

Proposals must be submitted before the **call deadline** *(see timetable section 4)*.

Proposals must be submitted **electronically** via the EuroHPC Peer-Review Platform, available at https://access.eurohpc-ju.europa.eu . Paper submissions are NOT possible.

Proposals (including annexes and supporting documents) must be submitted using the forms provided *inside* the Submission System (NOT the documents available on the Topic page — they are only for information).

Proposals must be **complete** and contain all the requested information and **all required annexes**:

- Application Form Part A — contains general information about the participants and is an online form. Application Form Part B — contains the technical description of the project and the summarised budget of the project.
- mandatory annexes :
    - CVs Europass format of core project team.
    - Additional annexes (e.g. letters of support), where relevant to the specific use case, should be included in Application Form Part B. These annexes will not count towards the page limit.

– ⚠ Please note that failing to provide all mandatory annexes in the requested format[1] by the proposal submission deadline may render your proposal inadmissible[2].

At proposal submission, you will have to confirm that you have the **mandate to act** for all applicants. Moreover, you will have to confirm that the information in the application is correct and complete and that all participants comply with the conditions for receiving EU funding *(especially eligibility, financial and operational capacity, exclusion, etc)*. Before signing the grant, each beneficiary will have to confirm this again by signing a declaration of honour (DoH). Proposals without full support will be rejected.

Your application must be **readable**, **accessible and printable** (please check carefully the layout of the documents uploaded).

Proposals are limited to maximum **10 pages** (Part B). Evaluators will not consider any additional pages. Shorter proposals are welcome.

You may be asked at a later stage for further documents *(for legal entity validation, financial capacity check, bank account validation, etc)*.

6. Eligibility

Eligible participants

Eligible entities receiving financial support must be established in Argentina, Brazil, Chile, Colombia, Costa Rica, Mexico, or Uruguay.

The following types of entities are eligible to receive financial support:

- **Higher education establishments**, whether public or private, recognised as such by the relevant national education system, including universities and other higher education institutions.
- **Research organisations**, whether public or private, established on a non-profit basis and having research and/or technological development as their main objective, including research institutes and eligible think tanks.
- **High-Performance Computing (HPC) and Artificial Intelligence (AI) centres**, whether public or private, including entities operating or providing access to HPC/AI infrastructure, facilities, expertise or related services, provided that their activities are relevant to the objectives of the call.

Consortium Composition

Single applicants or consortia may apply. Specific requirements regarding consortium composition may apply depending on the relevant use case or topic.

The consortium structure is relevant only for the proposal submission and evaluation stage. If a proposal is selected for funding, each participating entity will sign its own grant agreement and receive funding directly. Consequently, there is no coordinator role during the implementation phase with respect to the management or distribution of the grant budget.

Eligible actions and activities

[1] PDF unless another format is specified in the templates.

[2] See judgment of 18 December 2024, *Institut Jožef Stefan vs European Commission*. T-134/23, EU:T:2024:907, paragraphs 43, 47, 59.

Financial support should only be used to pay for personnel salaries, not for equipment, resources, or travel.

Geographic location (target countries)

Proposals must relate to activities taking place in the eligible countries *(see above)*.

Duration

Projects should normally range between 12 and 18 months.

Projects of longer duration may be accepted in duly justified cases. Extensions are possible, if duly justified and through an amendment.

Project budget

The requested grant amount per project must not exceed EUR 50.000 per project.

The grant awarded may be lower than the amount requested.

7. Evaluation and award procedure

The proposals will have to follow the **standard submission and evaluation procedure** (one-stage submission + one-step evaluation).

An **evaluation committee** (assisted by independent outside experts) will assess all applications. Proposals will first be checked for formal requirements (admissibility, and eligibility, *see sections 5 and 6*). Proposals found admissible and eligible will be evaluated (for each topic) against the award criteria *(see section 8)* and then ranked according to their scores.

For proposals with the same score (within a topic or budget envelope), a **priority order** will be determined according to the following approach:

Successively for every group of *ex aequo* proposals, starting with the highest scored group, and continuing in descending order:

1) The ex aequo proposals within the same topic will be prioritised according to the scores they have been awarded for the award criterion 'Excellence'. When these scores are equal, priority will be based on their scores for the criterion 'Impact'. When these scores are equal, priority will be based on their scores for the criterion 'Implementation'.

2) If this does not allow to determine the priority, a further prioritisation can be done by considering the overall project portfolio and the creation of positive synergies between projects, or other factors related to the objectives of the call. These factors will be documented in the panel report.

3) After that, the remainder of the available call budget will be used to fund projects across the different topics in order to ensure a balanced spread of the geographical and thematic coverage and while respecting to the maximum possible extent the order of merit based on the evaluation of the award criteria.

All proposals will be informed about the evaluation result (**evaluation result letter**). Successful proposals will be invited for grant preparation; the other ones will be put on the reserve list or rejected. That decision is final and the project will not enter into discussions concerning the evaluation results, and no appeals process will be provided.

⚠ No commitment for funding — Invitation to grant preparation does NOT constitute a formal commitment for funding. Various legal checks before grant award may still be needed.

Funding Grant preparation will involve a dialogue in order to fine-tune technical or financial aspects of the project and may require extra information from your side. It may also include adjustments to the proposal to address recommendations of the evaluation committee or other concerns. Full compliance will be a pre-condition for signing the grant.

8. Award criteria

The **award criteria** for this call are as follows:

1. Excellence:

- Relevance of the applicant's qualifications, experience and skills to the objectives and specific selection criteria of the selected joint use case. Quality and relevance of the proposed research contribution to the selected joint use case. Availability and accessibility of relevant data or other research resources, where applicable, and their potential contribution to the selected joint use case.

2. Impact:

- Relevance and expected contribution of the proposed outputs and deliverables to achieving the objectives of the joint use case. Potential for knowledge exchange, capacity building and strengthening LAC–Europe research collaboration, including opportunities for future cooperation.

3. Implementation:

- Clarity and feasibility of the proposed work plan, including the planned activities, timeline and collaboration arrangements with the European partner. Appropriateness of the requested salary funding in relation to the proposed activities, duration and level of researcher involvement.

Criteria 1 to 3 will carry a score ranging from 0 to 5 (steps of 0.5). Criterion 1 will have a weight of 2×, Criteria 2 and 3 a weight of 1× (leading to a maximum score of 20 points). A threshold score of 3 will apply to all criteria.

Proposals that pass the thresholds will be considered for funding — within the limits of the available budget (i.e. up to the budget ceiling). Other proposals will be rejected.

9. Set-up of the Funding Grant Agreements

If you pass evaluation, your project will be invited for funding grant preparation, where you will be asked to prepare the Funding Grant Agreement.

This Agreement will set the framework for your grant and its terms and conditions, in particular concerning deliverables, reporting and payments.

The Model Funding Grant Agreement that will be used as a basis (and all other relevant templates and guidance documents) will be made available as a supporting document. This agreement will concluded between the Third party selected to receive funding and the Barcelona Supercomputing Center – Centro Nacional de Supercomputación who is the Coordinator of the project entitled "EU-LAC Supercomputing Network for AI".

Additional supporting documents may be requested (e.g. to confirm the operational capacity of any applicant).

Third Parties will receive 100% funding of incurred, eligible direct costs necessary for the completion of activities; no indirect costs or overheads will be funded. The budget categories and cost eligibility rules are fixed in the Funding Grant Agreement. The only budget category eligible for funding is 'personnel costs'.

The project starting date and duration will be fixed in the Funding Grant Agreement . Project duration: between 12 and 18 months (*see section 6 above*).

The grant will be an actual cost grant. This means that it will reimburse ONLY certain types of costs (eligible costs) and costs that were *actually* incurred for your project (NOT the *budgeted* costs).

Grants may NOT produce a profit (i.e. surplus of revenues + EU grant over costs). For-profit organisations must declare their revenues and, if there is a profit, we will deduct it from the final grant amount.

The funding grant agreement will define the payment schedule, whereby payment will be made in multiple tranches, partly linked to the delivery of reports and outputs defined in the workplan and in the funding agreement itself. The principle followed will be that the participants will receive an advance payment at the start of the joint use case followed by subsequent instalments.

10. How to submit an application

All proposals must be submitted directly online via the EuroHPC Peer Review platform https://access.eurohpc-ju.europa.eu/calls/68. Paper applications are NOT accepted.

Submission is a **2-step process**:

a) **create a user account**

To use the Submission System (the only way to apply), all applicants (in case of consortium, only one account is enough) to create an account on https://access.eurohpc-ju.europa.eu .

b) **submit the proposal**

Submit your proposal in 3 parts, as follows:

- Part A includes administrative information about the applicant organisations and is an online form.
- Part B (description of the action) covers the technical content of the proposal. Download the mandatory word template from the Submission System, fill it in and upload it as a PDF file.
- Annexes (*see section 5*). Upload them as PDF file (single or multiple depending on the slots).

The proposal must keep the format unchanged and respect the **page limits** *(see section 5)*; excess pages will be disregarded.

Documents must be uploaded to the **right category** in the Submission System, otherwise the proposal may be considered incomplete and thus inadmissible.

The proposal must be submitted **before the call deadline** (*see section 4*). After this deadline, the system is closed and proposals can no longer be submitted.

Once the proposal is submitted, you will receive a **confirmation email** (with date and time of your application). If you do not receive this confirmation email, it means your proposal has NOT been submitted.

11. Help

As far as possible, ***please try to find the answers you need yourself***, in this and the other documentation (we have limited resources for handling direct enquiries):

– *Contact*

Questions should be sent to the following email address: RI@eurohpc-ju.europa.eu.

Please indicate clearly the reference of the call and topic to which your question relates *(see cover page)*.