EPICURE

From Code to Scale:

EPICURE's Services, Stories, and Training for EuroHPC Users

POP-3 Webinar • Free for All Attendees

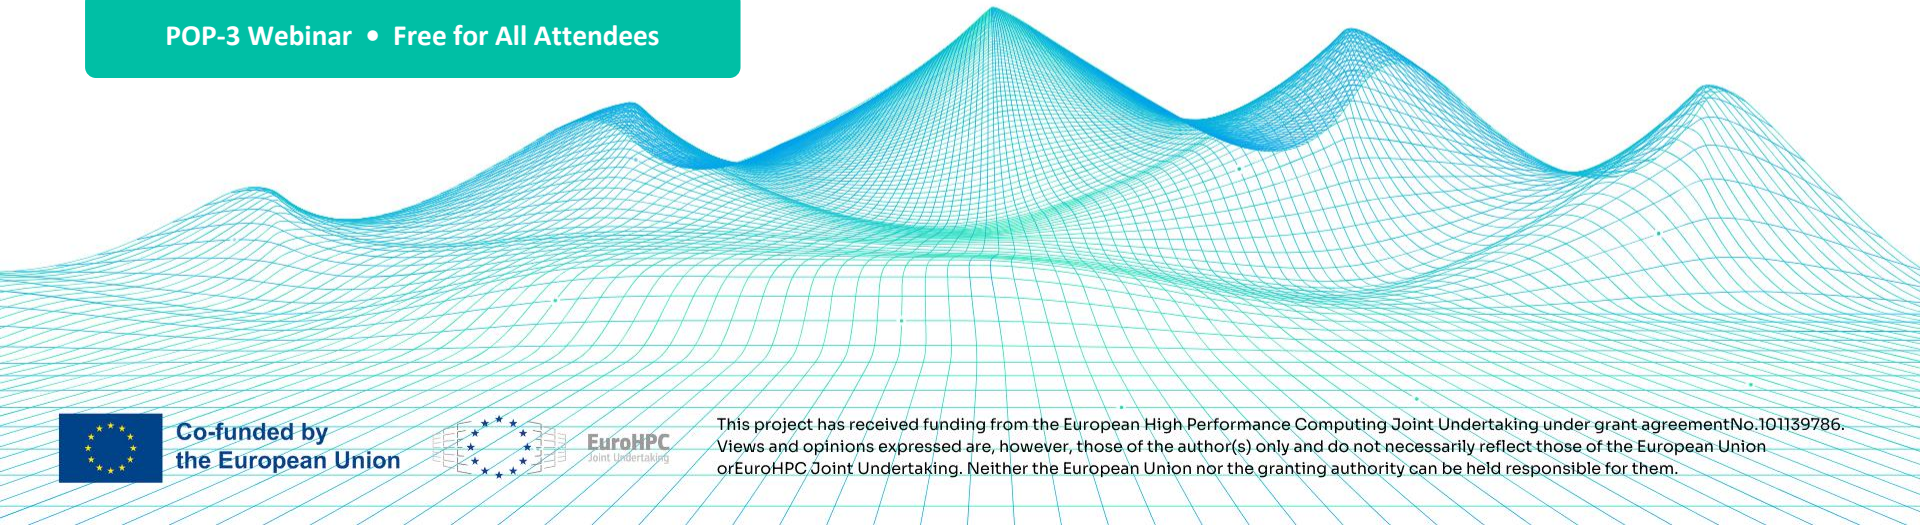

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreementNo.101139786. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union orEuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

Agenda

01 **What is EPICURE?**

Mission, consortium & the 4-year programme

02 **Support Services**

Level 2 & Level 3 — what we do, what we don't

03 **How to Apply**

Eligibility, the submission portal & what to expect

04 **Success Stories**

Real results across energy, climate, health & more

05 **Knowledge Base & Training**

Guides, best-practice docs, code snippets & on-demand materials

06 **Q & A**

Your questions answered live

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

What is EPICURE?

4-year project (Feb 2024 – Jan 2028)

- Coordinated by CSC with 17 partners from European HPC sites
- Provides a distributed network of Application Support Teams (ASTs) for Level 2 and Level 3 support
- Covers: Porting & Optimisation · Advanced Training · Best Practice Guides
- Collaboration at European and global level for pre/exascale applications
- Works with Centres of Excellence and National Competence Centres

Mission: to establish a European HPC Application Support Service and develop an educated, HPC-capable research community.

17
Partner Organisations

120+
Projects Supported

70+
Application Codes Optimised

FREE
For EuroHPC Allocation Holders

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Who Can Apply?

Supported — any project with EuroHPC resources:

- **Extreme Scale Access**
- **Regular Access**
- **Development Access**
- **Benchmarking Access**

Not Supported:

National projects using resources **not** coming from EuroHPC (e.g. resources from purely national allocations).

Tip: If you do not yet have a qualifying EuroHPC project, consider submitting a Development Access proposal and requesting EPICURE application support directly in the form — it's the easiest route in.

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Support Services Overview

Meet our Support Services

KNOW MORE

APPLY FOR EUROHPC ACCESS CALLS

Code enablement and scaling

Support for enabling and increase the scalability of user codes to EuroHPC supercomputers

Performance Analysis

Performance analysis for HPC codes

Benchmarking

Our service focuses on developing a benchmarking suite to evaluate the performance of EuroHPC machines.

Code refactoring

This service involves restructuring or rewriting parts of an application code to improve it maintainability but without changing its function.

Code optimization

Our service aims at improving the efficiency and performance of the software such that it consumes fewer resources

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Level 2 vs Level 3 Support

2nd Level Support | Code Porting, Enabling and Scaling

Work limited to 1-2 months with focus in compilation improvements, vectorization and scalability analysis

3rd Level Support | Code Optimization

Handling large-scale workloads with durations of 2 to 6 months, focused on performance improvements that require code modifications, such as inter-node optimizations, GPU porting and scalability improvements

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

What EPICURE Does NOT Provide

EPICURE does not provide Level 1 Support — that is provided by your local computing centre or EuroHPC helpdesks.

- ✕ General information (e.g. how to generate SSH keys)
- ✕ Login and access to HPC systems
- ✕ Setting up your environment / modules
- ✕ Preparing workflow or job scripts (e.g. SBATCH templates)
- ✕ Standard software installation not related to your project's code

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

How to Apply for EPICURE Support

1

Submit EuroHPC Access Proposal

Apply via the EuroHPC JU open calls at pracecalls.eu. Select Development, Regular, Extreme Scale, Benchmarking or AI access.

2

Request Application Support in the Form

In the web submission form, check:
"Does your proposal require assistance from an AST on the selected partition(s)?" → Select Yes.

3

Project Accepted → Support Begins

Once awarded, you are matched with an Application Support Team expert. No extra application needed — support is automatic for approved projects.

4

Collaborate & Achieve

Work closely with the AST on profiling, optimisation and scaling. Outcomes are documented and shared on the HPC Application Support Portal.

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

EPICURE Impact — First 24 Months

Driving forward the development of the European HPC Application Support Service

120+

Projects Supported

completed or ongoing

70+

Application Codes Optimised

incl. GROMACS, QuantumEspresso, PyTorch

28

AI / LLM Projects

model support & optimisation

18

Profiling & Benchmarking

performance analysis cases

24

Enabling & Porting Projects

compilation & optimisation

8

Workflow & Data / I/O Cases

containerisation, HPC workflows

Success Stories

ENERGY

Methane→H_2 Catalytic Reactor Simulation

EHPC-EXT-2023E01-036 · CINECA on Leonardo

Python implementation updated and replaced with more efficient libraries.

Result: 20× performance improvement

ENVIRONMENT & CLIMATE

North West European Shelf Ecosystem

EHPC-REG-2023R02-047 · IZUM on Vega

Process memory pinning via quiet option applied to multi-node runs.

Result: Up to 7% speedup on real-life workloads

HEALTH & MEDICINE

Heart Language Foundation Model

EHPC-DEV-2025D04-139 · Luxprovide on MeluXina

Distributed training pipeline optimised for multi-node multi-GPU environments: improved communication, memory usage and throughput.

Result: Significantly reduced training time

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Success Story Deep-Dive: Nuclear Physics Code

Project 114 · EHPC-DEV-2025D04-097 · LUMI & MareNostrum 5

Large scale fission properties with energy density functionals

Wouter Ryssens – Institute of Astronomy and Astrophysics – Université Libre de Bruxelles

Objectives

- Benchmark ScaLAPACK for large eigenvalue problems
- Investigate benefit of replacing with ELPA
- Refactor Fortran code for improved maintainability
- Expose code as Python module for ML integration

Support Given (WP2)

- Explore PyScalapack
- Install ELPA and pyELPA on LUMI
- Build PyEV: backend-agnostic wrapper (ScaLAPACK / ELPA)
- Extensive benchmarking across systems

Achievements

- ELPA 2–5× faster for small systems vs ScaLAPACK
- ELPA 10–20× faster for large systems
- Python wrappers enable ML-library integration
- Team convinced to adopt Modern Fortran → Python module approach

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

HPC Application Support Portal

A growing, open knowledge base for the entire European HPC community

How-to Guides

Step-by-step instructions for compiling, running, and tuning codes on all EuroHPC systems.

Best Practice Documents

Recommended workflows, compiler flags, MPI settings, and GPU configurations for major HPC applications.

Code Snippets

Reusable scripts for profiling, benchmarking, job submission, and common HPC patterns.

On-Demand Training Materials

Recorded lectures, hackathon content, workshop slides and exercises — available 24/7.

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Example: Extreme Scalability in EPSILOD

Projects 51 / 119 / 118 / 153 · EHPC-DEV · LUMI · WP2

A framework for generic iterative loop stencil application

Arturo Gonzalez-Escribano · Dpto. de Informática, Universidad de Valladolid

Objectives
- Code porting & scaling issues on LUMI GPU partition

Support (WP2)
- Compilation Celerity / Controllers
- Optimizations: hipHostMalloc, async faults

Achievements
- Code successfully running on LUMI

Lessons Learned
- hipcc/Cray runtime differences require careful testing
- Ask detailed questions at project start

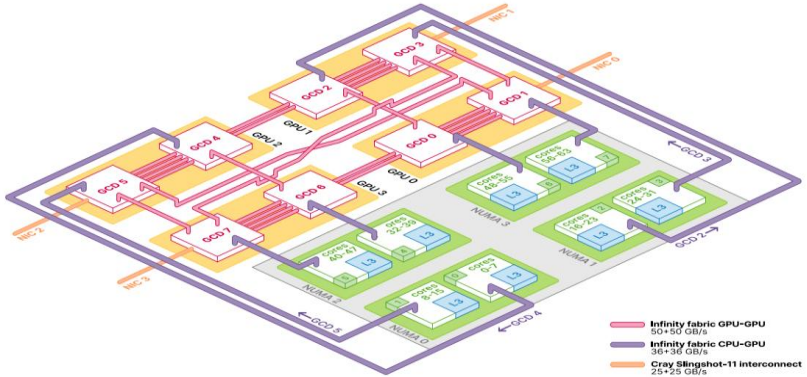

LUMI-G node: AMD MI250X GPU topology & Infinity Fabric

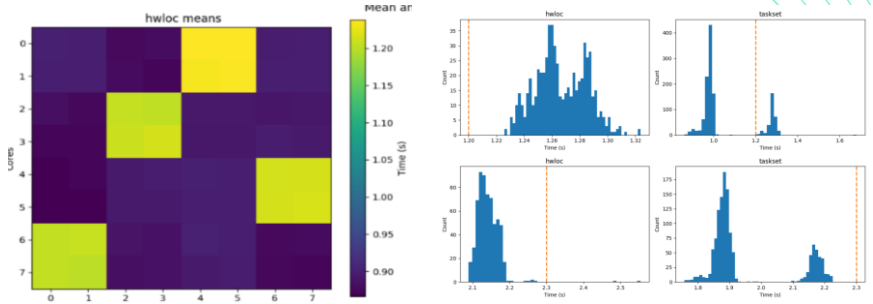

hwloc core-affinity heatmap (left) and timing distributions hwloc vs taskset (right)

Co-funded by the European Union
EuroHPC Joint Undertaking
This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Example: HIDALGO 2 — Wildfire Deployment 2024

Project 67 · EHPC-DEV-2024D09-058 · LUMI · WP2 / WP3

Specific MicroHH configuration for real-time wildfire atmosphere modelling

HIDALGO 2 project · Coupled WRF-SFIRE / µHH weather-fire simulation on LUMI

Objectives

- Run coupled WRF-SFIRE and MicroHH (µHH) on LUMI
- Support real-time wildfire spread simulation workflows

Support Given (WP2/WP3)

- EasyBuild recipe for WRF-SFIRE on LUMI
- Added MicroHH (µHH) EasyBuild module to LUMI Software Library
- Module now available to all LUMI users

Achievements & Lessons

- µHH EasyBuild added to LUMI Software Library
- Dependency chain for both codes fully resolved
- Lesson: document dependencies early
- Lesson: always include working

LUMI Software Library ✓

WRF-SFIRE — user-installable

EasyBuild recipe contributed by EPICURE

LUMI Software Library ✓

MicroHH (µHH) — user-installable

EasyBuild module added to LUMI Software Library

Both codes now user-installable via LUMI Software Library — outcome contributed by EPICURE support

Co-funded by the European Union

EuroHPC

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Example: CANONICS — Charge & Current on Cosmic Strings

Project 103 · EHPC-REG-2024R02-178 · Leonardo · WP2 / WP3

Distributed FFT acceleration for cosmic string network s

Carlos Martins · Instituto de Astrofísica e Ciências do Espaço, Universidade do P

Objectives

- Implement distributed 3D FFT on Leonardo using HeFFTe
- Maximise GPU bandwidth across A100 nodes

Support (WP2/3)

- HeFFTe installation & configuration on Leonardo
- Development tips and communication pattern guidance

Achievements

- Code running on Leonardo A100 GPUs
- A100 bandwidth: ~1,320 GB/s — near theoretical peak

Lessons

- PI should provide runnable code at project start
- A working minimal example accelerates support significantly

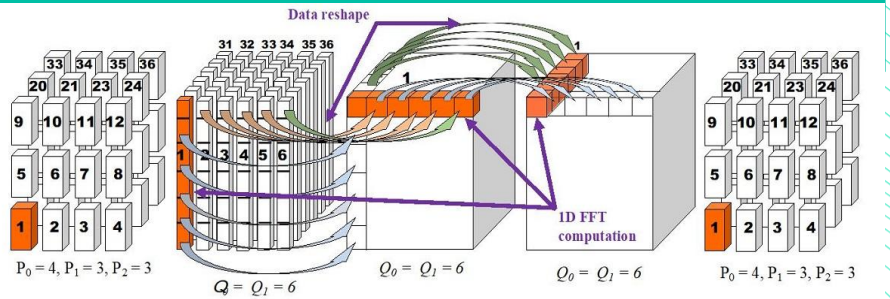

Distributed 3D FFT communication pattern (HeFFTe)

GPU	Type	Time (ms)	Bandwidth (GB/s)
A100	float3	1.184	1,319
A100	float4	1.462	1,336
MI250X	float3	1.347	1,160
MI250X	float4	1.687	1,158
P4000	float3	8.423	186
P4000	float4	10.18	192

GPU bandwidth comparison across EuroHPC systems (1D FFT kernel, float3/float4)

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

Interaction with POP3 Centre of Excellence

Start your analysis in POP3 — bring the performance bottlenecks identified to EPICURE for expert porting & optimisation support

1

Request a POP3 Performance Assessment

POP3 evaluates code performance & scaling, quantifies inefficiencies, and delivers expert recommendations — free of charge for all EuroHPC users.

2

Identify Problems: Scaling, Architecture Gaps & Optimisation Needs

POP3 reports pinpoint bottlenecks: parallelisation inefficiencies, memory issues, energy use, and readiness gaps for EuroHPC pre/exascale systems.

3

Bring Identified Problems to EPICURE for Level 2 / Level 3 Support

Apply for EPICURE support to address the issues — porting to new architectures, deep code optimisation, scalability work, and best-practice documentation.

Tip: POP3 and EPICURE are complementary services. POP3 diagnoses the problem — EPICURE implements the fix. Reference your POP3 report in your EPICURE support request.

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No.101139786.

EPICURE
Unlocking European-level HPC Support

Thank you!

pmo-epicure@postit.csc.fi

Follow us

Apply for support → pracecalls.eu

Explore guides → epicure-hpc.eu

Contact → pmo-epicure@postit.csc.fi

Co-funded by the European Union

EuroHPC Joint Undertaking

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreementNo.101139786. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union orEuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.