RÉPUBLIQUE FRANÇAISE
Liberté Égalité Fraternité

Forum TERATEC
20 - 21 OCTOBRE 2026
PARIS · PALAIS DES CONGRÈS

cea cnrs Inria

FRANCE 2030
PROGRAMME DE RECHERCHE
NUMÉRIQUE POUR L'EXASCALE

Integrating Supercomputers into Cross-Facilities Workflows (CWF)

ANR Numpex/Exa-AToW

François Bodin

October 21, 2026

University of Rennes

Introduction

- The role of supercomputers in science is expanding beyond traditional simulation, toward high-performance data processing, and compute agents supporting LLM-based scientific workflows.
- Supercomputers are becoming integral components of a digital continuum, itself embedded within a broader datasphere.
- Cybersecurity and operational constraints must be accounted for to ensure real-world deployability.
  - Usage is governed by ZRR regulations.
  - ZRR access requires a personal account.
  - Requires HPC system knowledge.
- Exa-AToW has designed a solution addressing both challenges and enable the mutualization of scientific applications within a community
  - Operate cross-facility workflows usable by non HPC expert users.
  - The deployment of scientific applications as MCP servers.

Cross-Facility Workflows

Cross-Facility Workflows

A Cross-Facility Workflow is a chain of processing tasks whose workload is distributed across multiple computing systems (e.g., data centers, cloud infrastructures, HPC centers, edge devices) that fall under distinct administrative authorities.

- Tasks are orchestrated and executed across heterogeneous infrastructures rather than within a single facility.
- The participating systems are governed by different organizations, each with their own policies, access controls, and operational constraints.
- Cross-Facility Workflow platforms are inherently collaborative systems of systems, requiring interoperability, trust, and coordination mechanisms across administrative boundaries.

What is a community HPC Application in exa-AToW Context

Community HPC Application

A Community HPC Application is a scientific application designed for use by a large, diverse research community, where computational resources and expertise are shared (mutualized) across users who are not necessarily proficient in HPC.

- The application serves multiple users or research groups, enabling mutualization of both computational infrastructure and domain expertise.
- Users are not required to have advanced HPC skills—the complexity of the underlying infrastructure is abstracted away from the end user.
- The application addresses needs that extend beyond a single developer or research team.
- Community HPC Applications are cross-facilities workflow building blocks.

Role Models in a community HPC Application

Community Developer

A Community HPC Application Developer is a scientist who actively contributes to the development of a Community HPC Application. This role is inherently tied to the community dimension — the developer contributes to a codebase that benefits multiple users, as opposed to developing a one-off, single-user tool.

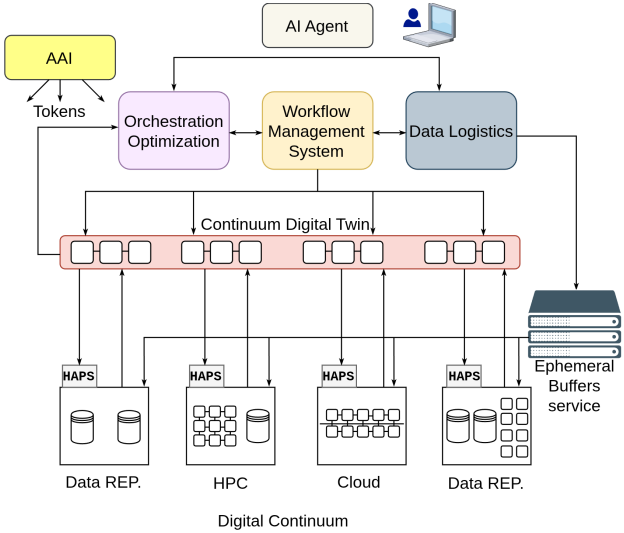
Community Maintainer

A Community HPC Application Maintainer is the individual responsible for deploying, configuring, and managing a Community HPC Application on behalf of a given research community. The Maintainer handles the practical, ongoing management of the application — including deployment, updates, and configuration — rather than the scientific development of its core functionality. The Maintainer holds an active account on the HPC system, granting them the technical access rights necessary to perform **these operational tasks**.

Community User

A Community HPC Application User is a scientist who holds an identity within a community-recognized federation and seeks to use a Community HPC Application to support their research. This notion is extended to **AI agent**.

Exa-AToW Architecture



- The **Continuum Digital Twin (CDT)** provides a uniform abstraction of the digital continuum, spanning HPC centers, data centers, cloud, and edge resources.
- The CDT acts as a discovery service, enabling users and workflows to locate available centers and Community HPC Applications.
- The CDT exposes rich descriptions of HPC centers, including their capabilities, resources, and exposed HAPS.
- It is built upon the Exa-AToW ontologies: *ImpactHPC*, *Job*, and *Workflow*.

Digital Twin Ontologies and PIDs

- Ontologies:
  - HPC Resource Ontology (High-Performance Computing) Environment Ontology: Covers physical and virtual resources in HPC, including compute nodes, storage, interconnects, and infrastructure.
  - Job Submission Ontology (in progress): This component focuses on the semantics of submitting computational jobs. It covers concepts like job parameters, resource requirements (e.g., CPU, memory, time), submission protocols, and execution targets.
  - Processor Indicator Estimator Ontology: Groups estimation tools, metrics, and models used to assess processor indicators such as power consumption, thermal footprint, and die size.
  - Workflow Ontology: This ontology formally defines the concepts and relationships related to scientific and computational workflows. It models the structure, dependencies, execution logic, and state transitions of complex multi-step processes.
- Persistent Identifiers (joint work with J. Sudre Chargé de mission PID - DGRI)
  - Organisations, people
  - National research infrastructures
  - Datasets
  - ...

`https://github.com/logistica-dev/exa-atow-ontologies/tree/main`

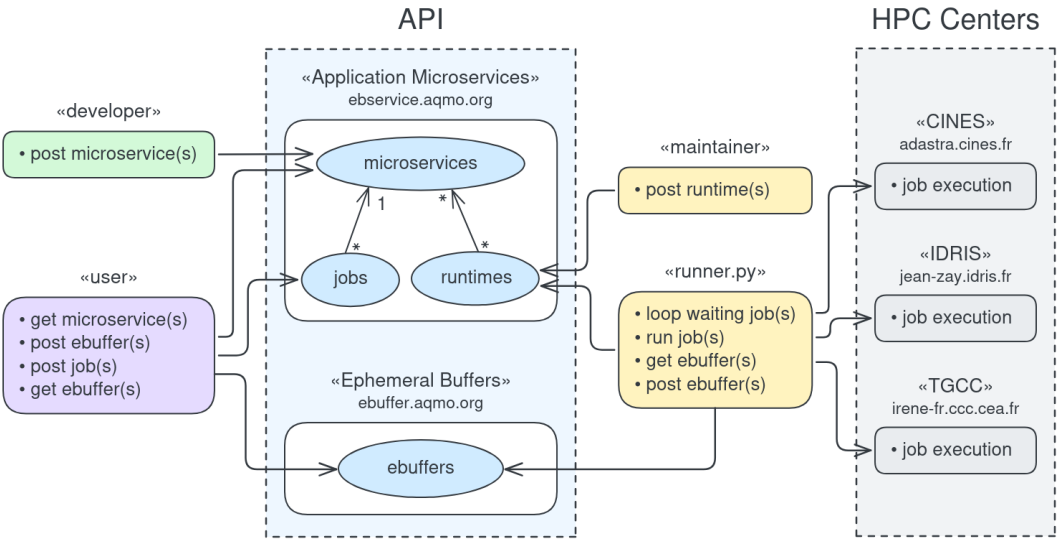
HPC APplication Services (HAPS)

HAPS

An HPC Application Service (HAPS) is a web service that enables Users to request, monitor, and retrieve the results of executions of a Community HPC Application, abstracting away the underlying complexity of the HPC infrastructure.

- Users interact with HAPS without needing HPC-specific expertise (job scheduling, resource allocation, batch systems, etc.)
- HAPS typically supports the full execution lifecycle — submission, monitoring, and results retrieval — not just the triggering of a run.
- HAPS relies on Ephemeral Buffer for transient data transfers.
- Access to HAPS relies on the User's identity within the community-recognized federation.

HAPS Architecture



HAPS in Practice (based on Gromacs)

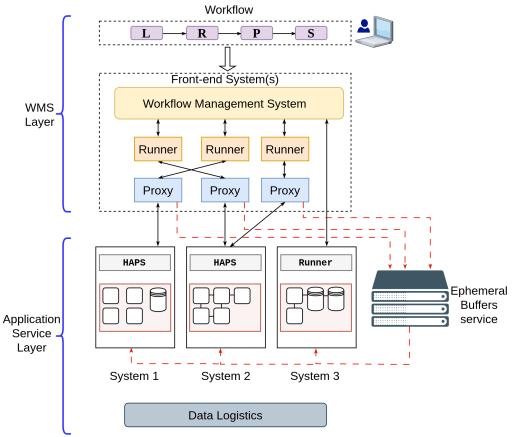
```
md_tpr = Path("./water.tpr")
ebin = api.buffers.create(
 reserved_size=md_tpr.stat().
    st_size,
 lifespan=LIFESPAN
)

ebin.fill(md_tpr)

SIZES = {
    "md.cpt": 10_000_000,
    "md.edr": 10_000_000,
    "md.gro": 10_000_000,
    "md.log": 10_000_000,
}
```

```
ebout = [
    api.buffers.create(
        reserved_size=SIZES[name],
        lifespan=LIFESPAN
    )
    for name in ms_gromacs.
    ebout_names
]
sched.submit_job(
    [],
    ebin=[ebin.uuid],
    ebout=[b.uuid for b in ebout],
    dry_run=False
)
sched.monitor()
```

Cross-Facility Workflow in Exa-AToW



- The application layer exposes HAPS and provides an Ephemeral Buffer (EB) service to support the associated data logistics.
- The architecture is agnostic to the Workflow Management System (WMS) used – currently supporting Galaxy and CWL-based tools.
- Proxy tasks run on the same system as the WMS and act as a bridge between the HAPS and the workflow tasks.
- Data can remain at the HAPS level, with only metadata exchanged at the WMS layer, avoiding unnecessary data transfers.
- Regular & HAPS-based tasks can be seamlessly mixed within the same workflow.

Conclusion

- Supercomputers are becoming key components of a broader digital continuum, requiring new models to expose scientific applications as accessible, secure services.
- Exa-AToW addresses this challenge through Community HPC Applications exposed as **HAPS**.
- HAPS are a technical tradeoffs between HPC governance and cross-facility workflow requirements.
- They provide a common framework for agentic systems, reproducibility, tracability, etc.
- A digital twin allows for a high level actionnable abstraction of the continuum.

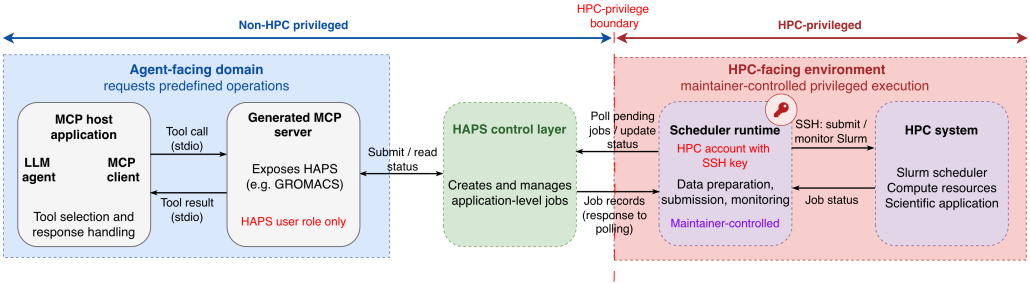
Perspectives

- HAPS as Model Context Protocol (MCP) services
- MCP Service Discovery
- Runtime and energy consumption prediction
- MCP Based Workflows
- Reproductibility and Tracability

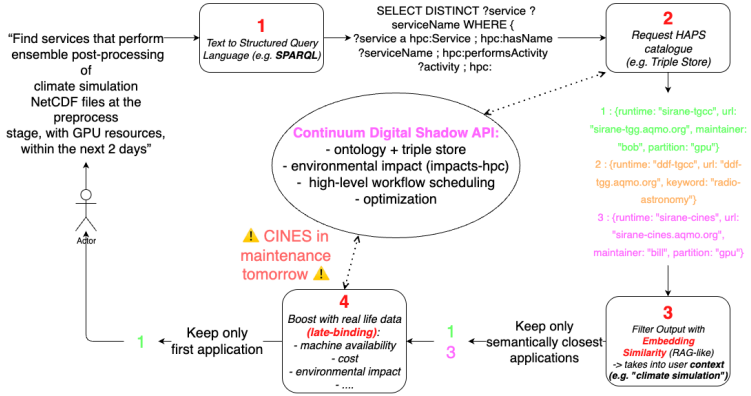
Making the Application Available Using MCP

HAPS MCP Service

A HAPS MCP Service is the packaging of a HAPS as a Model Context Protocol (MCP) server, exposing the Community HPC Application's execution capabilities as standardized tools consumable by AI agents and LLM-based applications.



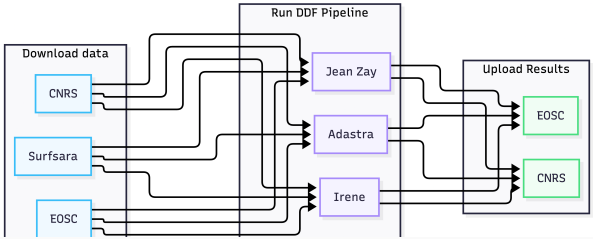
MCP Service Discovery



- Convert natural language request to structured request
- Request Digital Twin (Ontology) for HAPS
- Similarity with Request (embeddings)
- Late-binding with real-time data (machine availability, ...)

Runtime and energy consumption prediction

- The Problem: Resource Sizing
  - Non-expert users → Difficulty translating scientific parameters into SLURM requests.
  - Risk: Job failure (under-allocation) vs. Resource waste & priority loss (over-allocation).
- The Solution: Continuous Learning Predictor
  - Estimation of Runtime, Memory footprint, and Energy consumption.
  - Adaptability to variable application parameters and input data.
- The Goal: Optimizing cross-facility workflows
  - Task placement by identifying the best Execution Time vs. Energy trade-off.



MCP Based Workflows

- Specifying workflows and tasks in natural language
- Based on the Exa-AToW ontologies (Workflow, etc.) and CDT
- Cross-facilities workflow optimization[1]

- See also "LARA: Validation-Driven Agentic Supercomputer Workflows for Atomistic Modeling", William Dawson, Louis Beal, Yoann Curé, Giuseppe Fisicaro, Dorian Rolland, Luigi Genovese, CEA[2]

[1] https://arxiv.org/abs/2609.07275
[2] https://arxiv.org/pdf/2604.22571

Reproductibility and Tracability

Provenance : the Lightcone ASTRA specification approach

- Explicit scientific choices - ASTRA (Agentic Schema for Transparent Research Analysis), a YAML spec linking inputs and outputs to decision. Written by hand or via agent skills : lc-new, lc-from-paper, lc-from-code
- Traceable configurations: named "universes" (decision sets) link results to specific analysis configurations and make them comparables.
- Recorded execution: each output carries its command, environment, Git revision and content fingerprints
- Reusable research records: results and provenance persist with Workflow Run RO-Crate providing an interoperable publication view.

Specification ASTRA : `https://astra-spec.org/latest/specification/`

Lightcone Research : `https://github.com/LightconeResearch`

References

- Mathis Certenais, François Bodin, Benjamin de Zordo. Enabling Radio Astronomy to Use French Supercomputers. 2026 IEEE International Conference for High Performance Computing, Networking, Storage and Analysis Workshops (SC-W) - 21st Workshop on Workflows in Support of Large-Scale Science (WORKS26), Nov 2026, Chicago (Illinois), United States. ⟨hal-05784277⟩
- Mathis Certenais, Laurent Morin, François Bodin. A description of the radio astronomy data processing tool DDF Pipeline. 2026. ⟨hal-05519244⟩
- Marius Garénaux, François Bodin. How Much Does Submitting a Burst of Jobs Impact Waiting Time? An Analysis of the Jean Zay Workload Trace. 2026. ⟨hal-05783397⟩
- all Exa-AToW Documents `https://exatow.fenix.rudi-univ-rennes1.fr/` (login and passwd: exa-atow)