The International Post-Exascale Project (InPEx)

InPEx 2026, Niteroi, Brazil

Computing Continuum Sessions: Introduction

Gabriel Antoniu (Inria), Mario Acosta (BSC), Carlos Jaime Barrios Hernandez (UIS/INRIA),
Marcio Castro (UFSC), Nicola Ferrier (Northwestern University), Maria Girone (CERN),
Kento Sato (RIKEN), Hiroyuki Takizawa (Tohoku University)

InPEx 2026 - Computing Continuum Sessions

International Post-Exascale Initiative

24th April 2026

08:30-09:00. Welcome onsite

09:00-11:00. Subgroup session A2

Each subgroup works in parallel for 120 minutes.
Few participants will join remotely

1. HPC/AI convergence (session 1/2)
Chair: Mohamed WAHIB (RIKEN)
Co-chairs: Emmanuel Jeannot (INRIA), Artur Lorenzon (UFRGS), Philippe Navaux (UFRGS), Brian

1. Computing continuum (session 1/2)
Chair: Gabriel ANTONIU (INRIA)
Co-chairs: Mario ACOSTA (BSC), Carlos Jaime Barrios Hernandez (UIS/INRIA), Marcio Castro (UFSC), Nicola Ferrier (Northwestern University), Maria Girone (CERN), Kento Sato (RIKEN), Hiroyuki Takizawa (Tohoku University)

11:00-11:30. Coffee break

11:30-12:30. Presentation of subgroups A1 and A2 wrap-up and discussion

12:30-14:00. Lunch

14:00-16:00. Subgroup session B1

Each subgroup works in parallel for 120 minutes.

1. Computing continuum (2/2)
Chair: Gabriel Antoniu (INRIA)
Co-chairs: Mario ACOSTA (BSC), Carlos Jaime Barrios Hernandez (UIS/INRIA), Marcio Castro (UFSC), Nicola Ferrier (Northwestern University), Maria Girone (CERN), Kento Sato (RIKEN), Hiroyuki Takizawa (Tohoku University)

1. GenAI and software engineering (session 2/2)
Chair: Anshu DUBEY (ANL)
Co-chairs : Roberto Di Cosmo (INRIA), Jens Domke (RIKEN), Alfredo Goldman (USP), Antonio Tadeu Gomes (LNCC), Christophe Prudhomme (UNISTRA)

16:00-16:30. Coffee break

InPEx 2025 Recap – Computing Continuum Sessions

International Post-Exascale Initiative

Use Case presentations with associated challenges

- The Square Kilometer Array (SKA), a workflow-intensive project that would benefit significantly from operating across multiple infrastructures;
- High-energy physics (HEP) use cases;
- EsiWACE – a high-performance climate and weather benchmark;
- Urgent computing: facilitating trade-off management on the Edge – Cloud Computing Continuum for Urgent Science;
- Micro-meteorology control;
- Disaster digital twins

Guidelines: each use case presents the corresponding challenges/roadblockers and initial approaches to tackle them. To help identifying commonalities, all use-case representatives are requested to comment on the criteria identified.

- *Shared notes: https://tinyurl.com/inpex-continuum*

InPEx 2025 Recap – Computing Continuum - Discussion

International Post-Exascale Initiative

- Pick-up your top 3 most relevant challenges and provide more details on what is difficult
- Ongoing R&D efforts to address these challenges?
- What important strategic actions do you see to address these challenges?
- What are the technical milestones whose achievement would help to address these challenges?

InPEx 2025 Recap – Computing Continuum – Top challenges?

International Post-Exascale Initiative

InPEx workshop - April 2025 - Japan

Digital continuum and data management session

	SKA - DDF pipeline	HEP use cases	ESiWACE	Urgent computing	Micro-meteorology control	Disaster Digital Twin Project
Access to HPC resources (as a community)	X (exploitation of final data products offline)	X	X	X (Urgent computing needs to harness all available resources achieve its objectives)	X	
Interface to federation of resources (e.g. multiple EuroHPC centers)	X (exploitation of final data products offline)	X	X (seeks opportunity to participate with domain cases in the development o the ederated infrastructure)	X (End to end workflows require the federation of resources across the digital continuum, including HPC centers)		X
Co-design for the post-exascale systems	X (especially for offline component, but downscaled solutions could be reused for online component in order to maximize efficiency)	X	X (seeks opportunity to participate in the co-design of the benchmark to assess upcoming HW for the PostExascale)	X (co-design of the digital continuum, both architecturally and operationally)		
End-to-end workflow control, workflow management	X (for the online part)	X (need custom interfaces to submit pilot-based workflows)	X (web - Cloud- HPC)	X	Custom workflow control	X (the urgent workflow consists of 200+ tasks, but only a small number of tasks need to be executed on HPC systems
Data logistics	X (across the continuum, from real-time data stream to offline data analysis: several data topologies and formats as well as provenance tracking strategies)	X (need transient storage; limited connectivity; need standard data transfer protocols)	X (need standard data transfer protocols)	X (Data, both real-time streaming and at rest, is critical to urgent applications)	Real-time data from sensors are used for invoking tasks. Some data come every 5min, some come every hour. Heterogeneous tasks run with their deadline.	
Reproducibility	X (for the online part)		X (CMake, SPACK)	X (Provenance, explainability is critical)		
Hardware Heterogeneity	X (performance / efficiency issue for the online part, portability issue for the offline part)	X large interest in High Level Frameworks for heterogeneous computing (KOKKOS, Alpaka, SYCL, ...)	X	X	X Main simulation tasks currently run o	X (use of multiple data centers in emergency)
Energy Efficiency	X (across the continuum, various aspects to be considered incl. operational constraints -- ie power capping scenarios -- and TCO. currently working on full LCA)	X (studies in progress)	X (tools for monitoring the energy and ensure that we have a compromise with the effficiency of the applications) Carbon footprint concerns	X		
Scheduling/resource allocation across the continuum	X (across the continuum, from very short-term real-time stream processing close to the sensors -- complex scheduling strategy -- to longer term final data product reduction and analysis on federated facilities -- scheduling and resources allocation strategy)	X (mismatch in allocation models; short-term allocations vs. long-term scientific planning -- our "runs" are 3-5 years long)	X (for operational weather/climate services or long climate projections, how to allocate resources)	X(discover and aggregate currently available resources based on current needs)		X (reduce the negative effect of urgent job execution on system operation

InPEx 2025 Recap - Computing Continuum – Next actions

International Post-Exascale Initiative

Potential goal:

- Co-write a white paper explaining the needs of the science communities that needs to be executed on the computing continuum, listing challenges, priorities and promising approaches to tackle them

Follow-up actions planned

- Continue the discussions via online meetings
- Continue to contribute to the shared nodes *: https://tinyurl.com/inpex-continuum*
  - *Use case characterization for the additional use cases*
  - *Ongoing related R&D projects, strategic actions, impactful technical milestones*
- Please add your e-mail in the list at the top of the document to keep being updated for forther actions!

Some actions since the last workshop

- CSA SPE-EU European project on post-exascale roadmapping starting in April 2026 (Europe)
- Seminar on interoperability of digital infrastructures (National, France, October 2025)
- Transdigital conference – session on the computing continuum challenges (National, France, February 2026)

HORIZON-CL4-2025-03-DIGITAL-EMERGING-04: Post-exascale HPC (CSA)

- **Expected Outcome:**
- Delivery of a **high-quality roadmap** addressing the post-exascale HPC/AI research challenges for applications, algorithms, software, hardware and systems, including a strong emphasis on AI
- Contribution to the development of a competitive **European converged HPC/Quantum/AI ecosystem**, including AI Factories and future AI Gigafactories
- Interaction and **collaboration with similar international efforts**, ensuring alignment with AI-driven computing paradigms worldwide

InPEx 2026 - Computing Continuum – Plans

International Post-Exascale Initiative

Goals and actions:

- Updated view on ongoing efforts/projects
- Complete the picture: add other relevant use cases, additional challenges
- Discuss concrete operational steps to enable international collaborations in this area

Expected outcome:

- A solid updated vision on the aforementioned topics at an international level
- A basis for upcoming roadmapping documents and funding strategies
- Concrete context in Europe for a first joint white paper: the CSA SPE-EU project

Proposed contributions

- Updates on HEP use cases, SPECTRUM (Maria Girone)
- The ODISSEY project (Damien Gratadour)
- Updates on High-Performance Climate and Weather Benchmarking (Mario Acosta)
- New use cases: agentic workflows, agentic AI (Manish Parashar)
- Additional use cases and use cases from Latin America (Carlos Jaime Barrios Hernandez)
- *To be completed*