


HORIZON-EUROHPC-JU-2021-COE-01



The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D5.2
First Updated Services, Training and Portal Progress
Report



The EXCELLERAT P2 project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101092621. The JU receives support from the European Union's Horizon Europe research and innovation programme and Germany, Italy, Slovenia, Spain, Sweden and France.

| | | |
|----------------------------|-----------------------------|------------------|
| Work Package: | 5 | Centre Operation |
| Author(s): | Claudio Arlandini | Cineca |
| | Francesca Gebbia | Cineca |
| | Tina Crnigoj Marc | ARCTUR |
| | Maksym Deliyergiyev | USTUTT |
| Approved by | Executive Centre Management | 05.12.2025 |
| Reviewer | Matthias Meinke | RWTH |
| Reviewer | Anne-Bernard Bedouet | SICOS |
| Dissemination Level | Public | |

| Date | Author | Comments | Version | Status |
|------------|------------------|----------------------------------|---------|--------------|
| 2025-05-14 | All | Initial draft | V0.1 | Draft |
| 2025-05-23 | All | Addressing Anne-Bernard comments | V0.2 | Draft |
| 2025-05-29 | All | Addressing Matthias comments | V0.3 | Second Draft |
| 2025-06-06 | All | Addressing second review | V0.4 | Final draft |
| 2025-06-13 | All | Addressing PMT review | V1.0 | Final |
| 2025-12-02 | Francesca Gebbia | revised version for resubmission | V2.0 | Final |

List of abbreviations

| | |
|------------|---|
| AI | Artificial Intelligence |
| API | Application Programming Interface |
| ARM | Advanced RISC Machines |
| CEEC | Committee of European Economic Co-operation |
| CFD | Computational Fluid Dynamics |
| CFL | Ported Courant–Friedrichs–Lewy |
| CoE | Centre of Excellence |
| CPU | Central Processing Unit |
| C2C1 | Use Case Identification Number |
| DoA | Description of Action |
| EoCoE | Energy oriented Centre of Excellence |
| ESP | EXCELLERAT Service Portal |
| EuroHPC JU | European High Performance Computing Joint Undertaking |
| FA | Further Application |
| F4E | Fusion for Energy |
| GPU | Graphics Processing Unit |
| HPC | High-Performance Computing |
| ITER | International Thermonuclear Experimental Reactor |
| KPI | Key Performance Indicator |
| MAR | Mediterranean and Related |
| MRF | Multiple Reference Frames |
| M | Month |
| MS | Milestone |
| NCC | National Competence Centre |
| Q | Quarter |
| RANS | Reynolds Averaged Navier-Stokes |
| TBD | To be defined |
| UL | University of Ljubljana |
| Y | Year |
| WEST | W (tungsten) Environment in Steady-state Tokamak |
| WHPC | Women in High Performance Computing |
| WP | Work Package |

Executive Summary

This deliverable provides an overview of the progress made in Work Package (WP) 5 of the EXCELLERAT P2 project, which is responsible for managing the service offering, training activities, resource acquisition supporting the technical work packages, and onboarding assessment and oversight.

The report includes an analysis of the usage and performance of the EXCELLERAT Service Portal, summarising user engagement, recent updates, and ongoing maintenance activities.

It also offers a detailed account of the training initiatives carried out under the EXCELLERAT P2 framework, accompanied by an evaluation based on KPIs analysis.

In addition, the deliverable examines and confirms the access to EuroHPC resources for each of the codes currently under development by the technical work packages.

Finally, the report addresses the task dedicated to further applications and external engagement. A notable achievement in this task is the successful onboarding of a new use case.

Table of Content

| | | |
|-----|--|----|
| 1 | Introduction | 8 |
| 2 | Service provisioning..... | 9 |
| 2.1 | Work Performed in T5.1 | 9 |
| 2.2 | Service requests..... | 13 |
| 2.3 | Future Outlook | 14 |
| 3 | Training and Education | 15 |
| 3.1 | Training events..... | 15 |
| 3.2 | KPIs Analysis..... | 17 |
| 3.3 | Outlook to the future | 19 |
| 4 | Resource Acquisition | 20 |
| 4.1 | Initial Activities and Collaboration with CASTIEL 2 | 20 |
| 4.2 | Monitoring and Oversight of Resource Access | 21 |
| 4.3 | Ongoing Activities and Future Planning..... | 22 |
| 5 | Further Applications..... | 23 |
| 5.1 | Tokamak Digital Twin | 24 |
| 5.2 | Axial Turbine Stages – OpenFOAM CPU integration & GPU porting..... | 25 |
| 5.3 | Outlook and Future Planning | 28 |
| 6 | Conclusion..... | 29 |

Table of Figures

| | |
|---|----|
| Figure 1: EXCELLERAT Service Portal entry page | 10 |
| Figure 2: Service Subpage, showing the Application Software list..... | 11 |
| Figure 3: Application Software - AVBP subpage with description..... | 12 |
| Figure 4: Mixing plane interpolation | 26 |
| Figure 5: SPUMA speed up | 27 |
| Figure 6: Visual comparison of the velocity field between OpenFOAM v2412 (left) and SPUMA (right)..... | 27 |
| Figure 7: Torque comparison between SPUMA (yellow line) and OpenFOAM v2412 (blue line)..... | 28 |

Table of Tables

| | |
|---|----|
| Table 1: Conducted trainings for Task 5.2..... | 16 |
| Table 2: KPIs for training events | 17 |
| Table 3: Computational resources obtained until Y2..... | 22 |

1 Introduction

Work Package 5 – Centre Operation is responsible for the orchestration and management of the project's service offering. Its primary role is to ensure that external users (academic, industrial, or institutional) can access, benefit from, and contribute to the Centre of Excellences (CoE's) HPC services, tools, and knowledge. WP5 acts as the operational interface between the technical development (WP2, 3 and 4), business validation (WP6), and the end-user community.

WP5 is structured into four main tasks, each addressing a key functional area:

- Task 5.1 – Service Provisioning: Oversees the continuous development and management of the EXCELLERAT P2 Service Portal, building on the legacy of EXCELLERAT P1. This includes integrating new services aligned with user needs, updating the portal interface, and refining the user experience based on stakeholder feedback.
- Task 5.2 – Training and Education: Designs, delivers, and evaluates training modules derived from the project's technical outputs and tailored to identified skill gaps in the European HPC ecosystem. It ensures the delivery of high-quality, accessible training aligned with the project's dissemination and outreach objectives.
- Task 5.3 – Resource Acquisition: Facilitates the planning and acquisition of computational resources required by the technical work packages. This task acts as the liaison between EXCELLERAT P2 and EuroHPC, assisting partners in navigating resource access pathways.
- Task 5.4 – Further Applications: Manages the onboarding of new use cases proposed by external stakeholders. It supports integration into the EXCELLERAT P2 ecosystem, ensuring alignment with project workflows and providing technical assistance as needed.

The following sections will examine each of the aforementioned tasks individually, outlining the main progress and updates since the previous deliverable, and providing an outlook on future activities.

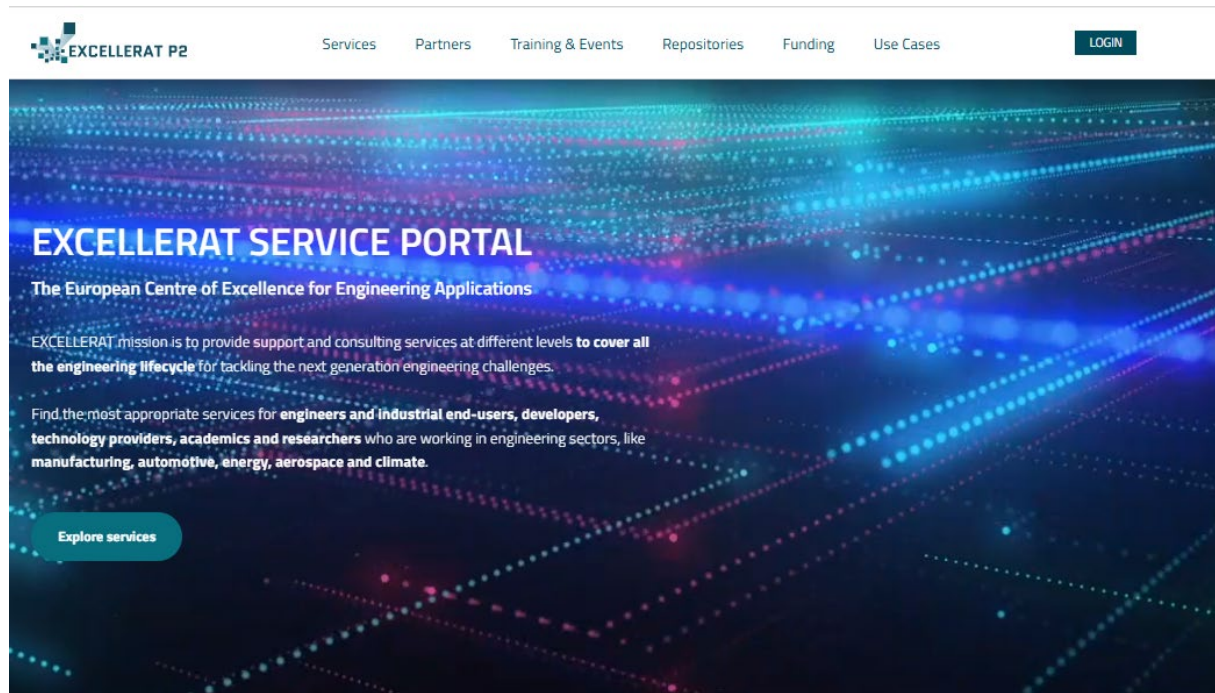
2 Service provisioning

2.1 Work Performed in T5.1

This section presents the work carried out within the framework of Task 5.1 – Service Provisioning, which includes all enhancements made to the Service Portal since its release in February 2024. This is followed by a section dedicated to the service requests received during the reporting period of this deliverable. Finally, the outlook on future developments of Task 5.1 is provided. In D5.1 “Services, Training and Portal Progress Report”, an introduction to the draft design of the new EXCELLERAT P2 Service Portal (ESP) was provided. In line with the Description of Action (DoA) and the WP5 plans, the ESP was developed and officially launched in early February 2024.

The ESP remains accessible at <https://services.excellerat.eu/> [3] and can also be reached via the EXCELLERAT P2 project homepage [1]. The portal is hosted and maintained by project partner Arctur.

The main menu of the new portal includes the following sections: Services, Partners, Training and Events, Repositories, Funding, Use Cases, and Login, as shown at the top of the Service Portal entry page in Figure 1.



Training and Events



19.05.2025 • Rimske Terme, Slovenia
Austrian-Slovenian HPC Meeting 2025 – ASHPC25

Events



20.05.2025 • Stuttgart, Germany
Automotive Testing Expo

Events



14.10.2025 • Barcelona, Spain
Advanced CFD and Turbulence Modelling targeting HPC

Excellerat trainings

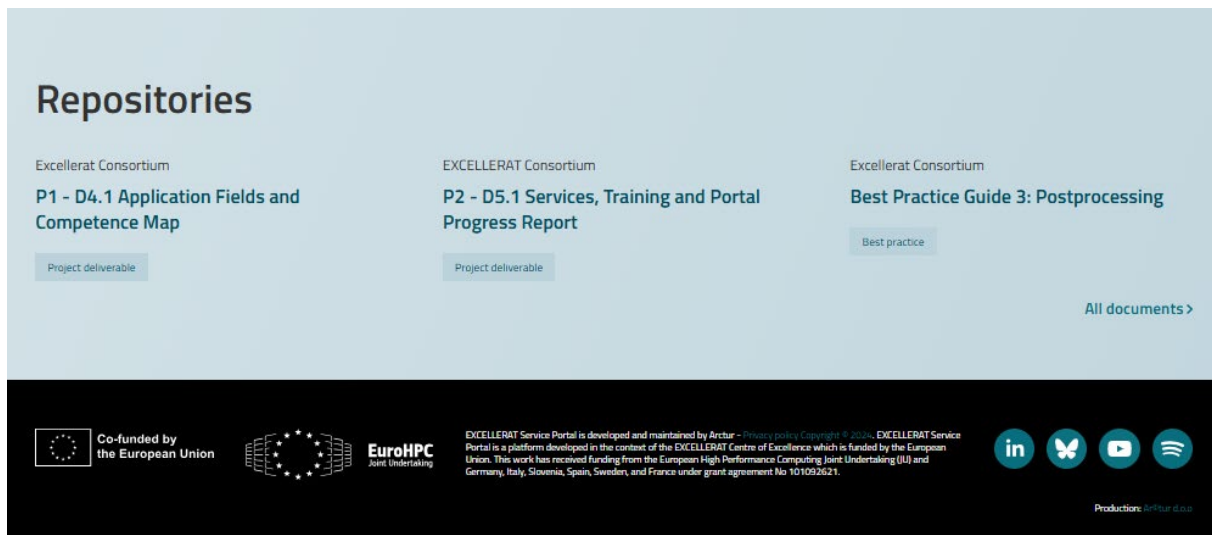


Figure 1: EXCELLERAT Service Portal entry page

The content provided on the ESP was further rewritten and adapted, to make it more user friendly, understandable and valuable by providing key information and details.

The services subpage (Figure 2) offers the following services by category:

- Application Software: support for AVBP, Alya, CODA, Neko, OpenFOAM, m-AIA, RaySect, Sod2D, L2G, STREAMS
- Tools: support for Data Management Platform, InSiDS, SimExplore, UQit, Vistle, Anubis & Marauder's Map
- Consulting: consulting on the following topics
 - Co-Design Service for Engineering Applications,
 - Data analytics in engineering,
 - Data management for large scale simulation result and input data,
 - Efficient and Modern Implementation of Exascale Ready Engineering Software,
 - Efficient Execution of Large Scale Engineering Simulation Workflows,
 - Holistic Testing and Validation for the Engineering Workflow,
 - Meshing and re-Meshing Techniques, Methodologies and Software,
 - Performance Engineering for the Complete Large Scale Engineering Workflow,
 - Strategies for Load-Balancing and Data-distribution,
 - Visualisation Methods and Tools for Large Scale Engineering Simulation Workflows.
- HPC Resources: support in using the HPC resources.
- Use-case Onboarding: possibility to submit proposals to address specific challenges.

Each category offers an overview of the services, accompanied by key details. Users can find more information about each service on its dedicated subpage (e.g. Figure 3). In the first year, four service requests were received, specifically for: Data Management Platform (tool), Anubis & Marauder's Map (tool), Co-Design Service for Engineering Applications (consulting), Use-case onboarding. A more detailed description of these requests is provided in Section 2.2. Project partners have internally discussed future actions to increase the number of requests, as outlined in the Future Outlook section.

The screenshot displays the 'Services' section of the EXCELLERAT CoE portal. At the top, a header reads 'Services' followed by a brief introduction: 'Tackling the next-generation engineering challenges could be easier by using EXCELLERAT services. EXCELLERAT CoE offers a variety of tailored services that could cover the whole engineering lifecycle. Find the appropriate service for your needs.' Below this is a navigation bar with five buttons: 'APPLICATION SOFTWARE', 'TOOL', 'CONSULTING', 'HPC RESOURCE', and 'USE CASE ONBOARDING'. The main content area is a grid of nine service cards, each featuring an icon, a title, a list of sectors, a list of capabilities, a short description, and a 'Learn more' button. The services listed are AVBP, Alya, CFD by ONERA, DLR and Airbus (CODA), FLEW, Neko, OpenFOAM, m-AIA, RaySect, and Sod2D.

| Service Name | Sectors | Capabilities | Description |
|-------------------------------------|--|--|--|
| AVBP | Automotive, Energy, Aerospace | Computational Fluid Dynamics, Turbulence modelling, Combustion | AVBP is a state-of-the-art large multi-physics solver developed with high performance computing in mind to tackled compressible reactive multi-phase flows. |
| Alya | Automotive, Energy, Aerospace | Computational Fluid Dynamics, Turbulence modelling, Aerodynamics | Alya is a parallel multi-physics / multi-scale simulation code, which has been developed to run efficiently in supercomputing environments, with optimal performance on heterogeneous architectures. |
| CFD by ONERA, DLR and Airbus (CODA) | Aerospace, (Wind) Energy | Computational Fluid Dynamics, Multi-Phase Flow, Adaptive Mesh Refinement | CODA is a Computational Fluid Dynamics software based on second-order finite-volume and higher-order Discontinuous-Galerkin discretizations on unstructured grids. |
| FLEW | Aerospace | Computational Fluid Dynamics, Visualization | FLEW, an in-house high-fidelity solver for Direct Numerical Simulation (DNS) of turbulent compressible flows over arbitrary shaped geometries. FLEW solves the Navier-Stokes equations in a generalized... |
| Neko | Automotive, Aerospace, Climate | Computational Fluid Dynamics | Neko is a portable framework for high-order spectral element flow simulations. Written in modern Fortran, Neko adopts an object-oriented approach, allowing multi-tier abstractions of the solver stack... |
| OpenFOAM | Magnetically Confined Nuclear Fusion, Automotive, Energy | Thermal Modelling, Visualization | A customized solver to solve the heat equation is using the OpenFOAM kernel, resulting in temperature distributions in solid bodies for user-defined heat fluxes. |
| m-AIA | All sectors | CFD, Multi-Physics Simulations, Multi-Phase Flow, Adaptive Mesh Refinement | m-AIA is a CFD related multi-physics simulation framework for aeroacoustics, aeroelastics or conjugate heat transfer, based on Cartesian meshes tailored for high-performance computing systems. |
| RaySect | Magnetically Confined Nuclear Fusion, Optics, Energy | Optical Simulation, Visualization | Ray-tracing is a package for the simulation of optical phenomena. |
| Sod2D | Aerospace, Automotive, Energy | Computational Fluid Dynamics | Sod2D is a new high-order code for turbulent scale resolving compressible and incompressible flow simulations. It has been designed from scratch to work optimally on GPU-based supercomputers. |

Figure 2: Service Subpage, showing the Application Software list

Further ESP improvements

After the release of the new portal, an additional round of user interviews was conducted. Six user interviews were held to gather feedback on the portal's usability. Users provided input on technical features, content, and design. The feedback was analysed, and suggestions for further improvements were identified.

Minor enhancements were suggested and subsequently implemented. These included:

- Improving text and background colour contrast to enhance readability
- Adding explanations for abbreviations such as FA and C2C1
- Including subtitles for figures
- Providing descriptions that clarify how each code differs from the others and what they may have in common
- Fixing broken or incorrect URLs

Additionally, users expressed a desire for more targeted content, such as detailed use cases and comprehensive guidelines for application software.

NEW Service Onboarded

This service offers an opportunity for potential users facing specific engineering challenges to collaborate with EXCELLERAT P2 partners in developing a suitable use case.

To apply, users must register on the EXCELLERAT P2 Portal, download the proposal template, complete it, and submit the filled-in proposal to the project's official email address. Once submitted, the proposal will be evaluated for feasibility by a team of experts. Users will be notified of the outcome via email within 10 to 30 working days.

In the second year, one use case onboarding was conducted, detailed information are provided in section 5 – Further Applications.

Services / Application software

AVBP

AVBP is a state-of-the-art large multi-physics solver developed with high performance computing in mind to tackle compressible reactive multi-phase flows.



About the software

AVBP focuses on compressible reactive and non-reactive flows. It is used for multiphase problems, turbulence modelling, plasma flows, wind turbines and turbomachinery. Designed with high-performance computing in mind, it is parallelized using MPI for distributed memory systems. At the node level, MPI3 RDMA communications allow for optimal performance combined with optimised vectorisation and cache blocking on CPUs. Accelerator offloading is possible via OpenACC pragmas with full solver support of NVIDIA systems for combustion simulation. The code uses classical multi-element grids (tetrahedras, hexahedras, prisms, pyramids, quads and triangles) with 3rd order in time and space numerical schemes.

Software details

AVBP is the state-of-the-art application for CFD of combustion in Europe and is used by both academia and industry. It is developed by around 40 people at CERFACS and many other researchers in Europe. The code has demonstrated excellent parallel performance in all current Intel, AMD and Arm CPU architectures as well as NVIDIA accelerators and more recently with prototypes on AMD GPUs.

Code license

Shared through a Limited Usage License agreement for scientific purposes. Commercial applications are to be discussed on a case-by-case basis.

Code URL

<https://www.cerfacs.fr/avbp7x/>

Programming language

Fortran, OpenACC

Availability

Private. Distributed on a case-by-case basis.




Figure 3: Application Software - AVBP subpage with description

To request services (sending inquiry) or download materials such as training materials, videos, use cases, users need to create a user account. To comply with the security measure, the registration is verified automatically through a two-step verification.

New Interest Group Subpage Development

A strategy for collaboration with the Interest Groups is currently being developed within WP6. To align with the needs of existing members and attract new ones, a dedicated subpage is being created. The Interest Group subpage will provide information on the benefits available to members through the support of EXCELLERAT P2 partners, as well as opportunities to promote their work or services via the portal, training sessions, and joint events. Member presentations will also be featured on the subpage. New members will be able to join by submitting a registration form.

Adding content to the WALDUR Platform

Waldur is a project co-developed by OpenNode and University of Tartu with contributions from different organisations. Waldur enables access and management of distributed HPC resources across Europe 30. It is a cutting-edge open-source multi-cloud marketplace designed to streamline the management of various cloud environments, including private clouds, public clouds, HPC centres, and IT infrastructure, all through a unified platform.

With Waldur, service providers can exercise resource governance and oversight, while offering end-customers a user-friendly self-service environment [5]].

At its core, the services from the National Competence Centres (NCCs) and CoEs are provided, meaning EXCELLERAT P2 CoE services could benefit from visibility through the additional platform, besides the EXCELLERAT P2 Portal. Moreover, it is part of the JU's strategy to establish a single-entry point for all NCCs and CoEs across Europe.

The aim of the “marketplace” solution to be set-up would be:

- an industrially-oriented tool for the NCCs to promote their services, products and expertise
- the marketplace to act as a single reference point for users to find expertise, services, codes and trainings
- to connect potential users of codes, training materials or customers for services
- to be based on the principles of yellow pages with domain-specific search capabilities and to evolve with relevant features.

EXCELLERAT P2 was the first CoE that joined the Portal. The migration process was aligned, and the migration of the services content was done through the Application Programming Interface (API). Furthermore, the final check of the services and branding were done on the Waldur platform Portal. These activities are still in early phase, meaning, we cannot provide any further statistics or details about the outcome. After setting up the whole services of the Waldur platform, additional promotional schemes shall be provided by Waldur, as well as the WP7 team, and further monitoring will be performed. Using the WALDUR platform does not exclude using the EXCELLERAT P2 Platform. However, partners need to carefully monitor and plan the future actions, and adapt activities accordingly.

2.2 Service requests

As previously mentioned, the services available on the Portal are organised into the following categories: support for tools and software, consulting topics, HPC resources and use-case onboarding.

Since the Portal's release, four requests have been received, specifically for the following services:

1. Data Management Platform (tool): this tool facilitates the access to HPC resources for industry, particularly for SMEs, and provides data transfer and HPC usage through a web front end.
2. Anubis & Marauder's Map (tool): this tool analyses HPC codebases using codemetrics to assess complexity, technical debt, and developer activity, highlighting challenges in legacy systems and providing tailored solutions for code structure and team interaction
3. Co-Design Service for Engineering Applications (consulting): it helps users prepare their engineering applications for future exascale systems by providing access to cutting-edge hardware, expert analysis of source code adaptation, and support for performance optimisation, benchmarking, and emerging technologies.
4. Use-case onboarding: this service provides an opportunity for potential users who are facing specific engineering challenges and would like to solve them with the help of EXCELLERAT partners to develop proper use case.

Requests for 1 and 2 primarily consisted of informational inquiries regarding the respective tools. The request related to item 4 lead to the successful onboarding of a new use-case: as the oversight and management of use-case onboarding fall under the responsibilities of Task 5.4, the result of this collaboration is described in detail in Section 5.2.

Regarding the consulting activity no. 3. the request involved seeking support to benchmark the NeoFOAM library (<https://github.com/exasim-project/NeoFOAM>) in order to perform hardware co-design on modern architectures such as ARM CPUs, NVIDIA GH200, and the MI300. NeoFOAM aims to modernise the OpenFOAM core components by providing full GPU support.

The outcome of this consulting service provided by EXCELLERAT partners lead to the following results:

- Improved vectorizability on Advanced RISC Machines (ARM) CPUs.
- Ported Courant–Friedrichs–Lewy (CFL) calculation to GPU.
- Performance on GH200 and MI300A shows room for improvement.
- Porting to MI300A omitting device data copies.
- NeoFOAM I/O to slash setup time for GPU-relevant case sizes.

2.3 Future Outlook

- The content on the ESP will be further updated, the portal will be maintained and further developed, according to the needs of the consortium.
- Terms and Conditions and Privacy policy for using Waldur portal are being developed.
- Activities on the Waldur platform will be monitored, and action plan needs to be established.
- Activities for Portal promotion will be defined, to raise the number of Portal visitors and services requests.

3 Training and Education

This section has the following structure: in sub-section 3.1 we list all the EXCELLERAT P2 training events organised in Quarter (Q) 4 2023 up to Q1 2025; in sub-section 3.2 we analyse the Key Performance Indicators (KPI) relative to the trainings; in the last sub-section we give an outlook on the future plans.

3.1 Training events

Building upon the structured approach established in D5.1, the Month (M)13-M30 report period includes a refined and expanded training program. This evolution aligns with EXCELLERAT P2's goal of optimizing training activities based on audience demand and feedback.

Many of the courses (Numerical Methods for Large Eddy Simulation, From Machine Learning to Deep Learning, and Scientific Visualisation with COVISE and Vistle) are continued editions, showcasing their sustained relevance and positive acceptance.

The Machine Learning course (2024-DL-HLRS) includes an improved curriculum incorporating BigWhoop for data compression and SimExplore for simulation data analysis.

The training structure in M13-M30 balances foundational courses with new emerging topics, such as Future Fusion Reactor Digital Twins, reinforcing EXCELLERAT P2's adaptability. Hands-on learning remains a priority, with continued emphasis on practical exercises (e.g., OpenFOAM programming).

Key developments in M13-M30 include the training efforts:

50. Expanded Collaboration and Interdisciplinary Training
 1. The Workshop on Future Fusion Reactor Digital Twins (September 2024) was delivered in collaboration with Plasma-PEPSC CoE, NCC Slovenia, and NCC Germany. Both EXCELLERAT P2 and Plasma-PEPSC contributed their expertise, ensuring a strong theoretical foundation and jointly developing interactive hands-on sessions to enhance practical understanding. This new training topic, addressing future challenges, introduces a forward-looking theme, integrating advanced simulation techniques in nuclear fusion, plasma physics, and High-Performance Computing (HPC). Plasma-PEPSC's role focused on demonstrating plasma physics codes that model behaviour in the scrape-off layer region and porting these codes to HPC infrastructure, scaling towards peta- and exa-scale computing. While, the EXCELLERAT P2 team showcased the development of ray-tracing algorithms for simulating light propagation, particularly electromagnetic waves. This is crucial for diagnostics and understanding plasma behaviour in the vacuum vessel of the future fusion reactors.
 2. The Scientific Visualisation with COVISE and Vistle training was conducted in collaboration with CASTIEL2 as part of the "Training Spring" initiative (Ref. D5.1), broadening outreach and ensuring alignment with European training strategies.
50. Expansion of numerical methods and Computational Fluid Dynamics (CFD) training: The NS-Large Eddy Simulation AVBP course was delivered twice (November 2023, November 2024), and the Numerical Methods for Large Eddy Simulation course continued in April 2024, reinforcing the focus on advanced CFD techniques.
51. Introduction of a major CFD and supercomputing school: The School on Computational Fluid Dynamics & Supercomputing (Third Edition) in July 2024 serves as a flagship event, emphasizing interdisciplinary applications in engineering.

52. Increased focus on Artificial Intelligence (AI) and data analytics: The recurring courses From Machine Learning to Deep Learning (March 2024) and Data Analytics for Engineering Data using Machine Learning (March 2025) highlight the growing importance of AI-driven approaches in HPC.
53. Hands-on Programming in OpenFOAM: The addition of Programming CFD in OpenFOAM (November 2024) meets the demand for open-source CFD training.

Table 1 provides a summary of the EXCELLERAT P2 training events that have taken place in Q2-Q3.

| Quarter | Date | Training event | Training category | Type | Number of participants |
|---------|------------------|--|-------------------|--------------------|------------------------|
| Q4-23 | Nov. 20-24 2023 | NS-Large Eddy Simulation AVBP https://cerfacs.fr/en/event/numerical-methods-for-large-eddy-simulation-using-avbp-2/ | CFD | on-site at CERFACS | 22 |
| Q1-24 | March 26-28 2024 | From Machine Learning to Deep Learning: A concise introduction [2024-DL-HLRS] https://www.hlrs.de/training/2024/dl-hlrs | DAT | online | 29 |
| Q2-24 | April 22-26 2024 | Numerical Methods for Large Eddy Simulation https://cerfacs.fr/en/event/numerical-methods-for-large-eddy-simulation-using-avbp-3/ | CFD | on-site at CERFACS | 17 |
| Q3-24 | July 14-19 2024 | "School on Computational Fluid Dynamics & SuperComputing(Third edition)" https://hpc-portal.eu/node/2116 | CFD | on-site | 34 |
| Q3-24 | Sep. 25-26 2024 | Workshop: Future Fusion Reactor Digital Twins [2024-FFR-DT] https://indico.ijs.si/event/2236/ | PAR, COM, VIS | hybrid at UL FS | 30 |
| Q4-24 | Oct. 8-9 2024 | Scientific Visualisation with COVISE and Vistle [2024-VIS1] https://www.hlrs.de/training/2024/vis | VIS | on-site at USTUTT | 13 |
| Q4-24 | Nov. 18-22 2024 | NS-Large Eddy Simulation AVBP https://cerfacs.fr/en/event/numerical-methods-for-large-eddy-simulation-using-avbp-4/ | CFD | on-site at CERFACS | 21 |
| Q4-24 | Nov. 25-27 2024 | Programming CFD in OpenFOAM https://www.events.cineca.it/en/hpc/programming-cfd-openfoam/bologna-20241125 | CFD, PAR | on-site at Cineca | 34 |
| Q1-25 | March 10-12 2025 | Data analytics for engineering data using machine learning [2025-ML4SIM] https://www.scai.fraunhofer.de/en/projects/EXCELLERAT/workshop.html | DAT | online | 41 |

Table 1: Conducted trainings for Task 5.2

3.2 KPIs Analysis

Concerning the KPIs collection, we start from those defined for EXCELLERAT P1 in the deliverable 5.1 “Initial Assessment of Training Needs and Services Building Plan”.

Table 2 gives an overview of the KPIs for training events in the 2023 and 2024 reporting period.

| KPI number | Description | Definition | Target | Score at 2023 | Score at 2024 |
|------------|--|---|--------------------------|------------------------|--|
| 5.2.1 | Number (total) of participants in trainings of EXCELLERAT P2 (per year). | $N_{\{total\ in\ year\}} = \sum_{i=1}^n N_{\{training, i\}}$ | 150/year | 22 + 26 + 18 + 14 = 80 | 29 + 22 + 34 + 30 + 13 + 34 + 21 = 183 |
| 5.2.2 | Balanced participation of males, females in EXCELLERAT P2 trainings. Gender ratio of the participants. | N_{female}/N_{total} or $R_{female_male_quotient}$ | 26% of female/non-binary | 27% | 15.56% |
| 5.2.3 | ACSI score: Overall training satisfaction as average per year (via survey/feedback sheet). | $N_{acsi_training_per_year}$ The scale is from 0 to 10. | 7.5 | 8.5 | 7.69 |
| 5.2.4 | The recommendation scale as average per year (via survey/feedback sheet). | $N_{acsi_recom_per_year}$ The scale is YES or NO. | 75% | 79% | 77.8% |

Table 2: KPIs for training events

All training KPIs were measured using an online survey. The evaluation was positive, but despite the advertising campaign carried out with the support of WP5, the number of course participants remained below the target value, 150 participants per year. By analysing the answers given by the participants, we will find relevant and interesting observations.

KPI 5.2.1 (“Number of participants”) bit the target value. In M1-12, participation fell short of the target, likely due to a limited number of training events (four courses) and their specific technical scope. In M13-30, participation more than doubled, surpassing the target, correlating with the expansion to seven training events, including the new School on Computational Fluid Dynamics & Supercomputing and the Workshop on Future Fusion Reactor Digital Twins. The increased number of courses, combined with diverse topics (CFD, AI, fusion reactor simulations, and OpenFOAM programming), likely attracted a broader audience.

KPI 5.2.2 (“Balanced participation of men”) Despite meeting the gender diversity target in M1-12, the percentage dropped significantly in M13-30 despite a higher total number of participants. Possible factors contributing to this decline include the predominance of CFD and HPC-focused training events, which historically attract a male-dominated audience. The introduction of new courses such as Future Fusion Reactor Digital Twins, Programming CFD in OpenFOAM and the CFD & Supercomputing School may have further skewed the gender ratio due to their technical nature.

KPI 5.2.3 (“Overall training satisfaction”) Despite an increase in training events, the slight drop in satisfaction in M13-30 with respect to in M1-12 suggests that some areas require improvement.

Possible reasons for the decline: Increased number of participants may have led to variability in individual experiences. More complex topics, such as Future Fusion Reactor Digital Twins, may have required additional theoretical background for participants. The introduction of hands-on programming (e.g., OpenFOAM training) may have been challenging for attendees with varying experience levels.

The post-event survey collected feedback from participants regarding the training courses, covering aspects such as content, organisation, hands-on exercises, theoretical depth, and logistics. Many answers indicating their willingness to attend future editions. The overall sentiment was highly positive, with notable appreciation for the expertise of trainers, the structure of the hands-on sessions, and the networking opportunities. However, some areas for improvement were highlighted, including a need for better pacing, additional theoretical coverage, and logistical refinements.

Importantly, the feedback results were shared with the instructors and training organisers to help guide improvements in future sessions. Key suggestions included adjusting the pacing, offering more theoretical background where needed, and refining logistical aspects of course delivery.

Positive aspects. Participants particularly valued:

50. Quality of training material and instruction: Trainers were consistently praised for their expertise, availability, and ability to explain complex concepts. Many appreciated the balance between theoretical and practical sessions.
51. Hands-on exercises: The variety and depth of hands-on tutorials were seen as beneficial in reinforcing learning. The “from scratch” exercises were particularly well received.
52. Networking opportunities: Participants highlighted the value of meeting experts, professors, and peers from diverse backgrounds, facilitating knowledge exchange.
53. Well-organised structure: The event was commended for its systematic approach, ensuring smooth transitions between lectures and practical sessions.
54. Logistical comfort: The training environment, including the venue setup, food, and availability of necessary computational resources, was appreciated.
55. Topic relevance: The course content was deemed cutting-edge and highly applicable to current research and industry challenges in CFD and HPC.

Areas for improvement. While the training was well received, participants suggested improvements in the following areas:

56. Pacing and structure: Some felt that certain days were too intense while others were more relaxed. A more balanced distribution of workload could enhance engagement.
57. More in-depth theoretical content: Several participants expressed a desire for deeper coverage of numerical methods, meshing, and turbulence models.
58. Better accessibility and logistics:
 1. Some found the entry process cumbersome due to security restrictions.
 2. Requests were made for digital rather than printed materials.
 3. More workstations could improve efficiency for hands-on exercises.
50. Inclusivity and representation: Multiple participants noted the low number of female attendees and suggested increasing efforts toward diversity and inclusivity in speaker selection and participant outreach.

51. More time for hands-on work: Some requested extended time for practical exercises to ensure deeper understanding.
52. Better flow of hands-on sessions: Some participants struggled to keep up with the hands-on exercises due to varying experience levels. More structured guidance and pacing adjustments were recommended.

Also, recommendations suggested extending the duration of the training to allow for more in-depth learning - extending practical sessions which should allow more time for hands-on exercises to reinforce learning; adjust session intensity to distribute workload more evenly; expand topics such as numerical methods, turbulence models, and equation-solving techniques; consider digitalizing materials, streamlining entry procedures, and increasing computing resources.

KPI 5.2.4 (“The recommendation scale”) The recommendation score remains strong and consistently above the target. The slight drop M13-30 reflects the trend seen in satisfaction scores, indicating that while the training remains well-received, refinements could enhance participant experience. Courses that received high engagement, such as Machine Learning, NS-Large Eddy Simulation, Programming CFD in OpenFOAM should continue, while new courses should integrate more participant-centred adjustments.

3.3 Outlook to the future

Looking forward, WP5 aims to address this imbalance through increased outreach and targeted engagement with diversity-focused HPC communities. For example, upcoming participation in the Women in High Performance Computing (WHPC) panel at ISC25 will include lightning talks highlighting the role of inclusive training initiatives and diverse participation in HPC. WP5 will also seek to strengthen collaboration with regional WHPC chapters, including the Mediterranean and Related (MAR) Chapter [9]] and the Central European Chapter [10]]. These communities actively promote gender diversity and inclusion through mentorship, education, and networking, offering valuable platforms to support the participation of women and underrepresented groups in future training activities.

Such partnerships will help ensure more balanced representation across future events and contribute to broader efforts to foster an inclusive HPC ecosystem.

To maintain alignment with our KPIs, we need to deliver approximately 7–8 training courses per year. The current schedule includes continued editions of well-established courses from previous years, such as:

- March 21 – April 4, 2024: NS-Large Eddy Simulation AVBP
- October 23 – October 24, 2024: Scientific Visualisation with COVISE and Vistle
- November 17 – November 21, 2024: NS-Large Eddy Simulation AVBP
- TBD: Data Analytics for Engineering Data using Machine Learning
- TBD: Programming CFD in OpenFOAM

In addition to these recurring courses, expanding the scope of training topics—including CFD, AI, fusion reactor simulations, and OpenFOAM programming—will help attract a broader audience. A key strategy to enhance topic diversity involves collaboration with other CoEs that work with similar computational tools. Potential partnerships include:

- CEEC and the Energy Oriented Centre of Excellence (EoCoE), both of which rely on the Alya code in their research and applications.
- HiDALGO2, which extensively uses OpenFOAM for computational fluid dynamics.

Unfortunately, the joint training event with Plasma-PEPSC CoE will not be repeated in Year (Y) 3 due to scheduling conflicts for both CoEs in the upcoming fall. However, we remain open to future collaborations when feasible.

To further enrich our training portfolio, we invite all EXCELLERAT P2 partners to reach out to WP5 if they are interested in organizing training courses in Y3. Collaborative efforts will be key in maintaining high-quality training while meeting community needs and KPI goals.

4 Resource Acquisition

As defined in the Grant Agreement, Task 5.3 “Resource Acquisition” is dedicated to facilitating the internal provisioning of HPC services and computational resources required by the technical Work Packages within EXCELLERAT P2. More broadly, this task serves as a key interface between the Centre of Excellence and the wider European HPC infrastructure, with a particular emphasis on the EuroHPC pre-exascale systems.

The main objectives of Task 5.3 include:

- Maintaining an up-to-date overview of the available HPC services relevant to the Consortium Partners from CoE EXCELLERAT P2.
- Offering a centralised access point for Consortium Partners from CoE EXCELLERAT P2 to request and obtain specific HPC resources.
- Providing technical support and guidance, including relevant documentation, to ensure effective utilisation of resources.
- Coordinating with HPC centres to ensure that requests from consortium members are efficiently handled and matched with the appropriate service delivery mechanisms.

The following paragraphs outline the key responsibilities, strategic actions, and outcomes of Task 5.3, including its collaboration with CASTIEL 2, its response to evolving access policies, and the ongoing effort to monitor and support technical partners in securing HPC resources. Particular attention is given to the operational measures introduced to ensure efficient and equitable access to EuroHPC systems, as well as the alignment of EXCELLERAT P2’s objectives with the broader goals of exascale computing in Europe. In general, no major complications have been encountered by the technical work packages in accessing EuroHPC resources

4.1 Initial Activities and Collaboration with CASTIEL 2

In the initial phase of the project, Task 5.3 focused on mapping the HPC resource requirements of technical partners to support the development and execution of their respective use cases. This effort aimed to provide a clear understanding of the computational needs across the full duration of EXCELLERAT P2.

During the first year, the project partners were also awaiting the implementation of a dedicated access scheme to be provided by the EuroHPC JU. In parallel, Task 5.3 initiated collaboration with CASTIEL 2, the Coordination and Support Action responsible for gathering HPC resource requirements across all CoEs. CASTIEL 2 proposed a differentiated approach to streamline access to JU systems, consisting of:

- Flat-rate or base-rate access: a predefined allocation of resources per code, enabling rapid and straightforward benchmarking and development access across all relevant systems.
- Individualised but accelerated access: A more tailored approach for regular and large-scale computational needs, with faster processing of access requests.

The detailed estimates of computational needs, collected code by code from the technical partners, were shared with CASTIEL 2 and have also been documented in WP5 previous deliverable D5.1 “Services, Training and Portal Progress Report”. Interested readers are encouraged to refer to those reports for further information.

4.2 Monitoring and Oversight of Resource Access

At the beginning of year 2, EXCELLERAT P2 received an official communication indicating that CoEs are not eligible to directly apply for Strategic Access to EuroHPC systems. This information was promptly disseminated to all technical partners. In response, Task 5.3 reiterated its commitment to supporting partners in applying to EuroHPC open calls and navigating the available access pathways.

Fortunately, the partners involved in EXCELLERAT P2 are experienced users of JU resources. Nonetheless, in the absence of a dedicated access scheme, Task 5.3 strongly recommended careful planning to ensure timely and sufficient access to HPC resources for upcoming large-scale simulation campaigns.

In addition to its coordination role, Task 5.3 has taken on a new responsibility: monitoring the actual usage and access to HPC resources by the technical EXCELLERAT P2 partners. This is being carried out through regular surveys, allowing the task team to supervise and support proper planning and timely access.

The results of the surveys conducted up to and including year 2 are summarised in Table 3.

| Partner | Use-case | Access scheme | Response date | Computational time obtained (node-hours) | Cluster | Computational time (node-hours) consumed by Y2 |
|---------|----------|-------------------------------|----------------|--|----------------|--|
| DLR | UC-1 | Internal resources | N.A. | N.A. | CARA | 133000 |
| DLR | UC-1 | Internal resources | N.A. | N.A. | CARO | 28000 |
| CERFACS | UC-2 | Customer provision | N.A. | Customer Limit | Topaze | 322000 |
| CERFACS | UC-2 | Inner provision | N.A. | Workload dependent | Kraken | 404000 |
| CERFACS | UC-2 | EuroHPC Benchmark access | N.A. | 10860 | Leonardo | N.A. |
| CERFACS | UC-2 | Discretionary allocation | N.A. | 819 | ADASTRA | N.A. |
| CERFACS | UC-2 | N.A. | N.A. | N.A. | LUMI-G | All |
| CERFACS | UC-2 | N.A. | N.A. | N.A. | AWS Graviton 3 | All |
| RWTH | UC-3 | EuroHPC Benchmark Access Call | September 2023 | 5000 (CPU) | Vega CPU | Significant share used for benchmarking |
| RWTH | UC-3 | EuroHPC Benchmark Access Call | October 2023 | 5000 (CPU) | MeluXina CPU | Significant share used for benchmarking |

| | | | | | | |
|--------------|------|---------------------------------------|-------------------------|-----------------------------|------------------------------------|--|
| RWTH | UC-3 | EuroHPC Bechmark Access Call | November 2023 | 7000 (CPU) | Discoverer | Significant share used for benchmarking |
| RWTH | UC-3 | EuroHPC Bechmark Access Call | May 2024 | 7000 (CPU) | Karolina CPU | Significant share used for benchmarking |
| RWTH | UC-3 | Internal re- sources | N.A. | N.A. | Hawk | N.A. |
| BSC | UC-4 | EHPC-BEN- 2023B08- 001 | August 2023 | 1000 (GPU) | Karolina | All |
| BSC | UC-4 | EHPC-BEN- 2024B02- 007 | February 2024 | 3500 (GPU) | Leonardo | All |
| BSC | UC-4 | Internal re- sources | N.A. | N.A. | Mare Nos- trum 5 | N.A. |
| KTH | UC-5 | EuroHPC Development Access Call | June 2023 | 3500 (GPU) + 10000 (CPU) | Leonardo- booster and LUMI-G | All |
| KTH | UC-5 | NAISS LUMI Swe- den | December 2024 | 230000 (CPU) | LUMI-C | Just acquired |
| Cineca | UC-6 | Inner provi- sion | N.A. | N.A. | Leonardo- booster | 15000 |
| Cineca | UC-6 | Inner provi- sion | N.A. | N.A. | Leonardo- DCGP | 1000 |
| Sapienza UNI | UC-6 | ISCRA B | | 60000 | Leonardo- booster | 30000 |
| SapienzaUNI | UC-6 | EuroHPC JU | March 2024 | 50000 | MeluXina | 35000 |
| Sapienza UNI | UC-6 | EuroHPC JU | Waiting for response | N.A. | LUMI-G | N.A. |
| UL | FA-1 | SLING | November 2023 | 5000 (CPU) | Vega CPU | Significant share used for benchmarking |
| UL | FA-1 | Internal re- sources | N.A. | N.A. | HPC FS | 1250 (CPU) |

Table 3: Computational resources obtained until Y2

From the data presented, it is evident that all technical partners have successfully accessed the necessary EuroHPC resources. All codes involved in EXCELLERAT P2 have been able to carry out development and/or production runs on these platforms.

Given that one of the core missions of EXCELLERAT P2 is to enable and support the transition toward exascale computing, it is particularly significant that five codes within the project framework have already been executed on EuroHPC's pre-exascale systems. These codes include: AVPB (UC-2), SOD2D (UC-4), Neko (UC-5), STREAMS (UC-6); which have run on Leonardo, LUMI and Mare Nostrum 5.

4.3 Ongoing Activities and Future Planning

An updated overview of the computational resources obtained and consumed by each technical partner will be required by mid-year 3 (presumably end of June 2025) and by the end of Y3. To support continued access and usage planning, Task 5.3 has also launched a follow-up survey aimed at gathering updated forecasts from technical partners regarding their future

computational needs. This activity is designed to promote strategic resource planning and to increase the number of codes that are actively run on pre-exascale EuroHPC systems. The ultimate goal remains to ensure that EXCELLERAT P2 partners are well-positioned to leverage Europe's most advanced HPC platforms and contribute meaningfully to the development of exascale-ready applications.

5 Further Applications

This report outlines the current progress of Task 5.4, titled “Further Applications,” within the EXCELLERAT P2 project. The core responsibility of this task is to manage the integration of new users and partners into the EXCELLERAT P2 ecosystem. This involves guiding them through the onboarding process and ensuring their alignment with EXCELLERAT P2's established workflows, particularly in areas such as requirements analysis, application development, and task coordination.

The onboarding workflow and associated procedures have been thoroughly detailed in the previous deliverable D5.1 “Services, Training and Portal Progress”. As such, this information will not be reiterated in the present document. Readers who require a comprehensive overview of the onboarding process are encouraged to consult the aforementioned deliverable.

A key use case currently under observation within Task 5.4 is the development of a digital twin for the first wall of the ITER tokamak. This effort aims to optimise this critical component of the fusion reactor. ITER (International Thermonuclear Experimental Reactor) is a major international initiative focused on demonstrating the feasibility of nuclear fusion as a sustainable energy source. The tokamak itself is a doughnut-shaped device designed to contain high-temperature plasma using powerful magnetic fields.

The main achievement for this task during the second year of the project was the successful onboarding of a new use case. This use case, submitted by De Pretto Industrie, focuses on Axial Turbine Stages – OpenFOAM CPU integration and GPU porting. Following the onboarding procedure structured by Task 5.4, it was accepted for development by Cineca. The expected duration is six months, with completion planned during Year 3. A full description of the use case and its current progress is provided in the following sections.

In addition to onboarding and monitoring activities, Task 5.4 collaborates closely with Task 7.2 to facilitate the transfer of technical expertise related to software tools and codes that are not yet part of EXCELLERAT P2's official offering. This ensures that the broader technical work packages can benefit from these new capabilities. As part of community-building efforts, collaboration between WP6 and WP7 has led to increased visibility of EuroHPC resources, particularly with the inclusion of LUMI in the EXCELLERAT P2 Service Portal.

Lastly, Task 5.4 continues to oversee the evolution of major use cases by maintaining regular contact with their respective owners. This ongoing communication allows the team to address emerging challenges proactively and ensure the smooth advancement of each application.

In the next two sections we will focus on a short update regarding the Tokamak Digital Twin and then a detailed description of the new onboard use-case “Axial Turbine Stages – OpenFOAM CPU integration and GPU porting”. After that, we will provide the outlook for this task.

5.1 Tokamak Digital Twin

The University of Ljubljana (UL) aims to develop a digital twin of the first wall of a tokamak fusion reactor, contributing to improved design, operation, and real-time control of fusion energy systems. This use case supports the broader goal of making fusion a sustainable energy source, as promoted by Fusion for Energy (F4E).

The digital twin will simulate thermal and optical behaviours using advanced modelling techniques such as field line tracing, Finite Element and Finite Volume Methods (via OpenFOAM), and ray-tracing algorithms (via Raysect). These simulations will help interpret camera data, accurately determine surface temperatures, and predict hotspots or stray light, supporting both safety and performance optimisation.

Experimental validation will rely on data from ITER and WEST tokamaks. Additionally, the use case explores AI and hardware accelerators (e.g., FPGAs) for solving inverse problems in real-time control scenarios.

UL started with benchmark simulations on Slovenia's VEGA supercomputer, eventually scaling up to exascale systems. Although not part of EXCELLERAT Phase 1, UL joined Phase 2 through structured onboarding and integration processes led by Task 5.4, contributing its expertise to advance state-of-the-art digital twin technology in fusion research.

A full analysis of the progress of the FA-1 use case is available in Deliverable D2.17 (Public dissemination level). For the sake of completeness, we provide a brief overview of the main achievements.

As an initial milestone (MS12), UL presented a small-scale simulation focused on a block of a temporary limiter inside the ITER reactor. This early demonstration successfully validated the full simulation workflow, using L2G for field line tracing and Raysect for optical simulations – both initially parallelised only within a single node.

Since then, significant progress has been made. The simulation now incorporates the magnetic ripple effect – variations in the magnetic field due to coil misalignment or current changes – even though the magnetic field is theoretically uniform in the toroidal direction. The geometric complexity has also increased, one of ITER's most heavily loaded components, which includes around 6,300 castellated cubes and a mesh of approximately 9 million triangles. The thermal model has been refined to account for multiple material properties (e.g., tungsten, CuCrZr, stainless steel) and active cooling, enabling more accurate simulations of heat dissipation.

In terms of optical modelling, surface radiation has been simulated using the Stefan-Boltzmann law. To address the challenge of integrating different input/output formats across the three main codes used in the workflow, the UL team has also developed a lightweight Python library. This tool simplifies case setup and job submission on HPC systems, making the workflow more accessible.

The next steps include enhancing scalability and completing detailed simulations for WEST and ITER reactors to produce synthetic diagnostics. These simulations, incorporating complex geometries and materials, will generate virtual infra-red camera signals for model validation and experimental comparison. Post-processing will play a central role, involving the transformation of simulation outputs into inputs for subsequent stages, the generation of synthetic imaging data, and optimisation of data handling. In addition, the use case foresees the integration of AI-driven approaches to address inverse problems – specifically, the reconstruction of surface temperature distributions from real camera images in highly reflective environments. Continued optimisation across all simulation components (field-line tracing, thermal modelling, and optical ray-tracing) will prepare the workflow for efficient execution on future exascale HPC systems.

5.2 Axial Turbine Stages – OpenFOAM CPU integration & GPU porting

This use case focuses on **compressible turbomachinery OpenFOAM** simulations, executed on both **CPUs and GPUs**, unlocking the full potential of **modern heterogeneous HPC** systems. Steam turbines simulations are characterised by several complexities, such as stator-rotor interaction, compressible high-speed flows, intricate turbulent flow field and steam properties evaluation.

The main motivations that characterise this use case are summarised below:

- OpenFOAM is widely regarded as the most used open-source CFD software.
- Specific implementations tailored to turbomachinery are often available only as third-party separate contributions.
- Experienced developers are needed to integrate these contributions and add new implementations.
- Cineca is currently putting a consistent effort into OpenFOAM GPU porting activities.

Work Package 1 (95% completed) - CPU execution

The first work package has been characterised by the CPU simulations of a real industrial geometry proposed by De Pretto Industrie. OpenFOAM currently offers most of the numerical capabilities required to simulate these machines, including compressible solvers, sliding mesh, Multiple Rotating Frames (MRF), and interfaces with interpolation schemes. Still, the following missing/incomplete features were identified:

- **Mixing plane interfaces** (partial implementation available as third party).
- **Partially overlapping interfaces** (partial implementation available as third party).
- **Domain specific boundary conditions** (missing)
- **Modelling steam thermophysical properties** (few attempts in scientific literature).

These missing features were integrated or implemented into the code and tested on the final test case to verify accuracy and stability when simulating compressible industrial geometries. In particular, the mixing plane interface was integrated from the work done by OpenFOAM Turbomachinery Special Interest Group [6]. The interpolation of the temperature using the typical circumferential average at the stator rotor interface can be observed in Figure 4.

Furthermore, **partial overlapping interfaces** were integrated from the ICSFoam library [7]. A new implementation of a velocity boundary condition to be used in conjunction with a user-defined value of total pressure has been proposed and named **turboPressureInletOutletVelocity**. Finally, the adoption of IAPWS97 steam thermophysical properties is currently being tested.




Figure 4: Mixing plane interpolation

The test case has been executed on both Galileo100 (G100) and Leonardo DCGP HPC systems. The simulation is a steady RANS (Reynolds Averaged Navier Stokes) analysis with a wall time that ranges from 5 to 8 hours on 2 nodes of G100, where 1 G100 node = 2 x CPU Intel CascadeLake 8260, with 24 cores each). The simulations converge smoothly at the testing conditions provided by the industrial end-user.

Work Package 2 (25% completed) - GPU porting

Recently, Cineca has undertaken efforts to adapt OpenFOAM to execute on Graphics Processing Units (GPUs), which resulted in the development of the SPUMA framework. This adaptation leverages the enhanced memory throughput and parallel computing capabilities of GPUs to enable more demanding and high-fidelity simulations with reduced hardware requirements compared to conventional CPU implementation. Notably, GPU tendency to work efficiently at higher cell count per GPU can decrease the domain decomposition burden in large-scale unsteady turbomachinery analyses, allowing cases that previously required hundreds of CPU cores to be executed on a single GPU with non-negligible speedup. This shift allows for more accessible and economically viable CFD workflows for industrial users.

SPUMA has been recently tested on the well-known DrivAer test case consisting of 236 M cells [8]], resulting in an equivalence between GPU and CPU cores of 1 to 220/300. This is shown in the following graph, where several analyses run on Leonardo DCGP (200,400,800,1600 cores) are compared to two simulations run on Leonardo Booster (8, 16 GPUs).




Figure 5: SPUMA speed up

The main goal of work package 2 is working towards the GPU simulation of a compressible axial turbine stage. To do so, three missing SPUMA features were identified:

- **MRF**
Implements fluid equations for multiple reference frames with the velocity defined in the inertial reference frame.
- **Thermophysical library**
Provides fluid thermophysical properties during the simulation.
- **Cyclic AMI interfaces**
Allow for interpolation of quantities across interfaces and non-conformal periodic boundaries.

The GPU porting of these features is ongoing. Currently, the MRF activities have been completed, and Cineca is testing the implementation on a simplified axial turbine test case provided by the OpenFOAM Turbomachinery Special Interest Group.

The validation of the results of the GPU porting are represented in Figures 6 and 7, where a visual representation of the velocity fields and a graph containing the Torque trends are shown respectively.




Figure 6: Visual comparison of the velocity field between OpenFOAM v2412 (left) and SPUMA (right)




Figure 7: Torque comparison between SPUMA (yellow line) and OpenFOAM v2412 (blue line)

Currently, the porting of AMI interfaces is being performed by the Cineca team. Lastly, the thermophysical properties library will be adapted to run on accelerators, followed by complete performance benchmarks.

5.3 Outlook and Future Planning

This task will remain active in assessing and overseeing the onboarding activities. It will continue to monitor the progress and successful completion of the use cases that have been onboarded so far, ensuring alignment with project objectives and identifying any areas requiring further support or improvement.

6 Conclusion

During the reporting period, WP5 has delivered steady progress across all its core responsibilities.

The EXCELLERAT Service Portal has been regularly updated and maintained, with new content and landing pages introduced to enhance usability and visibility of the services offered. Since its release, four requests have been received: two for questions regarding specific tools (Data Management Platform and Anubis & Marauder's map), one for the consulting service "Co-Design Service for Engineering Applications", and one request for use-case onboarding.

In the area of training and education, WP5 has achieved overall positive results with respect to the defined KPIs, that assess participation, satisfaction, and inclusivity (the latter is the only one that did not reach the target value). A diverse range of training activities was successfully delivered, covering topics relevant to the project's target communities. Participation levels, satisfaction rates, and feedback indicate that the training offerings are aligned with community needs.

Regarding resource acquisition, no major issues were encountered in supporting the technical work packages. All codes currently under development have successfully accessed EuroHPC resources, ensuring continuity in simulation campaigns and technical developments.

Finally, WP5 has made tangible progress in onboarding new external applications. As mentioned, one new use case has been successfully integrated into the project framework, demonstrating the functionality and readiness of the onboarding mechanism.

Continued efforts will focus on strengthening service visibility, expanding training opportunities, and fostering external collaboration through strategic onboarding and resource support.

References

- [1] EXCELLERAT project, <https://www.excellerat.eu>
- [2] EXCELLERAT P2 Grant Agreement project 101092621
- [3] <https://services.excellerat.eu/viewcode/7>
- [4] <https://docs.waldur.com/latest/about/governance/>
- [5] <https://interoperable-europe.ec.europa.eu/collection/open-source-observatory-osor/waldur>
- [6] <https://sourceforge.net/p/turbowg/wiki/Home/>
- [7] <https://github.com/stefanoOliani/ICSFoam>
- [8] https://develop.openfoam.com/committees/hpc/-/tree/develop/incompressible/simpl-Foam/occDrivAerStaticMesh?ref_type=heads
- [9] https://www.bsc.es/mar_whpc
- [10] <https://www.linkedin.com/company/central-european-women-in-hpc/>