International PostExascale

Workshop Series

InPEx 2025 workshop – April 14-17, Japan

Digital continuum and data management
InPEx working groups results
Session Co-leads: Gabriel Antoniu (Inria), Manish Parashar (U Utah), Kentaro Sano (RIKEN)

InPEx - Digital Continuum and Data Management – Session 1 Recap

International Post-Exascale Initiative

- Use Case presentation with associated challenges
  - [SKA] Direction-Dependent Facet (DDF) calibration Use-Case (astronomy/radioinferometry-based imaging)
    - *Mathis Certenais (U Rennes), Damien Gratadour (Obs Paris) et al (Inria)*
  - [HEP] High-Energy Physics use cases for HPC
    - *Maria Girone et al. – CERN*
  - [ESiWACE] A Workflow for HPCW - The High-Performance Climate & Weather Benchmark
    - *Mario Acosta et al. – BSC*
  - [Urgent computing] Facilitating trade-off management on the Continuum for Urgent Science
    - *Manish Parashar (U Utah), Daniel Balouek (Inria)*

- *Shared notes: https://tinyurl.com/inpex-2025-continuum*

Additional use cases

Micro-meteorology control project
- Prof. Ryo Onishi, Science Tokyo
  - https://www.turb.gsic.titech.ac.jp/en/project/project.html

Disaster Digital Twin Project
- Prof. Shunichi Koshimura, Tohoku University

Using High-performance Imaging and Advanced AI Analytics to Inspect Road Infrastructure,
- Peng Chen (RIKEN R-CCS) and Mohamed Wahib (RIKEN R-CCS)

Discussion

- **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 - Digital Continuum and Data Management – Use Case Characterization

International Post-Exascale Initiative

InPEx workshop - April 2025 - Japan

Digital continuum and data management session

					Other use cases?	
	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)		
Portability of the benchmarks and codes	X (exploitation of final data products offline)	X (porting to GPUs or new accelerators)	X (best efficiency possible in different CPU machines and porting to GPUs or new accelerators)	X (essential as workflows will have to run on available resources, which may not be know until runtime).		
Deployement issues of the workflows and applications	X (across the continuum: each sub-systems handle a portion of the overall pipeline)	X (having CVMFS as a baseline solution available at all centers would be largely beneficial [either full installation or possible at user level])	X (need automatization; on-demand workflow submission)	X		
Metrics for different deployement scenarios			X (need to collect metrics in real time)	X (Evaluation of cost/benefit at runtime is essential)		
Semantics and quality of data				X		
Resources provisionning	X (mostly for exploitation of final data products offline)	X (an optimal utilization would require longer provisioning - e.g. 72 hours)		X		
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
Multi-tenancies	X (exploitation of final data products offline)			X		

InPEx - Digital Continuum and Data Management – Outcome

International Post-Exascale Initiative

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	X (use of multiple data centers in emergency)
Cybersecurity and Privacy	X (mostly for exploitation of final data products offline)	X (need federated authentication/identity management compatible with the WLCG trust model -- which is adhereing the AARC Blueprint)		X		
Multiscale System-of-systems Integration	X (intercontinental experiment, from data production at the edge on two different continents to a global user communtiy accessing federated resources at regional level)			X		
Sustainability	X (across the continuum, ~50 years lifetime of the infrastructure)	X (studies in progress)	X (how to be suistanable beyond the lifecycle of a scientific project)			
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
Programming models		X (SYCL, Alpaka, Kokkos)	X (Maintain and optimize Fortran)	X		
Edge Application Requirements		X (need to deploy domain-specific federation services)		X (Edge integration is critical)		
Reliability issues	X (especially close to the edge, data streaming and reduction)	X (our standard requirements are at 99% availability, which are not currently met by HPC centers; but the requirements are mostly for storage which is less relevant at HPC centers)	X(availability and storage are relevant for operational cases, something less relevant for HPC centers)	X (support dynamic execution)		X (urgent simulation must be executed for sure)
Tradeoff management	?			X (Making cost/benefit evaluations is essential)		X (use of publicly-shared academic computing resources to minimize the operation cost)
Scale	X (cost / operations / maintenance issue for the online part, performance / deployment / usability issue for the offline part)			X (Yes, but scalability is not the key goal)		
Real-time data streaming at high rates	X (cope with buffer-based strategy and limited buffer capacity -- 24h max)	X (while not stricly a requirement, we have utilization patterns in which data is streamed to the centers instead of using pre-allocations or "lazy downloads")	X (not in HPCW use case but in DestinationEarth, the digital twin date is streamed directly to applications used by the users)	X		
Managing QoS under constraints, incertainty				X		X (hard deadline contraint)

InPEx - Digital Continuum and Data Management – 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 - Digital Continuum and Data Management – Ongoing R&D Efforts?

International Post-Exascale Initiative

- Interface to federation of resources
  - JP (RIKEN): Working on Data Acquisition System for AI for Science by Globus
  - JP(U Osaka): Working on data aggregation infra. named ONION to make data searchable; high-speed data transfer service on campus
  - EU: Destination Earth studies on federation challenges
  - JP (Tohoku): meta-scheduling
- Co-design for the post-exascale systems
  - EU: DARE and Center of Excellences as ESiWACE3
  - JP (RIKEN): FugakuNEXT
- End-to-end workflow control, workflow management
  - EU: ODISSEE project, CoEs like ESiWACE3
- Reproducibility
  - JP (U Osaka) : Working on Provenance management system
  - EU (Inria): E2Clab project - https://e2clab.gitlabpages.inria.fr/e2clab/

InPEx - Digital Continuum and Data Management – Strategic actions?

International Post-Exascale Initiative

- EU: EuroHPC projects comes with no computational resources. Communities can benefit if EuroHPC resources are provided along projects and CoEs granted.
- JP (RIKEN); Ecosystem and community building around Globus (collaboration with US).
- EU: Work in close collaboration with the EuroHPC Federation Platform (EFP) in 2025. This corresponds to the development of the EFP first release (development phase), foreseen to be ready end of Q1 2026. Domain communities do need to participate and provide feedback (co-design).
- JP (U Osaka): conformation to OpenScience
- EU: Collaborate with the community CoEs to ensure that we have an useul co-desing while the new hardware is developed
- JP (RIKEN): Selection of system software gueranteeing continuous/sustainable development (e.g., use of OSS / proprietary software developed by sustainable vendors)

InPEx - Digital Continuum and Data Management – Technical milestones?

International Post-Exascale Initiative

- JP (RIKEN): Integration of Globus with other data sharing service because different organizations use different data sharing services; JP (U Osaka): meta-data management of aggregated data
- EU: Interoperable workflow control / management tools
- JP (RIKEN): Stardadization of workflow management system; there are many workflow systems (https://workflows.community/systems) which burdens system support
- Technology to monitor (and optimize) energy efficiency across the continuum (extensions of existing solutions like e.g. EAR from EAS or EnergyScopium from DENERGIUM taking into account the full continuum of systems). Something portable and comparable across platorms and applications
- JP (RIKEN): Dedicated use of computational resources for 24/7 to deploy AI services for inference (e.g., code generation) which does not align with operation by a scheduler

InPEx - Digital Continuum and Data Management – 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

Next actions

- Continue the discussions via online meetings
- Continue to contribute to the shared nodes *: https://tinyurl.com/inpex-2025-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!