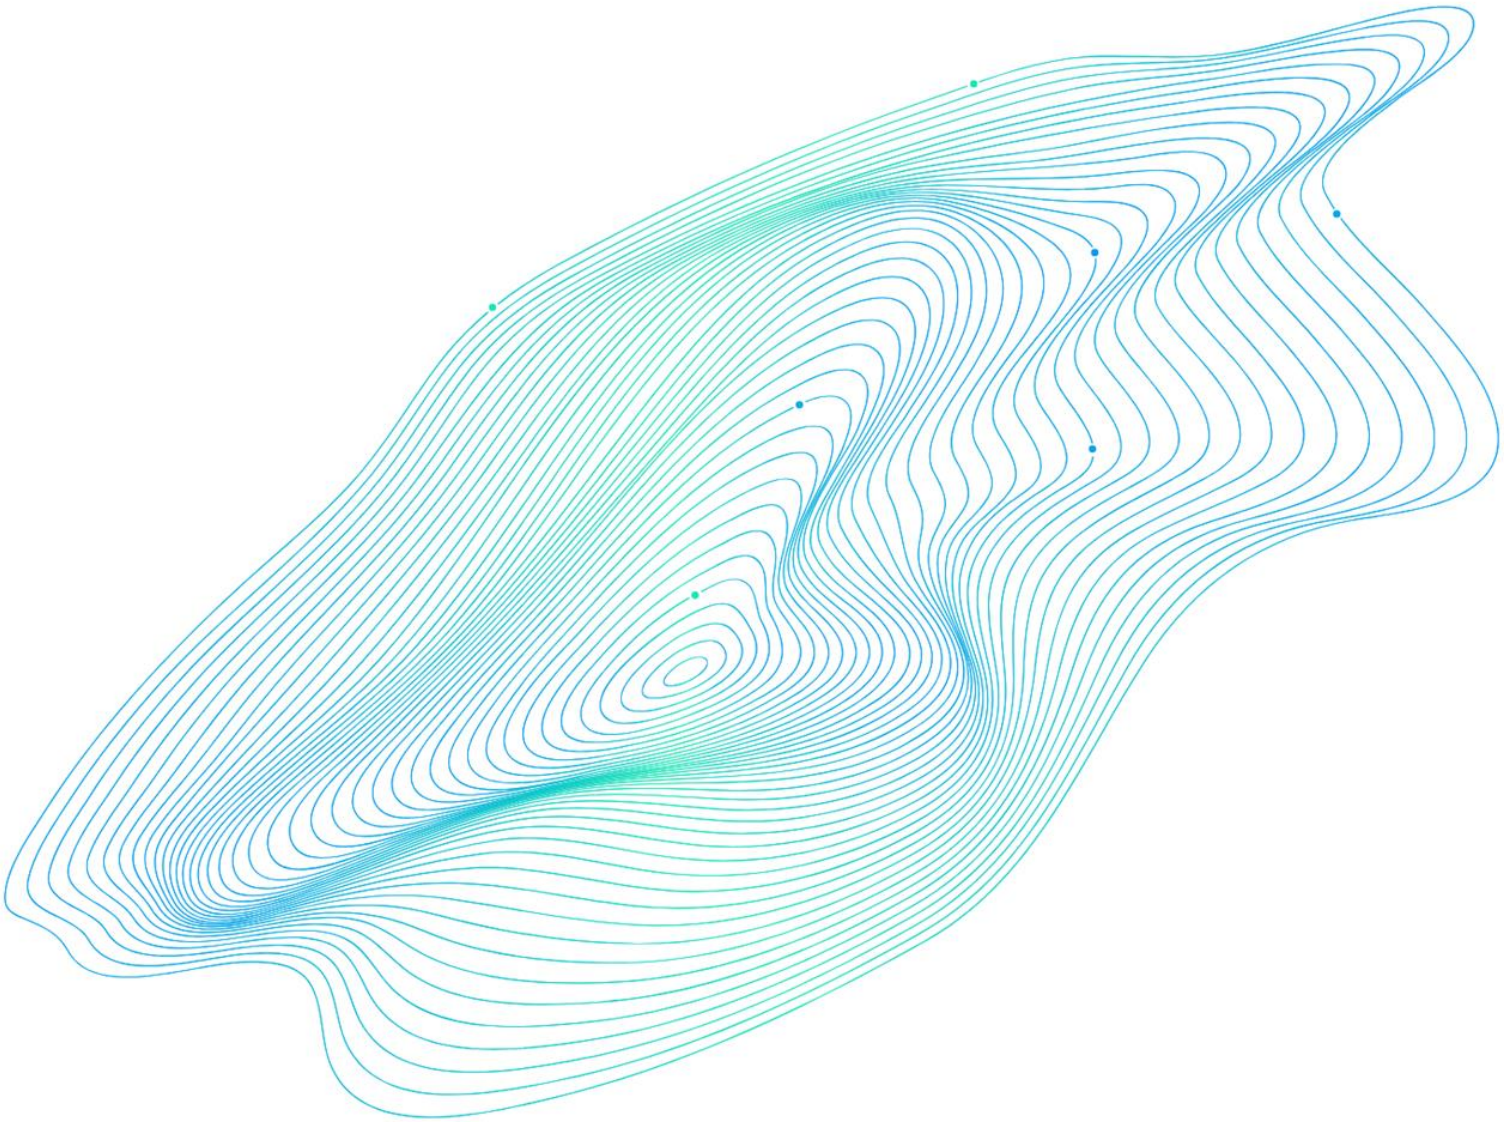
EPICURE

Unlocking European-level HPC Support

HIGH-LEVEL SPECIALISED APPLICATION SUPPORT SERVICE IN HIGH-PERFORMANCE COMPUTING (HPC)



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. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

EPICURE

DELIVERABLE 5.2

Training activities report year 2



Project Title	High-level specialised application support service in High-Performance Computing (HPC)
Project Acronym	EPICURE
Project Number	101139786
Type of Action	DIGITAL JU Simple Grants
Topic	DIGITAL-EUROHPC-JU-2022-APPSUPPORT-01-01
Starting Date of Project	2024-02-01
Ending Date of Project	2028-01-31
Duration of the Project	48 months
Website	epicure-hpc.eu

Work Package	WP 5 | Advanced Training
Task	5.3 | Communication, Coordination, and Reporting
Lead Authors	António Sousa (INESC TEC) Cláudia Brito (INESC TEC)
Contributors	Marcos Gouveia (INESC TEC)
Peer Reviewers	Alberto Lanzanova (CSC) Klemens Noga (Cyfronet)
Version	1.0
Due Date	31/01/2026
Submission Date	29/01/2026

Dissemination Level

X	PU: Public
	SEN: Sensitive – limited under the conditions of the Grant Agreement
	EU-RES. Classified Information: RESTREINT UE (Commission Decision 2005/444/EC)
	EU-CON. Classified Information: CONFIDENTIEL UE (Commission Decision 2005/444/EC)
	EU-SEC. Classified Information: SECRET UE (Commission Decision 2005/444/EC)

Co-funded by the European Union
EuroHPC Joint Undertaking

Version History

Revision	Date	Editors	Comments
0.1	19/12/2025	António Sousa (INESC TEC)	Table of Contents
0.2	03/01/2026	António Sousa (INESC TEC) Cláudia Brito (INESC TEC)	First Complete draft of the deliverable
0.3	09/01/2026	Alberto Lanzanova (CSC) Klemens Noga (Cyfronet)	Review conducted by the appointed reviewers
0.4	18/01/2026	António Sousa (INESC TEC) Cláudia Brito (INESC TEC)	Reviewers input consolidated and second draft sent to the Project Management Board (PMB)
0.5	22/01/2026	Virvatuli Uusimäki (CSC) Alberto Lanzanova (CSC) Joao Barbosa (IT4I)	Project Management Board feedback on the deliverable received
0.6	25/01/2026	António Sousa (INESC TEC) Cláudia Brito (INESC TEC)	Project Management Board feedback consolidated and final version of the deliverable made available to the Project Management Office
1.0	29/01/2026	Alberto Lanzanova (CSC) Virvatuli Uusimäki (CSC)	Final quality check performed by the Project Management Office and submission of the deliverable

Declaration on the use of AI

This report has been prepared with the support of AI tools which were used to assist in drafting text and checking language. All content has been thoroughly reviewed, fact-checked, and edited by the authors to ensure accuracy and alignment with the objectives of the report.

Co-funded by the European Union
EuroHPC Joint Undertaking

Glossary of Terms

Item	Description
AI	Artificial Intelligence
ARM	Advanced RISC Machine
AST	Application Support Team
CoE	Centre of Excellence
CPU	Central Processing Unit
EESSI	European Environment for Scientific Software Installation
ESM	EPICURE Support Manager
FPGA	Field Programmable Gate Array
GPU	Graphical Processing Unit
HPC	High-Performance Computing
MPI	Message Passing Interface
NCC	National Competence Centre
NCCL	NVIDIA Collective Communication Library
PCI	Peripheral Component Interconnect
PI	Principal Investigator
WP	Work Package

Co-funded by the European Union
EuroHPC Joint Undertaking

Executive Summary

This document serves as a report of the activities organised by EPICURE's Advanced Training Work Package between months 13 and 24 of the project. During the second year of the project, the primary objective was to implement the developed training plan.

From the 11 proposed webinars in the training plan, 5 were planned and organised[1]. This year, two hackathons were organised: one in-person and the other in a hybrid format. The hackathons have been organised in the "Bring Your Own Code" hackathon style.

The first EPICURE hackathon of the year occurred in February 2025 in Braga, Portugal. This hackathon had as its central theme "HPC in ARM architecture", where the participants learned about the basics of the use of ARM processors in HPC, besides getting the chance to run their code on the ARM partition of Deucalion.

The Second hackathon took place in December 2025 in Stockholm, Sweden, and was a hybrid event, with some participants and mentors on-site, while the majority participated online. This was a successful experience, which allowed us to gather participants and mentors who would have been very difficult to recruit on-site otherwise.

The webinars have been recorded. Materials and videos are freely available on the European HPC Application support portal[2]. Even though we have not yet reached the desired level of participation in the post-event feedback form, which makes it challenging to assess the effectiveness of the training program, the overall outcome of this year's training activities has been positive.

[1] A sixth webinar was scheduled for 2025 with the participation of the EuroHPC JU Peer Review Team, but it was cancelled due to the low number of registrations. Webinars not organised in Year 2 were postponed to Year 3, as summarised in the training plan.
[2] European HPC Application Support Portal: https://eurohpcsupport.eu/knowledge-base/training-material/

Co-funded by the European Union
EuroHPC Joint Undertaking

Table of Contents

1. INTRODUCTION 8
2. DEVELOPMENTS IN 2025 9
2.1 EPICURE METHODOLOGY IN PRACTICE 9
2.2 SCOPE AND STRUCTURE OF THE EPICURE ADVANCED TRAINING PLAN 11
2.3 ADVANCED TRAINING AT THE EUROPEAN HPC APPLICATION SUPPORT PORTAL 12
3. WEBINARS 14
3.1 CONTAINERIZATION IN HPC ENVIRONMENTS 14
3.2 UNLOCKING EXASCALE: A FIRST LOOK AT JUPITER 15
3.3 CPU-BASED THREADING MODELS APPLICATION IN PYTHON FOR HPC COMPUTING 17
3.4 PYTHON ON GPUS, HOW HARD CAN THAT BE? 18
3.5 HOW TO PROFILE YOUR PYTHON CODE 20
4. HACKATHONS 22
4.1 HPC IN ARM ARCHITECTURE 22
4.2 EPICURE HACKATHON: PYTHON MEETS HPC 27
5. CONCLUSIONS AND FUTURE WORK 31
APPENDIX A.YEAR 2 TRAINING PLAN 32
1. INTRODUCTION 2
1.1 TRAINING PLAN DEFINITION GUIDELINES 2
1.2 TARGET AUDIENCE 3
1.3 SUGGESTED TOPICS 3
1.4 SELECTION OF TRAINING EVENTS 4
2. PROJECT YEAR 2 TRAINING PLAN 6
3. CONCLUSIONS AND FUTURE WORK 12
APPENDIX B.YEAR 3 TRAINING PLAN 1
1. INTRODUCTION 3
2. PROJECT YEAR 3 TRAINING PLAN 4
2.1 TARGET AUDIENCES 4
2.2 YEAR 3 TRAINING PLAN 4
3. CONCLUSIONS AND FUTURE WORK 25

Co-funded by the European Union
EuroHPC Joint Undertaking

1. Introduction

The EPICURE training program, for which the Advanced Training Work Package (WP5) is responsible, is closely linked to the user support activities provided by the Application Support Teams (ASTs). While more technical work packages, such as WPs 2, 3, and 4, assist researchers who have been awarded computing time via a EuroHPC JU call, WP5's mission is to create an efficient training plan that empowers users by increasing their knowledge of High-Performance Computing concepts. The goal is that, in the long run, the training provided in the context of EPICURE will bidirectionally improve the high-level support experience. Since the training program also has application support specialists in mind, both parties involved in user support can benefit from it, facilitating effective communication during discussions and making user support more efficient.

During this period, the Advanced Training WP organised two types of training events: Webinars and Hackathons. Webinars are live, online events where an invited speaker delivers a presentation on a specific topic within their area of expertise. In a Hackathon, several participants are selected to participate in multiple-day, in-person or hybrid events where they can learn about a particular subject and, most importantly, bring their own code. They receive the support of mentors who help them profile, port, and optimise their application.

The remaining of this document is structured as follows. Section 2 describes the developments achieved during the reporting period. Sections 3 and 4 report the webinars and hackathons organised during the reporting period. Finally, Section 5 presents the conclusions regarding the outcome of the training activities, the difficulties encountered during implementation of the training plan, and a mitigation plan for the future. Two appendices complement this deliverable, containing the detailed training plans for Year 2 (2025) and Year 3 (2026).

Co-funded by the European Union
EuroHPC Joint Undertaking

2.Developments in 2025

Deliverable D5.1 introduced a structured methodology to support the systematic definition of advanced training activities within EPICURE. The methodology combines qualitative and quantitative inputs from multiple sources, namely project feedback, Application Support Team (AST) experience, project technical meetings, and information collected through the EPICURE Support Manager (ESM), to ensure that training activities are closely aligned with user needs, project maturity, and emerging trends in High-Performance Computing (HPC).

This methodology was first applied in preparing the 2025 training plan (Appendix A), providing a coherent framework for translating heterogeneous inputs into a consistent, balanced set of webinars, hackathons, and workshops. Despite the limited amount of historical data available at that stage of the project, the methodology proved effective in guiding topic selection, defining training formats, and establishing a regular cadence of activities. The experience gained in 2025 further validated the approach and enabled its application more effectively and with greater confidence in preparing the 2026 training plan, benefiting from richer feedback (i.e., from webinar and hackathon attendees), a larger portfolio of supported projects, and more mature interactions with ASTs and users.

2.1 EPICURE methodology in practice

In 2025, the methodology was applied more systematically by exploiting data collected through the EPICURE Support Manager (ESM). Information extracted from ESM, e.g., requests, project interactions, and technical discussions, was analysed and summarised into a word cloud (Figure 1) representing the most recurrent technologies, methodologies, and challenges faced by supported projects. This word cloud, combined with inputs and experience from EPICURE's team members and collected feedback from the events, was used to help in the internal discussions that led to the definition of the 2026 training plan.

Co-funded by the European Union
EuroHPC Joint Undertaking

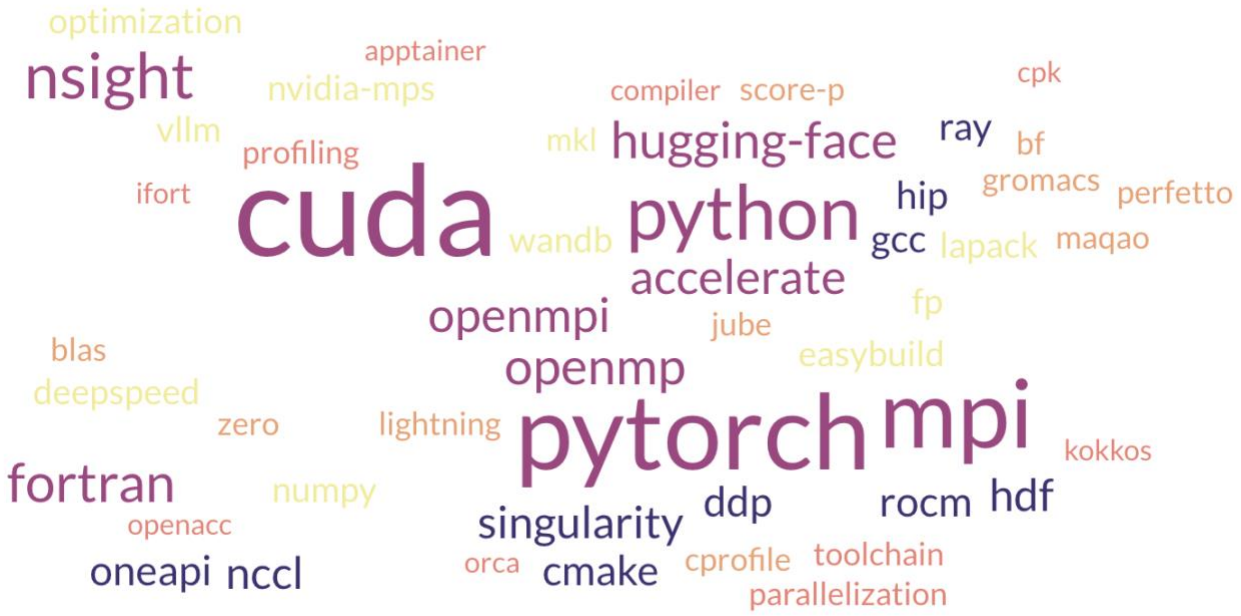


Figure 1 - Word cloud

Based on this analysis and discussions, three main training tracks were identified to structure the training plan:

1. Tools, covering core HPC software infrastructure, performance analysis, optimisation, and workflow support;
2. Scientific tools, focusing on domain-specific applications and libraries widely used by EPICURE-supported projects;
3. AI for Science, addressing the growing convergence between AI methodologies and large-scale scientific computing.

The knowledge derived from the word cloud, combined with input from the EPICURE team members, was then used to define advanced training subjects to be delivered through webinars, hackathons, and workshops. In addition, recognising that EPICURE internal technical meetings are also key moments for sharing progress and results from supported projects, WP5 decided to actively participate in these meetings. This participation enables the identification of particularly relevant themes that are well-suited for broader dissemination through dedicated webinars. Once identified, these dissemination-oriented webinars are dynamically added to the training plan. This methodology allows for a more project-centric approach, enabling AST teams to also be fully involved in the training plan initiatives

Co-funded by the European Union
EuroHPC Joint Undertaking

2.2 Scope and Structure of the EPICURE Advanced Training Plan

Building on the portfolio of EPICURE-supported projects, alongside the Year 2 training plan activities not yet delivered, current trends in High-Performance Computing (HPC), input from EPICURE team members, medium-term technological perspectives and users' feedback, WP5 has designed a comprehensive advanced training plan for Year 3. This plan is structured around three complementary training tracks: (i) Tools, (ii) Scientific Tools, and (iii) AI for Science. These tracks will be delivered through a combination of webinars, workshops, and hackathons, ensuring both regular knowledge dissemination and in-depth, hands-on experience.

The training tracks serve as the backbone of the program, maintaining a regular cadence throughout the year to sustain user engagement and collaboration with Application Support Teams (ASTs). Webinars will provide balanced coverage of both horizontal HPC skills and domain-specific topics, addressing essential tools and methodologies such as fault tolerance, software stack management, performance portability, I/O optimisation, profiling, and optimisation techniques. Simultaneously, there will be a strong emphasis on scientific applications and domain-specific tools—including Fortran, OpenFOAM, GROMACS, and AlphaFold—to reflect the needs and interests of active EPICURE-supported projects. A dedicated AI for Science track will further address the growing convergence between HPC and AI. Topics will include distributed PyTorch training, Hugging Face ecosystems, DeepSpeed, large-scale inference with vLLM, and AI-driven scientific workflows.

Complementing the webinars, the training plan will include two hackathons organised under a "bring your own code" framework and delivered in a hybrid format. This approach is informed by lessons learned from the EPICURE Hackathon Python Meets HPC, held in December 2025, which demonstrated that hybrid events significantly increase accessibility and participation without compromising technical depth, as long as projects and objectives are clearly defined. The experience also highlighted the importance of actively involving EPICURE-supported projects and application support projects, as their participation enhances the relevance of the challenges addressed and facilitates effective knowledge transfer between users, developers, and support teams. Building on these lessons, one hackathon will focus on code optimisation and portability for heterogeneous HPC systems, as the growing adoption of ARM-based and accelerator-rich architectures continues. The second hackathon will target the effective use and optimisation of scientific and domain-specific tools. Together, these events promote peer learning, mentoring, and the translation of training outcomes into tangible improvements in the performance, scalability, and sustainability of HPC applications.

The program will be further enhanced through in-depth workshops and tutorials that provide extended, hands-on training on strategically important topics. These include preparing applications for the exascale era and addressing scalability, performance, portability, and system-level challenges. Workshops will also be dedicated to exploiting the unique features of the MeluXina supercomputer, with particular emphasis on FPGA-based acceleration. This will enable participants to explore alternative computing paradigms and energy-efficient acceleration technologies,

Co-funded by the European Union
EuroHPC Joint Undertaking

alongside tutorials on distributed deep learning workflows and reproducible bioinformatics pipelines using Nextflow.

Overall, the EPICURE Year 3 training plan (presented in Appendix B) reflects a mature and balanced approach to advanced HPC training. By integrating regular dissemination activities with targeted, project-driven events and aligning training topics with user demand, technological advancements, and the EPICURE support ecosystem, the program contributes to building sustainable skills, improving application performance and portability, and strengthening the European HPC user and support communities.

2.3 Advanced Training at the European HPC Application Support Portal

The European HPC Application Support Portal (accessible at https://eurohpcsupport.eu) plays a central role in extending and amplifying the impact of EPICURE's advanced training activities within the EuroHPC JU ecosystem. Developed under the EPICURE project, the portal serves as a centralised hub for HPC support and training materials for both new and advanced users of EuroHPC JU supercomputing systems.

At its core, the portal consolidates and disseminates knowledge resources, best-practice guides, case studies, how-to documentation, and training materials that are directly linked to the skills and competencies targeted by EPICURE's advanced training programme. These resources include recordings, slide decks, and supplementary content from webinars, hackathons, and workshops delivered as part of the Advanced Training Work Package, making them persistently accessible to the broader EuroHPC JU community beyond the live events themselves.

Concretely, the portal supports the EPICURE advanced training in several strategic ways:

- Knowledge Continuity: It functions as a long-term repository of training materials and best practices, allowing participants to revisit content and enabling users who could not attend live events to benefit from the same materials.
- Skill-Building Ecosystem: By integrating training outputs with guidance on accessing EuroHPC JU systems, architecture overviews, and support services available through EPICURE, the portal bridges the gap between theoretical training and practical application. Users can immediately apply what they have learned to their own HPC allocation projects, fostering a more effective transfer of skills.
- Alignment with Project Support: The portal's training Section is closely aligned with the EPICURE support pipeline; as projects gain access to EuroHPC JU systems and request support, they are guided toward relevant training materials that match their needs, whether that be performance analysis, optimisation methodologies, or tools for code scaling. This integrated experience enhances the overall value of both support and training activities.
- Community Empowerment and Outreach: By aggregating resources in a user-friendly and searchable format, the portal enables broader dissemination of advanced training content across scientific disciplines, geographies, and experience levels. This supports the EuroHPC

Co-funded by the European Union
EuroHPC Joint Undertaking

JU's broader mission to build a strong, collaborative European HPC community with deep technical capabilities and cross-institutional knowledge sharing.

In summary, the EuroHPC Support Portal is not merely a repository but a strategic enabler of the EPICURE advanced training programme. It ensures that training outputs are reusable, discoverable, and connected to practical support mechanisms, significantly enhancing the reach and longevity of EPICURE's training impact across the European HPC user community.

Co-funded by the European Union
EuroHPC Joint Undertaking

3.Webinars

3.1 Containerization in HPC Environments

3.1.1 Date and Time

This webinar took place on March 31st 2025, from 11:00 a.m. to 12:00 p.m. (CET). The event recording can be found on the Training page of the European HPC Application Support Portal.
Link: https://eurohpcsupport.eu/training-material/containerization-in-hpc-environments/

3.1.2 Description

The webinar on basic containerisation for HPC environments provided an overview of key tools and techniques. It covered how to use Singularity/Apptainer containers from various perspectives, including parallelisation with MPI and CUDA for GPU-accelerated workloads, as well as management and environment systems such as Conda for creating smaller, lightweight environments. The session emphasised practical strategies for optimising HPC workflows using containerised solutions.

3.1.3 Speakers

Teo Prica is an HPC Development Engineer at IZUM (HPC Vega). His interests include Virtualisation and System Orchestration, Containerization, Scientific Software in HPC, and Optimisation with a focus on Machine Learning. He is a member of the HPC Federation Working Group (EGI), the EuroHPC Containers Forum, and member of the Slovenian National Supercomputing Network (SLING) consortium. Additionally, he is a contributor to European and National projects, including EPICURE, interTwin, SMASH, Castiel2, and EUMaster4HPC. He was part of the Deployment of the 1st Operational EuroHPC JU system and the Engineering team of EuroHPC JU Vega, providing high-level SLING User support, Technical Support, and Maintenance for HPC Vega.

3.1.4 Relevance

Containerisation is widely used in modern software engineering, relying on operating system and application virtualisation. A container is a sandbox in which a user can install a base operating system and then prepare the container to run a specific application by installing the necessary dependencies. The result is a single file that can be shared among users, regardless of their machines.

Containerisation is a crucial concept closely related to portability, particularly in high-performance computing. Software distributed in a container is more straightforward to run across multiple HPC systems, even if their architectures and environments differ.

This training activity is considered extremely important for users who want to run a containerised, parallel application and are unfamiliar with how to bind the container's MPI and GPU acceleration to the HPC systems' job submission system. Moreover, containerisation is also very important in AI workflows, as these workflows can generate many small files, quickly drain the user's file quota, and

Co-funded by the European Union
EuroHPC Joint Undertaking

make I/O systems less efficient. If all the small files generated, for example, by Python environments, are confined in a single container file, this problem is avoided. As such, this activity suits user support specialists in both Level 2 and Level 3 roles, as well as users with EuroHPC JU projects, particularly those awarded resources in AI calls.

3.1.5 Summary

- Introduction and outline of the talk.
- Why the rise of AI and the future of AI Factories create the need for containerisation.
- The challenges of producing HPC scientific software and the issue of dependencies.
- What is a container? How is it different from a virtual machine?
- Advantages of using containers for HPC.
- Apptainer and Singularity containers.
- Building and executing a container.
- Managing files and data. Parallel I/O performance.
- Python environments and containers.
- Using containers with different interconnects (InfiniBand, Slingshot, ...).
- Launching a contained application using SLURM and MPI/CUDA.
- Containers at HPC Vega.

3.1.6 Participant Data

- Number of registrations: 35
- Number of participants: 21
- Attendance: 60%
- Gender: Man (81%); Woman (14%); Prefer not to say (5%)
- Type of organisation: Academia (43%); R&D Institution (33%); Company (20%); Other (4%)
- Country of affiliation: 10 different countries.

3.2 Unlocking Exascale: A First Look at Jupiter

3.2.1 Date and Time

This webinar took place on April 30th 2025, from 11:00 a.m. to 12:00 p.m. (CET). The event recording can be found on the Training page of the European HPC Application Support Portal.
Link: https://eurohpcsupport.eu/training-material/unlocking-exascale-a-first-look-at-jupiter/

3.2.2 Description

The JUPITER Exascale supercomputer, Europe's first Exascale system, will soon be ready for large-scale computations within the upcoming allocation period at the Jülich Supercomputing Centre. This webinar provided an overview of JUPITER's system architecture, technical specification, and pathways to gain access to the machine's available compute resources. Furthermore, a performance comparison summary was presented between the JUWELS Booster production system and JEDI, the preparatory system with an identical hardware configuration to the JUPITER Booster. The

Co-funded by the European Union
EuroHPC Joint Undertaking

webinar also presented firsthand experiences in porting a plasma physics simulation code to an ARM-architecture-based system, such as JEDI/JUPITER, building software modules required by the application via EasyBuild, and profiling.

3.2.3 Speakers

Jolanta Zjupa is a research staff member at the Jülich Supercomputing Centre (JSC) in Germany, where she has been focusing on application optimisation, parallelisation, HPC support, and training since 2022. With a background in theoretical and computational astrophysics, she brings a unique perspective to her work in High-Performance Computing. Jolanta holds a PhD in Astrophysics from the University of Heidelberg and has worked as a Postdoctoral Researcher at renowned institutions in France and Austria.

Junxian Chew has been a research staff member of JSC in Germany since 2023. He currently focuses on HPC support, parallel I/O, training, and porting applications. Having a background in aerospace engineering, then pivoting to HPC study of plasma physics in the nuclear fusion topic, he brings a solution-oriented mindset to the field of computational physics in an HPC environment. He holds a PhD in physics from Ruhr-Universität Bochum.

3.2.4 Relevance

When Jupiter enters the production phase, it will officially be the first exascale machine in the EuroHPC JU ecosystem. Naturally, it is expected that many EuroHPC JU project applications will request computing resources from Jupiter; therefore, the pool of future users must be informed about what Jupiter is, as well as its hardware and software configurations. Since ARM processors will exclusively power Jupiter, it is also essential that users are aware of the amount of work necessary to port code from x86 to ARM in a scalable and optimised manner. Since the second EPICURE Hackathon focused on using Deucalion's ARM processors, this webinar complemented the information shared during the Hackathon, while also making clear the differences between Deucalion's ARM partition and JUPITER's architecture.

3.2.5 Summary

- Past, current, and future JSC supercomputing systems.
- JUPITER Booster partition.
- Current progress of HW and SW installation.
- JSC HPC Tools available on JUPITER.
- JSC Support Infrastructure.
- Getting early access to JUPITER: national and European calls.
- JUPITER and EPICURE support services.
- Job monitoring in JUPITER with LLview.
- EasyBuild stack for ARM partitions.
- GH200 unified memory: System vs Managed memory.
- ARM vs x86 architecture.

Co-funded by the European Union
EuroHPC Joint Undertaking

3.2.6 Participant Data

- Number of registrations: 93
- Number of participants: 52
- Attendance: 56%
- Gender: Man (77%); Woman (15%); Prefer not to say (8%)
- Type of organisation: Academia (62%); R&D Institution (23%); Company (10%); Other (5%)
- Country of affiliation: 14 different countries.

3.3 CPU-based Threading Models Application in Python for HPC Computing

3.3.1 Date and Time

This webinar took place on November 21st, 2025, from 3:00 to 4:00 p.m. (CET). The recording of the event can be found on the Training page of the European HPC Application support portal.
Link: https://eurohpcsupport.eu/training-material/webinar-cpu-based-threading-models-application-in-python-for-hpc-computing/

3.3.2 Description

The webinar covered three main approaches to overcome Python's Global Interpreter Lock (GIL) limitation: using compiled C/C++/Fortran code with OpenMP or TBB for maximum performance, implementing PyOMP with Numba for HPC users who prefer Python development, and advanced LLVM integration with GIL-disabled Python for professional developers. Each approach offers different performance gains, ranging from 6x to over 10x speedup compared to sequential Python, with varying skill requirements, from moderate Python knowledge to expert-level understanding of LLVM IR. The webinar was designed for HPC practitioners, researchers, and developers who need to implement parallel computing solutions in Python applications.

3.3.3 Speakers

Veselin Kolev (Discoverer Supercomputer) has a research scientist background. He holds a PhD in Chemical Physics. He has extensive knowledge and experience in designing and applying efficient parallel molecular simulation codes. His most recent scientific position was at the University of Southern California, where he was a member of Professor Ariel Warshel's group (2013 Nobel Laureate in Chemistry). With more than 20 years of Python programming experience (starting with Python 2.0), he has a strong background in scientific computing, data visualisation, data management, and managing high-performance computing (HPC) systems with Python. Currently, he works as an HPC systems engineer, developing both commercial and scientific software. Additionally, he participates in several defence projects for developing intelligent autonomous weapon systems with autonomous decision-making abilities.

Co-funded by the European Union
EuroHPC Joint Undertaking

3.3.4 Relevance

Python is a key enabling technology for many EPICURE-supported projects, including scientific computing, data analytics, and AI-driven workflows. However, the performance of Python applications on multi-core CPU architectures is often limited by the Global Interpreter Lock (GIL), which restricts the accurate parallel execution of threads in standard Python implementations. This webinar explicitly addresses this challenge by presenting CPU-based threading models and practical strategies to mitigate or overcome the limitations imposed by the GIL in HPC contexts.
By covering approaches such as multiprocessing, native extensions, thread-parallel libraries, and hybrid programming models, the webinar equipped participants with the knowledge required to design Python applications that effectively exploit modern multi-core CPUs, which is particularly relevant for EPICURE-supported HPC environments. The focus on overcoming GIL-related limitations directly contributes to improved resource utilisation, performance portability, and sustainability of Python-based applications deployed on European HPC systems. The webinar enabled users to make informed architectural decisions when integrating Python into HPC production workflows and when selecting appropriate parallelisation strategies.

3.3.5 Summary

- Short introduction
- Overview of the CPython direct interaction with multi-threading libraries
- Run Python code as multi-threaded using just-in-time (JIT) compilation
- Disable GIT to use Python as a multi-threaded application
- How to handle data in Python to facilitate multi-threading executions
- Configure SLURM jobs to handle multi-threading applications in Python

3.3.6 Participant Data

- Number of registrations: 93
- Number of participants: 62
- Attendance: 68%
- Gender: Man (77.4%); Woman (20.4); Prefer not to say (2.2%)
- Type of organisation: Academia (58%); R&D Institution (19%); Company (16%); Other (7%)
- Country of affiliation: 30 different countries.

3.4 Python on GPUs, how hard can that be?

3.4.1 Date and Time

This webinar took place on November 28th 2025, from 3:00 to 4:00 p.m. (CET). The recording of the event can be found on the Training page of the European HPC Application support portal.
Link: https://eurohpcsupport.eu/training-material/webinar-python-on-gpus-how-hard-can-that-be/

Co-funded by the European Union
EuroHPC Joint Undertaking

3.4.2 Description

Numba is a package that will do the heavy lifting for you when you want to offload computations to GPUs. In this webinar, users were able to learn how to use it, as well as about potential performance issues, and how to remedy those. As these performance optimisations are not specific to Numba, the learning outcomes are transferable to GPU programming in general.

3.4.3 Speakers

Geert Jan Bex has been working for the Vlaams Supercomputing Centrum (VSC) since 2009. Besides user support, his focus is mainly on training. Over the years, he has developed and led training sessions on scientific computing, parallel programming, GPU programming, and machine learning.

3.4.4 Relevance

The webinar on Numba was highly relevant for high-performance computing (HPC) and data-intensive applications, as it addressed a key challenge in Python-based scientific workflows: achieving near-native performance while preserving programmer productivity. Numba enables just-in-time (JIT) compilation of Python code to optimise machine instructions, allowing computational kernels to efficiently exploit modern CPU architectures without requiring a complete rewrite in low-level languages. This capability is particularly valuable in HPC environments, where performance, scalability, and portability are essential.

By focusing on the practical application of Numba, including parallel execution, vectorisation, and interoperability with NumPy, the webinar equipped participants with concrete techniques to accelerate existing Python applications and integrate them into production HPC workflows. The knowledge transferred through this activity supports better utilisation of HPC resources, reduces development time, and promotes best practices in performance-aware scientific software development, thereby contributing to more efficient and sustainable use of HPC infrastructures.

3.4.5 Summary

- Introduction and motivation
- GPU architecture and features
- Moving data between the host and device
- Linear algebra on GPUs
- Writing your own kernels
- Data science with Rapids
- Machine learning on GPUs
- Wrap up

3.4.6 Participant Data

- Number of registrations: 65
- Number of participants: 49

Co-funded by the European Union
EuroHPC Joint Undertaking

- Attendance: 75%
- Gender: Man (76.9%); Woman (21.5); Prefer not to say (1.5%)
- Type of organisation: Academia (55.4%); R&D Institution (18.5%); Company (20%); Other (6.2%)
- Country of affiliation: 24 different countries.

3.5 How to profile your Python code

3.5.1 Date and Time

This webinar took place on December 2nd 2025 from 3 to 4 pm (CET). The recording of the event can be found on the Training page of the European HPC Application support portal.
Link: https://eurohpcsupport.eu/training-material/webinar-how-to-profile-your-python-code/

3.5.2 Description

In this webinar, the goal was to demonstrate how to profile users' Python code, starting from a high-level view and progressing to a code-level analysis. This involved identifying bottlenecks at the execution level, then drilling down into functions and individual lines to pinpoint specific issues and determine the necessary fixes. The webinar exemplified how the same approach scales seamlessly from small scripts to large-scale applications and training loops.

3.5.3 Speakers

Ziyad AlBanoby studied at Alexandria University in Alexandria, Egypt, earning his degree in Computer and Communications Engineering. He specialises in CUDA software optimisation and porting, with a focus on algorithms that utilise the GPU. He worked for two years at Brightskies Inc. and, in 2024, joined the High-Level Support Team at the Barcelona Supercomputing Centre.

3.5.4 Relevance

Code profiling is highly relevant for high-performance computing (HPC) and data-intensive application development, as it addresses a fundamental requirement for achieving efficient and scalable software: performance analysis and optimisation. Python is widely used in scientific computing and research workflows, but performance bottlenecks are often non-obvious without systematic profiling. This webinar provided essential knowledge on identifying computational hotspots, memory inefficiencies, and scalability limitations, forming the foundation for effective optimisation strategies.
By introducing practical profiling techniques and tools applicable to Python-based applications, the webinar enabled participants to make data-driven decisions when improving code performance and resource utilisation. The skills acquired support better integration of Python applications into production HPC environments, leading to improved scalability, reduced execution time, and more efficient use of computational resources. Overall, the webinar contributed to capacity building and promoted best practices in performance-aware scientific software development.

Co-funded by the European Union
EuroHPC Joint Undertaking

3.5.5 Summary

- Introduction and motivation
- Code optimisation
- Profiling
  - Profiling Granularity Hierarchy
  - Tool selection
  - Examples
  - Visualisation
- Profilers
  - Line profilers
  - Timeline profilers
  - Programmatic profilers

3.5.6 Participant Data

- Number of registrations: 52
- Number of participants: 39
- Attendance: 75%
- Gender: Man (71%); Woman (27); Prefer not to say (2%)
- Type of organisation: Academia (59.6%); R&D Institution (11.5%); Company (23.1%); Other (5.8%)
- Country of affiliation: 21 different countries.

Co-funded by the European Union
EuroHPC Joint Undertaking

4.Hackathons

4.1 HPC in ARM architecture

This training event started at 14:00 (WET) of February 4th and ended at 14:00 of February 7th, 2025. Link: https://eurohpcsupport.eu/training-material/assessing-the-performance-and-scalability-of-armprocessors-for-industrial-applications/

4.1.1 Overview

The second EPICURE hackathon was organised by INESC TEC and took place in Braga, Portugal, from February 4th to 7th, 2025. This training event focused on using processors with the ARM architecture for High-Performance Computing. Currently, Deucalion is the only EuroHPC JU system that provides users with access to ARM CPU nodes.

Considering that the future exascale machines Alice Recoque and Jupiter will be equipped with ARM processors, this hackathon served as an entry point for people who want to learn how to port their code from x86 to ARM, with all the optimisation tasks and code source changes it may require. The previously summarised webinar “Unlocking Exascale: A First Look at Jupiter” (see section 3.2 of this document) was a natural continuation of this hackathon and provided users with additional information on this specific architecture.

This second hackathon closely followed the format of the previous one, with expert lectures during the first days combined with practical sessions based on the “Bring Your Own Code” format for the remainder of the event.

4.1.2 Description

The second EPICURE Hackathon focused on utilising HPC in CPUs with the ARM architecture and took place at INESC TEC’s hub at the University of Minho in Braga, Portugal, from February 4th to 7th, 2025.

The event featured lectures by experts on utilising the ARM architecture in the context of HPC, as well as practical sessions. The “Bring Your Code” hackathon type was adopted: participants were required to bring the code they were using for their research and that they wished to port or optimise to run on ARM CPUs.

Participants could register for this event individually or as part of a team and develop their code together. The participants had access to the ARM partition of the EuroHPC supercomputer Deucalion.

Co-funded by the European Union
EuroHPC Joint Undertaking

4.1.3 Agenda

Day 1 | February 4, 2025

14:00 – 14:15 Welcome Reception*

14:15 – 14:30 EPICURE HPC Application Support for EuroHPC Users | **Speaker:** António Luís Sousa (INESC TEC, UMinho and MACC)

14:30 – 16:00 Introduction to the A64FX Processor | **Speaker:** John Wagner (Fujitsu) | ***ONLINE SESSION***

16:00 – 16:15 Coffee Break*

17:00 – 18:00 Visit to Deucalion

19:30 Social Dinner*

Day 2 | February 5, 2025

09:30 – 10:30 Porting Code From x86 to ARM | **Speakers:** Bernardo Malaca (CNCA and Deucalion Supercomputer) and Ricardo Fonseca (ISCTE)

10:30 – 11:00 Coffee Break*

11:00 – 11:30 How to use Deucalion | **Speaker:** Rui Silva (CNCA and Deucalion Supercomputer)

11:30 – 12:30 AI in ARM Architecture | **Speaker:** Alícia Oliveira (INESC TEC)

12:30 – 14:00 Lunch*

14:00 – 15:00 Assessing the Performance and Scalability of ARM Processors for Industrial Applications | Gabriel Marcos Magalhães (INESC TEC e PIEP)

15:10 – 15:45 Installing software in ARM can be EESSI | **Speaker:** Kenneth Hoste (Ghent University and MultiXscale) | ***ONLINE SESSION***

15:45 – 16:00 Coffee Break*

16:00 – 18:00 Lab Session

Day 3 | February 6, 2025

09:30 – 10:30 Lab Session

10:30 – 11:00 Coffee Break*

11:00 – 12:30 Lab Session

12:30 – 14:00 Lunch*

14:00 – 15:45 Lab Session

15:45 – 16:00 Coffee Break*

16:00 – 18:00 Lab Session

Day 4 | February 7, 2025

09:30 – 10:30 Lab Session

10:30 – 11:00 Coffee Break*

11:00 – 12:00 Project Presentations

12:00 – 12:15 Closing Remarks

12:15 – 14:00 Lunch*

Figure 2 – Agenda of the Second Hackathon

Co-funded by the European Union
EuroHPC Joint Undertaking

4.1.4 Lectures

The list of lectures and a summary of them are presented as follows:

- **EPICURE HPC application support for EuroHPC users** (António Sousa, INESC TEC, U. Minho, MACC):
  - What is the EPICURE project, and how can researchers apply for user support?
- **Introduction to the A64FX Processor** (John Wagner, Fujitsu):
  - What is the A64FX Processor?
  - Fujitsu History of HPC Systems.
  - Technical information of the A64FX.
  - A64FX Performance.
  - Single Instruction Multiple Data (SIMD) Performance.
  - Low power consumption technology.
  - Fujitsu Software Stack.
  - The future of Fujitsu ARM processors: Monaka.
  - Deucalion's Architecture.
- **An introduction to ARM Scalable Vector Extensions** (Ricardo Fonseca, ISCTE) and **Porting x86 code to ARM** (Bernardo Malaca, CNCA):
  - Improving performance with SIMD instructions.
  - What are the Scalable Vector Extensions (SVE)?
  - Programming with Arm C Language Extensions (ACLE).
  - SVE Datatypes, Predicates and intrinsic functions.
  - SVE Scatter and Gather operations.
  - Migrating x86 AVX code to ARM SVE.
  - ZPIC: An educational Particle-in-Cell code that supports SVE.
  - Porting ZPIC to ARM.
  - Non-uniformity of ARM architecture.
  - A comparison between GCC and Fujitsu compilers for ARM.
  - The simd library.
  - A success case of porting: Osiris.
  - Manual vectorization techniques: unrolling, interleaving.
  - Performance comparison of other software.
- **How to use Deucalion** (Rui Silva, CNCA)
  - First contact with Deucalion.
  - SSH access.
  - Using modules.
  - Job scheduling.
  - Compiling programs in the ARM partition.
  - Profiling code using Score-P.

  This session was conducted using a worksheet with guided exercises for the participants to practice and experiment with Deucalion's system.
- **AI in ARM architecture** (Alícia Oliveira, INESC TEC).
  - What is Artificial Intelligence?
  - Fields of AI.
  - Training and Inference.
  - What is PyTorch?

Co-funded by the European Union
EuroHPC Joint Undertaking

  - A typical AI workflow in PyTorch.
  - Large language models (LLMs).
  - AI using Deucalion's ARM partition.
  - Scalable Distributed Training and Low Power Consumption.
  - Live Demos: Training a neural network in ARM, performing inference.
  - Comparing the performance of LLMs in x86 and ARM architectures.
- **Assessing the Performance and Scalability of ARM Processors for Industrial Applications** (Gabriel Marcos Magalhães):
  - What is CFD (Computational Fluid Dynamics)?
  - Domain discretization.
  - What is OpenFOAM?
  - Examples of simulations done with OpenFOAM.
  - Using OpenFOAM in Deucalion: exaFOAM project.
  - Challenges of OpenFOAM in ARM architecture.
  - OpenFOAM matrix format.
  - OpenFOAM compilation tests on ARM.
  - Performance and energy efficiency comparison between x86 and ARM architectures.
  - Future Work: Improving OpenFOAM's performance on ARM.
- **Installing software in ARM can be EESSI** (Kenneth Hoste, Ghent University and MultiXscale):
  - The MultiXscale CoE.
  - A landscape of scientific computing.
  - Optimized scientific software installations.
  - European Environment for Scientific Software Installations (EESSI).
  - Getting access to EESSI.
  - Live Demo: Using EESSI in Deucalion's ARM Partition.

The slides used by the lecturers in their respective talks are publicly available on the Training page of the European HPC Application support portal.
Link: https://eurohpcsupport.eu/training-material/assessing-the-performance-and-scalability-of-armprocessors-for-industrial-applications/

4.1.5 Lab Sessions

During the hands-on phase of the Hackathon, participants were allowed to work on their own projects while having access to the x86 and ARM partitions of Deucalion. With the possibility to submit jobs to both partitions, participants were able to compare their software's performance after porting it to the ARM architecture. The thematic range of the projects brought by the participants went from particle physics simulations to solving optimisation problems in an industrial pipeline. Examples of these projects include:

- Porting SIESTA to ARM architecture.
- Discovering new materials using Time-Dependent Ab-initio Density Functional Theory.
- Computational Fluid Dynamics for Industrial Applications.
- Optimisation algorithms applied to industrial management problems (job-shop).
- 3D cutting, packing, and vehicle routing algorithms.
- Video and Audio processing models using machine learning.

Co-funded by the European Union
EuroHPC Joint Undertaking

- Green Function Calculation using Lattice QCD.

In lab sessions, participants worked on their project either in groups or individually, with the constant supervision of experienced mentors.

4.1.6 Participant Data and Feedback

A total of 22 participants joined the hackathon. The analysis of the registration forms yielded the following data:

- **Gender:** Male (86%), Female (9%), Prefer not to say (5%).
- **Nationality:** Portugal (85%), Bulgaria/Iran/Mexico (5% each).
- **Positions:** Post Doctorate (32%), Master Student (23%), Other (23%), PhD Student (14%), Non-academic (4%), Industry (4%).

At the end of the event, participants were asked to fill out a feedback form to evaluate their experience during the Hackathon. The form consists of 8 fields: 4 questions where the participant had to rate their level of satisfaction on a scale of 1 to 10, with 10 being the highest level of satisfaction. The rest of the form consisted of four open-ended questions. A total of 12 participants completed the feedback form.

Q1: How do you rate the event regarding the usefulness of the information provided during the hackathon?

- Options: Very Good (84%); Good (8%); Fair (8%); Bad (0%);

Q2: How do you rate the event regarding the adequacy/relevance of topics covered?

- Options: Very Good (75%); Good (17%); Fair (8%); Bad (0%)

Q3: How do you rate the hackathon regarding its format?

- Options: Very Good (75%); Good (16%); Fair (8%); Bad (0%);

Q4: Was the session up to your expectations?

- Options: Very Good (75%); Good (17%); Fair (8%); Bad (0%);

Q5: How will you apply the knowledge from this hackathon in your work/research? (Open question)

Some of the answers to this question include: "Running Notebooks for Machine Learning", "Submitting jobs", "learned about using profilers", "optimise my code with SVE vectorisation", "learned how to compile code in an ARM processor in an HPC environment", "improved my expertise in parallel applications".

Q6: What theme/topic would you like to see covered during our training events? (Open question)

Some of the answers to this question include: "open-source tools for optimisation", "how to apply to HPC grants", "a more comprehensive lecture on profiling", "good practices in programming with floating point operations", and "SVE vectorisation with Python".

Co-funded by the European Union
EuroHPC Joint Undertaking

Q7: Could you please share a short testimonial about your experience at this hackathon? (Open question)
Some of the answers to this question include: "This Hackathon allowed me to get in touch with supercomputers and ARM architecture", "Better than expected. I improved my knowledge in HPC", "I had a great experience and would be thrilled to join another hackathon.", "I was pleasantly surprised by the relevance of the topics presented and the helpfulness of all the staff", "The visit to Deucalion was really interesting, and the breaks were good for socialising".

Q8: Other comments and suggestions? (Open question)
Some of the answers to this question include: "Do organise further events like this!", "It was a great experience. Got a lot of help from the staff".

4.1.7 Significant Outcomes

During the event, some participants inquired about how they could access computational resources to run their models and simulations. It was explained in detail how they can submit applications to EuroHPC JU calls. The participants who showed interest were also informed about how to obtain help from the EPICURE Application Support Teams. This indicates that training events like the EPICURE Hackathons not only transfer technical knowledge to participants, but also serve as a means to inform them about the EuroHPC JU ecosystem, as well as the work performed by EPICURE ASTs.

From a scientific perspective, the participation of a postdoctoral researcher led to the successful development of an optimised EasyBuild recipe, enabling the installation of the software SIESTA on ARM machines[3]. The work of this participant, with the help of mentors from the Hackathon, was extended beyond the event's duration and eventually led to the integration of his recipe into the EESSI ecosystem. As a member of the MaX Centre of Excellence, this contribution was the result of synergies within the EuroHPC community, combining MaX, EPICURE, and the MultiXscale project, which was responsible for the development of EasyBuild and EESSI.

4.2EPICURE Hackathon: Python Meets HPC

This training event started at 10:00 (CET) on December 10th and ended at 14:00 on December 12th, 2025.
Link: https://epicure-hpc.eu/2025/09/30/epicure-hackathon-python-meets-hpc/

4.2.1 Overview

The third EPICURE hackathon, organised by KTH and INESC TEC (as the Work Package Leader), took place in Stockholm, Sweden, from December 10th to 12th 2025. The Initial focus of the event was on Python and HPC, and it was planned to follow a different format compared to the previous ones, as it focused on group work and allowed for online participation.

[3] SIESTA 5.2.2 : From running on ARM in DEUCALION to global release through EESSI: https://max-centre.eu/siesta-5-2-2-from-running-on-arm-in-deucalion-to-global-release-through-eessi/

Co-funded by the European Union
EuroHPC Joint Undertaking

During the registration and preparation phases, due to the low number of registrations, the organising team, with the concordance of the PMB, decided to i) accept HPC projects, whether related to Python or not, ii) move the event to hybrid format with participants on-site and participants online, and iii) invite members of support projects to participate in the hackathon.
The event started with a welcome session and a brief presentation of the EPICURE project, followed by a day of group work. The second day began with a work status update the plan for the working day. On the third day, at first, it was checked that everything went as planned and there were no impediments to continue working. After that, and the teams were asked to prepare presentations of the developed work. The Hackathon concluded with a final presentation as the last session of the event, where each group presented its work and the advancement made during the event.

4.2.2 Description

"The Python Meets HPC" Hackathon focused on utilising Python in High-Performance Computing (HPC) and took place from December 10th to 12th 2025, hosted at KTH Royal Institute of Technology, Stockholm, with participants both on-site and online. For a broader approach, the event focused not only on Python code but on any code that needed optimisation and mentorship from experts.

This event brought together researchers, developers, and HPC practitioners to tackle real-world HPC challenges, optimise codes, and share best practices for performance and scalability.

4.2.3 Agenda

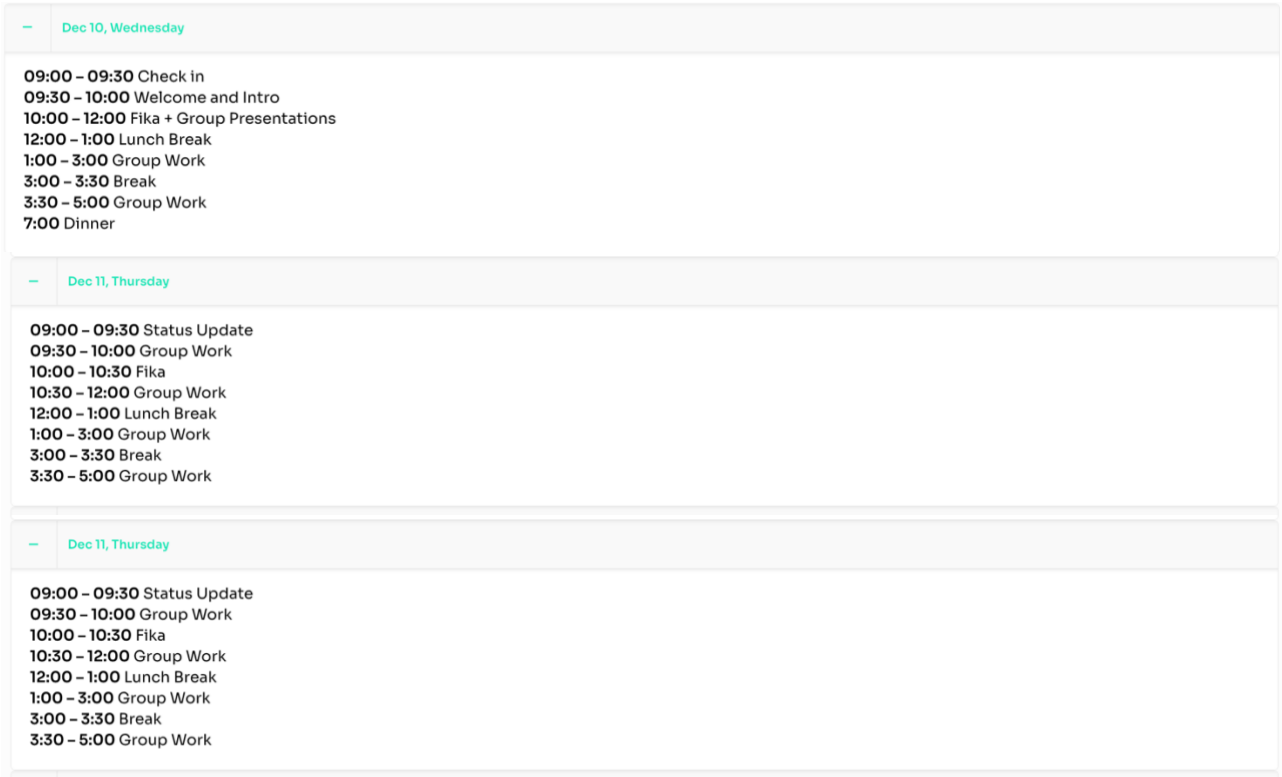


Figure 3 - Agenda of the Third Hackathon

Co-funded by the European Union
EuroHPC Joint Undertaking

4.2.4 Group Work

During the hands-on phase of the Hackathon, participants were given the opportunity to work on their own projects in collaboration with mentors from the EPICURE AST teams. All participants already had accounts in one or more of the EuroHPC JU supercomputers. However, as the hackathon progressed, one of the teams had difficulties running the code on the computers they had access to, and an account in a similar supercomputer was created to try to figure out if it was a software problem or a hardware issue.

12 mentors were assigned to work on the code brought by the 8 attending participants. The projects being worked on were:

- PyWake
- Air Quality Dispersion Model
- NEMA
- Boltz2
- DiscoFlow
- Solver for a nonlinear eigenvalue equation

The work developed during the event ranged from initial code parallelisation to optimisations of existing code and profiling.

4.2.5 Participant Data and Feedback

A total of 8 participants joined the hackathon. The analysis of the registration forms yielded the following data:

- **Gender:** Male (75%), Female (25%).
- **Nationality:** Portugal (25%), Belgium, Greece, Lithuania, Germany, Indonesia, Japan (12.5% each).
- **Positions:** Post Doctorate (25%), Other (12.5%), PhD Student (25%), Non-academic (25%), Industry (12.5%).

At the end of the event, participants were asked to fill out a feedback form to evaluate their experience during the Hackathon. The form is composed of 7 questions. 3 questions where the participant had to give a rating of Very Good, Good, Fair or Bad, one Yes/No questions and 3 open questions. In total, 2 participants filled out the feedback form.

Q1: How do you rate the event regarding the usefulness of the information provided during the hackathon?

- Options: Very Good (100%) / Good (0%) / Fair (0%) / Bad (0%).

Q2: How do you rate the event regarding the adequacy/relevance of topics covered?

- Options: Very Good (100%) / Good (0%) / Fair (0%) / Bad (0%).

Q3: How do you rate the hackathon regarding its format?

- Options: Very Good (100%) / Good (0%) / Fair (0%) / Bad (0%).

Q4: Was the session up to your expectations?

- Options: Yes (100%) / No (0%).

Co-funded by the European Union
EuroHPC Joint Undertaking

Q5: What theme/topic would you like to see covered during our training events? (Open question)
The suggestions were " how to write GPU kernels (in C++) in order to interface them with a code written in another language" and "At some point I would like to get a better understanding of CUDA, that I only use through python libraries now".

Q6: Could you please share a short testimonial about your experience at this hackathon? (Open question)
Some interesting quotes from the answers are:

- "I expected the Hackathon would revolve around certain general, typical HPC problems, while I was looking to solve a specific hiccup in our development. It was a nice surprise that, on the contrary, this fitted the format of the Hackathon perfectly."
- "The profiling, switching to clusters and setting up a good test environment would for sure have taken a lot more time and motivational hurdles when done independently. Roughly we went through the things I had hoped to address: memory profiling, solving the problem and testing, where quickly moving to different machines was especially easy with Epicure support"
- "I had great supervision with all my mentors, we communicated very nicely and we managed to run through the code -what it does and how- very quickly. They also provided feedback that would be impossible to obtain otherwise."

Q7: Other comments and suggestions? (Open question)
There were no answers to this question

4.2.6 Significant Outcomes

This event served as a practical implementation of one of the recommendations highlighted in the project review regarding the organisation of hybrid hackathons. Although the event was not originally designed as a hybrid initiative, it was successfully adapted to support both on-site and online participation, providing valuable insights for future activities.

While the level of interaction differs from an entirely in-person format, the experience demonstrates that, when projects are sufficiently well-defined and mature, hybrid hackathons can still enable significant progress. This was observed across both individual and team-based projects, with tangible outcomes achieved in both cases. The results confirmed that hybrid participation can effectively support focused development and problem-solving activities.

Moreover, the inclusion of online participation made it possible to involve mentors who would otherwise not have been available due to financial or logistical constraints. This significantly enriched the technical support provided during the event, highlighting the added value of hybrid formats in increasing accessibility, expertise, and overall impact.

Co-funded by the European Union
EuroHPC Joint Undertaking

5.Conclusions and Future Work

The training activities carried out during the second year of the EPICURE project successfully implemented part of the training plan defined for WP5, demonstrating the relevance and impact of targeted advanced training in High-Performance Computing (HPC). Through a balanced combination of webinars and hackathons, EPICURE addressed key technical topics aligned with user needs, including Python performance, threading models, GPU acceleration, profiling, containerisation, and emerging architectures such as ARM and exascale systems. Overall participation levels, diversity of backgrounds, and geographic reach confirm the strong interest of the community in these activities and validate the chosen training formats.

The organisation of both in-person and hybrid hackathons proved particularly effective in translating theoretical knowledge into practical outcomes. The hybrid hackathon, although not originally designed as such, successfully implemented one of the project review's recommendations and demonstrated that meaningful progress can be achieved when projects are sufficiently mature and well-defined. The possibility of remote participation also enabled the involvement of highly qualified mentors who would otherwise not have been available due to financial or logistical constraints, thereby increasing the quality and reach of the technical support provided.

Despite these positive outcomes, some challenges remain, in particular the need to increase the response rate to post-event feedback questionnaires. Although efforts have been made to encourage participant engagement in providing feedback, the level of responses obtained has not yet reached the desired target. This limitation restricts a more comprehensive and quantitative assessment of the effectiveness and impact of the training activities, highlighting the need for improved feedback collection strategies in future events.

Future work will focus on refining the training strategy based on the experience gained during the first two years. This includes improving feedback collection mechanisms, further consolidating the hybrid event format, and strengthening the alignment between training activities and the evolving needs of EPICURE-supported projects. Additional emphasis will be placed on hands-on, project-driven events, closer integration with Application Support Teams, and continued coverage of emerging HPC technologies to ensure that EPICURE remains responsive, impactful, and aligned with the European HPC ecosystem.

Co-funded by the European Union
EuroHPC Joint Undertaking

Appendix A. Year 2 Training plan

Executive Summary

This training plan outlines a series of webinars designed to strengthen skills in High-Performance Computing (HPC) for the AST team, and EUROHPC users. The topics address current technological trends, emerging hardware platforms, and best practices relevant to both AST teams, developers and researchers. The objective is to provide participants with a compact yet impactful introduction to each subject, ensuring they are well-prepared to leverage modern HPC infrastructures, including EuroHPC JU systems.

By focusing on GPU programming, software management, performance optimization, and resilience, the training plan equips participants with essential knowledge for the exascale era. The structure allows broad participation, as the short webinar format is accessible and flexible. The sessions will be supported by practical demonstrations, ensuring applicability to real-world research and development challenges.

Co-funded by the European Union
EuroHPC Joint Undertaking

Table of Contents

1. INTRODUCTION 2

1.1 TRAINING PLAN DEFINITION GUIDELINES 2
1.2 TARGET AUDIENCE 3
1.3 SUGGESTED TOPICS 3
1.4 SELECTION OF TRAINING EVENTS 4

2. PROJECT YEAR 2 TRAINING PLAN 6

3. CONCLUSIONS AND FUTURE WORK 12

Co-funded by the European Union
EuroHPC Joint Undertaking

1. Introduction

As proposed in the project's description, the training events to be organized by the project will be of three distinct types: i) webinars; ii) hackathons; iii) workshops. This document describes the process associated with the selection and delivery of each event.

This training plan proposes for the second year of the project a series of webinars covering key topics in High-Performance Computing (HPC). The sessions are designed to provide participants with both theoretical knowledge and practical insights, with justifications highlighting the relevance of each subject to current and future HPC practices.

1.1 Training Plan definition guidelines

The WP5 presented in Deliverable 5.1 a methodology to be used in the selection of the training events to be included in the Projects' training plan. This Section refers to the proposed methodology for the selection of the theme of each training event. By systematically applying a set of selection criteria and considering the opinions of the different members of the Advanced Training WP, it is guaranteed that the theme selected will best serve the target audience while fulfilling the purpose of EPICURE's training program.

The steps involved in organising a training event can be divided into three main processes: **Topic Suggestion**, **Selection of Training Events**, and **Event Organization**. Figure 1 depicts this process graphically. The steps involved in the selection will be described next in more detail.



Figure 4 - Diagram representing the different steps in the topic selection process

Co-funded by the European Union
EuroHPC Joint Undertaking

1.2 Target Audience

In order for a Training plan to be effective, it must identify its target audiences. The EPICURE training plan, delivered through webinars, workshops, and hackathons, is designed to engage and support the following key communities:

- **Industry Users and SMEs**: Organisations adopting HPC for innovation in domains such as energy, manufacturing, finance, and health. Training will lower entry barriers, demonstrate use cases, and accelerate adoption.
- **Researchers and Academics**: Primary users of HPC systems for simulations, modelling, AI, and data-intensive science. Training will help them scale applications, improve performance, and adopt best practices.
- **Emerging Researchers** (Students, PhD Candidates, Junior Developers): Future HPC professionals. Training introduces HPC fundamentals and offers hands-on experience in collaborative problem-solving.
- **Application Support Teams** (ASTs) and Hosting Centres: Providers of advanced user support across EuroHPC sites. Training ensures consistency, deep technical expertise, and cross-centre collaboration.

Each of the training formats proposed aims at a specific goal:

- **Webinars** – broad awareness, access to expertise, strategic overviews.
- **Workshops** – focused technical deep dives and skill building.
- **Hackathons** – applied, collaborative environments fostering innovation and knowledge transfer.

1.3 Suggested Topics

Considering that HPC comprises so many subareas of interest to the HPC user community, a pre-selection of a particular set of topics viewed as of interest to the EPICURE community and that fit the goal of the Advanced Training WP must be made. Several sources can be considered when creating this preliminary topic list, such as:

- **Suggestions from the EPICURE community:** Every member of the EPICURE community is welcome to suggest topics they deem appropriate for a training activity. Because the EPICURE member network's knowledge base is so vast, with people with different areas of expertise working on different WPs, it is essential to listen to the input from inside the project. During the WP5 General meetings, everyone attending can intervene and suggest topics for new training activities and potential speakers for those activities (including themselves, if they wish).
- **End-of-project PI feedback forms:** After the user support activity for a particular project ceases, the principal investigator of the project is asked to fill in a feedback form where they can evaluate their experience. There, they can also indicate what

Co-funded by the European Union
EuroHPC Joint Undertaking

difficulties they had during the process. By analysing this document, information is collected that can help steer the training program to address those difficulties and provide better experiences for future projects the PI may apply to. Empowering these users through training activities is one of the main objectives of the EPICURE project.

- **Feedback from Application Support Teams:** In addition to the previous point, ASTs can evaluate the user support experience from their perspective. Throughout the project, they can list issues that may have arisen during support. Compiling these reports can help user support activity on future projects. If there are any recurrent doubts or issues with a particular concept, a training activity on the subject may fill in those gaps and improve the user support experience and the support specialists' know-how.
- **Past training event feedback forms:** After each training event, webinar, or hackathon, all participants will be contacted to fill in a form evaluating their experience. Every form has a section where participants can suggest other themes they want addressed in a training event.
- **Emergent Topics:** Nowadays, technology is advancing rapidly, and HPC is one of the scientific areas in which this advancement is well-known. As such, whenever new hardware emerges, especially hardware that may be integrated into a EuroHPC JU supercomputer in the future, it is relevant to create an event where an expert on this new technology can present it to the community, focusing on how it can be helpful to a large part of the user community. This is also true for software that can interest a large percentage of the user base.
- **Success Stories:** When the user support service for a particular project ceases, the project PI will use the results obtained from their computations to advance their research. If a PI is willing, they can deliver a talk on how the user support provided by the EPICURE ASTs was essential to their research. These success stories would show how the EPICURE application support can help research in different areas, encouraging new people from the scientific community to apply to EuroHPC projects and ask for EPICURE support.

1.4 Selection of Training Events

1.4.1 Filtering process

After a list of topics is compiled considering the sources described in Section 2.2, the list must be analysed for activities that may not be suited for the EPICURE Advanced Training program. For that purpose, two main criteria are considered:

- **Avoid duplication of training activities:** Besides EPICURE, other EuroHPC JU initiatives have their own working groups that organize training activities for their target audience. To avoid duplication of efforts and expand the range of topics

Co-funded by the European Union
EuroHPC Joint Undertaking

covered in training sessions, it is necessary to check if the topic has already been covered in a training activity organized by other programs. In some cases, if the topic is considered very relevant to the EPICURE ecosystem, organizing a training activity already offered by another program can be considered, but preferably offering a new approach to the topic more directed to EPICURE's interests.

- **Avoiding topics that are beginner-oriented:** In an Advanced Training program like EPICURE's, it is not expected to see training sessions on fundamental HPC skills, e.g., learning to program in C, SSH key generation, Linux environment, etc. Although these skills must be taught to users, the Advanced Training program of EPICURE aims to go further. Also, there are other training programs, namely the ones organized by NCCs, where these more essential tools are taught. This criterion complements the "no duplication" of activities presented in the previous point.

After this process, the list of topics grows shorter, and the topics that have been filtered out can be suggested to other partners like Castiel 2 or other training programs if it is considered that there is a lack of material on the topic available to users.

1.4.2 Approval of training actions

Once a reasonable number of topics and speaker combinations are available, a WP5 General Meeting is convened. The shortlist for future training events is presented to the members attending during that meeting. Each idea is put to a vote, and if no one opposes it, the topic is approved for a future training activity.

Co-funded by the European Union
EuroHPC Joint Undertaking

2.Project Year 2 Training Plan

At the end of the first year of the project, in WP5, members started preparing the training plan using the proposed methodology. The training plan resulted mainly from the inputs of the project members, as the number of projects finished was low, and the PI's project feedback form did not include suggestions. The main source of suggestions for themes comes from the EPICURE community, mainly members of the AST, who also collaborate with WP5.

During the WP5 general meeting, two themes clearly emerged as the most referenced: Python in HPC and AI-related issues, and as such, it was decided to propose a series of webinars related to Python, which will be complemented by a hackathon on Python for HPC to consolidate the lessons learned in the webinars. Regarding the IA-related, several webinars were proposed, ranging from GPU programming to checkpointing and IO optimization. Additionally, several of interest for the community have also been identified and added to the plan.

The following Table 1 presents the plan for Year 2 of the EPICURE Project.

Topic	Objective	Format	Instructor	Justification	Date
Containers in HPC	Introduce the use of container technologies (e.g., Singularity, Apptainer, Docker) in high-performance computing to improve portability, reproducibility, and ease of deployment of scientific applications.	1-hour webinar	Teo Prica	The use of containers in HPC simplifies portability, reproducibility, and the management of complex scientific applications, ensuring that researchers and developers can run their codes efficiently and consistently across different infrastructures.	March 2025

Co-funded by the European Union
EuroHPC Joint Undertaking

Topic	Objective	Format	Instructor	Justification	Date
EuroHPC's Jupiter – Presentation	Present the architecture and opportunities of Europe's first exascale system.	1-hour webinar	Jolanta Zjupa & Junxian Chew	Awareness of Jupiter prepares researchers to leverage exascale computing early, aligning with EuroHPC JU strategy	April 2025
OpenMP for Python	Enable parallelism in Python using OpenMP via Numba or Cython.	1-hour webinar		Many researchers rely on Python but face performance limitations. Combining OpenMP with Python enables multi-core CPU utilization, bridging the gap between usability and high performance.	Q4 2025
Profiling Python Code	Demonstrate profiling tools for Python in HPC contexts.	1-hour webinar		Performance tuning is critical in HPC. Profiling empowers researchers to identify hotspots, optimize code, and make best use of HPC hardware, particularly in compute- and data-intensive workflows.	Q4 2025

Co-funded by the European Union
EuroHPC Joint Undertaking

Topic	Objective	Format	Instructor	Justification	Date
CUDA for Python	Accelerate Python codes using Numba and CuPy.	1-hour webinar		Python dominates in AI, ML, and data science. Bridging Python with CUDA allows researchers to harness GPUs without abandoning familiar tools, accelerating adoption of HPC among non-traditional users.	Q4 2025
Programming AMD GPUs	Overview of HIP and ROCm for GPU programming.	1-hour webinar		EuroHPC JU systems like LUMI use AMD GPUs; researchers and users must adapt their codes to run on non-NVIDIA accelerators. This training ensures portability and full use of European HPC resources.	Q4 2025 /Q1 2026
Multi-GPU Programming with NCCL	Introduce NCCL for scalable GPU communication.	1-hour webinar		Many modern AI/HPC workloads require multiple GPUs. NCCL is a standard for high-bandwidth, low-latency communication between GPUs. Training ensures participants can scale applications beyond single devices	Q4 2025 /Q1 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Topic	Objective	Format	Instructor	Justification	Date
Fault Tolerance / Checkpoint-Restart	Techniques to improve resilience in HPC workloads.	1-hour webinar		As systems scale toward exascale, hardware failures become more frequent. Fault tolerance and checkpointing are essential to avoid loss of large-scale computations, reducing costs and wasted compute hours.	Q4 2025 /Q1 2026
Efficient I/O	Discuss strategies to optimize I/O with parallel file systems.	1-hour webinar		I/O bottlenecks can limit application scalability. Learning efficient strategies (MPI-IO, HDF5, parallel file systems) helps researchers achieve faster runtimes and more reproducible workflows.	Q4 2025 /Q1 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Topic	Objective	Format	Instructor	Justification	Date
Spack Software Manager	Introduce Spack for building and managing software stacks in HPC.	1-hour webinar		Scientific software stacks are complex and often need multiple versions. Spack is one of the building tools for reproducible, portable builds in HPC. Training ensures users and admins can manage software efficiently. There have been already webinars on EasyBuild the other reference toll	Q4 2025 /Q1 2026
Kokkos Library	Introduction to performance portability with Kokkos.	1-hour webinar		HPC applications must run across CPUs, NVIDIA GPUs, AMD GPUs, and emerging architectures. Kokkos allows a single codebase to target multiple platforms, future-proofing scientific applications.	Q4 2025 /Q1 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Topic	Objective	Format	Instructor	Justification	Date
hackathon on Python for HPC	Hands-on consolidation of the Python webinars	2/3 days hackathon @ KTH		To provide participants with a collaborative, hands-on opportunity to apply and integrate the skills acquired during the training webinars, fostering innovation, teamwork, and practical problem-solving in real HPC environments.	Q4 2025

Table 1 - EPICURE training plan Year 2

Co-funded by the European Union
EuroHPC Joint Undertaking

3.Conclusions and Future Work

The Training plan presented in this section results from the first run of the proposed methodology for the selection of the training events. Being done in an early stage of the project, the available inputs are limited.

It will enhance participants' skills and knowledge in HPC programming, optimization, and application deployment, while broadening participation by engaging both advanced users and newcomers, such as students and early-career researchers. Through a mix of webinars, workshops, and hackathons, it will foster collaboration and knowledge transfer across academia, industry, and application support teams, thereby promoting a consistent and sustainable HPC culture in Europe.

The Lessons that will be learned from the implementation of this training plan, as well as the number of projects executed during the first 24 months of the project, will surely be precious for the definition of the training plans for the years 3 and 4 of the project, where we expect to have 6 webinars, 1 workshop, and 2 hackathons per year.

Co-funded by the European Union
EuroHPC Joint Undertaking

Appendix B. Year 3 Training plan Executive Summary

This training plan outlines a series of webinars, workshops and hackathons designed to strengthen skills in High-Performance Computing (HPC) for the AST team, and EUROHPC users. The topics address current technological trends, emerging hardware platforms, and best practices relevant to both HPC users (i.e., developers and researchers) and the AST teams. The objective is to provide participants with a compact yet impactful introduction to each subject, ensuring they are well-prepared to leverage modern HPC infrastructures, including EuroHPC JU systems.

By focusing on GPU programming, software management, performance optimisation, and resilience, the program equips participants with essential knowledge for the exascale era. The structure allows broad participation, as the short webinar format is accessible and flexible. The sessions will be supported by practical demonstrations, ensuring applicability to real-world research and development challenges.

Co-funded by the European Union
EuroHPC Joint Undertaking

Table of Contents

1. INTRODUCTION 3

2. PROJECT YEAR 3 TRAINING PLAN 4

2.1 TARGET AUDIENCES 4
2.2 YEAR 3 TRAINING PLAN 4

3. CONCLUSIONS AND FUTURE WORK 25

Co-funded by the European Union
EuroHPC Joint Undertaking

1. Introduction

As proposed in the project's description, the training events to be organised by the project will be of three distinct types: i) webinars; ii) hackathons; iii) workshops. This document outlines the process involved in selecting each event.

This training plan proposes for the third year of the project a series of webinars, workshops and hackathons covering key topics in High-Performance Computing (HPC). The sessions are designed to provide participants with both theoretical knowledge and practical insights, with justifications highlighting the relevance of each subject to current and future HPC practices.

The Year 3 Training Plan of EPICURE builds on the foundations established in previous year, continuing to strengthen skills development, knowledge sharing, and community building within the project. The plan is designed to address the evolving needs of partners and participants, fostering effective knowledge transfer, international collaboration, and increased visibility of EPICURE outcomes.

Throughout the year, the training programme foresees at least one webinar per month, scheduled for the last Friday of each month, except in January, July, and August. This regular cadence ensures continuous engagement and timely dissemination of relevant topics aligned with EPICURE's technical and strategic priorities in advanced computing.
In addition, the plan includes in-person workshops hosted at the Barcelona Supercomputing Centre (BSC) and LuxProvide (during their SCynergy[4] event), offering opportunities for intensive training, hands-on interaction, and direct engagement with leading European computing infrastructures. The programme also foresees two hackathons, one in Portugal and one in Poland, organised in a hybrid format, with a strong emphasis on on-site participation while maintaining remote access to maximise inclusiveness and impact.

Furthermore, the training plan remains flexible to accommodate additional EPICURE dissemination webinars whenever appropriate. These may include sessions dedicated to presenting project work and results, showcasing outcomes from hackathons, or addressing emerging topics of interest to the community, thereby reinforcing the visibility and overall impact of EPICURE.

[4] SCynergy event: https://www.scynergy.events/

Co-funded by the European Union
EuroHPC Joint Undertaking

2.Project Year 3 Training Plan

This Section describes the training plan proposed for the third year of EPICURE. It starts by recalling the target audiences and the training events formats, before presenting the trainings proposed.

2.1 Target audiences

To be effective a Training plan must identify its target audiences and formats. The EPICURE training plan, delivered through webinars, workshops, and hackathons, is designed to engage and support the following key communities:

- **Industry Users and SMEs**: Organisations adopting HPC for innovation in domains such as energy, manufacturing, finance, and health. Training will lower entry barriers, demonstrate use cases, and accelerate adoption.
- **Researchers and Academics**: Primary users of HPC systems for simulations, modelling, AI, and data-intensive science. Training will help them scale applications, improve performance, and adopt best practices.
- **Emerging Researchers** (Students, PhD Candidates, Junior Developers): Future HPC professionals. Training introduces HPC fundamentals and offers hands-on experience in collaborative problem-solving.
- **Application Support Teams** (ASTs) and Hosting Centres: Providers of advanced user support across EuroHPC sites. Training ensures consistency, deep technical expertise, and cross-centre collaboration.

Each of the training formats proposed aims at a specific goal:

- **Webinars** – broad awareness, access to expertise, strategic overviews.
- **Workshops** – focused technical deep dives and skill building.
- **Hackathons** – applied, collaborative environments fostering innovation and knowledge transfer

2.2Year 3 Training Plan

In the fourth quarter of the second year of the project, WP5 members started preparing the training plan for the third year, using the methodology proposed in Deliverable 5.1.
The Year 2 training plan could not be fully implemented. Development of the new plan started with a review of which undelivered activities should be integrated into the Year 3 plan. It was agreed that several of these activities should indeed be carried over, while others may be incorporated later if considered appropriate.

The current Year 3 training plan is largely shaped by the events that remain to be delivered, as well as by feedback gathered from previous initiatives, such as webinars and hackathons,

Co-funded by the European Union
EuroHPC Joint Undertaking

from both participants and internal AST teams. It also reflects lessons learned from supported projects and input from project members.
During the WP5 general meeting in November 2025, several themes clearly emerged as the most referenced in the support projects (see Figure 1):

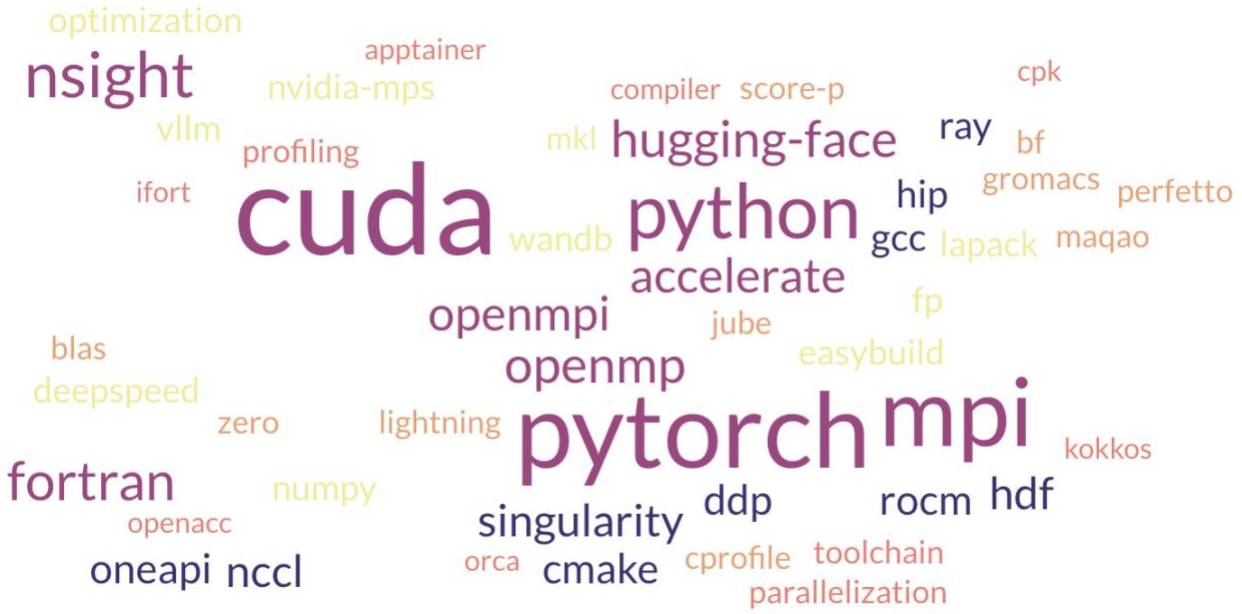


Figure 5 - Word cloud

Python in HPC and AI-related issues were the most referenced topics, and as such, it was decided to propose a series of webinars related to the Python language. Regarding AI-related subjects, it was decided to propose several webinars, covering topics from AI frameworks and GPU programming to checkpointing and I/O optimisation.

There were also references to profiling and optimisation tools, as well as references to domain-specific tools. In this respect, the plan will explore the tools during webinars, while one of the hackathons will focus on domain-specific tools.

The following Table 1 presents the training plan for Year 3 of the EPICURE Project. It is structured using a set of clearly defined columns that ensure coherence, transparency, and traceability in the selection of training activities. Each activity is assigned to a Training Track, representing a high-level thematic area derived from recurring needs identified across EPICURE-supported projects, Application Support Team (AST) experience, and feedback from previous training events. Within each track, a clearly defined Topic specifies the technical focus of the activity, enabling participants to easily assess its relevance and ensuring that the overall programme addresses both current and emerging challenges in High-Performance Computing.

Co-funded by the European Union
EuroHPC Joint Undertaking

For each activity, the Objective describes the expected learning outcomes, supporting alignment between content, delivery, and evaluation. The chosen Format (webinar, workshop, or hackathon) aligns with the activity types defined in the Grant Agreement and was selected according to the complexity of the topic and the required level of interaction, balancing efficient knowledge dissemination with hands-on experience. The Justification column provides the rationale for the inclusion of each activity, linking it to identified user needs, project experience, and strategic developments in the HPC landscape, thereby ensuring the training plan remains demand-driven and aligned with WP5 objectives. Finally, the Date indicates the planned timing of each activity, following a balanced scheduling approach while retaining sufficient flexibility to adapt the programme during execution.

Co-funded by the European Union
EuroHPC Joint Undertaking

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
Tools	Fault Tolerance / Checkpoint-Restart	Techniques to improve resilience in HPC workloads.	1-hour webinar	As systems scale toward exascale, hardware failures become more frequent. Fault tolerance and checkpointing are essential to avoid loss of large-scale computations, reducing costs and wasted compute hours.	Q1 2026
Tools	Spack Software Manager	Introduce Spack for building and managing software stacks in HPC.	1-hour webinar	Scientific software stacks are complex and often need multiple versions. Spack is one of the building tools for reproducible, portable builds in HPC. Training ensures users and admins can manage software efficiently. There have been already webinars on EasyBuild, the other reference tool.	Q2 2026
Tools	Kokkos Library	Introduction to performance portability with Kokkos.	1-hour webinar	HPC applications must run across CPUs, NVIDIA GPUs, AMD GPUs, and emerging architectures. Kokkos allows a single codebase to target multiple platforms, future-proofing scientific applications.	Q2 2026
Tools	Efficient I/O	Discuss strategies to optimize I/O with parallel file systems.	1-hour webinar	I/O bottlenecks can limit application scalability. Learning efficient strategies (MPI-IO, HDF5, parallel file systems)	Q3 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
				helps researchers achieve faster runtimes and more reproducible workflows.	
Tools	Profiling	introduce the fundamental principles and practical techniques of performance profiling in High-Performance Computing (HPC)	1-hour webinar	Performance profiling is a critical skill in HPC, as application efficiency directly impacts time-to-solution, energy consumption, and the effective use of shared computing resources. Many HPC users rely on intuition or ad-hoc tuning, often overlooking systematic profiling approaches that can reveal scalability issues, load imbalance, memory inefficiencies, or communication overheads. By promoting best practices in profiling, this webinar supports more sustainable and efficient use of HPC infrastructures, improves application portability across architectures, and aligns with the growing emphasis on performance, energy efficiency, and reproducibility in modern scientific computing.	Q4 2026
Tools	Optimization	provide an overview of key optimization strategies in High-Performance Computing (HPC)	1-hour webinar	Optimization is essential in HPC to fully exploit advanced computing infrastructures and to achieve efficient, scalable, and sustainable scientific and industrial workloads. As HPC systems grow in complexity, incorporating deep memory hierarchies, accelerators, and large-scale parallel applications that are not explicitly	Q4 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
				optimized often fail to deliver expected performance gains. By promoting optimization methodologies, this webinar contributes to better application efficiency, improved use of shared HPC resources, and increased return on investment in advanced computing infrastructures.	
Scientific Tools	Fortran	Stress the role of Fortran in High-Performance Computing (HPC), highlight its relevance for developing efficient, scalable, and portable scientific applications.	1-hour webinar	Fortran continues to be a cornerstone language in HPC, underpinning many large-scale scientific and engineering codes due to its performance-oriented design, mature compiler ecosystem, and strong support for numerical computing and parallelism. Despite the emergence of newer languages and frameworks, Fortran remains critical for both maintaining legacy codes and developing modern, high-performance applications using features such as coarrays, OpenMP, MPI, and GPU offloading. By strengthening Fortran literacy, this webinar supports the sustainability, performance, and longevity of HPC software ecosystems.	Q4 2026
Scientific Tools	Alphafold	Introduction to the use of AlphaFold in High-Performance Computing (HPC) environments, highlighting how large-scale computing	1-hour webinar	AlphaFold has transformed structural biology by enabling highly accurate protein structure prediction, but its computational demands pose significant challenges when applied at scale. HPC systems play a crucial role in supporting large datasets, accelerated inference, and	Q2 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		infrastructures enable efficient and scalable protein structure prediction.		high-throughput studies required in modern bioinformatics and life sciences research. By exposing researchers and HPC practitioners to the computational aspects of AlphaFold, this webinar promotes efficient use of advanced computing infrastructures, fosters interdisciplinary collaboration between HPC and life sciences communities, and supports reproducible, scalable scientific discovery.	
Scientific Tools	OpenFoam	To gain an overview of OpenFOAM's parallel execution model, typical HPC workflows, and key considerations for running and scaling OpenFOAM applications on modern multi-core and distributed-memory systems. It also aims to provide practical insight into leveraging HPC resources for large-scale CFD problems.	1-hour webinar	OpenFOAM is one of the most widely used open-source CFD frameworks in scientific and industrial applications, yet its effective use often depends on access to and understanding of HPC systems. Large meshes, complex physics, and long simulation times make HPC essential for achieving accurate results within reasonable timeframes. By connecting CFD concepts with HPC execution practices, this webinar supports more efficient use of advanced computing resources, improves simulation scalability and reliability, and helps bridge the gap between application-domain users and HPC infrastructures.	Q1 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
Scientific Tools	Gromacs	To gain a high-level understanding of GROMACS' execution model, including parallelisation strategies and accelerator usage, as well as typical HPC workflows for running molecular dynamics simulations at scale. To familiarise attendees with best practices for leveraging HPC resources in biomolecular simulation studies.	1-hour webinar	GROMACS is one of the most widely used molecular dynamics packages in computational chemistry, biophysics, and related life sciences, where realistic simulations often require significant computational power. Large system sizes, long time scales, and ensemble simulations make HPC essential for achieving scientifically meaningful results within reasonable timeframes. By connecting molecular dynamics applications with HPC execution practices, this webinar promotes more efficient use of advanced computing resources, improves scalability and reproducibility, and strengthens the interaction between life sciences researchers and HPC infrastructures.	Q4 2026
AI4Science	PyTorch	To gain a high-level understanding of PyTorch's execution model in distributed settings, including data and model parallelism,	1-hour webinar	PyTorch has become a dominant framework for deep learning across research and industry, yet scaling workloads beyond a single node introduces challenges related to parallelism, communication, resource management, and reproducibility. HPC systems provide the compute, interconnects, and storage required to address these challenges, but effective use requires an	March 2026?

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		accelerator usage, and typical HPC workflows. To familiarise attendees with best practices for running PyTorch workloads on multi-node, multi-GPU HPC systems.		understanding of both the framework and the underlying infrastructure. By bridging machine learning workflows and HPC practices, this webinar supports more efficient use of advanced computing resources, accelerates AI-driven research, and promotes convergence between AI and HPC communities.	
AI4Science	Hugging-face	To gain a high-level understanding of how Hugging Face libraries and model hubs integrate with HPC resources, including multi-GPU and multi-node systems, and how typical workflows are adapted for batch scheduling and distributed execution.	1-hour webinar	Hugging Face has become a central platform for developing and sharing state-of-the-art NLP and multimodal models, but practical adoption at scale often requires substantial compute, fast interconnects, and high-performance storage capabilities provided by HPC systems. Scaling model training and inference introduces challenges related to distributed execution, resource allocation, and reproducibility that are not addressed by single-node setups. By bridging modern AI tooling and HPC practices, this webinar supports efficient use of advanced computing resources and accelerates AI-driven research and innovation.	Q2 2026
AI4Science	DeepSpeed	To know core concepts behind DeepSpeed, such as memory	1-hour webinar	As state-of-the-art AI models continue to grow in size and computational demand, traditional single-node or naïve distributed training approaches become impractical.	Q3 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		optimisation, model parallelism, and efficient distributed training. To know how DeepSpeed core concepts map to typical HPC environments using multi-node, multi-GPU. To understand when and why DeepSpeed is used, how it integrates with existing deep-learning workflows, and the main considerations for deploying it on HPC systems.		DeepSpeed, originally developed at Microsoft, addresses these challenges by enabling efficient large-scale training through advanced optimisation techniques that significantly reduce memory usage and improve performance. For HPC centres and users, DeepSpeed is particularly relevant because it aligns well with the characteristics of modern supercomputing systems, including GPU accelerators, high-speed interconnects, and batch-scheduled multi-node execution. This webinar helps bridge the gap between cutting-edge AI frameworks and production HPC infrastructures, empowering users to run larger models, reduce resource waste, and achieve better scalability on shared HPC systems.	
AI4Science	vLLM	To know key design principles behind vLLM: i) efficient memory management, ii) continuous batching, iii) high-throughput GPU utilization.	1-hour webinar	With the rapid adoption of large language models across scientific, engineering, and data-driven domains, inference has become a major computational workload in HPC centres, often rivaling or exceeding training demands. Efficiently serving large models at scale requires careful management of GPU memory, parallelism, and throughput. These challenges are	Q4 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		To know how the key design principles enable scalable, low-latency inference on single-node and multi-node HPC systems.		particularly pronounced in batch-scheduled, multi-user HPC environments. vLLM directly addresses these challenges by providing a purpose-built inference engine that maximizes hardware utilization while reducing memory fragmentation and idle time. This webinar is therefore highly relevant for researchers, engineers, and HPC practitioners who need to deploy LLM-based services or workflows on supercomputing systems.	
Tools	Code optimisation for heterogeneous HPC environments	The objective of this hackathon is to challenge participants to optimise real-world scientific and engineering codes for heterogeneous High-Performance Computing (HPC) environments that combine multiple CPU architectures, GPUs, and other accelerators. Participants will analyse performance	Hackathon @ INESCTEC	Modern HPC systems are increasingly heterogeneous, integrating multi-core CPUs with GPUs and other specialized accelerators to meet growing performance and energy-efficiency demands. While these architectures offer significant computational power, achieving optimal performance requires specialized knowledge of programming models, memory hierarchies, and performance analysis tools. This hackathon provides a practical, collaborative environment where participants can apply theoretical knowledge to concrete optimization challenges, bridging the gap between algorithm design and efficient execution on heterogeneous systems. By focusing on hands-on code optimization, the event addresses a critical skills gap in the HPC ecosystem and directly supports the	April 27-29

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		bottlenecks, apply optimisation techniques (e.g. parallelisation, memory optimisation, vectorisation, and accelerator offloading), and adapt applications to efficiently exploit heterogeneous architectures. By the end of the hackathon, teams will have improved the performance, scalability, and portability of their code while gaining hands-on experience with modern heterogeneous HPC systems and optimisation workflows.		training of researchers, engineers, and students to effectively use current and future supercomputing infrastructures. The hackathon format further encourages knowledge exchange, peer learning, and cross-disciplinary collaboration, fostering best practices in performance-aware and portable HPC software development.	
Scientific Tools	HPC Domain Specific	To enable efficient use of domain-specific HPC tools.	Hackathon @ Cyfronet	Domain-specific HPC tools play a critical role in modern scientific research and industrial innovation, encapsulating advanced algorithms and domain	Q3/4 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
	tools efficient usage	To improve application performance and scalability. To promote optimal utilisation of HPC resources. To foster cross-domain knowledge exchange. To strengthen interaction between domain scientists and HPC experts. To encourage portability and reproducibility. To identify dissemination-ready outcomes.		expertise. Despite their maturity, these tools are often used sub-optimally on HPC systems due to limited user familiarity with performance-critical configuration options, resource mapping, and system-specific constraints. A hackathon-based training format provides an effective mechanism to address these gaps through hands-on experimentation with real applications and datasets. By focusing on efficiency and best practices, participants gain practical insights that are difficult to achieve through conventional lectures or isolated tutorials. The collaborative nature of the hackathon fosters peer learning and cross-domain exchange, revealing common performance patterns and optimisation strategies applicable across multiple scientific fields. This approach supports more efficient, sustainable, and productive use of HPC infrastructures. Finally, the hackathon acts as a feedback and knowledge-gathering mechanism, helping to identify emerging training needs and high-impact use cases that can inform future advanced training activities within the overall training programme.	
Tools	high scalability for large	The objective of this workshop is to equip participants with the	Workshop @ BSC	As HPC systems continue to grow in size and complexity, application scalability has become a critical factor in determining scientific and industrial productivity. Even	March /April 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
	scale applications	knowledge and practical insights required to design, analyse, and optimise highly scalable applications for large-scale High-Performance Computing (HPC) systems. The workshop will address key scalability challenges, including strong and weak scaling, communication overhead, load imbalance, and I/O bottlenecks, and will explore strategies to efficiently exploit thousands to millions of processing elements. By the end of the workshop, participants will understand the principles and best		well-optimised codes may fail to deliver expected performance at scale due to communication costs, synchronisation, memory contention, or inefficient parallel algorithms. This workshop addresses these challenges by focusing on scalability-aware application design and performance analysis rather than isolated node-level optimisations. It provides participants with a systematic understanding of how applications behave at large scale and how to identify and mitigate scalability limiters. The topic is particularly relevant for researchers and engineers targeting current and next-generation supercomputers, where efficient scaling is essential to fully exploit available resources, reduce time-to-solution, and improve energy efficiency. By fostering a scalability-first mindset, the workshop directly supports the development of robust, future-proof large-scale HPC applications.	

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		practices needed to achieve high scalability for large-scale applications on modern supercomputing platforms.			
AI4Science	Distributing PyTorch Model Training	To Lauch the training of PyTorch training on an HPC. To learn how to profile, optimize and parallelize it.	Workshop @ scynergy 2026	The increasing scale and complexity of modern deep learning models have made distributed training a fundamental requirement rather than an advanced optimisation. As model sizes, dataset volumes, and training times continue to grow, single-node training quickly becomes impractical, particularly in research and production environments that rely on shared accelerator-based infrastructures. This workshop addresses the critical need to understand how to effectively distribute PyTorch model training across multiple GPUs and nodes. It focuses on the practical challenges of parallelism strategies, communication overhead, synchronization, and fault tolerance that arise in real HPC and large-scale computing environments. By demystifying distributed training concepts and connecting them to concrete PyTorch mechanisms, the workshop helps participants	April 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
				move beyond trial-and-error approaches toward scalable solutions. The topic is especially relevant for users of modern HPC systems, where efficient use of GPU resources and high-speed interconnects is essential to achieve acceptable time-to-solution and cost efficiency. By emphasizing practical deployment considerations alongside core concepts, this workshop supports researchers and engineers in transitioning from single-GPU prototypes to robust, scalable training workflows capable of exploiting large-scale computing infrastructures.	
Scientific Tools	Bioinformatics workflow with Nextflow	To explain how to configure Nextflow to take advantage of MeluXina platform. To use Nextflow in a bioinformatics with a genome analysis example.	Workshop @ scynergy 2026	The increasing complexity of bioinformatics analyses, driven by high-throughput sequencing technologies and large-scale genomic datasets, has made workflow automation, scalability, and reproducibility essential requirements rather than optional features. Manual or ad hoc pipeline execution is no longer adequate for ensuring reliable results or for efficiently exploiting modern HPC infrastructures. This tutorial addresses these challenges by focusing on Nextflow, a workflow management framework specifically designed for building portable, scalable, and reproducible scientific workflows. By demonstrating how to configure Nextflow to take full advantage of the MeluXina	April 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
				supercomputing platform, the tutorial bridges the gap between workflow design and efficient execution on large-scale HPC systems. Applying the concepts to a concrete genome analysis use case allows participants to understand how abstract workflow definitions translate into real bioinformatics pipelines, including data staging, task parallelism, and resource management. The tutorial is particularly relevant for researchers and practitioners in bioinformatics and computational biology who need to run complex analyses reliably at scale, while ensuring reproducibility across computing environments and over time.	
Tools	Rust on accelerators (GPU + FPGA)	To introduce Rust as a systems programming language for accelerators. To enable accelerator programming using Rust-based ecosystems. To develop safe and high-performance accelerator kernels.	Workshop @ scynergy 2026	Accelerators such as GPUs and FPGAs are central to modern HPC systems, enabling high performance and energy efficiency for a wide range of scientific and industrial workloads. However, traditional accelerator programming models often rely on low-level languages and frameworks that increase development complexity and the risk of programming errors. Rust offers a compelling alternative by combining systems-level performance with strong guarantees on memory safety and concurrency. Its growing ecosystem for accelerator programming makes it an increasingly relevant language for developing reliable and	April 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		To explore heterogeneous CPU–GPU–FPGA execution models. To improve portability and maintainability of accelerator code. To expose performance analysis and optimisation techniques. To build capacity for next-generation HPC software development.		maintainable HPC applications targeting heterogeneous architectures. A dedicated training activity on Rust on accelerators (GPU + FPGA) addresses a clear skills gap by introducing participants to emerging tools and programming paradigms that complement existing accelerator technologies. Through hands-on examples and real use cases, participants gain practical insight into how Rust can be effectively integrated into accelerator-based workflows. By focusing on both GPUs and FPGAs, the training reflects the heterogeneity of current and future HPC systems, encouraging portable and sustainable software design. This activity supports long-term capacity building in advanced computing, equipping participants with skills aligned with next-generation HPC architectures and software ecosystems.	
Tools	Accelerating HPC Workloads on Intel FPGA Clusters with SYCL	To provide participants with hands-on experience in accelerating High-Performance Computing (HPC) workloads on Intel FPGA clusters	Tutorial @ ISC	Since its release in 2021, the EuroHPC Supercomputer MeluXina is the only HPC platform with a dedicated FPGA partition. Each of the 20 MeluXina FPGA computer nodes comprises two BittWare 520N-MX FPGAs based on the Intel Stratix 10 FPGA chip. Designed for compute acceleration, the 520N-MX are PCIe boards featuring	June 2026

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		using SYCL, a modern C++-based heterogeneous programming model. To learn the fundamentals of SYCL for FPGA development, demonstrate how to offload compute-intensive kernels to FPGAs, and explore performance and energy-efficiency considerations in multi-node FPGA-enabled HPC environments. Be able to identify suitable HPC workloads for FPGA acceleration, develop SYCL-based kernels targeting Intel FPGAs, and understand the practical steps required to deploy and		Intel's Stratix 10 MX2100 FPGA with integrated HBM2 memory. The size and speed of HBM2 (16GB at up to 512GB/s) enables acceleration of memory-bound applications. Programming with high abstraction C/C++ with SYCL is now possible.	

Co-funded by the European Union
EuroHPC Joint Undertaking

Training Track	Topic	Objective	Format	Justification	Date
		scale these applications on FPGA clusters.			
Scientific Tools	Operational GeoAI for Scalable Satellite Analytics	To introduce scalable GeoAI workflows for Earth Observation. To implement an end-to-end GeoAI pipeline for object detection. To integrate heterogeneous geospatial data sources. To address core research challenges in GeoAI. To operate GeoAI workflows on a production HPC system. To promote reproducibility and experimental rigour.	Tutorial @ GeoAI 2026 Tutorial @ scynergy 2026	Earth Observation (EO) has become a cornerstone of both scientific research and operational services in the space industry. This workshop provides a hands-on, research-oriented introduction to scalable and reproducible GeoAI workflows for satellite imagery, using the MeluXina supercomputer as a concrete execution platform. Participants will implement an end-to-end GeoAI pipeline for large-scale object detection from Sentinel-2 imagery, integrating raster, vector, and auxiliary data sources. The tutorial emphasizes challenges central to current GeoAI research—multimodal learning, spatial generalization, benchmarking, and efficiency.	June 2026

Table 2 - EPICURE Year 3 Training Plan

Co-funded by the European Union
EuroHPC Joint Undertaking

3.Conclusions and Future Work

The training plan for the third year of the project was developed in accordance with the methodology defined in Deliverable 5.1. Unlike the initial iterations of the training programme, this plan benefits from experience accumulated over previous project years, knowledge gained from the execution of supported projects, and feedback and suggestions from Application Support Teams.

Based on these inputs, the resulting training plan reflects a more mature and informed understanding of user needs, emerging technical challenges, and strategic priorities within the EuroHPC JU ecosystem. We consider this plan ambitious, both in its scope and in the breadth of topics covered. Its full implementation requires, in practice, organising more than one training event per month, combining regular webinars with additional workshops, hackathons, and dissemination-oriented sessions.

At the same time, the training plan is conceived as a living document. It remains flexible and can be adjusted throughout its execution, allowing new training activities to be added whenever relevant topics emerge that would benefit from broader dissemination through dedicated webinars. This adaptability ensures that the training programme remains responsive to evolving project needs and technological developments, while maximising its overall impact.

Co-funded by the European Union
EuroHPC Joint Undertaking