Call for Papers

AHPC³: The 2nd Workshop on Accelerated HPC in the Cloud-Edge Continuum

Affiliated with 13th IEEE International Conference on Cloud Engineering (IC2E)

Website: http://ahpc3.di.unipi.it/

Important Dates

Date	Milestone
June 3, 2025	Submission of paper abstracts
June 16, 2025	Submission of short/regular papers
July 21, 2025	Notification to authors
July 31, 2025	Submission of camera-ready

Today, we are witnessing an increasing demand for high-performance computing infrastructures. Modern applications are becoming ever more dependent on the processing of computationally intensive workloads. Historically, HPC infrastructures were utilized by a limited number of application domains. However, in the current era, a diverse range of domains and applications are increasingly requesting HPC infrastructures, driven by the high availability of large volumes of data and the need for AI/ML and federated learning across distributed Cloud-Edge-HPC environments.

In the past, HPC infrastructures significantly differed from cloud infrastructures for their distinct software and hardware requirements and their on-premises nature.

However, such distinctions have become increasingly blurred, driven by the proliferation of data-intensive, AI/ML, and federated learning applications. Modern cloud and edge environments are progressively integrating capabilities once exclusive to HPC systems in terms of performance capabilities and hardware specifications.

This workshop aims to explore the intersection of high-performance computing and modern Cloud-Edge Continuum architectures. The focus will be on achieving HPC by

relying on Cloud-Edge architectures. In this year's edition, a particular emphasis will be placed on promoting energy-efficient execution, green computing practices, and intelligent resource management to reduce environmental impact. Additionally, topics of interest are related to the challenges of adapting, tailoring, and orchestrating technologies typically exploited in cloud and edge contexts, including serverless computing, microservices, and fine-grained task offloading, to support efficient, scalable, and resilient high-performance computing applications.

We welcome original contributions on novel programming paradigms, runtime systems, and orchestration frameworks that empower high-performance and federated learning applications across cloud-edge and HPC architectures. Submissions may address efficient network communication, scalable data management, robust fault-tolerance mechanisms, end-to-end security, and energy-aware execution. Case studies and tool papers on orchestrating data-intensive workflows, implementing comprehensive monitoring solutions, and optimizing sustainable execution in distributed and hybrid environments are also encouraged.

The **2nd International Workshop on Accelerated HPC in the Cloud-Edge Continuum (AHPC³)** aims to bring together cloud, edge computing, and HPC experts from academia and industry to identify new challenges, discuss novel systems, methods, and approaches for Hybrid and accelerated HPC cloud-edge infrastructures and architectures, and promote this vision toward academia and industry stakeholders.

Topics of interest

Topics of interest for the workshop include but are not limited to the following ones:

- Adaptation of cloud-edge technologies and methodologies for HPC (e.g., serverless, microservices, task offloading)
- Cloud-edge computing architectures for HPC (e.g., resource federation)
- Lightweight virtualisation tools, execution environments and scheduling techniques
- Orchestration, deployment techniques and algorithms for high-performance workflows in Cloud-Edge environments
- Programming paradigms for high-performance Cloud-Edge computing
- Communication, data and resource management for Cloud-Edge computing
- Fault tolerance, reliability and security in the Cloud-Edge continuum
- Data-intensive workloads and tools
- Tools and techniques for monitoring HPC Cloud-Edge applications
- Sustainable and energy-efficient computing for HPC in the cloud-edge continuum

- Accelerated computing architectures for cloud-edge environments (e.g., GPU-based execution, hardware accelerators for AI/ML)
- Network-aware optimizations and intelligent data movement for performance and energy efficiency
- Federated resource-sharing mechanisms for hybrid HPC environments
- Federated learning in the Cloud/Edge Continuum
- Novel and efficient techniques and algorithms for Cloud/Edge orchestration
- ML/AI based algorithms to support Cloud/Edge computing

Submissions and attendance

http://ahpc3.di.unipi.it/submissions.html

Accepted papers will be published in the conference Proceedings. Submitted papers must be original work that has not appeared in and is not under consideration for another conference/workshop or a journal. Every submitted paper will be reviewed by at least three members of the Program Committee. Reviewing will be single-blind. Authors are invited to submit papers in the <u>IEEE Conference Proceedings format</u> style, which should be **at least 5 pages** and **should not exceed 9 pages** including everything.

Please note that registering on the submission site with a title and meaningful abstract by the earliest deadline is required to enable the actual paper submission.

The authors must be prepared to sign a copyright transfer statement. At least one author of each accepted paper must register for the workshop by the early date, to be indicated by the organisers, and ***must*** present the paper.

Special Issue

The workshop chairs will invite authors of outstanding research papers to submit an extended version of their work to a Journal Special Issue. We are in contact to propose a Special Issue in a high-quality journal. Information will appear on the AHPC3 website.

Invited Speaker

We are pleased to announce that **Prof. Guillaume Pierre**, a leading researcher in distributed systems and professor of Computer Science at **Université de Rennes**, **France**, will be joining AHPC³ as our invited keynote speaker.

In his keynote, Prof. Pierre will explore the challenges of supporting data-intensive workloads in heterogeneous edge infrastructures, with a particular focus on operating under strong energy constraints.

Organisers

- Luca Ferrucci, University of Pisa, luca.ferrucci@unipi.it, General Chair
- Stefano Forti, University of Pisa, stefano.forti@unipi.it, General Chair
- Valerio Besozzi, University of Pisa, valerio.besozzi@phd.unipi.it, Program Chair
- Jacopo Massa, University of Pisa, jacopo.massa@di.unipi.it, Program Chair

Program Committee

- Jörn Altmann, Seoul National University
- Hojjat Baghban, Chang Gung University
- Roberto Casadei, University of Bologna
- Emanuele Carlini, ISTI-CNR
- Marcin Copik, ETH Zürich
- Massimo Coppola, ISTI-CNR
- Patrizio Dazzi, University of Pisa
- Maria Fazio, University of Messina
- Carlos Guerrero, University of Balearic Islands
- SongHee Kang, Seoul National University
- Hanna Kavalionak, ISTI-CNR
- Isaac Lera, University of Balearic Islands
- Matteo Mordacchini, IIT-CNR
- Paolo Palazzari, ENEA
- Paul Rourab, Siksha 'O' Anusandhan University
- Jocelyn Sérot, Université Clermont Auvergne
- Konstantinos Tserpes, NTUA
- Paolo Trunfio, University of Calabria