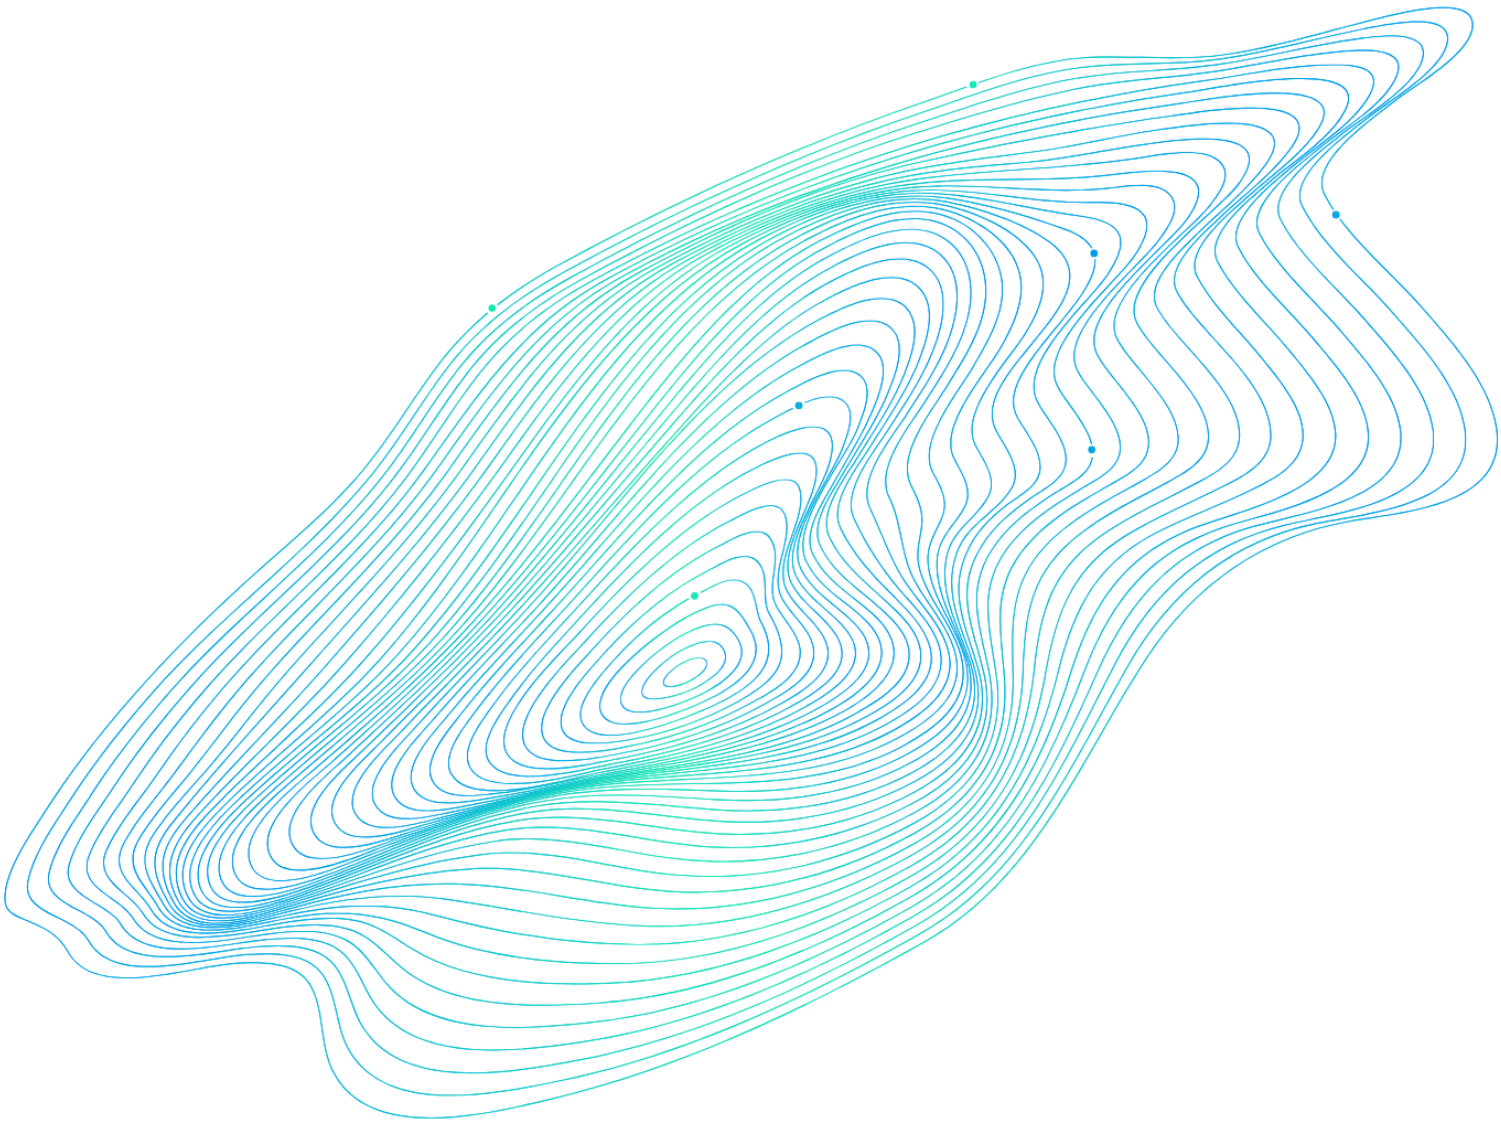
EPICURE

Unlocking European-level HPC Support

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



Co-funded by the European Union

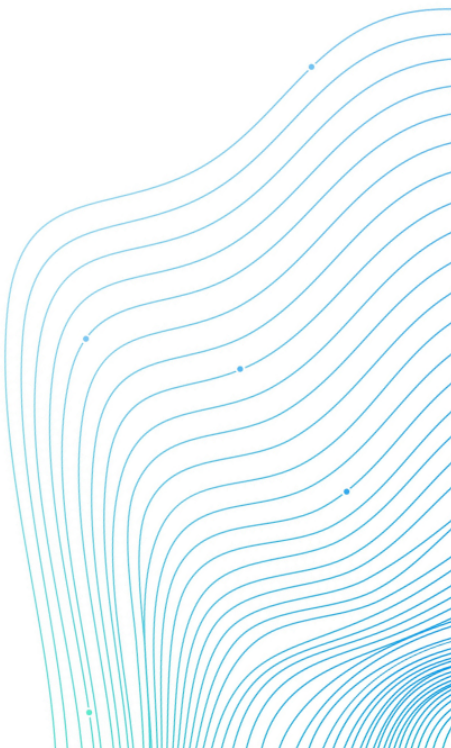
EuroHPC Joint Undertaking

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 5.1

Training activities report year 1



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	Marcos Gouveia (INESC TEC)
Contributors	Marcos Gouveia (INESC TEC), André Pereira (INESC TEC), Andrew Emerson (CINECA)
Peer Reviewers	João Barbosa (IT4I@VSB); António Sousa (INESC TEC)
Version	1.0
Due Date	2025-01-31
Submission Date	2025-01-31

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)

Version History

Revision	Date	Editors	Comments
0.1	10/12/2024	Marcos Gouveia	Table of Contents
0.2	06/01/2025	Marcos Gouveia	First Draft
0.3	16/01/2025	Marcos Gouveia	Corrections after internal review
0.4	22/01/2025	Marcos Gouveia	Updated info on sections 3.5 and 4.2
1.0	31/01/2025	Eeva Harjula	File completed and ready to be submitted

Glossary of Terms

Item	Description
AI	Artificial Inteligence
ARM	Advanced RISC Machine
AST	Application Support Team
CoE	Centre of Excellence
CPU	Central Processing Unit
EESSI	European Environment for Scientific Software Installation
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

Executive Summary

This document serves as a report on the training activities organised by EPICURE's Advanced Training Work Package during the first twelve months of the project. During the project's first year, it was fundamental to establish a transparent methodology to select which training activities should be organised. This implies considering the opinions of the different members of the Advanced Training Work Package, the feedback from the technical work packages intimately related to the user support activity, and studying the training program landscape already present in the EuroHPC ecosystem. Related to this last point, an effort was made to actively collaborate with the Castiel 2 project that organises training activities under the National Competence Centres and Centres of Excellence.

Nine EPICURE training activities have been planned and organised: six webinars, two in-person hackathons, and one workshop. Of these nine activities, four webinars, one hackathon, and a workshop have already been executed, while two webinars and one hackathon will occur soon. The executed webinars covered different topics related to HPC: in situ visualisation, SYCL programming for GPUs, the EESSI software stack, and the EasyBuild package manager. The following webinars will focus on extreme-scale climate modelling and how to use containers in HPC systems. In October, the first EPICURE Hackathon took place in CINECA, Bologna, and it gave 25 participants a chance to learn about GPU programming while using the Leonardo GPU partition, taking advantage of the "Bring your Own Code" hackathon style. The feedback regarding the hackathon was very positive and encouraging.

The second EPICURE hackathon is planned for the beginning of February 2025 in Braga, Portugal. This hackathon will have the same format as the former, and the central theme will be "HPC in ARM architecture", where the participants will learn 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.

Finally, a three-hour, hands-on workshop about FPGA programming was delivered by the EPICURE team at LuxProvide, to many attendees. Most of these events have been recorded and are freely available for anyone to watch on EPICURE's webpage, while the slides used by the speakers are also available. Overall, the outcome of this first year of training activities is positive. However, the lack of post-event feedback provided by the participants has been very scarce, making it difficult to assess the efficacy of the training program. Mitigation strategies to increase the number of participants who fill out the feedback forms have already been discussed and studied and will be implemented for future events.

Table of Contents

1. Introduction 8
2. Training Event Selection Procedure 9
2.1 Overview 9
2.2 Suggested Topics 9
2.3 Selection of Training Events 10
2.4 Organising the Event 11
3. Webinars 13
3.1 Scientific Visualisation at the Argonne Leadership Computing Facility 13
3.2 Heterogeneous Computing with SYCL and oneAPI on CINECA's Leonardo 14
3.3 Streaming Optimised Scientific Software: An Introduction to EESSI 15
3.4 Building Software with Ease: an Introduction to EasyBuild 17
3.5 Computational Challenges in Advancing Earth Sciences 19
3.6 Containerization in HPC Environments 21
4. Hackathons 22
4.1 GPU Hackathon 22
4.2 HPC in ARM architecture 26
5. Workshops 29
5.1 Introduction to FPGA computing for the HPC ecosystem on MeluXina 29
6. Conclusions and Future Work 31

Table of Figures

Figure 1 Diagram representing the different steps in the topic selection process. 9
Figure 2 Complete agenda of EPICURE's GPU Hackathon in CINECA, Italy. Images taken from the EPICURE website page for this event. 23
Figure 3 Preliminary agenda for EPICURE's HPC in ARM architecture Hackathon. Images taken from the EPICURE website page for this event. 27

1. Introduction

The EPICURE training program, for which the Advanced Training Work Package (WP5) is responsible, is tightly connected with the user support activities provided by the Application Support Teams (ASTs). While more technical work packages like WPs 2, 3, and 4 assist researchers awarded with computing time via a EuroHPC call, WP5's mission is to create an efficient training plan that empowers users by increasing their expertise in High-Performance Computing. The goal is to provide the training in the context of EPICURE and, in the long run, bidirectionally improve the high-level support experience. Since the training program may also target the application support specialists, both parties involved in user support can take advantage of the same training program, facilitating their communication during discussions and making user support more efficient. This document serves as a report of the training activities organised during the first twelve months of EPICURE. In this period, nine training events have been planned, and so far, six of them have been executed. The remaining three events have been organised and will occur until the end of February 2025.

The first step in organising the EPICURE training program plan was to define a methodology for selecting the topics to include in the program. Section 2 of this document contains a detailed explanation of the workflow used for this selection.

The Advanced Training WP has organised three types of training events so far: webinars, hackathons, and workshops. Webinars are live, online events where an invited speaker gives a presentation about a specific topic in which they are experts. In a hackathon, several participants are selected to participate in multiple-day, in-person events to learn about a particular subject. Most importantly, they can bring a code of their choice and have the support of mentors who will help them profile, port, and optimise it. Besides webinars and hackathons, the EPICURE WP5 has also organised a workshop. Workshops have a format similar to webinars, but they are typically longer and follow a hands-on approach with practical examples and demonstrations. The webinars, hackathons, and workshops organised in the period this document covers are described, respectively, in Sections 3, 4, and 5.

Finally, Section 6 presents conclusions about the training activities' outcomes, the difficulties encountered in implementing the training plan, and a mitigation plan for the future.

2. Training Event Selection Procedure

2.1 Overview

Since WP5's mission is to organise training events covering more advanced topics on HPC, the theme for each event must undergo a rigorous selection process. 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 Organizing the Event. Figure 1 depicts this process graphically. This Section will describe each of these steps in more detail.

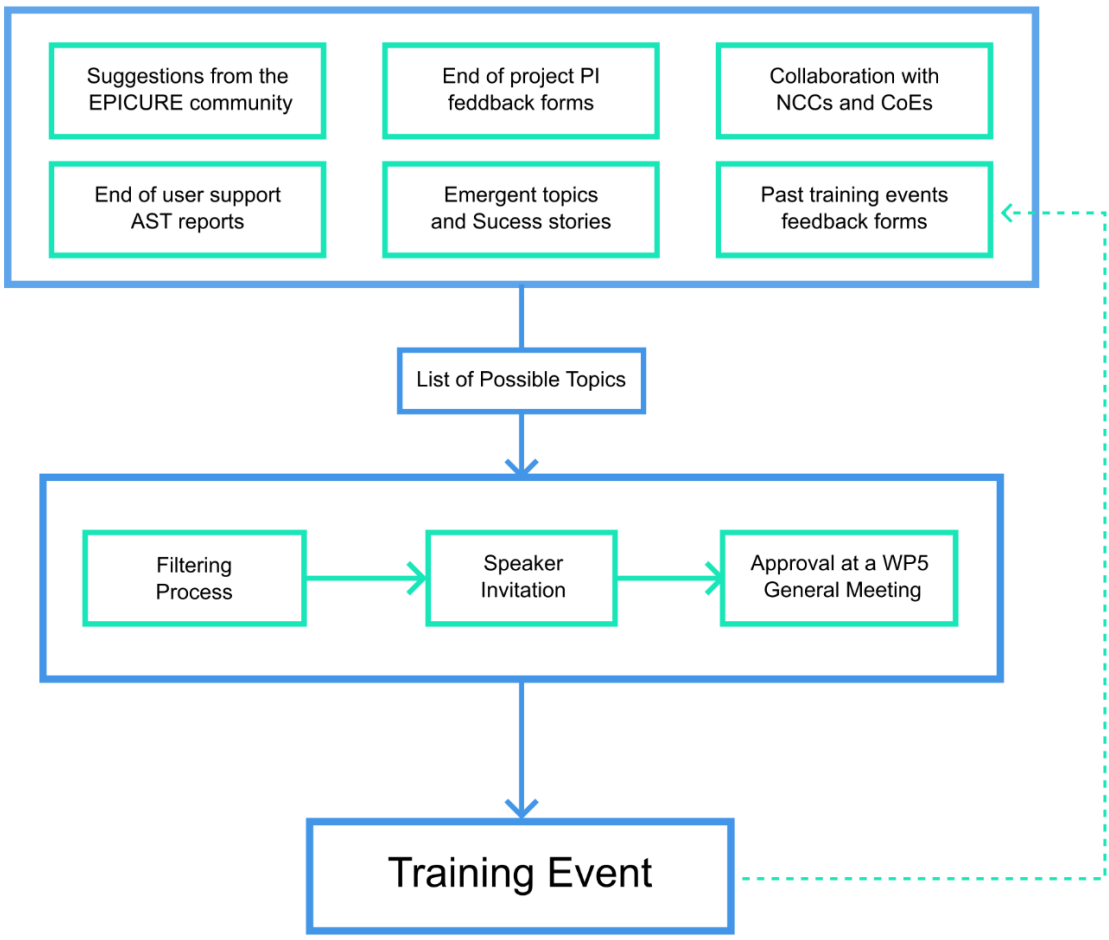


Figure 1 Diagram representing the different steps in the topic selection process.

2.2 Suggested Topics

Considering that HPC comprises so many subareas of interest to the EuroHPC 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 to 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 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 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.

2.3 Selection of Training Events

2.3.1 Filtering process

After a list of topics is compiled considering the sources described in Section 2.2, it is 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 programs have their own working groups that organise training activities for their target audience. To avoid duplication of efforts and expand the range of topics covered in training sessions, it is necessary to check if the topic has already been covered in a training activity organised by other programs. In some cases, if the

topic is considered very relevant to the EPICURE ecosystem, organising 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, this program aims to go more in-depth. Also, there are other training programs, namely the ones organised by NCCs, where these 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 issues 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.

2.3.2 Finding speakers

Before the topic is approved for a training event, an expert or team of experts who can lead a training session must be found. The first place to check for a suitable speaker is inside our own EPICURE community since it includes many experts in an extensive range of areas of knowledge. Second, in the interest of collaboration between different EuroHPC programs, people from Centres of Excellence (CoE) whose experience can increase the quality of the training activity, particularly in the area of interest covered by the CoE, are also considered. If needed, the speaker can be found outside the EuroHPC community, especially in HPC topics that may not be explored in the European community.

2.3.3 Approval of training actions

Once a reasonable number of topics and speaker combinations are available, a WP5 General Meeting is summoned. 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.

2.4 Organising the Event

2.4.1 Meeting with the speaker(s)

After a training event has been approved, the speaker(s) who have agreed to lead the activity are contacted to discuss details about the organisation of the event, the most important being:

- **Date and time:** The choice of a date for the activity must consider the speaker(s) availability, the possible overlap with important events of the HPC community (Supercomputing, ISC, etc.), and the different time zones involved.
- **Who will be the speakers:** When organising a training event, it is up to the invited person to decide if they want to ask other people to participate if they find that they can contribute to the quality of the event.
- **The content of the activity:** It is important to understand what the speaker wishes to show in this presentation. In some cases, suggestions may be made to the speakers to slightly adapt their plan for the event to cover some details surrounding the central theme and better suit the interests of the EPICURE community.
- **Information for event communication:** In this meeting, essential information that will help create a webpage for the training event is requested, namely the title of the event, the names and pictures of the speakers, a short bio for each of them, and an abstract describing what will be covered in the session.

2.4.2 Participant Registration

With the information obtained in the meeting with the trainers, event promotion can start and the registration process by which people can show their interest in attending the session. The registration form is designed to acquire essential data about the participants in the event, which will later be used for statistical analysis. The registration form includes the following fields:

- Name.
- Email.
- Gender
    - Woman
    - Man
    - Other
    - Prefer not to say.
- Main affiliation.
- Position at your affiliation organisation.
- Type of organisation
    - Academia
    - R&D institution or interface centre
    - Company
    - Other.
- Country of the affiliation organisation

In addition to this data, it is also asked if the participants want to subscribe to EPICURE's newsletter and warn them that the event will be recorded and made available online for future viewing.

Ideally, the registrations for an event will be open a month before the day of the session.

After completing the registration form, participants will receive an email confirming their registration. The email will also include the link for the online session (in the case of online live events) and other information specific to each activity.

2.4.3 Post-Event

After an event, the session recording is made public (with the speakers' consent). It can be accessed either directly from EPICURE's YouTube channel or on the event page on the EPICURE website. In the spirit of Open Science, it is crucial for the material from past training events to be available for people who could not attend the live session and for future users.

After each event, everyone who registered for it will receive a form to provide feedback and give their opinion about how EPICURE's training plan can be improved.

3. Webinars

This section will describe each webinar organised in the first twelve months. The description and speaker bios will be the same as those provided by the speakers and presented on the EPICURE website's events page: https://epicure-hpc.eu/events/.

3.1 Scientific Visualisation at the Argonne Leadership Computing Facility

3.1.1 Date and Time

This online webinar occurred on July 10th, 2024, from 16:00 to 17:00 CET.

3.1.2 Description

This talk will delve into the history of scientific visualisation at the Argonne Leadership Computing Facility (ALCF), highlighting successful collaborations with science teams. We will showcase exemplary science problems tackled at the facility and discuss our efforts in modernising visualisation tools to harness the power of exascale computing, enabling groundbreaking insights and discoveries.

3.1.3 Speakers

Silvio Rizzi: Silvio Rizzi is a Computer Scientist at ALCF. His research interests are large-scale and in situ scientific visualisation and analysis, display technologies, and immersive visualisation environments. He has an MS in Electrical and Computer Engineering and a PhD in Industrial Engineering and Operations Research, both from the University of Illinois-Chicago.

Joseph Insley: Joseph A. Insley leads the Visualization and Data Analysis team at ALCF. He develops scalable methods for large-scale data analysis and visualisation and investigates immersive virtual and augmented reality environments. He holds an MFA and MS in Computer Science from the University of Illinois at Chicago, and has published over 70 peer-reviewed papers, receiving several awards, including the R&D100 award, Best Paper Awards, and an ACM Gordon Bell Finalist (Honorable Mention).

Victor Mateevitsi: Victor A. Mateevitsi is a Computer Scientist at ALCF, specialising in visualisation, virtual reality, and human-computer interaction. With a background in computer science and experience at Pixar, DreamWorks, and Microsoft Research, he brings a unique perspective to the field. He holds a PhD in Computer Science from the University of Illinois at Chicago.

3.1.4 Relevance

Nowadays, the simulation of a complex system generates a massive amount of data that can quickly fill several disks on an HPC cluster. Also, since these simulations can run for an extended time, it is essential to check if the simulation is going in the right direction. As such, it is critical to visualise the data in real time. In situ visualisation is a technique that allows the user to overlap simulation and data visualisation and analysis and to establish a way to redesign the simulation in real time. This topic was selected for a webinar since it is an emerging technology that may benefit users with EuroHPC projects running lengthy and data-intensive simulations, particularly users with Extreme Scale projects.

3.1.5 Summary

The presentation was split between the three speakers, covering the following themes:

- A short introduction to the Argonne Leadership Computing Facility, including the contribution of the institution to research in the North American Continent, group members, and the evolution of the HPC systems hosted at ALCF.
- Examples of scientific visualisation projects where the ALCF was involved.

- What is In Situ Visualization, and why is it necessary?
- ASCENT – An in-house developed software for In Situ Visualization.
- Use cases of ASCENT.
- A peek into the future of In Situ Visualization.
- Q&A session.

3.1.6 Participant Data and Feedback

Sixty-one people registered for this event, while 49 of those participated. The registration forms provided the following metrics.

- **Attendance:** 80%
- **Gender:** Male (47%), Female (24%), Prefer not to say (29%).
- **Country of Affiliation:** 17 different countries.
- **Type of Organization:** Academia (39%), R&D Institutions (29%), Other (27%), Industry (6%).

3.2 Heterogeneous Computing with SYCL and oneAPI on CINECA's Leonardo

3.2.1 Date and Time

This online webinar occurred on October 10th, 2024, from 9:30 to 11:30 CET.

3.2.2 Description

In this webinar, attendees will learn the fundamentals of the SYCL programming model, exploring its key concepts and capabilities, and how to leverage oneAPI tools to bridge SYCL code with proprietary programming models such as CUDA and HIP, and deploy on a variety of hardware platforms.

The webinar also features a detailed example of a successful CUDA to SYCL migration, showcasing the practical steps and benefits of this approach. With ample time for questions and discussion, attendees will leave with a solid understanding of SYCL and the confidence to apply it to their own HPC projects.

Finally, the webinar also introduces developers to the SYCL oneAPI Hackaton at CINECA, a three-stage event (this webinar, an online learning course, and a final on-site Hackaton) where attendees will learn how to effectively write efficient heterogeneous SYCL code, and how to use oneAPI tools to deploy it on CINECA's Leonardo, a Tier-0 EuroHPC supercomputer ranked #7 on the Top500 list of the world's most powerful supercomputers.

3.2.3 Speaker

Biagio Cosenza is a tenure-track Assistant Professor in the Department of Computer Science at the University of Salerno (Italy), an Intel oneAPI Certified Instructor, and a member of the Khronos SYCL Working Group. He joined the University of Salerno in August 2019 through a national brain gain program (Attraction and International Mobility) and received the Abilitazione for Italian Associate Professorship. From 2015 to 2019, he was Senior Research at the TU Berlin (Germany), where he was Principal Investigator for the DFG project Celerity and received the Habilitation from the Faculty IV. From 2011 to 2015, he was Postdoctoral researcher at the University of Innsbruck (Austria), where he contributed to the Insieme Compiler project and the DK-Plus multidisciplinary platform for Scientific Computing. His research is currently funded by the European HPC Joint Undertaking (LIGATE project), the Italian Ministry of Research (LibreRT project, PRIN 2022), and several industrial projects. Cosenza's main research interests are in the field of high-performance computing, in particular with respect to programming models, compiler technology, optimisation and tuning.

3.2.4 Relevance

SYCL (pronounced “sickle”) is a programming model for heterogeneous programming. Developed by the Khronos Group, it is sometimes called Intel Data Parallel C++ (DPC++). With small source code changes, SYCL allows users to write code that works for CPUs, GPUs, FPGAs, and other types of accelerators. One of the main objectives of SYCL is to enable users to write GPU kernels that are architecture agnostic, such that the user can run the same code on LUMI’s AMD GPUs and Leonardo’s NVIDIA A100s.

One of the positive features of SYCL is that beginners can use it as a first approach to heterogeneous programming while also providing experienced users with advanced options that can make a difference when optimising the performance of more complex codes.

Since portability, acceleration, and optimisation are essential concepts in HPC, this training activity can give participants the basic tools and concepts needed to write portable and versatile code such that they can run the same code in multiple EuroHPC supercomputers with different architectures.

3.2.5 Summary

The outline of the presentation was as follows:

1. Introduction to heterogeneous computing.
2. Introduction to SYCL.
    a. oneAPI setup in CINECA
    b. Basic: queue and parallelism, device, range, task scheduling, and memory access modes.
    c. Advanced: kernel reduction, subgroups, and group algorithms, atomics.
    d. SYCL and the oneAPI ecosystem.
3. CUDA to SYCL.
4. Online learning on the CodeReckons platform.
5. Q&A session.

During the presentation, the speaker used a hybrid format to present the theoretical aspect of heterogeneous, some code examples to introduce SYCL concepts and hands-on exercises. For the demonstrative moments, the speaker used Leonardo to run their examples.

3.2.6 Participant Data and Feedback

Note: This event had minimal information about the participants.

- Registrations: 35
- Participants: 25
- Attendance: 71%

None of the participants responded to the feedback form.

3.3 Streaming Optimised Scientific Software: An Introduction to EESSI

3.3.1 Date and Time

This online webinar occurred on November 15th, 2024, from 14:00 to 15:30 CET.

3.3.2 Description

What if you could avoid having to install a broad range of scientific software from scratch on every HPC cluster, laptop, or cloud instance you use or maintain, without compromising on performance?

Installing scientific software for supercomputers is known to be a tedious and time-consuming task. Especially as the HPC user community becomes more diverse, computational science expands rapidly, and the diversity

of system architectures increases, the application software stack continues to deepen. Simultaneously, we see a surge in interest in cloud computing for scientific computing. Delivering optimised software installations and providing access to these installations in a reliable, user-friendly, and reproducible way is a highly non-trivial task that affects application developers, HPC user support teams, and the users themselves.

This webinar aims to address these challenges by providing the attendees with the knowledge to stream optimised scientific software installations. For this, we will introduce the European Environment for Scientific Software Installations (EESSI), a collaboration between various European HPC sites & industry partners, with the common goal of creating a shared repository of scientific software installations that can be used on a variety of systems, regardless of which flavour/version of Linux distribution or processor architecture is used, or whether it's a full size HPC cluster, a cloud environment, or a personal workstation.

We will cover the design and usage EESSI, different ways to accessing EESSI, how to add software to EESSI, and highlight some more advanced features including support for NVIDIA GPUs and facilitating the deployment of pre-release builds of scientific software. We will also show how to engage with the community and contribute to the project.

3.3.3 Speakers

Kenneth Hoste: Kenneth Hoste holds a Masters (2005) and PhD (2010) in Computer Science from Ghent University (Belgium). His dissertation topic was "Analysis, Estimation and Optimization of Computer System Performance Using Machine Learning". Since October 2010, he is a member of the HPC team at Ghent University where he is mainly responsible for user support & training. He is the lead developer of EasyBuild since 2011, a software build and installation framework for (scientific) software on HPC systems, an active contributor in the European Environment for Scientific Software Installations (EESSI) project since 2020, and a partner in the MultiXscale EuroHPC Centre-of-Excellence since 2023. He enjoys (co-) organising workshops and meetups, like the EasyBuild User Meetings (yearly since 2016) and the HPC, Big Data, and Data Science devroom at FOSDEM (since 2014).

Alan O'Cais: Dr. Alan O'Cais is currently the technical manager of MultiXscale, a EuroHPC centre of Excellence. He previously acted as the Software Manager of E-CAM, an EU Centre of Excellence in computing applications, while at Juelich Supercomputing Centre (JSC). He has previously managed the user software and application environment on the multi-PFlop JURECA hybrid cluster system at JSC. Between 2010 and 2015 he was primarily focused within the LinkSCEEM-2 project helping to develop a Virtual Research Community in Computational Science in the Eastern Mediterranean region. He has held research positions at Trinity College and the University of Adelaide. In 2008 he joined the Cyprus Institute as Scientific Coordinator of the Computation-based Science and Technology Research Centre (CaSToRC). He has worked at the Juelich Supercomputing Centre from 2010 to 2021 and at the University of Barcelona (with an affiliation to CECAM) from 2022.

Pedro Santos Neves: Pedro Santos Neves joined the HPC team at the University of Groningen (The Netherlands) in 2023. While his role focuses on user support, training, and development he is also a contributor to the EESSI project and the MultiXscale EuroHPC Center of Excellence. Previously, he worked as a scientific programmer at the the Groningen Institute for Evolutionary Life Sciences, where he helped researchers design, develop, and maintain scientific software. His background is in Evolutionary Biology in which he holds a Master's degree (University of Lisbon, 2016) and PhD (University of Groningen, 2024). His doctoral research revolved around the use and development of stochastic models of island community evolution.

3.3.4 Relevance

The European Environment for Scientific Software Installations (EESSI) initiative simplifies and standardises software installation in HPC environments. EuroHPC's MultiXscale Centre of Excellence funds EESSI's development. As such, this webinar results from a collaboration between the EPICURE project and the CoE network under the umbrella of Castiel 2.

Using EESSI, researchers can focus on their scientific work instead of spending a non-productive amount of time trying to compile some applications and their dependencies and worrying about optimising the installation while taking full advantage of the machine's architecture. Considering that this operation is often the focus of EPICURE's Level 2 support, knowing about EESSI can lead to a more productive user support activity.

The EESSI project also encourages open science. Its active community, which anyone can join, is helpful for discussing problems, sharing ideas, and contributing to the EESSI code.

This training event occurred a week before the Supercomputing 24 meeting, where the EESSI project was awarded the "HPCwire 2024 Reader's Choice Award for Best HPC Programming Tool or Technology" prize, showing that EESSI is well recognized not only in Europe but globally.

3.3.5 Summary

The outline of the presentation was as follows:

- The MultiXscale EuroHPC Centre-of-Excelence in a Nutshell.
- Motivation for and high-level introduction to EESSI.
- Hands-on live demo: getting access to and using EESSI.
- More in-depth overview of the layered design of EESSI.
- Overview of available software and supported hardware.
- Workflow for contributing additional software to EESSI.
- Status on GPU support (including a live demo).
- Use cases enabled by EESSI: Continuous Integration, training, and collaboration with developers.
- Q&A session.

The speakers constructed this webinar presentation to encourage the audience's strong participation. During the webinar, there were a few moments when the speakers introduced competitive, interactive quizzes that participants could answer by accessing a link on their computer or smartphone. In addition to these interactive moments, the speakers resorted mainly to hands-on demonstrations rather than following a traditional theoretical approach.

3.3.6 Participant Data and Feedback

Fifty-six people registered for this event, while 28 of those participated. The registration forms provided the following metrics.

- **Attendance:** 50%
- **Gender:** Male (84%), Female (16%).
- **Country of Affiliation:** 13 different countries.
- **Type of Organization:** Academia (57%), R&D Institutions (33%), Company (10%)

None of the participants completed the post-event feedback form

3.4 Building Software with Ease: an Introduction to EasyBuild

3.4.1 Date and Time

This online webinar occurred on December 13th, 2024, from 14:00 to 15:30 CET.

3.4.2 Description

Maintaining a collection of software installations for a diverse user base can be tedious, repetitive, error-prone, and time-consuming.

Because most end-user software packages for an HPC environment are not readily available in existing OS package managers, they require significant extra effort from the user support team.

In this presentation, we present EasyBuild (https://easybuild.io), a software installation framework written in Python that aims to support the various installation procedures used by the vast collection of software packages that are typically installed in an HPC environment.

It is built on top of existing tools, provides support for well-established installation procedures, and makes sharing implementations of installation procedures very easy. Installing software packages that are supported can be done by issuing a single command, even if dependencies are not installed yet. Hence, it simplifies the task of HPC site support teams, and even allows end-users to keep their software installations consistent and up to date.

We will also show how you can build and install software with EasyBuild on top of EESSI, the European Environment for Scientific Software Installations (EESSI, https://eessi.io).

3.4.3 Speakers

Kenneth Hoste: Kenneth Hoste holds a Masters (2005) and PhD (2010) in Computer Science from Ghent University (Belgium). His dissertation topic was “Analysis, Estimation and Optimization of Computer System Performance Using Machine Learning”. Since October 2010, he is a member of the HPC team at Ghent University where he is mainly responsible for user support & training. He is the lead developer of EasyBuild since 2011, a software build and installation framework for (scientific) software on HPC systems, an active contributor in the European Environment for Scientific Software Installations (EESSI) project since 2020, and a partner in the MultiXscale EuroHPC Centre-of-Excellence since 2023. He enjoys (co-) organising workshops and meetups, like the EasyBuild User Meetings (yearly since 2016) and the HPC, Big Data, and Data Science devroom at FOSDEM (since 2014).

Pedro Santos Neves: Pedro Santos Neves joined the HPC team at the University of Groningen (The Netherlands) in 2023. While his role focuses on user support, training, and development he is also a contributor to the EESSI project and the MultiXscale EuroHPC Center of Excellence. Previously, he worked as a scientific programmer at the the Groningen Institute for Evolutionary Life Sciences, where he helped researchers design, develop, and maintain scientific software. His background is in Evolutionary Biology in which he holds a Master’s degree (University of Lisbon, 2016) and PhD (University of Groningen, 2024). His doctoral research revolved around the use and development of stochastic models of island community evolution.

Caspar van Leeuwen: Caspar van Leeuwen holds a PhD (2017) in Medical Physics from the Amsterdam UMC. After his PhD, he joined the HPC team at SURF, which hosts the Dutch National Supercomputer. There, he was responsible for managing the software environment for several years. Since 2020, he is an EasyBuild maintainer and an active contributor in the European Environment for Scientific Software Installations (EESSI) project. Since 2023, he leads one of the work packages of the MultiXscale EuroHPC Center-of-Excellence.

3.4.4 Relevance

This webinar on EasyBuild serves as a complement to the previous training event on EESSI. EESSI aims to create a stack that can deliver optimised software to systems with varying architectures. EasyBuild is the tool behind EESSI’s software creation, and it relies on writing easyconfig files. These files describe everything that the user intends to use to compile the desired application, such as what compiler toolchain to use (compilers and corresponding versions) and the source code that will be compiled, and it also enables the user to select which versions of the several dependencies they want to use.

Along with Spack, EasyBuild is the most used software installation tool in EuroHPC supercomputers, and as such, users of those machines need to know how they can build software that is not yet present in the system. Knowing how to build software using EasyBuild is also essential for AST members, particularly in Level 2

support. A combination of EESSI and EasyBuild can be a very powerful tool as it simplifies the installation of optimised software while allowing advanced users total control of the compilation process.

3.4.5 Summary

The outline of the presentation was as follows:

- What is EasyBuild? + Terminology.
- Installation and configuration of EasyBuild.
- Basic Usage of EasyBuild.
- Troubleshooting + a sneak peek at EasyBuild 5.0.
- Adding support for additional software.
- Advanced topics: hooks and beyond.
- Using EasyBuild on top of EESSI.
- Outlook to EasyBuild 5.0.
- Q&A

3.4.6 Participant Data and Feedback

Sixty-four people registered for this event, while 43 of those participated. The registration forms provided the following metrics.

- **Attendance:** 67%
- **Gender:** Male (74%), Female (18%), Other (5%), I prefer not to say (3%).
- **Country of Affiliation:** 16 different countries.
- **Type of Organization:** Academia (53%), R&D Institutions (24%), Company (18%), Other (6%)

None of the participants completed the post-event feedback form.

3.5 Computational Challenges in Advancing Earth Sciences

This training event has been organised but has not yet been executed.

3.5.1 Date and Time

This online webinar is scheduled to occur on January 17th, 2025, from 14:00 to 15:30 CET.

3.5.2 Description

Earth sciences are crucial for understanding and addressing challenges like climate change, extreme weather, and environmental sustainability. Advances in computational science have revolutionised the field, enabling the development of sophisticated models and workflows that simulate the complex processes shaping our planet.

This talk will explore the computational challenges faced in various domains of Earth sciences. Topics include building and running advanced climate and weather models, analysing and processing massive observational and simulated datasets, and applying high-performance computing to areas such as atmospheric modelling, ocean dynamics, and ecosystem simulations. These challenges highlight the intricate balance of scientific inquiry and computational innovation required to push the boundaries of Earth system research.

3.5.3 Speakers

Oriol Tintó Prims holds a PhD in Earth Sciences from the Universitat Autònoma de Barcelona and the Barcelona Supercomputing Center, where he developed advanced methods to optimise the computational performance of Earth science models on HPC infrastructure.

After contributing to the Waves to Weather project in Germany from 2021 to 2023, he rejoined the Barcelona Supercomputing Center in 2024, where he currently holds an AI4S fellowship.

His expertise includes mixed-precision computing, data compression techniques, and performance optimisation, with contributions to major European projects like EERIE and ESiWACE3, driving innovation in computational Earth sciences.

3.5.4 Relevance

The extensive work done at BSC in climate modelling is of great importance to scientific knowledge in general and society. Their work in climate prediction, atmospheric and air quality modelling, and studying extreme events such as hurricanes is essential for the well-being of populations and industries directly related to agriculture.

Studying such complex phenomena requires computing very heavy calculations, made possible only through HPC systems such as MareNostrum 5. Any application used to study these systems must be optimised to scale well using hundreds of computing nodes and accelerators such as GPUs. The work involved in optimising extreme-scale applications is very complex since every aspect of the computation process must be considered, such as intra and internode communication, kernel launching on accelerators, intensive I/O operations, etc.

This training event aims to provide a glimpse into the struggles related to extreme-scale operations. This subject is essential for users who want to scale their code to run with a high node count and for support specialists who may need to deal with these problems when assisting Extreme-Scale projects.

3.5.5 Summary

The outline of the presentation was as follows:

- Why is it essential to model climate?
- Climate change and global warming.
- Climate as a chaotic system.
- The Computational Earth Sciences Group at BSC.
- How to bring legacy code to HPC systems.
- Data management.
- The role of AI in climate modelling.
- Visualisation techniques.

3.5.6 Participant Data and Feedback

Thirty-seven registered for this event, while twenty-three of those participated. The registration forms provided the following metrics.

- **Attendance:** 62%
- **Gender:** Male (83%), Female (17%)
- **Country of Affiliation:** 12 different countries.

Type of Organization: Academia (47%), R&D Institutions (33%), Company (10%), Other (10%)

3.6 Containerization in HPC Environments

This training event has been organised but has not yet been executed.

3.6.1 Date and Time

To be confirmed. The speaker will be available only after the second half of February.

3.6.2 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 aspects, such as parallelisation with MPI and CUDA for GPU-accelerated workloads and including management and environment systems such as Conda for managing smaller, lightweight environments. The session emphasised practical strategies to optimise HPC workflows with containerised solutions.

3.6.3 Speakers

Teo Prica is an HPC Development Engineer at IZUM (HPC Vega). His interests include Virtualization and System Orchestration, Containerization, Scientific Software in HPC, and Optimization with a focus on Machine Learning. He is a member of HPC Federation Working Group (EGI), EuroHPC Containers Forum, a member of Slovenian National Supercomputing Network (SLING) consortium, Leonardo High Level Support Team, and contributions within European and National projects such as EPICURE, interTwin, SMASH, Castiel2, and EUMaster4HPC. He was part of the Deployment of 1st Operational EuroHPC JU system and Engineering team of EuroHPC Vega and provides High-level SLING User support and Technical Support and Maintenance for HPC Vega.

3.6.4 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 machine.

Containerisation is a crucial concept closely related to portability, particularly in high-performance computing. Software distributed in a container is more straightforward to run in multiple HPC systems, even if their architecture and environment are different.

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 submitting system. Moreover, containerisation is also very important in AI workflows, as it can generate many small files, quickly draining the user's file quota while making I/O systems less efficient. If all the small files generated, for example, by Python environments, are all inside a single container file, this problem is avoided. As such, this activity suits user support specialists in Level 2 and Level 3 work and users with EuroHPC projects, particularly those awarded in AI calls.

4. Hackathons

4.1 GPU Hackathon

4.1.1 Overview

The EPICURE team at CINECA, Bologna, organised the first EPICURE hackathon with the support of the WP5 team at INESC TEC. The topic chosen for the hackathon was GPU programming, taking advantage of Leonardo's GPU partition, which is equipped with approximately 3,500 GPU nodes and 4 A100 64GB GPUs each. During the lab sessions, the hackathon participants were provided temporary training accounts to test their code.

Over the years, GPU acceleration has become very popular among HPC users, particularly since the rise in AI research. As such, organising a training event around GPU computing would make the Hackathon popular, particularly in the EuroHPC community. In addition to the "Bring your Code" modality of the hackathon, participants could also choose to participate in challenges created by the event's organising team.

The maximum number of participants was 25, and they could participate individually or as a team. The place chosen for the event was the CINECA building located in Casalecchio di Reno on the outskirts of Bologna. However, a few days before the start of the event, there was a need to change the location of the event to the INFN CINECA building, where Leonardo is hosted. This was due to the heavy rain that affected the region of Bologna the previous week, causing floods in different areas of the city and its surroundings, resulting in electrical problems in the main CINECA building and making it impossible for the event to occur in that location. As such, the event was relocated to the INFN CINECA building.

Although the maximum number of registrations was reached, only 23 people participated in the Hackathon out of the 25.

4.1.2 Description

Taken from the Event page in EPICURE's webpage.

"The EPICURE GPU Hackathon, the first one organised by EPICURE, will focus on GPU programming, and it will take place in CINECA, Bologna, Italy, from 28-31 October 2024.

This event will include a series of lectures by experts on GPU computing, and hands-on practice sessions and the last day, there will be a "Bring your own code" session where you can bring code you have been working on and have some assistance from GPU programming experts. Preferably, you will register as a two-person team (you can also register individually) so that you can work in pairs during hands-on sessions, the "Bring your own code" session and other challenges.

For this event, you will have access to the Leonardo Booster cluster, where you will be able to test your code on its state-of-the-art NVIDIA GPUs."

4.1.3 Agenda

Day 1 | October 28

9.00 – 9.15	Welcome and registration
9.15 – 9.30	EPICURE HPC application support for EuroHPC users
9.30 – 10.30	Introduction to GPU computing
10.30 – 11.00	Coffee break
11.00 – 12.30	Introduction to CUDA
12.30 – 14.00	Lunch
14.00 – 15.30	OpenAcc programming and profiling with Nsight
15.30 – 16.00	Coffee break
16.00 – 17.30	OpenAcc programming and profiling with Nsight

The 1st day will serve as an introduction to GPU computing. As such, participants who are comfortable with GPU computing may join the hackathon only on the 2nd day

Day 2 | October 29

9.00 – 10.30	Advanced GPU optimisation and profiling
10.30 – 11.00	Coffee break
11.00 – 12.30	Advanced GPU optimisation and profiling
12.30 – 14.00	Lunch
14.00 – 17.30	Lab session (with your code or a provided code)

Day 3 | October 30

9.00 – 12.30	Lab session
12.30 – 14.00	Lunch
14.00 -17.30	Lab session

Day 4 | October 31

9.00 – 12.30	Lab session
12.30 – 14.00	Lunch
14.00 – 17.30	Lab session

Figure 2 The complete agenda for EPICURE's GPU Hackathon in CINECA, Italy. Images taken from the EPICURE website page for this event.

4.1.4 Lectures

The expositive part of the Hackathon was composed of seven sessions, each presenting essential concepts that would later be useful for participants during the Lab Sessions. The list of lectures and a summary for each of them is given as follows:

- **EPICURE HPC application support for EuroHPC users** (Andrew Emerson, CINECA):
  - What is the EPICURE project, and how can researchers apply for user support?
- **Introduction to GPU computing** (Nitin Shukla, CINECA):
  - Introduction to heterogeneous computing.

  - CPUs vs GPUs: What is the difference?
  - GPU Architecture. Example: Leonardo's NVIDIA A100 GPU.
  - GPU as an accelerator.
  - Different ways to program GPUs are Libraries, CUDA, OpenACC, and SYCL.
- **Introduction to CUDA** (Nitin Shukla, CINECA):
  - CUDA programming and execution model.
  - Launching a CUDA kernel.
  - Concepts of Thread, Block and Grid.
  - Data Movement.
  - Compiling and Running CUDA code.
  - Memory Management
  - Example: Matrix multiplication
- **OpenACC programming and profiling with Nsight I & II** (Laura Bellentani, CINECA):
  - Why OpenACC?
  - OpenACC Syntax: parallel, reduction, sequential, private,…
  - Compiling an OpenACC application.
  - Implicit and Explicit Data Management
  - Gang, worker, and vector.
  - CUDA - OpenACC interoperability.
  - OpenACC profiling with Nsight Systems.
- **Advanced GPU optimisation and profiling I** (Sergio Orlandini, CINECA):
  - Leonardo Booster modules: InfiniBand, PCI, and NVLink connections and bandwidth.
  - Asynchronous programming.
  - CUDA Streams and concurrent kernel execution.
  - Multi-GPU parallelism with OpenACC, MPI-aware CUDA, and NCCL.
- **Advanced GPU optimisation and profiling II** (Ilya Zhukov, Jülich Supercomputing Centre):
  - Tuning Tools Overview: Score-P, Cube, Vampir and Scalasca.
  - Performance Engineering principles.
  - How to use Score-P, Cube, Vampir and Scalasca.
  - How to combine multiple tools for detailed profiling.
- **How to use Leonardo** (Moreno Guernelli, CINECA):
  - Leonardo's Architecture.
  - Accessing Leonardo.
  - Login Nodes, Accounting, and Filesystem.
  - Software Modules.
  - Job Scheduling.

On the first day of the Hackathon, there was time for a guided tour of the Leonardo supercomputer, allowing the participants to see the system and its infrastructure.

4.1.5 Lab Sessions

4.1.5.1 Bring Your Own Code:

Of the 23 participants, 14 participated in the "Bring your own code" lab sessions. Some examples of project ideas brought by the participants:

- Particle Finite Element Method (PFEM) applied to Computational Fluid Dynamics.
- Finite-difference time-domain (FDTF) written in the C programming language, including a module for Fast Fourier Transform algorithms.
- Material Point Method/Particle in Cell (MPM/PIC) code for simulating compressible fluids.
- Time-Dependent Density Functional Theory (TDDFT) code used to calculate the spectra of chemical species.

- An automatic machine learning pipeline for earthquake pick detection that can identify P and S-wave arrival times.
- Quantum computer emulator powered by tensor networks. Implementation of multi-GPU emulation.
- Numerical solution of Stochastic Differential Equations (SDEs) for the propagation of cosmic rays in the heliosphere.
- Understanding the NCCL communication library for a multi-GPU implementation of the FFTXlib library used in Quantum Espresso.

It is important to note that at least two participating groups already had access to the Leonardo cluster via EuroHPC development and benchmark calls.

4.1.5.2 Proposed Challenges:

The participants who chose to work on the challenges proposed by the mentors had three different challenges to choose from:

- Calculating the Mandelbrot fractal set using OpenACC.
- Solving the Poisson equation using the Jacobi method using CUDA.
- Solving a Computational Fluid Dynamics problem using OpenACC for single and multi-GPU implementations.

At the beginning of the hackathon, the participants were asked to present which objectives they would like to reach during lab sessions. During the hackathon, participants had access to a HackMD repository where they could document their progress. On the last day of the hackathon, participants were allowed to present their work during the event and to show if the objectives they had set in the beginning were accomplished. After watching the final presentations and considering the work done by the different groups, the team that showed the most improvement was awarded a prize of an EPICURE mug for each of the team members.

4.1.6 Participant Data and Feedback

Analysing the registration forms of the 25 participants, we can extract the following metrics:

- **Gender:** Male (100%) / Female (0%).
- **Format chosen:** Bring your Own code (60%) / Challenges (40%).
- **Nationality:** Italian (76%), Indian/Iranian/Danish/German/Mexican/Spanish (4% each).
- **Position:** PhD Student (32%), Postdoctoral Researcher (16%), Other (16%), Undergraduate (12%), Industry (12%), Master Student (8%), Non-Academic (4%)

After the Hackathon, participants were contacted to give their opinions about the event. The feedback form sent to them contained the following questions, and their response to each one was:

- **Q1:** How do you rate the event regarding the usefulness of the information provided during the hackathon? Options: Very Good (80%) / Good (20%) / Fair (0%) / Bad (0%).
- **Q2:** How do you rate the event regarding the adequacy/relevance of topics covered? Options: Very Good (80%) / Good (20%) / Fair (0%) / Bad (0%).
- **Q3:** How do you rate the hackathon regarding its format? Options: Very Good (70%) / Good (20%) / Fair (10%) / Bad (0%).
- **Q4:** Was the session up to your expectations? Options: Yes (100%) / No (0%).
- **Q5:** What theme/topic would you like to see covered during our training events? Open answer.
- **Q6:** Could you please share a short testimonial about your experience at this hackathon? Open answer.
- **Q7:** Other comments? Open answer.

The following list of suggested topics for future training events was extracted from the answers to Q5:

- CUDA programming in Python and PyTorch profiling.

- Half precision in scientific computing.
- NVSHMEM and NCCL.
- Hybrid MPI – OpenACC programming.
- CUDA kernel optimisation for AI training/inference and distributed AI in multi-node, multi-GPU machines.
- Cybersecurity.
- Best practices for HPC, including parallel computing and GPU usage.

From the testimonials gathered in Q6 and the comments in Q7, it is possible to extract a list of positive remarks and a list of improvements for future events given by the participants:

Positive comments:

- The teaching material was very helpful, and the lectures were well-paced.
- It is a good opportunity to meet people with the same interests and exchange ideas.
- The support from the mentors was invaluable.
- The hackathon covers the essential material for beginners and allows us to learn more advanced concepts to apply to our case studies.

Improvements:

- The number of mentors should be higher.
- Having access to the training material ahead of the event.
- Include prizes in the hackathon.

4.2 HPC in ARM architecture

This training event has been organised but has not yet been executed.

4.2.1 Overview

The second EPICURE hackathon will be organised by INESC TEC and take place in Braga, Portugal, between the 4th and 7th of February 2025. This training event will focus on using processors with the ARM architecture for High-Performance Computing. EuroHPC's Deucalion is the only HPC infrastructure to which users can apply via EuroHPC access calls and have access to nodes with ARM CPUs. Considering that the future exascale machines Alice and Jupiter will be equipped with ARM processors, this hackathon will serve 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.

This second hackathon will closely follow the previous one's format, 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.

4.2.2 Description

"The second EPICURE Hackathon will focus on using HPC in CPUs with the ARM architecture, and it will take place at the INESC TEC's hub at the University of Minho, Braga, Portugal, from 4-7 February 2025.

The event will include lectures by experts on using ARM architecture in the context of HPC and practical sessions. We will adopt the "Bring Your Code" hackathon type, where the participants will bring code they are using for their research and wish to port or optimise so they can run on ARM CPUs.

You can register for this event individually or as part of a team with whom you will work to develop your code.

The participants will have access to the ARM partition of the EuroHPC supercomputer Deucalion."

4.2.3 Preliminary 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 ARM Architecture | **Speaker:** TBU
16:00 – 16:15 Coffee Break

16:15 – 18:00 Visit to Deucalion

20:00 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 – 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 How to Use Deucalion | **Speaker:** TBU

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 3 Preliminary agenda for EPICURE's HPC in ARM architecture Hackathon. Images taken from the EPICURE website page for this event.

4.2.4 Lectures

The list of lectures and a summary of 4 of them is presented as follows:

- **EPICURE HPC application support for EuroHPC users** (António Sousa):
  - What is the EPICURE project, and how can researchers apply for user support?
- **Introduction to ARM architecture** (TBD):
  - Overview of processor architectures.
  - Origin of ARM processors.
  - What is RISC?
  - Components of ARM architecture: Registers, Instruction set, memory management, …
  - ARM vs x86.
  - Advantages and Limitations of ARM.
  - Emergence in the world of HPC
  - Fugaku and the A64FX
  - Technical aspects of ARM processors for HPC.
- **Porting x86 code to ARM** (Bernardo Malaca and Ricardo Fonseca):
  - Benefits of ARM in HPC.
  - Future consequences of sticking to x86.
  - Overview of the porting process.
  - Challenges in porting x86 to ARM.
  - SVE vectorisation.
  - Osiris: A successful case of porting.
- **AI in ARM architecture** (Alícia Oliveira).
  - Why run AI models in ARM processors?
  - Training and Inference in ARM.
  - Acceleration with GPUs, NPUs, and TPUs.
  - AI workflows in Deucalion's ARM partition.
  - Deucalion's OpenOnDemand platform.
- **Assessing the Performance and Scalability of ARM Processors for Industrial Applications** (Gabriel Marcos Magalhães):
  - Quick introduction to OpenFoam.
  - OpenFoam in ARM architecture.
  - Performance: ARM vs x86.

5. Workshops

5.1 Introduction to FPGA computing for the HPC ecosystem on MeluXina

5.1.1 Date and Time

This workshop occurred on November 28th, 2024, from 9:30 to 12:30 CET.

5.1.2 Description

Taken from the Event page in EPICURE's webpage:

"In this session, participants will explore how to program High-Performance FPGA cards using the Intel DPC++ compiler, a SYCL implementation, to facilitate heterogeneous device programming.

Field-Programmable Gate Arrays (FPGAs) are integrated circuits that can be configured by users post-manufacturing. They consist of a network of programmable logic blocks and reconfigurable interconnects, enabling the creation of custom digital circuits. Traditionally, FPGAs have been programmed using Hardware Description Languages (HDLs) such as Verilog and VHDL. However, designing high-performance accelerators with these languages demands extensive knowledge and experience in hardware design.

By leveraging higher-level abstractions like SYCL and OpenCL—C/C++-based programming models familiar to software developers—developers can generate hardware kernels through an offline compiler. This significantly simplifies FPGA programming, reducing development time compared to HDL, which typically involves more complex coding, simulation, and debugging processes.

Following an introduction to Meluxina's FPGA cards and an overview of SYCL programming using the OneAPI software development toolkit, participants will engage in a live demonstration led by the presenter. This demonstration will feature the coding of a brief example: "How to code quantum circuits on FPGA".

This session is open to every FPGA enthusiast and person interested in learning from SYCL programming."

5.1.3 Speaker

Emmanuel Kieffer is an HPC (High-Performance Computing) Software Engineer in the Supercomputing Applications Team at LuxProvide, where he contributes to the development and optimisation of software for MeluXina, the EuroHPC supercomputer based in Bissen, Luxembourg.

Before joining LuxProvide, Emmanuel was an HPC Research Scientist at the University of Luxembourg, focusing on large-scale optimisation strategies leveraging high-performance computing. His expertise spans heterogeneous programming, code optimisation, and software acceleration, strongly emphasising improving computational performance and efficiency in HPC environments.

5.1.4 Relevance

Luxembourg's MeluXina stands out from other infrastructures in the EuroHPC consortium since it is the only supercomputer with a dedicated partition equipped with FPGAs. This type of accelerator can be used in scientific simulations, although its usage is more common in data-intensive tasks such as bioinformatics, computational finance, and high-energy physics. These areas have in common the need for real-time data acquisition and analysis, a task that can be optimised using FPGAs and that can be part, for example, of an AI application's workflow.

In the past, programming an FPGA required using specific programming environments, but nowadays, users can program one of these devices using C++ and SYCL.

This training activity exposes attendees to FPGA programming in the context of HPC, which is rare in HPC training programs. Since this session is primarily practical, including hands-on moments and live demos, it is helpful for people who have access to a system with an FPGA, such as MeluXina but have never tried to use these devices. The use cases of FPGA programming can interest researchers with projects coming from AI calls from both the scientific field and industry-oriented research, who can also apply for EuroHPC projects hosted at LuxProvide.

The speaker included an example of using an FPGA to emulate a quantum circuit in the event agenda, appealing to researchers from the increasingly popular field of quantum computing. Thus, this event can also be of interest to researchers who, in the future, will use the quantum computers inserted in the EuroHPC consortium.

5.1.5 Summary

The outline of the workshop was as follows:

- What is an FPGA, and what are its applications?
- Introduction to MeluXina's FPGAs.
- How to program an FPGA: HDL vs HLS.
- HLS and Kernel-based programming: from OpenCL to SYCL.
- Intel LLVM SYCL oneAPI DPC++ for FPGAs.
- Compiling and running SYCL programs.
- Quantum Circuit emulation with FPGAs.
- Q&A session.

This presentation started with a small introductory, more theoretical session about FPGAs. However, the speaker executed most of the presentation using coding examples that they compiled and ran in real time using the MeluXina account.

5.1.6 Participant Data and Feedback

Sixty-six people registered for this event, while 34 of those participated. The registration forms provided the following metrics.

- **Attendance:** 52%
- **Gender:** Male (86%), Female (14%).
- **Country of Affiliation:** 10 different countries.
- **Type of Organization:** Academia (93%), Company (7%)

None of the participants filled the post-event feedback form.

6. Conclusions and Future Work

Overall, EPICURE's training program's first year of activity can be considered successful. However, due to logistical and organisational issues that had to be resolved, training activities did not begin at the start of the project. Thus, the first training event occurred on July 10th, 2024, approximately five months after the project started. However, now that the training selection workflow has been defined and the training program's popularity is rising, the training program is expected to be improved during the project's second year, considering the experience acquired during the first year.

Considering that WP2 and WP3 just finished the first round of user support activity, the feedback of ASTs and users, which constitutes valuable information for developing the training plan, is still scarce. During the following year, as the ASTs' activity produces more reports, the Advanced Training plan will try to study the shortcomings identified during the user support activity and create training events that will address those same problems.

One of the EPICURE project's objectives is to create productive collaborations with NCCs and CoEs, especially in training. During the first year, there were opportunities to present EPICURE's Advanced Training plan to the Castiel 2 community, and the possibility of collaboration between these two EuroHPC initiatives was discussed. This possibility is expected to continue to develop in the following years.

The disparity between the number of people who register for a training event and the number of actual participants will be the main challenge for the Advanced Training WP for the following year. During the first year, the ratio of registrations and participation was always around 50%. This value of attendance can be explained as the combination of different factors: people who register forget they have done so; at the time of registration, the person had free time at the hour of the event, which was eventually occupied. To mitigate this problem, some strategies have already been suggested, such as creating a calendar event for each training session, sending extra reminders about the event one or two days before the scheduled day, shortening the registration period to a maximum of two weeks before the event and querying a sample of past participants to find the most suitable hour and day of the week when to schedule an event to maximise attendance.

Another difficulty identified during this year's training activities is that a minimal number, if any, of participants filled out the post-event feedback form. Considering that the audience's feedback is such an essential source of information that can help increase the quality of the training program and the organisation of future training events, this issue must be solved as quickly as possible. To increase the number of feedback forms received after an event, it is believed that offering a certificate of participation to participants who attend the training event and fill in the form can help. Also, offering EPICURE merchandise (mug, shirt, notebook) to a random person taken from the pool of people who submitted the form may encourage more people to provide feedback.

One objective for future in-person hackathons is to gather participants from various countries. This would increase the range of cultural and scientific ideas at the meeting, provide the participants with a better social and scientific experience, and encourage collaboration. To satisfy this requirement, registrations must open at least five to six months before the event so that people from other countries can plan their trip and accommodation and, in some cases, apply for the necessary visas.