EPICURE

Unlocking European-level HPC Support

SUPPORTING USERS RESEARCH

academia & industry

EuroHPC JU Supercompting Ecosystem

Roadmap for the next 10 minutes

The problem: fragmented, ad-hoc HPC support

The model: what EPICURE is

The flow: how support works end-to-end

Evidence: who it helps + early impact

What’s next + ways to engage

Q&A

Consortium

CSC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CINECA

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

KTH VETENSKAP OCH KONST

IZUM
Institute of Information Science, Maribor, Slovenia

LUX PROVIDE
NEXT-LEVEL INSIGHTS WITH SUPERCOMPUTING

INESC TEC

CYFRONET

DTU Technical University of Denmark

Jožef Stefan Institute Ljubljana, Slovenia

Universiteit Antwerpen

SOFIA TECH PARK

INES

GENCI
HPC at the service of knowledge

JÜLICH Forschungszentrum

EPICURE: the three pillars

01

Access — one simple entry point that routes requests to the right experts.

02

Delivery — cross-site teams led by a named lead; time-bounded, outcome-focused.

03

Reuse — we package results (containers, playbooks, lessons) so others benefit.

EPICURE: consortium

Our research organizations

- CSC – IT Center for Science +
- Cineca Consorzio Unteruniversitario +
- Institut Informacijskih Znanosti +
- INESC TEC +
- Universiteit Antwerpen +
- Academic Computer Centre CYFRONET AGH +
- Forschungszentrum Julich Gmbh +
- Institut Jozef Stefan +
- Barcelona Supercomputing Centre +
- VSB – Technical University of Ostrava +
- Luxprovide SA +
- Sofia Tech Park JSC +
- Kungliga Tekniska Hoegskolan +
- Danish e-Infrastructure Consortium +
- Grand Equipement National de Calcul Intensif +
- Centre Informatique National de L'Enseignement Superieur +

Our supercomputers

- LUMI +
- LEONARDO +

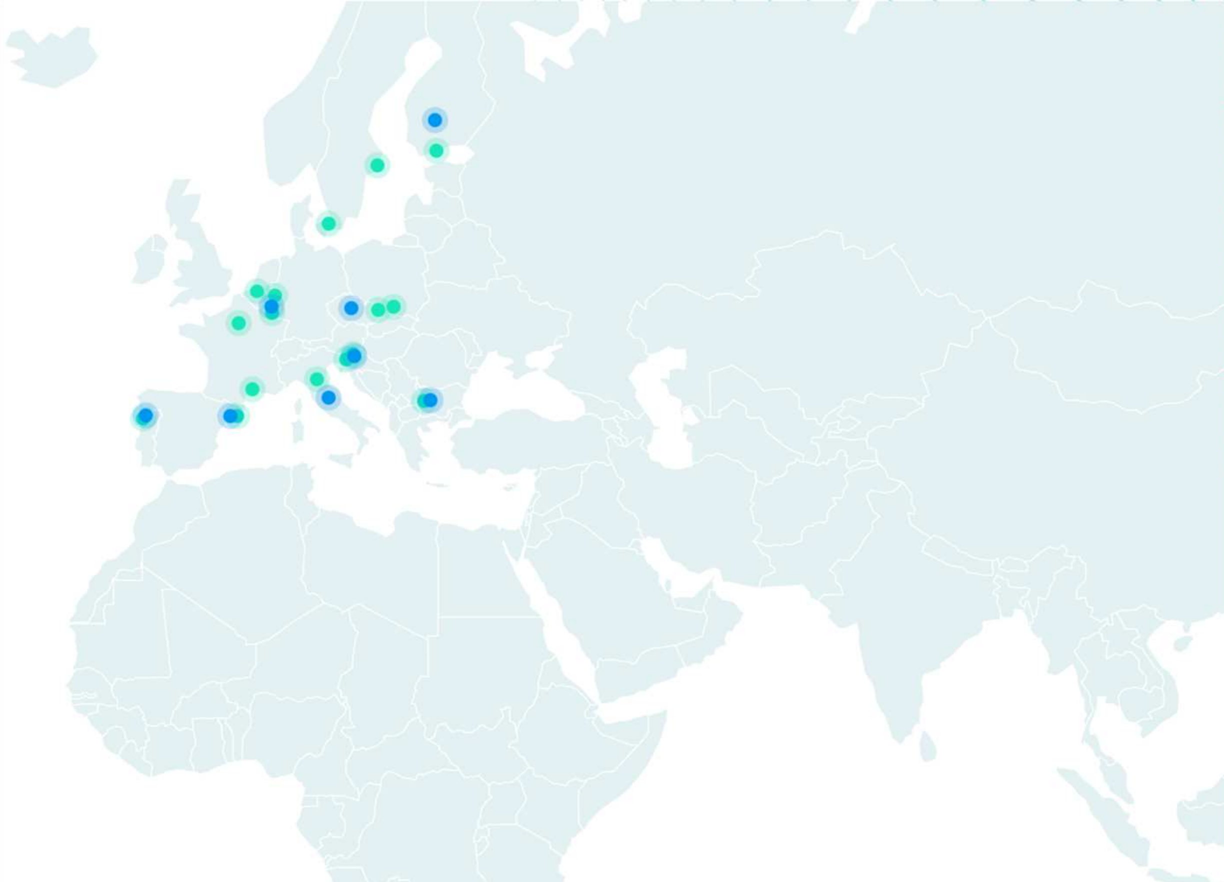

Who we serve

Academia: researchers, PIs, students, research software engineers.

Industry: startups, SMEs, large R&D teams.

Public labs and national centers collaborating across sites.

Cross-disciplinary communities needing scalable HPC expertise.

Operating principles

User-first scoping and time-bounded engagements.

Reuse by default: document once, apply widely.

Cross-site collaboration with clear roles (lead, contributors, artefact owner).

Transparent metrics and continuous feedback.

How EPICURE support flows

Start: project enters via a simple support form / single entry point.

Triage: route to the right level (L2/L3) and partner site.

Assign: designate a lead expert and assemble a cross-site team.

Execute: iterate with regular check-ins; track progress centrally.

Share: document outcomes, lessons, and reusable assets.

Application Enablement

AI applications and LLM models support and optimization

28

Enabling, Compilation, porting and Optimization Projects

24

Profile, Performance Analysis and Benchmarking

18

HPC Workflow and Job Management, Containerizatio n, Data and I/O

8

Example: From methane and iron nanoparticles - catalytic process inside a reactor

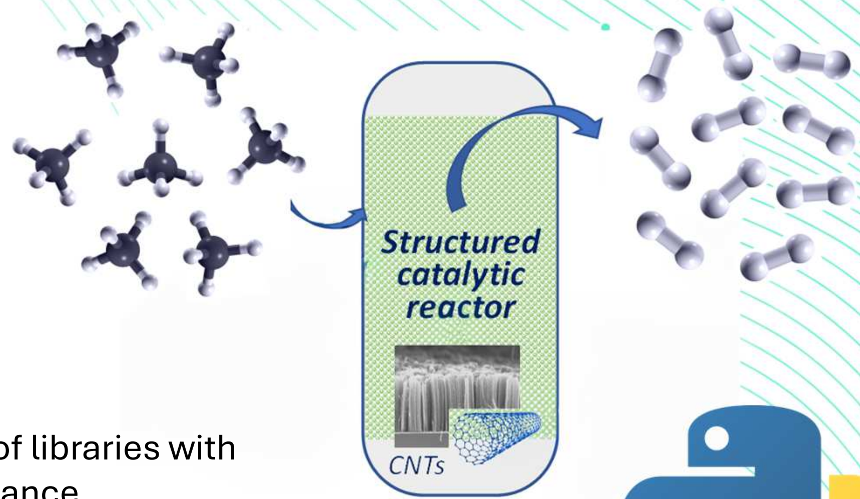

- Bottlenecks in the workflow's pre- and post-processing functions
- Update of the Python environment to later version
- Implementation/extension of the unit tests

Simply the update of the python implementation and the replacement of libraries with more efficient ones produced a huge (20x) improvement of the performance.

Interface	Old Rbf/old environment	old Rbf/new environment	new Rbf/new environment
Ag(111)/Cu(111)	54s	12s	11s
Al(111)/Cu(111)	1h18m	3m38s	2m12s
Ti(001)/Mg(001)	> 17h	43m49s	28m27s

Interface	serial	parallel
Ti(001)/V(110)	2m12s	23s
Ti(001)/Fe(110)	38m45s	5m33s
Ti(001)/Cr(110)	1h0m1s	9m4s

Example: DEAREL: Large-scale Deep rEinforcement leARning for activE fLow control in wings

- Performed in-depth tracing analysis (Nsight) to identify bottlenecks and synchronization issues
- Optimization targeted several kernels (OpenACC Fortran)
  - Memory access patterns (stride, indirect access)
  - Cache utilization
  - Reduction of atomic operations
  - Improved register utilization
  - Loop optimization
- Significant improvement for parallel construct in save_hdf5_restartfile resulting in 1400x speedup

Kernel Name	Time Before	Time After	Speedup
visc_dissipationRate	140.08 ms	72.75 ms	× 1.93
save_hdf5_restartfile	542.27 ms	370.53μs	× 1463
eval_laplacian_diag	30.51 ms	3.41 ms	× 8.95
full_diffuision_ijk_incomp	1.21 ms	1.18 ms	× 1.03
smart_visc_spectral_incomp	2.63 ms	785.41 μs	× 3.35
eval_gradient	701.12 μs	486.08 μs	× 1.44
full_convec_ijk_incomp	2.53 ms	1.40 ms	× 1.81

Early impact

Time-to-engagement reduced from X weeks to Y days for intake sprints.

Performance gains of A--B% on key applications after optimisation.

Reusable artefacts (containers, scripts, playbooks) adopted across sites.

Portfolio Snapshot (End of M18)

Project status (n=136)

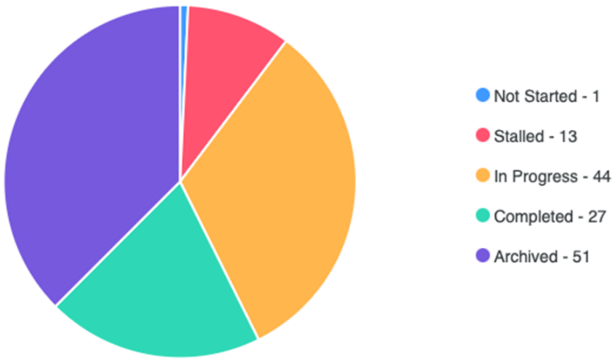

Stalled reasons (n=13)

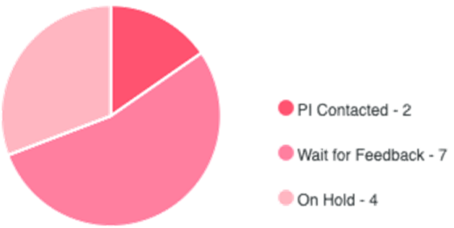

Archived reasons (n=51)

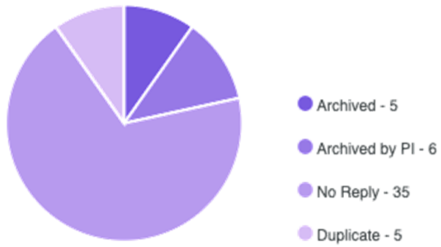

ESM – Assign to Lead Partner or Rejected Avg. time (Days) [N=48]

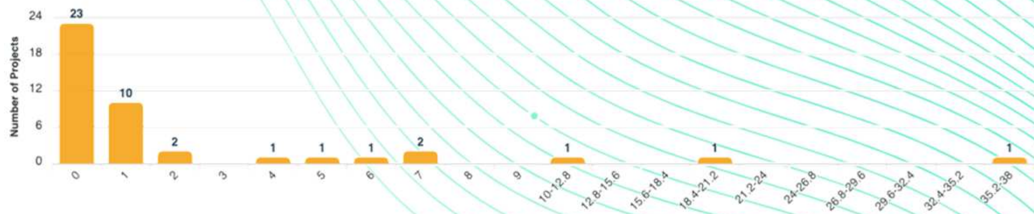

ESM - User Support Received

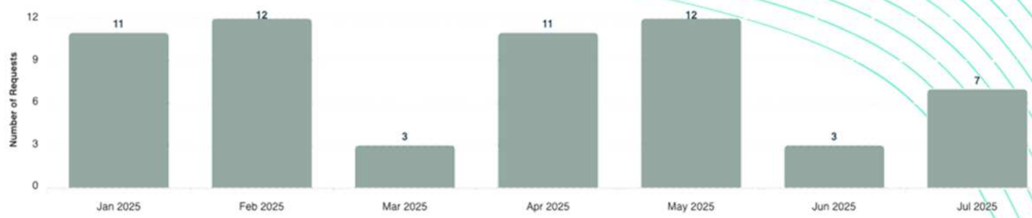

ESM - Person-months per project

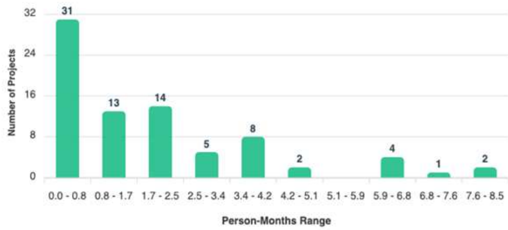

ESM - Project duration (days)

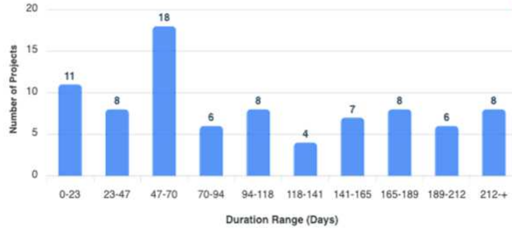

User Feedback Metrics

Helpfulness 4.7/5

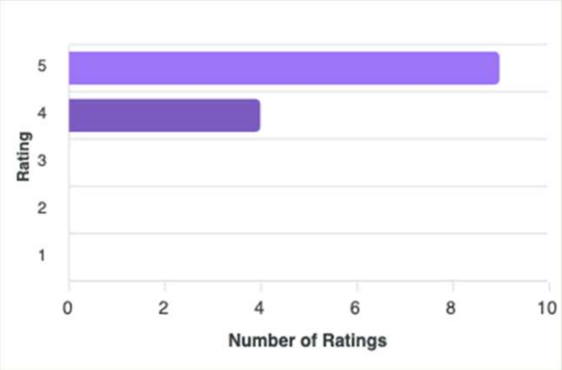

Recommendation 4.8/5

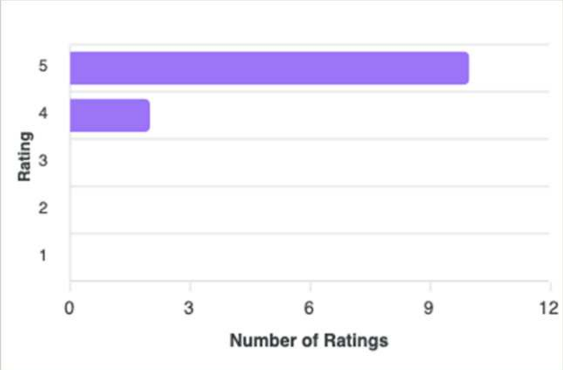

Staff availability 4.7/5

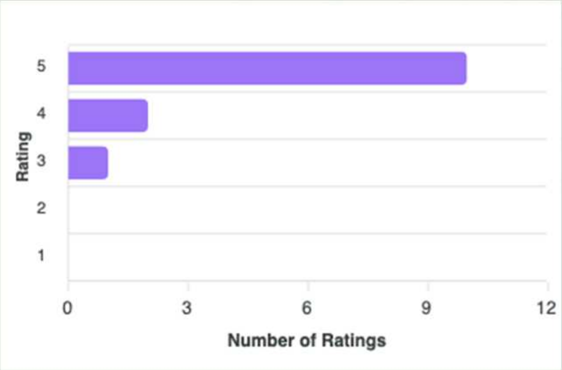

What's next & how to engage

GROW THE EXPERT POOL; LIGHTWEIGHT ONBOARDING AND MATCHMAKING.

PUBLISH A PUBLIC CATALOGUE OF REUSABLE ARTEFACTS AND PLAYBOOKS.

DEEPER INTEGRATION WITH EUROHPC CALLS AND SITE SUPPORT CHANNELS.

CLOSE THE LOOP: SIMPLE USER FEEDBACK, SHARED METRICS DASHBOARD.

We’re Hiring — High Level Support Team

IT4Innovations National Supercomputing Center

What you’ll do

- Tackle complex HPC user cases: profiling, scaling, and code optimization.
- Co-design and tune workflows on national/EuroHPC systems.
- Mentor researchers and deliver hands-on trainings & best practices.
- Collaborate technical publications with domain scientists

What we’re looking for

- Solid C/C++/Fortran and parallel programming (MPI, OpenMP; GPU: CUDA/SYCL/OpenACC a plus).
- Performance analysis & tuning tools (e.g., VTune, Arm MAP, TAU, Nsight).

Why IT4Innovations

- Work directly on cutting-edge Tier-1/Tier-0 HPC systems.
- Impact national and EU research across disciplines.
- Collaborative, international team with real user impact.z

João Barbosa
joao.barbosa@vsb.cz

IT4Innovations National Supercomputing Center
VSB – Technical University of Ostrava
Studentská 6231/1B
708 00 Ostrava-Poruba, Czech Republic
www.it4i.cz

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER