H2020-INFRAEDI-2018-2020

The European Centre of Excellence for Engineering Applications

Project Number: 823691

D5.7
Services, Training and Portal Report

The EXCELLERAT project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823691

Workpackage:	WP5	Centre implementation
Author(s):	Claudio Arlandini	CINECA
	Lorenzo Zanon	USTUTT
	Francesco Salvadore	CINECA
	Antonio Memmolo	CINECA
Approved by	Executive Centre Management	
Reviewer	Thomas Gehrold	DLR
Reviewer	Niclas Jansson	KTH
Dissemination Level	Public	

Date	Author	Comments	Version	Status
2022-04-29	C. Arlandini	Initial draft	V0.1	Draft
2022-05-05	All	Revised draft	V0.2	Draft
2022-05-18	All	Fully revised draft after first review	V0.5	Draft
2022-05-25	C. Arlandini	Final	V1.0	Final

List of abbreviations

ACSI	*American Customer Satisfaction Index*
CFD	*Computational Fluid Dynamics*
CoE	*Centre of Excellence*
CSA	*Coordination and Support Action*
CUDA	*Compute Unified Device Architecture*
DIMA	Department of Mechanical and Aerospace Engineering
ERPP	*Exact Regularized Point Particle*
ESP	*EXCELLERAT Service Portal*
FLUMACS	*FLUid dynamics of MAcro and micro Complex Systems*
HPC	*High Performance Computing*
KPI	*Key Performance Indicator*
PMO	*Project Management Office*
PRACE	*Partnership for Advanced Computing in Europe*
SZE	*Széchenyi István Egyetem University of Györ*
TWG	*Thematic Working Group*
WP	*Work Package*

Executive Summary

The objective of WP5 “Centre Implementation" is to ensure the management of the services offered by EXCELLERAT and, in particular, to support the implementation of the services for external users whose requirements have been defined in the technical work packages (WP2, 3 and 4) and validated in relation to the business model of WP6. WP5 also manages services, users’ access, and related issues, both for internal services and for services towards the users.

This includes:

- Services for end-users as designed in WPs 2,3,4 with a particular focus on the industrial users,
- Training and education services,
- Centre of Excellence (CoE) internal management services, and
- Support for new application codes.

Another key aspect is the management of the provisioning of HPC resources (CPU hours, storage, etc.), including the integration of infrastructure services, for EXCELLERAT activities, to be coordinated with the provision services of the HPC centres.

The main output of this work package is the provision of a single-entry point to its ecosystem of services, thus, building up an online access point – hereinafter referred to as the EXCELLERAT Service Portal (ESP) distinct but integrated with the project website, hereinafter referred to as EXCELLERAT Website (EW) – accessible at:

https://services.excellerat.eu

This deliverable “*Services, Training and Portal Report*” presents the outcomes of the final year of WP5 activities for all its five tasks and in particular for the implementation of the above-mentioned portal from a user point of view, while its companion D5.6 *"Portal Release"* describes the portal from a technical point of view.

The determination of the final KPIs as defined in D5.1 “*Initial Assessment of Training Needs and Services Building Plan*” to support a successful implementation is also presented.

Table of Contents

Table of Figures

Table of Tables

1 Introduction

This deliverable "*Services, Training and Portal Report*" presents the outcomes of the final year of WP5 activities for all its five tasks and in particular for the implementation of the above-mentioned portal from a user point of view, while its companion D5.6 *"Portal Release"* describes the portal from a technical point of view.

After an Introduction section, Section 2 "*Service portal Content Review*" presents the main areas of the final version of the service portal, acting as a user guide. This is the result of work on Tasks 5.1 and 5.2.

Section 3 "*Further applications*" describes Task 5.3 activities related to the interaction with the Interest Groups and onboarding of new applications in the project, describing the two completed (specifically, Annex 1 presents the report for the PAinG-Flow code activity) and all other collaborations started.

Section 4 "*Training*" describes Task 5.4 activities related to the analysis of skill gaps related to the project topics and the training events organized by the consortium to address these gaps, supplemented by material in Annex 2, while Section 5 "*HPC Provisioning*" accounts for Task 5.5 effort in providing the HPC resources on which the CoE activities are based, with further data in Annex 3.

The document is concluded by Section 6, where we draw some conclusions.

2 Service Portal Content Review

We provide here a brief overview of the content currently available on the EXCELLERAT Service Portal (ESP) [1]. Instead, we refer to deliverable D5.6 *"Portal Release"* for technology updates related to the portal platform. We remind you that the main portal content available to users, from the content management system perspective, is summarized in Figure 1.

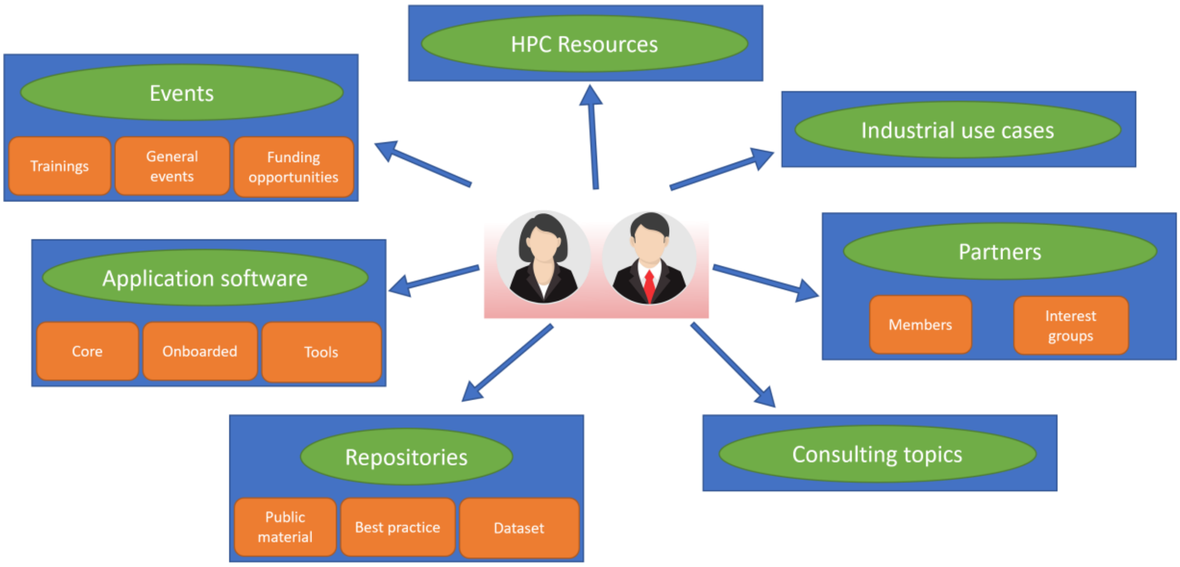

Figure 1: Hierarchy of contents managed by the ESP portal and accessible for external users. The content types are represented in green, while the content subtypes are represented in orange (where present).

The different content types basically correspond to different database tables that coordinate the saving of the different metadata associated with each content. For the Events and Repositories, in addition to the metadata, there is the possibility of storing material through the uploader integrated in the portal. In terms of presentation, each type of content typically has:

- a page for listing and searching for content based on keywords
- a specific page for each content
- a page for creating and updating content, accessible with administrator rights.

These pages are created ad-hoc depending on the specifics of each content.

2.1 HPC Resources

The portal currently presents 9 HPC clusters belonging to the CoE partners. An excerpt of the portal page presenting the list is provided in Figure 2, left. For each portal, basic information is provided, including description, web-url, an image and a list of related tags (Figure 2, right).

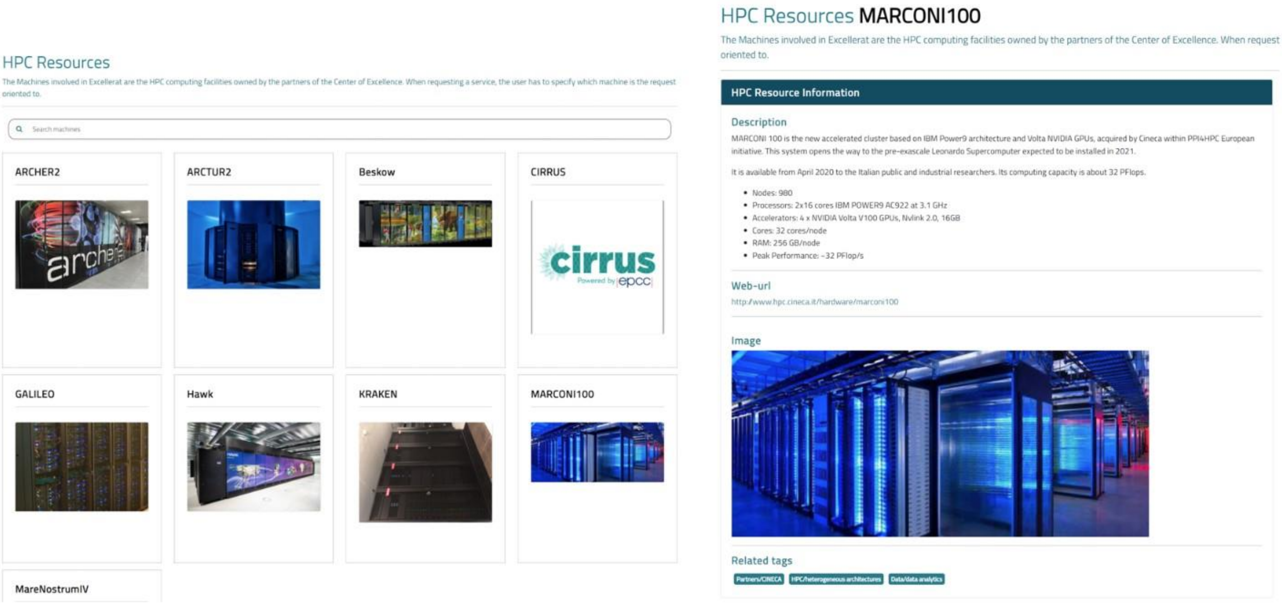

Figure 2: EXCELLERAT Service Portal: HPC resources search page (left) and information on the Cineca Marconi100 cluster (right).

2.2 Partners & Interest groups

The Partners & Interest group section contains the members involved in the CoE, in the two different possibilities of project partner or interest group. The partner page features 13 members (Figure 3, left) along with a map of their geo-located locations. It is possible to specify multiple locations for the same partner to support multi-location institutions. For each partner, you can access the information page containing a description, an image, addresses, web-url and related tags (Figure 3, right).

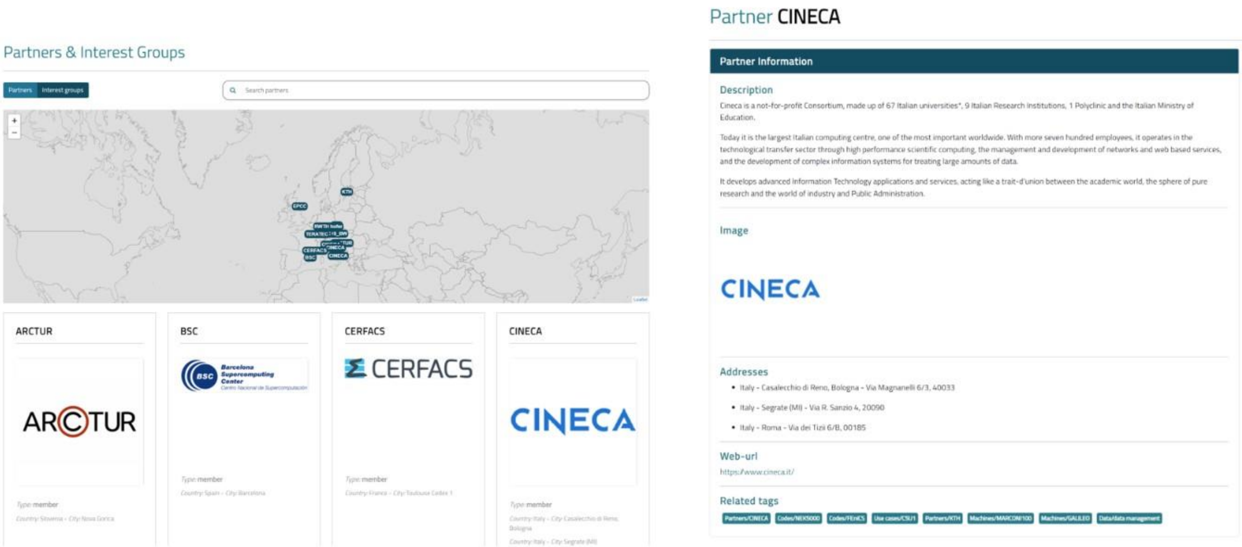

Figure 3: EXCELLERAT Service Portal: Partners & Interest groups search page (left) and information on Cineca Partner (right).

From the point of view of the management of the portal, the presentation of the Interest groups which are currently 19 is very similar. A button at the top of the page allows you to switch between Partners and Interest groups.

2.3 Industrial use cases

As far as industrial use cases are concerned, there are currently 8 industrial use-cases in the portal (Figure 4).

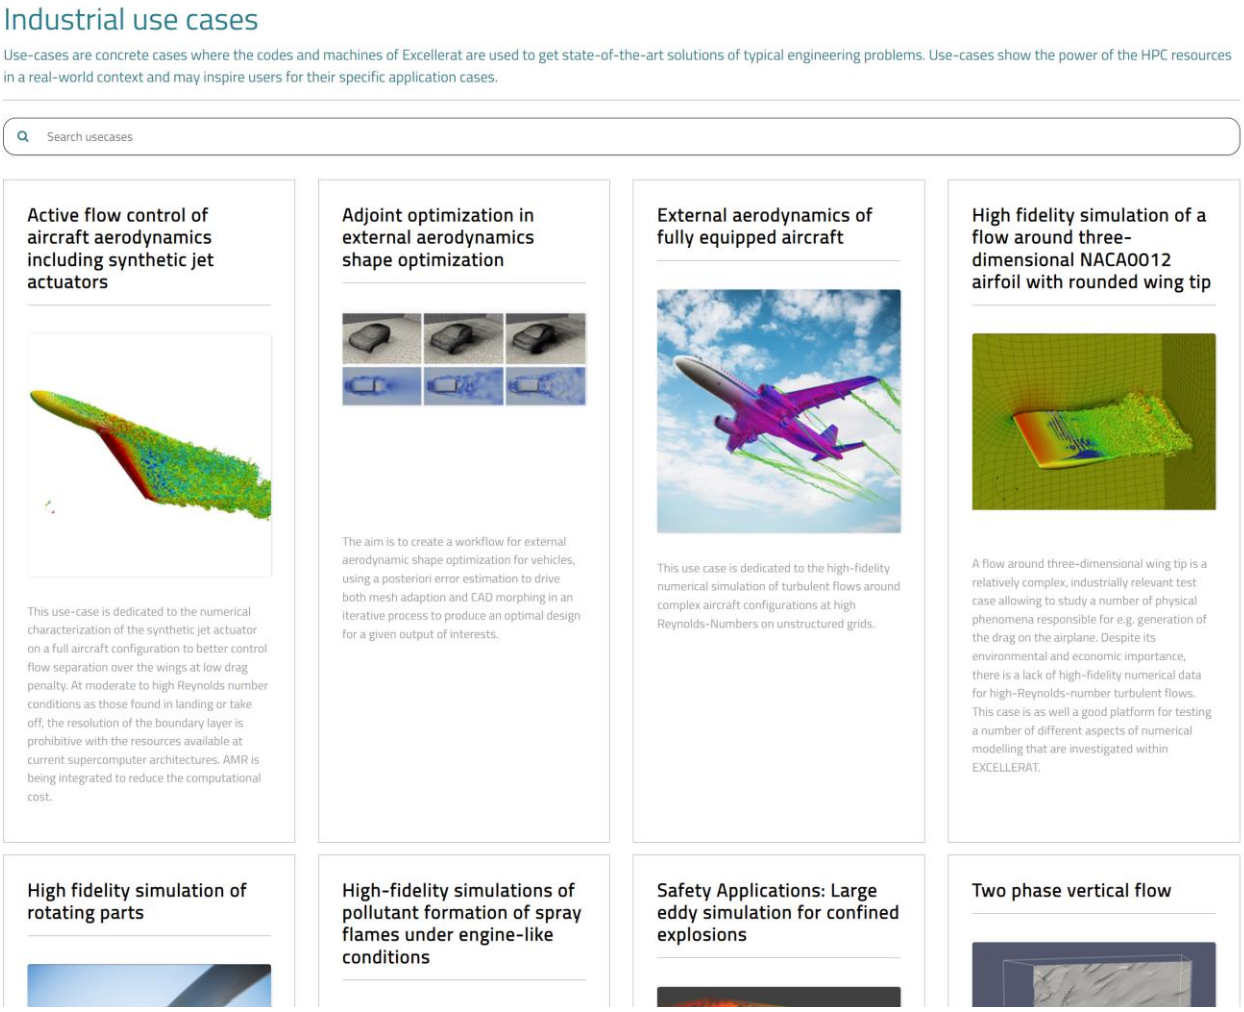

Figure 4: EXCELLERAT Service Portal: search page of industrial use cases.

For each use case it is possible to access a descriptive page that contains abstract, image, tags and a description from two different points of view, namely engineer and developer. The user will always find both descriptions available but only the description relating to the dashboard of origin (through which the page was accessed) will be directly legible (open). The descriptions present different and complementary technical details separately to give the user the opportunity to quickly navigate to the most relevant content with respect to their profile and interests. In Figure 5 there is an example of display relating to the use-case C1U1 *"High fidelity simulation of a flow around three-dimensional NACA0012 airfoil with rounded wing tip*". In this figure both descriptions are open, by way of example.

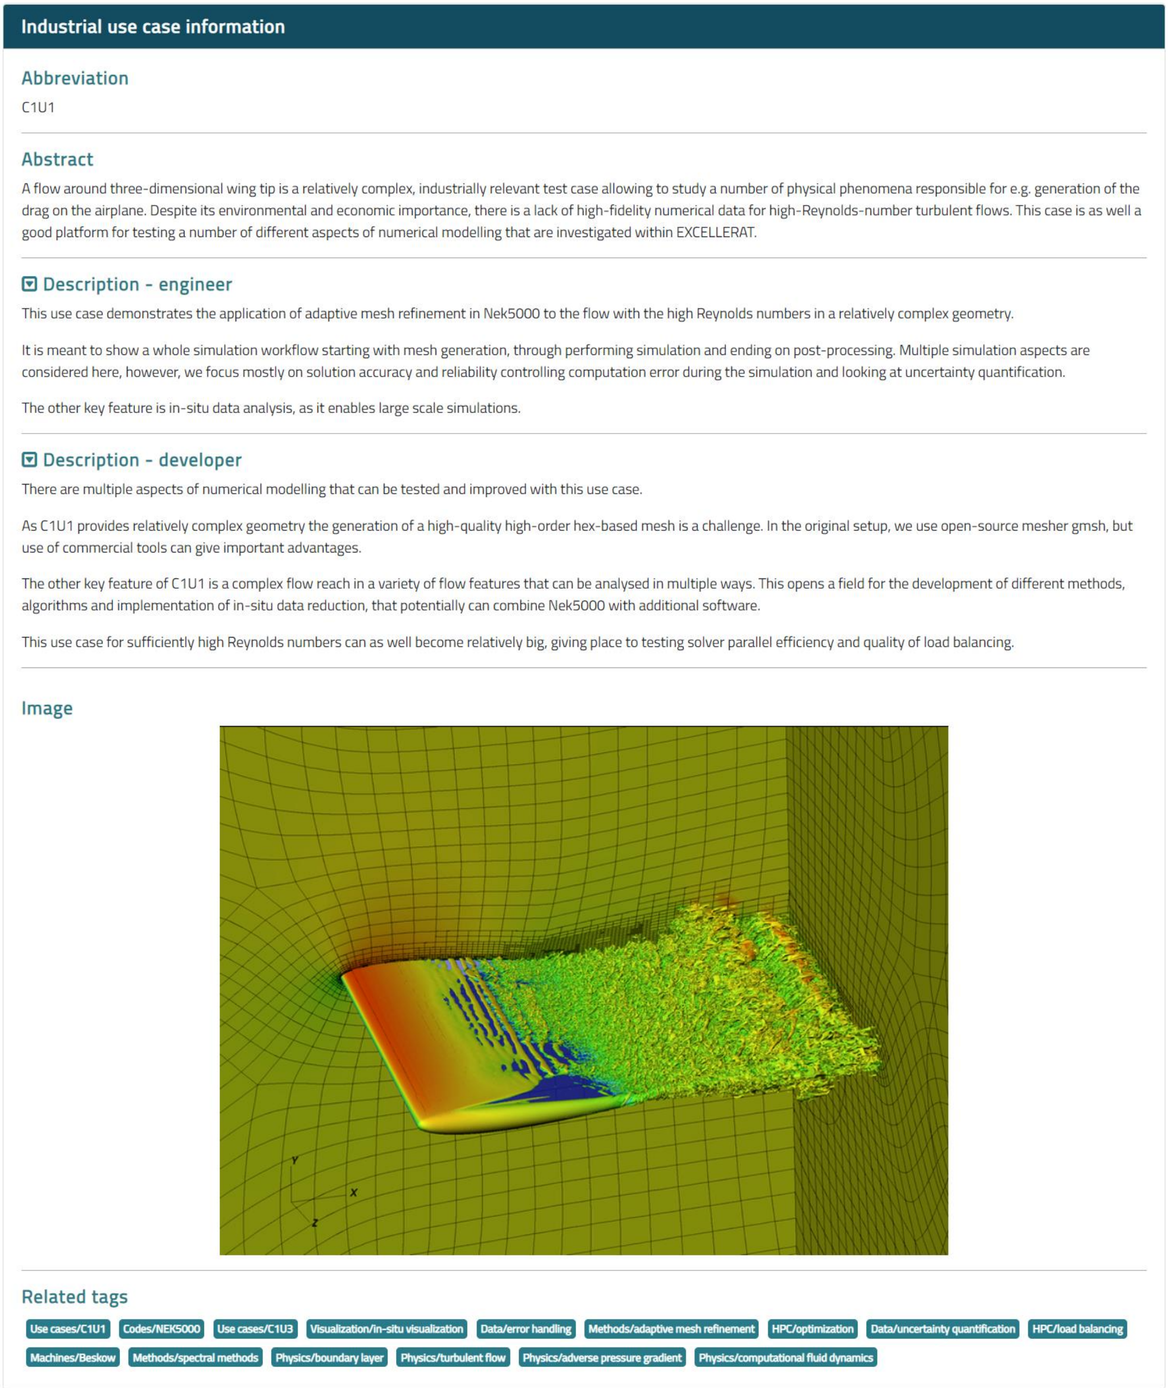

Figure 5: EXCELLERAT Service Portal: example of page showing the details relating to the use-case C1U1.

2.4 Application software

The application software presented in the portal are divided into three categories that can be selected via a button at the top of the page. The first category contains the six EXCELLERAT core-codes. The second category contains the onboarded codes during EXCELLERAT support activities towards external codes: these are currently two codes relating to the two onboarding projects carried out during the project period (see Section 3.2 of this deliverable). Finally, the third category contains the six software tools which, while not representing core codes for simulating problems, are auxiliary software (data management, visualization, etc.) that support various components of the simulation workflow and are part of the EXCELLERAT development. As an example, in Figure 6, the page listing the available tools is shown:

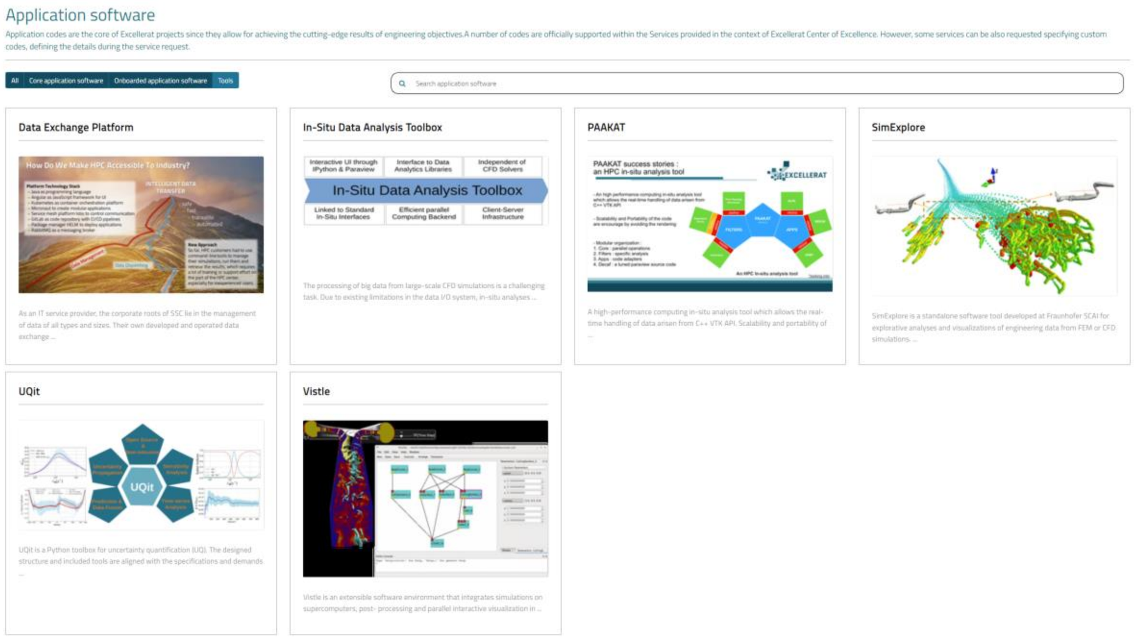

Figure 6: EXCELLERAT Service Portal: application software search page. The three sub-categories (core-codes, onboarded codes and tools) can be selected with the buttons at the top left.

An application software page contains details such as abstract, availability, license, web-url, code-url, image, and tags. It also contains two types of description: the first entitled "The engineering challenge" is aimed at the engineer user and is activated automatically when coming from the engineer dashboard, while the second descriptive section is entitled "Application software code details" and is activated automatically when coming from the developer dashboard. An example of an application software page (AVBP)[1] is provided in Figure 7.

[1] https://services.excellerat.eu/viewcode/5

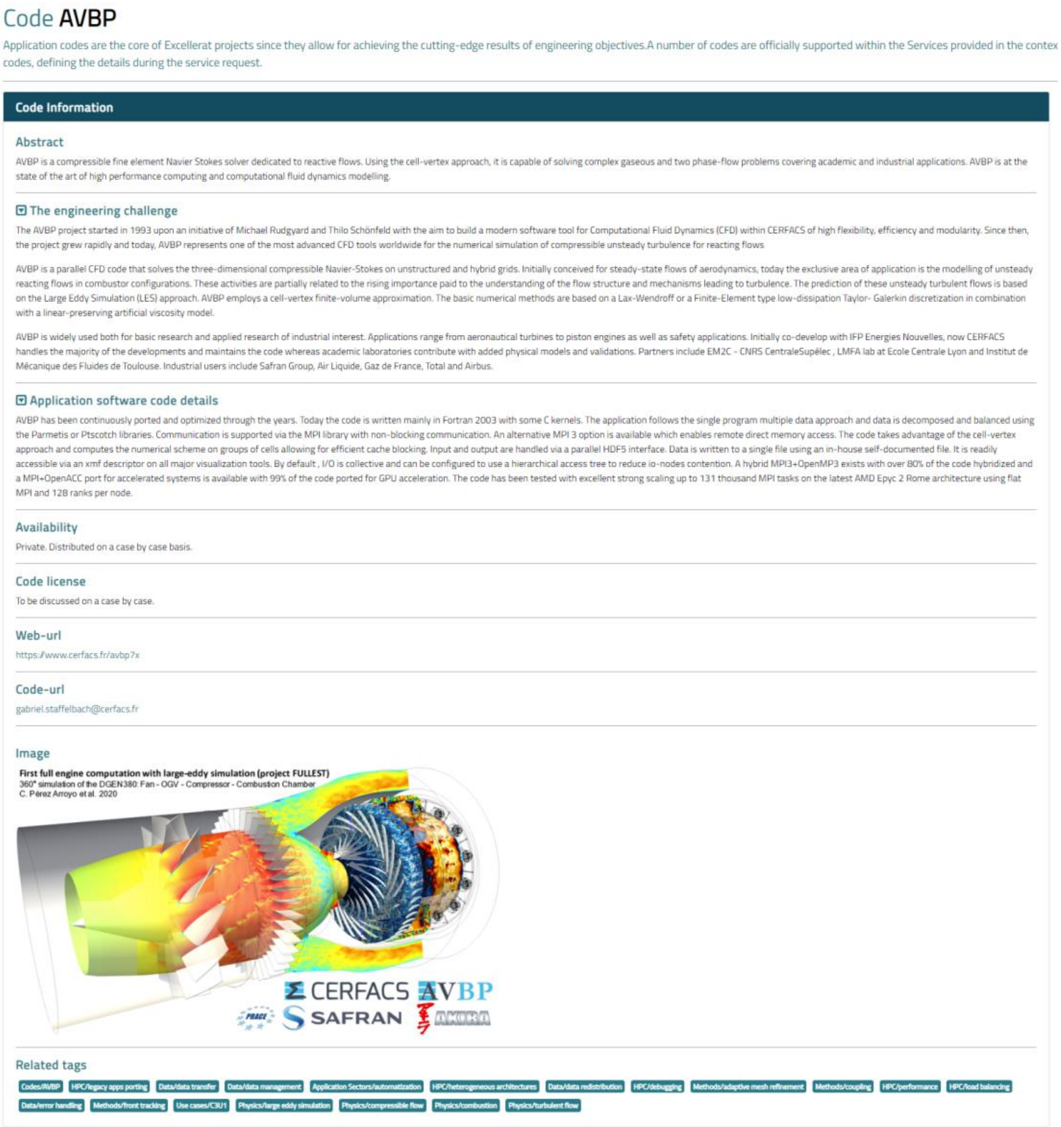

Code AVBP

Application codes are the core of Excellerat projects since they allow for achieving the cutting-edge results of engineering objectives.A number of codes are officially supported within the Services provided in the contex codes, defining the details during the service request.

Code Information

Abstract

AVBP is a compressible fine element Navier Stokes solver dedicated to reactive flows. Using the cell-vertex approach, it is capable of solving complex gaseous and two phase-flow problems covering academic and industrial applications. AVBP is at the state of the art of high performance computing and computational fluid dynamics modelling.

The engineering challenge

The AVBP project started in 1993 upon an initiative of Michael Rudgyard and Thilo Schönfeld with the aim to build a modern software tool for Computational Fluid Dynamics (CFD) within CERFACS of high flexibility, efficiency and modularity. Since then, the project grew rapidly and today, AVBP represents one of the most advanced CFD tools worldwide for the numerical simulation of compressible unsteady turbulence for reacting flows

AVBP is a parallel CFD code that solves the three-dimensional compressible Navier-Stokes on unstructured and hybrid grids. Initially conceived for steady-state flows of aerodynamics, today the exclusive area of application is the modelling of unsteady reacting flows in combustor configurations. These activities are partially related to the rising importance paid to the understanding of the flow structure and mechanisms leading to turbulence. The prediction of these unsteady turbulent flows is based on the Large Eddy Simulation (LES) approach. AVBP employs a cell-vertex finite-volume approximation. The basic numerical methods are based on a Lax-Wendroff or a Finite-Element type low-dissipation Taylor- Galerkin discretization in combination with a linear-preserving artificial viscosity model.

AVBP is widely used both for basic research and applied research of industrial interest. Applications range from aeronautical turbines to piston engines as well as safety applications. Initially co-develop with IFP Energies Nouvelles, now CERFACS handles the majority of the developments and maintains the code whereas academic laboratories contribute with added physical models and validations. Partners include EM2C - CNRS CentraleSupélec , LMFA lab at Ecole Centrale Lyon and Institut de Mécanique des Fluides de Toulouse. Industrial users include Safran Group, Air Liquide, Gaz de France, Total and Airbus.

Application software code details

AVBP has been continuously ported and optimized through the years. Today the code is written mainly in Fortran 2003 with some C kernels. The application follows the single program multiple data approach and data is decomposed and balanced using the Parmetis or Ptscotch libraries. Communication is supported via the MPI library with non-blocking communication. An alternative MPI 3 option is available which enables remote direct memory access. The code takes advantage of the cell-vertex approach and computes the numerical scheme on groups of cells allowing for efficient cache blocking. Input and output are handled via a parallel HDF5 interface. Data is written to a single file using an in-house self-documented file. It is readily accessible via an xmf descriptor on all major visualization tools. By default , I/O is collective and can be configured to use a hierarchical access tree to reduce io-nodes contention. A hybrid MPI3+OpenMP3 exists with over 80% of the code hybridized and a MPI+OpenACC port for accelerated systems is available with 99% of the code ported for GPU acceleration. The code has been tested with excellent strong scaling up to 131 thousand MPI tasks on the latest AMD Epyc 2 Rome architecture using flat MPI and 128 ranks per node.

Availability

Private. Distributed on a case by case basis.

Code license

To be discussed on a case by case.

Web-url

https://www.cerfacs.fr/avbp7x

Code-url

gabriel.staffelbach@cerfacs.fr

Image

First full engine computation with large-eddy simulation (project FULLEST)
360° simulation of the DGEN380: Fan - OGV - Compressor - Combustion Chamber
C. Pérez Arroyo et al. 2020

CERFACS AVBP
PRACE SAFRAN AKIRA

Related tags

Codes/AVBP, HPC/legacy apps porting, Data/data transfer, Data/data management, Application Sectors/automatization, HPC/heterogeneous architectures, Data/data redistribution, HPC/debugging, Methods/adaptive mesh refinement, Methods/coupling, HPC/performance, HPC/load balancing, Data/error handling, Methods/front tracking, Use cases/C3U1, Physics/large eddy simulation, Physics/compressible flow, Physics/combustion, Physics/turbulent flow

Figure 7: EXCELLERAT Service Portal: example of page showing the details of the AVBP code.

2.5 Events

The events section[2] contains three distinct types. The first category contains general purpose events, which can be conferences or conventions of various kinds. The second category are Training events. The third category are Funding-related events. For the latter case, these are real events or even simply calls for projects related to funding. An example of event availability for 2022 is exemplified in Figure 8.

For each event there is the possibility to specify the type of event, date, time, description, web-url, presenters of the event, availability of participation, image. In addition, it is possible to upload various kinds of material. The material is then available for registered users (no matter with what permissions). The constraint of being a registered user is a strategy, proposed and discussed by the EXCELLERAT partners, to try to promote the creation of a more informed user base for the portal.

[2] https://services.excellerat.eu/searchevents

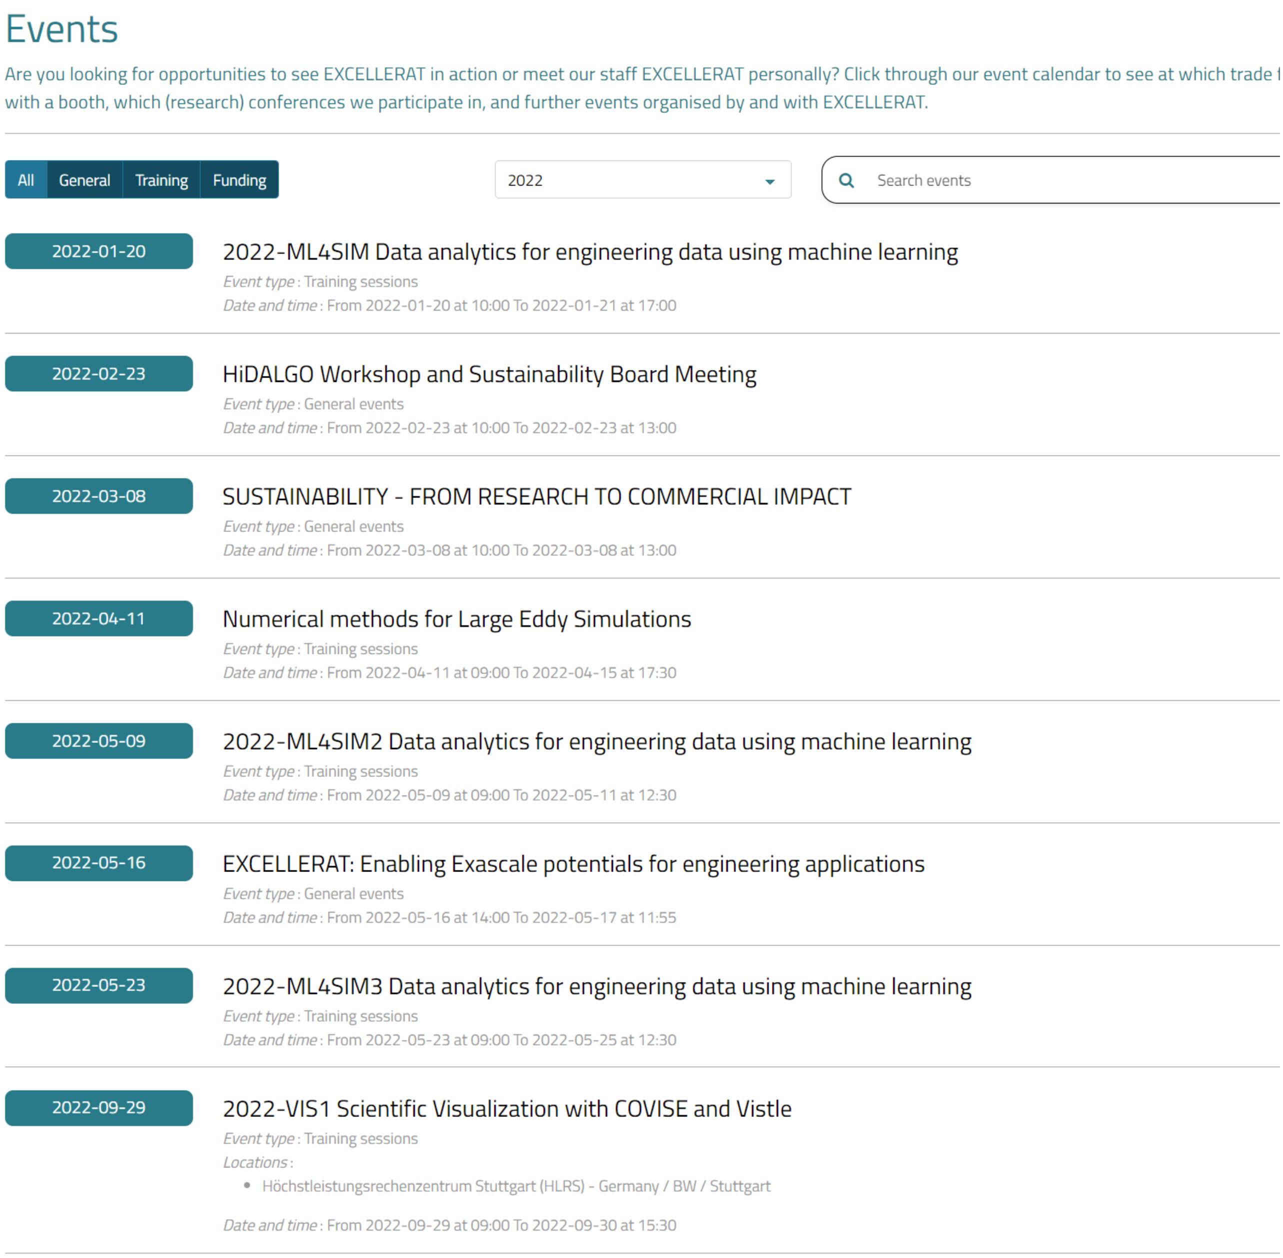

Figure 8: EXCELLERAT Service Portal: event search page. The three sub-categories general events, training events and funding opportunities can be selected using the buttons on the top left.

In addition to these categories, the events page allows you to access external training events, i.e. training of EXCELLERAT partners or training of particularly significant stakeholders in the training field such as PRACE and EUROHPC. For further discussion on training, please refer to the Section 4 of this deliverable.

2.6 Repositories

The repositories are of three types. The first category is called "Public material", the second category "Best practices", and the last are the datasets. There are currently 16 contents, none of which belong to the dataset category which has been prepared for later use. A repository is characterized by a description, a web-url, a license, an image and some tags. In addition, similarly to the events, it is possible to upload material, via the uploader integrated in the portal, which then becomes available for registered users of the portal itself. As an example, Figure 9 shows the page describing a "Best practice" type repository, in which the material download part is visible in particular.

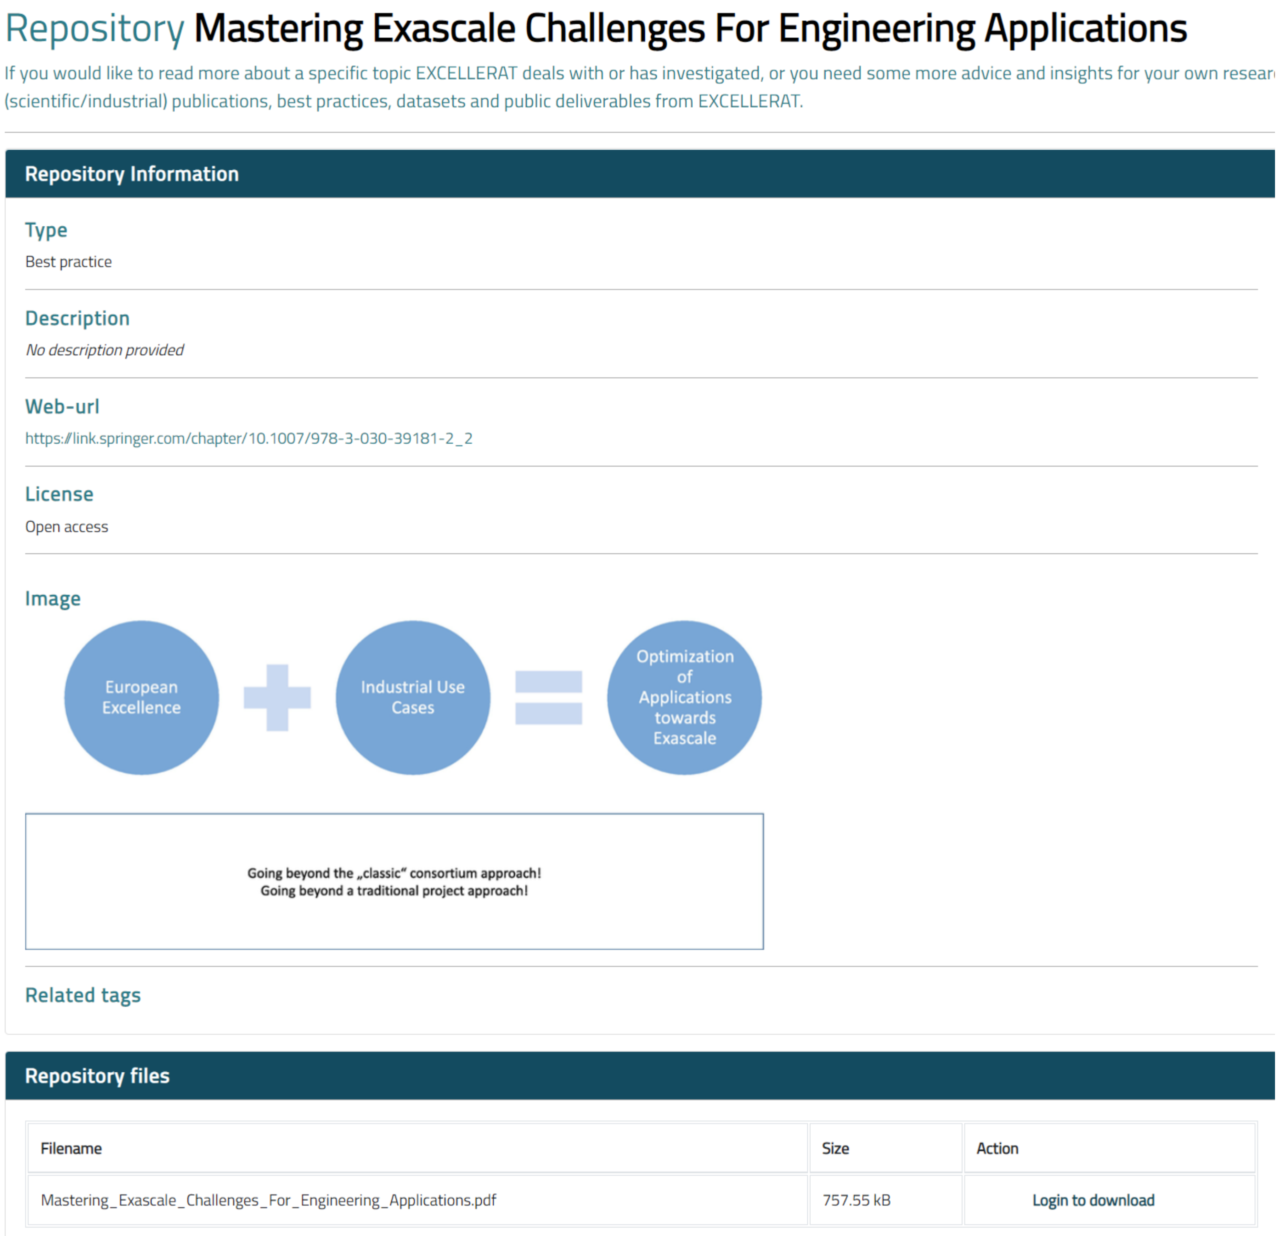

Figure 9: EXCELLERAT Service Portal: example of page showing the details of a best practice repository.

2.7 Consulting topics

For some of the contents already presented, such as industrial use cases, application software, and events, it is possible to use support request features that channel the user towards an EXCELLERAT service request among the different types allowed. Consulting topics, on the other hand, represent a form of content that is more directly and explicitly service-oriented for possible EXCELLERAT users. The consulting topics in fact summarize possible issues and needs related to the trend of exascale computing oriented towards engineering applications.

The portal currently collects 15 consulting topics[3]. In addition to the list page and the possibility to search for a consulting topic through free search, it is also possible to filter using the tags defined in the various consulting topics. An excerpt from the consulting topics list and search page is shown in Figure 10.

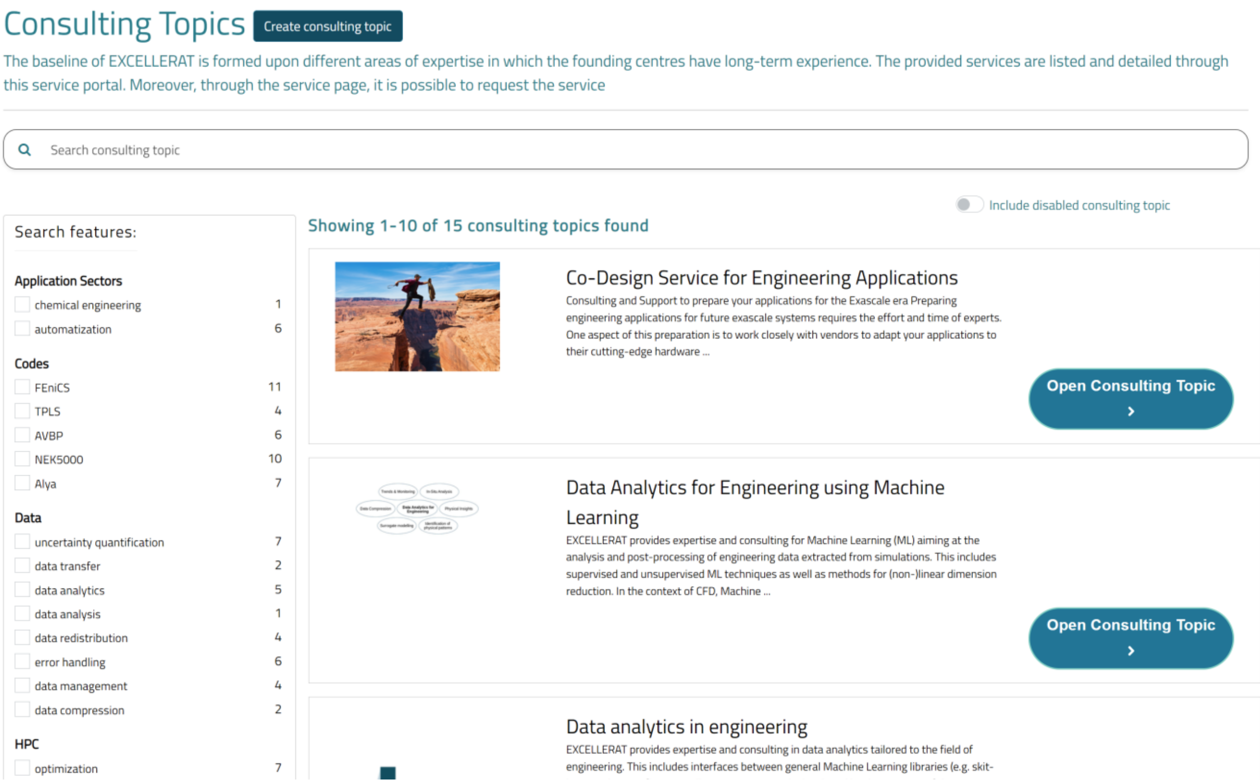

Figure 10: EXCELLERAT Service Portal: consulting topics search page. In addition to the free search at the top, on the left you can see the availability of filters based on tags and usable for structured search.

For each consulting topics it is possible to specify a description in markdown that can describe in more or less detail the topics and strategies dealt with in the consulting topic in question.

2.8 KPI Analysis

Concerning the KPIs collection, we will start from those defined in D5.1 *"Initial Assessment of Training Needs and Services Building Plan"*. In that report, a distinction was made between the KPIs for internal services, defined by Task 5.1, and those for external services, defined by Task T5.2.

Table 1 gives an overview of the KPIs for internal services.

[3] https://services.excellerat.eu/searchtopics

KPI number	Description	Definition	Target	Score at M40
5.1.1	Percentage of the number of internal services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	92% (6,5 implemented / 7 high priority)
5.1.2	Number of internal active users	*N_active_monthly*	≥15 for the first year after MS4 completion, with a 20% average increase later	51
5.1.3	ACSI score: user satisfaction for internal services, annually	N_{ACSI}*_int_services*	≥80	61

Table 1: KPIs for Task 5.1

As regards the ACSI score, it was measured thanks to an online survey. The evaluation, although positive, still does not reach the desired target. Analyzing the answers provided by the participants, we find relevant and interesting observations. In particular, most notes are specific for some tools:

- OpenProject is not enough user-friendly
- Etherpad does not work sufficiently well for collaborative editing
- MediaWiki would be easier to use with the WYSIWYG extension

While other notes are more general:

- Tools require more integration among each other
- Guidelines are missing so as to keep things tidy and uniform and allow for easier search and find

In the follow-up of the project, if funded, these points will be addressed to improve the experience of using internal services.

Table 2 gives an overview of the KPIs for external services. Despite the promotional campaign executed with the support of WP7, the number of external users remains under the target. Since the event is close to the submission time of the Deliverable, we are not able to account in the table the new registrations due to the interest generated by the final EXCELLERAT workshop.

KPI number	Description	Definition	Target	Score at M40
5.2.1	Number of external	*N_ext_users*	≥20 at the end of the first year after MS4 completion, ≥30 at the end of the project	18
5.2.2	Number of requested services (for the present implementation release)	*N_services*	>1	3
5.2.3	Percentage of the number of external services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	100%
5.2.4	Uptime	*Time of portal availability /Time from first production release*100*	>70% for the first year after MS4 completion >85% for the second, >95% later	95%
5.2.5	ACSI score: user satisfaction for external service, annually	N_{ACSI}*_ext_services*	≥80	80

Table 2: KPIs for Task 5.2

3 Further Applications

In D5.1 *"Initial Assessment of Training Needs and Services Building Plan"*, several ways have been foreseen for EXCELLERAT to attract new applications and, thus, potential new end-users and customers. Among the main ones, the progress in the interaction with the Interest Groups since D5.5 *"Updated Services, Training and Portal Report"* is described in detail in Section 3.1. Another aspect regards the Service Portal development, already object of Section 2, whereas the particular issue of the onboarding workflow in the Portal is described in Section 3.2.1 below.

The onboarding of further codes (Section 3.2) has technical and scientific advantages for both EXCELLERAT and for the external codes involved. These can benefit from the EXCELLERAT achievements, and at the same time act as proof-of-concept of the transferability of the centre's results. The onboarding also allowed to create success stories for outreach and dissemination, and to show that there is indeed a need of EXCELLERAT services from the community (therefore creating an exploitation channel for the post-funding phase).

3.1 Status and Activities Performed: Interest Groups

To recap the introduction to the Interest Groups given in D5.5, the EXCELLERAT Interest Groups are recruited to communicate, monitor, and validate the project's overall goals and its industrial, as well as technological relevance for the user communities. The first contact through "Letters of Support" in Y1 marked the kick-off of the interaction. Next, with an official invitation, the Interest Groups have been asked to fill in and sign a consent form, requesting their agreement to the following: to become member of the EXCELLERAT Interest Groups, to appear in EXCELLERAT dissemination material, and to subscribe to the Interest Group mailing list (each a separate choice).

The Interest Groups have been so far involved in two virtual meetings. A first meeting hosted by HLRS on November 8th, 2019, was dedicated to showcase the EXCELLERAT activities to the Interest Groups. A short overview of the project was given, and three applications were presented by the respective developers: the data exchange platform SWAN, visualisation with Vistle, and the BigWhoop compression library. During a second web-seminar on June 26, 2020, the Interest Groups (two new ones, in particular) were introduced to the activities of the CoE and guided through the refurbished Service Portal.

In discussion with the members of the Interest Groups, topics for a last workshop were identified. For a wider reach, these were finally incorporated in the programme of the two-days online conference "*EXCELLERAT: Enabling Exascale potentials for engineering applications*"[4] on May 16-17th, 2022. Participation of representatives of at least two Interest Groups has been already confirmed, among which Leonardo Labs. This final conference is dedicated to showcase the impact, the innovations and tools resulting from EXCELLERAT. It will highlight the industrial perspective of the EXCELLERAT technical achievements, such as the potential towards exascale computing, scalability, accelerators and GPU computing, adaptive mesh refinement, data management, visualisation, and emerging technologies.

The Interest Groups that consented to be publicly mentioned are listed in Table 3, together with their attendance to the online meetings in Y1 and Y2. These Interest Groups also appear in both the EXCELLERAT Website (Figure 11) and in the Service Portal, in the latter with a short description and tags connecting their areas of interest to the other sections of the Portal. As an update to D5.5, Constelcom and "Falquez, Pantle and Pritz GbR" (Nuberisim) joined as EXCELLERAT Interest Groups in 2021. As reported in the Service Portal, the goal of

[4] https://services.excellerat.eu/viewevent/61

Constelcom is “*linking providers and consumers in its unique, secure and web accessible HPC ecosystem*”, while Nuberisim provides “*consulting, research and development services for fluid machinery and fluid devices*” and access to the “*NUBERISIM Cloud Platform*”.

Figure 11: Interest Groups in the EXCELLERAT Website [https://www.excellerat.eu/interest-groups/].

We can mention several activities that EXCELLERAT undertook together with the Interest Groups, such as the EXCELLERAT HPC Podcast (see D5.5), training (see Section 4.2) and the onboarding applications (Section 3.2): For instance, the Interest Groups received a prioritised invitation to the ML4SIM workshops in 2022. Further, AMD’s involvement in EXCELLERAT was strengthened through three online workshops provided online at USTUTT in January and June 2021, as well as through discussions towards the acceleration of the onboarding application code STREAmS. Finally, the online round table on a test case present in both EXCELLERAT and exaFOAM also took place in the frame of the Interest Groups, since the Interest Group “ESI Group” acts as the Principal Investigator of exaFOAM.

A videoconference, on January 12th, 2022, has been dedicated to synchronise WP5, 6, 7 on the EXCELLERAT sustainability strategy, including the cooperation with the Interest Groups. Some of the proposals included further engaging the Interest Groups in the EXCELLERAT community, e.g. through an optimised sharing of EXCELLERAT resources (blogs, whitepapers, use cases, events), while, on the other hand, allowing the Interest Groups to share their services on the Service Portal or through the EXCELLERAT HPC Podcast. The Interest Groups were already part of a series of interviews initiated by WP6 (Task 6.1) to complete the EXCELLERAT market assessment (see D6.5 *“Final Market Assessment”*). In order to further involve them in the EXCELLERAT outreach (WP7), “exclusive” Interest Groups content will be added to the EXCELLERAT newsletter (see D7.4 *"Updated Dissemination, Communication, Collaboration, Community Building and Standardization Report"*).

The Interest Groups presence on the Service Portal has been refurbished by T5.3 during 2021 (e.g. through requesting high-quality logos and an updated description). In February 2022, they have been invited to request an account on the Service Portal. Such an invitation was sent at a later stage of the project, since several sections of the Service Portal needed previous internal agreement, checking, and updates. The advantages of such an account have been highlighted in the invitation: Logged-in users can access EXCELLERAT training material and repositories and contact directly EXCELLERAT support staff to handle service requests. By creating an account, the Interest Groups got the option to subscribe to the EXCELLERAT newsletter to receive the last news first-hand and in advance, e.g. about EXCELLERAT success stories.

Finally, in March 2022, the current Interest Groups have been requested to support the application process of EXCELLERAT Phase 2 (through sending back a Letter of Support), in order to pursue the cooperation and to ensure the alignment of EXCELLERAT Phase 2 with the Interest Groups' and the consortium stakeholders' needs.

Interest Group	Typology	Meeting 1	Meeting 2
ANSYS Germany GmBH	Code-Developer/ISV		
Constelcom	Code-Developer/ISV		
Dassault Aviation	Code-Developer/ISV		
ESI Group	Code-Developer/ISV		
Fraunhofer IGD	Scientific Expert		
Nuberisim	Scientific Expert		
Politecnico di Milano	Scientific Expert	Yes	Yes
STAI	Scientific Expert		
University of Ljubljiana	Scientific Expert		
University of Rome Tor Vergata	Scientific Expert	Yes	
AMD	Technology Provider	Yes	
ARM	Technology Provider	Yes	
ATOS	Technology Provider	Yes	Yes
CRAY	Technology Provider	Yes	Yes
HPE	Technology Provider		
Intel	Technology Provider	Yes	
Leonardo	Technology Provider		Yes
NVIDIA	Technology Provider		Yes
Westinghouse	Technology Provider	Yes	

Table 3: EXCELLERAT Interest Groups and their attendance to the online meetings in Y1 and Y2.

3.2 Status and Activities Performed: Onboarding Applications

The EXCELLERAT onboarding requires the completion of the workflow steps described in the next section. In particular, two points have been deemed essential to initiate an onboarding: On the one hand, the application must be consistent with the scope of the project, namely an engineering code with potential for high performance and scalability (possibly exascaling). On the other hand, at least one EXCELLERAT partner must have available effort within the project (or as in-kind contribution) to work on it.

Two applications successfully went through the onboarding process:

- **PAinG-Flow**[5] [2] is a project of FLUMACS (FLUid dynamics of MAcro and micro Complex Systems), Department of Mechanical and Aerospace Engineering (DIMA) at the Sapienza University of Rome. EXCELLERAT support has been granted by HPC specialists at CINECA. The collaboration has been accepted in February 2020 (prior to the finalised onboarding procedure on the Service Portal), and a final report has been sent in October 2021 (see Annex 1: EXCELLERAT Onboarding Final report: PAinG-Flow). In short, the project aims at implementing the Exact Regularized Point Particle method (ERPP) to exert the particle/fluid momentum coupling in multiphase turbulent flows on a code able to run on supercomputers with GPUs.
- **HiDALGO-UAP** originates from the partner CoE HiDALGO [3] within the Urban Air Pollution Pilot (UAP), CFD Module, led by the Szechenyi Istvan University (SZE) in Győr, Hungary. After the initial contacts (see D5.5), HiDALGO went through the onboarding process in March 2021 and the collaboration was formally accepted. An overview of the UAP project has been provided for the Service Portal[6], while the final results have been summed up in a success story [4] which was published in December 2021.

  The central part of the UAP application is an HPC-framework for simulating the air flow in cities by taking into account real 3D geographical information of the landscape, applying a CFD simulation on a highly resolved mesh. Boundary conditions are derived from weather forecasts and re-analysis data (**Errore. L'origine riferimento non è stata trovata.**). The primary aim of the EXCELLERAT onboarding has been to improve UAP's performance on HPC systems. The OpenFOAM CFD module for pollution dispersion simulation has been used to discretise the problem.

[5] https://services.excellerat.eu/viewcode/11
[6] https://services.excellerat.eu/viewcode/12

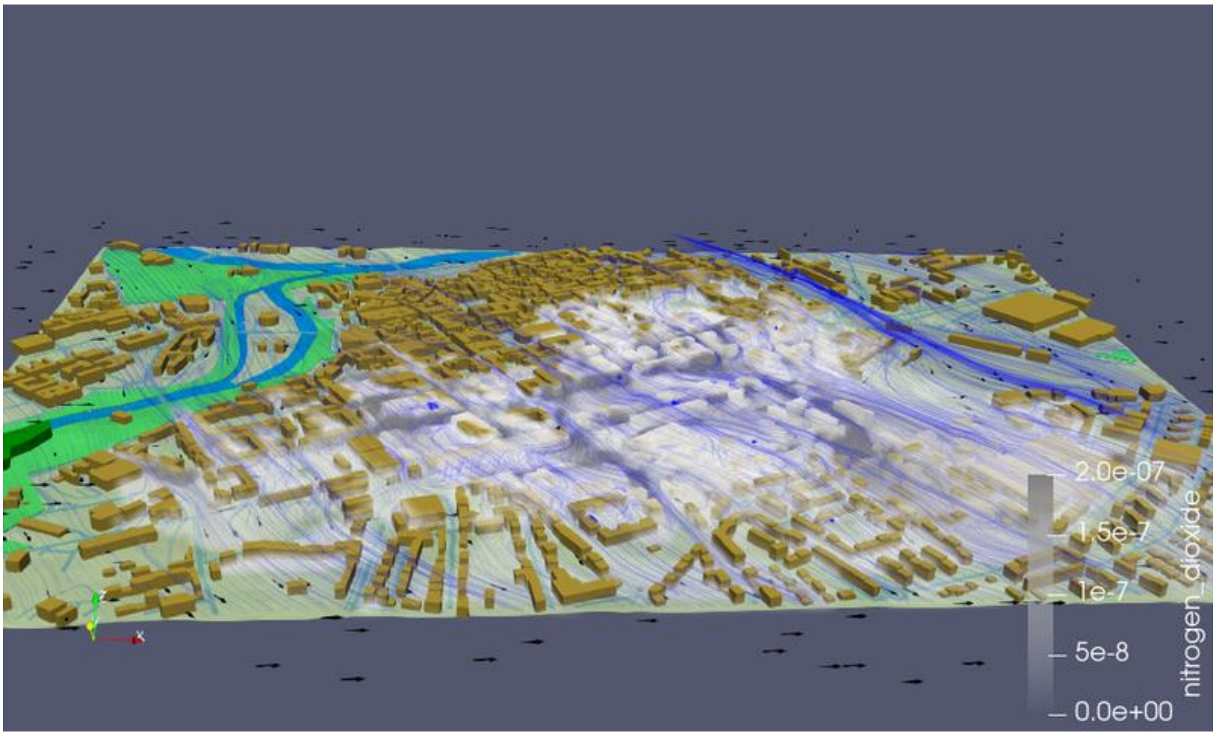

Figure 12: 3D snapshot of results of a city sample model with a synthetic wind profile as boundary condition and pollution from simulated random traffic as source (courtesy László Környei, SZE)

One further application submitted an onboarding request:

- **STREAmS (**Supersonic TuRbulEnt Accelerated navier stokes Solver) is an open-source project at the Department of Mechanical and Aerospace Engineering (DIMA), Sapienza University of Rome. The onboarding request has been submitted in November 2021. EXCELLERAT resources to actively support this application will be allocated during the next funding period (if granted).
The motivation for an onboarding lies in the aim of porting the existing open-source compressible fluid dynamics code to AMD GPU-based architectures (or more general heterogeneous HPC platforms). Another goal is the development of advanced features for simulation of flows in complex geometries, as immersed boundary capabilities to be implemented in a specific private plugin. Through EXCELLERAT, the STREAmS team has already been liaised to AMD personnel to discuss options and obstacles in porting to GPUs the existing Fortran code. The reader can consult [5] for further information.

Even though only a few applications either initiated or completed the onboarding workflow, T5.3 monitored collaborations and interactions between EXCELLERAT and other external codes and applications beyond the onboarding:

- EXCELLERAT T5.3 hosted the **OpenFOAM Thematic Working Group** (TWG) bi-monthly conference calls from the beginning of the project through October 2020 (see also D5.5). The TWG web-page is part of the internal EXCELLERAT Wiki. The purpose of the calls lied in exchanging information and best practices regarding the use of OpenFOAM in the context of an HPC-Cloud. During the calls, OpenFOAM members presented challenges and open problems that were discussed together with EXCELLERAT experts (e.g. the complex fluids use-case [6] already mentioned in D5.5).
- The collaboration between EXCELLERAT and the EuroHPC JU project **exaFOAM** [7] is connected to the EXCELLERAT Interest Group ESI, the Principal Investigator of

exaFOAM. A round table between exaFOAM and EXCELLERAT took place on December 16, 2021 to share different approaches about the common use-case DrivAer [8] In particular, CINECA presented their specific approach using the EXCELLERAT flagship code FEniCS.

- The Department of Mathematics and Computational Sciences of the Széchenyi István Egyetem University of Györ (SZE) [9] requested the consultancy of EXCELLERAT experts about their 3D compressible Fluid-Solver code, numerically focused on POD-DEIM reduced order models (see[7] for a demo). The request encompassed benchmarking possibilities, application to industrial use cases and further development with respect to acceleration on GPGPU. A round table took place on October 18, 2021 with the participation of CINECA, DLR, RWTH and USTUTT. After a presentation by SZE, feedback and insight were provided, especially from the perspective of Nek5000, CODA and the use-cases in acoustics at RWTH (see AIA [10]).
- The tool SimExplore has been developed by Fraunhofer SCAI in cooperation with EXCELLERAT as an EXCELLERAT extension of a Fraunhofer project. It has been therefore listed in the Tools page of the Service Portal[8]. SimExplore has been the object of extensive demonstrations during the Fraunhofer training workshops ML4SIM (see Section 4.2).

3.2.1 EXCELLERAT Portal Onboarding Workflow

The structural requirements for the Service Portal in order to attract external codes have been listed in D5.1, while details about their implementation were provided in D5.3 *"Services, Training and Portal Report"* and D5.5. As mentioned in D5.5, the onboarding workflow is a trade-off between the initial idea (in D5.1) and a feasible workflow compatible with the evolution of the Service Portal. The final implementation has been approved by the PMO, and the workflow is summed up here below.

After successful registration to the Service Portal, specifying an organisation or team, the applying user creates a project by compiling and submitting a form containing name and description of the project and linking to the corresponding owner organisation.

After submission, the second step is carried out by an evaluation team composed by a PMO representative and leaders of WP 2, 3, 4, and 5. General evaluation criteria include a feasibility checklist to pre-screen the application, and at least one partner having effort and interest to work on the submitted code. The evaluation might involve the request of additional information, e.g. via tele-conferences, and should be concluded within a specified timeframe. All communication with the applicant will happen through the "Service requests" feature, which functions as a request-ticket tracker[9]. During evaluation, staff members (i.e., EXCELLERAT partners) can both send staff-confidential messages and communicate with the applicant through "Service requests".

In case of positive evaluation, both the partner selected for collaboration and the applicant will receive a message through "Service requests". In particular, the applicant will also receive a template to provide information about the application software. This will be displayed on a dedicated section in the Portal[10] The applicant will also decide which results and advances they would like to post on the Service Portal during the joint work. Within a specified timeframe,

[7] https://www.youtube.com/watch?v=uTHnPNdGmWY
[8] https://services.excellerat.eu/viewcode/13
[9] https://services.excellerat.eu/searchrespsis
[10] https://services.excellerat.eu/searchcodes/onboarded

the designated EXCELLERAT partner will produce a workplan with a time-schedule and submit it to the applicant.

In case of negative evaluation, the applicant will receive a message on their Service Portal account (visible at “Service requests”), containing an explanation for the rejection.

Finally, the applicant will be also requested to evaluate the application procedure and, for successful applications, the onboarding process, through a survey prepared by T5.3 (KPIs 5.3.4 and 5.3.3 in the next section).

3.3 KPI Analysis

In D5.1, four functional and quality-related KPIs have been proposed for T5.3, which can be seen in Table 4: KPIs for Task 5.3. The quality-related KPIs refer to the ACSI Score (American Customer Satisfaction Index), defined in Section 2.3 of D5.1.

KPIs for this task have been measured after the first Portal release (MS4 in M14, see D5.3). In this final deliverable, we therefore present the total numbers from M15 to now (M40): number of submitted and approved applications in M15-M40 (target 9 at M42, corresponding to 1 per quarter), and averaged satisfaction of both candidates and successful applicants.

The below-target amount of sent and selected applications can be justified with the delay in finalising the onboarding workflow, not in production until the second portal release at M24. Nonetheless, collaborations between EXCELLERAT and external applications have taken place also outside the formal onboarding framework, as described in Section 3.2. The satisfaction KPIs, on the other hand, lie inarguably above the target.

KPI number	Description	Definition	Target at M42	Score at M40
5.3.1	Number of external entities sending an application (quarterly)	*N_app_sent_quarterly*	≥9 [in total]	3
5.3.2	Number of selected applications to interact with (quarterly)	*N_app_interacting_quarterly*	≥9 [in total]	2
5.3.3	ACSI score: Satisfaction of the on-boarding process for a service (e.g. via survey/feedback sheet) as average per year	N_{ACSI}*_onboarding*	≥80 [total average]	84,73
5.3.4	ACSI score: Satisfaction of the processed applicants (e.g. via survey / feedback sheet) as average per year	N_{ACSI}*_applications_per_year*	≥80 [total average]	100

Table 4: KPIs for Task 5.3

4 Training

This section reports on the training activities performed since the previous deliverable D5.5 (M24-M40), and on those scheduled until and beyond the end of the project at M42. Regarding the latter activities, it is assumed that the content will have been developed with EXCELLERAT effort until M42, while the course itself delivered at a later point in a different framework. Also in the last year of the project, most activities up to a few exceptions have been carried out online.

As in D5.1, D5.3 and D5.5, the first sub-section 4.1 is dedicated to the assessment of EXCELLERAT training. Then, the training activities since M24 are summed up (Section 4.2), and finally the KPIs for the entire funding period are reported (Section 4.3).

4.1 Status of Assessment of Training Activities

The two past assessments of the EXCELLERAT training activities have been carried out in **January 2019** (D5.1) and **July 2020** (D5.5). They encompassed events organized, held, or advertised by the organisations participating as members of EXCELLERAT. A regular assessment has the goal of maintaining an updated map of the existing initiatives in training and education, in order to also highlight the emerging needs for structuring and evaluating the EXCELLERAT training roadmap and the next initiatives.

As already mentioned in D5.5, the Coordination and Support Action for the European National Competence Centres CASTIEL [11] conducted a thorough training assessment among the 33 National Competence Centres of the network. Since most EXCELLERAT partners are part of the NCC network, the last EXCELLERAT assessment could be done in synergy with CASTIEL. A detailed description of the methodology and the results can be found in Annex 2: Assessment of Training Activities 2022, while here only the main outcome has been summed up.

The analysis considered three categories (Scientific Domain, Technical Domain, HPC Topic of each training activity). The EXCELLERAT training offer has been compared with the **baseline of eight EXCELLERAT centres** (training providers) for the joint periods 09/2019-08/2020 and 09/2020-08/2021 (Figure 13, Figure 14, Figure 15). All mentions of individual labels (i.e., each option in the three categories) have been collected for both the centres' and for EXCELLERAT training activities. The latter have been then computed separately and put into perspective by relating them to the overall sum. A summary of the main outcome follows:

- **Scientific Domain** (Figure 13): The EXCELLERAT offer shows a great amount of non-domain-specific training (comparable to the baseline) and is significantly above the baseline in the domains of engineering and mathematics. This was expected, based on the engineering vocation of EXCELLERAT. Keeping in mind that all applying domains of a training course had to be selected during the evaluation, it does not surprise mathematical aspects dominate, since they are also a prerequisite or a necessary introductory step in most EXCELLERAT courses (e.g., in a CFD course). Moreover, it is expected that the option "non-domain specific" would be selected whenever the specific topic is lacking from the list. In fact, topics of EXCELLERAT training are often specialised e.g. on a particular code or library, and cannot be found in this taxonomy.
- **Technical Domain** (Figure 14): The EXCELLERAT offer positions itself above the baseline with respect to numerical libraries and methods, performance analysis, scientific programming, software engineering and visualisation. These topics are all pertinent to EXCELLERAT applications, and EXCELLERAT is expected to fill in training gaps in these fields. On the other hand, several EXCELLERAT partners offer training on generic AI, Data Science and HPC topics, which justifies the EXCELLERAT ranking below the baseline in these domains.

- **HPC Topics** (Figure 15): EXCELLERAT outperforms the baseline in the topics of data visualisation, accelerators, HPC on exascale architectures, parallel libraries, performance engineering and analysis: These are all core-topics of the EXCELLERAT mission. Among the categories where EXCELLERAT is below the average, we can mention "parallel algorithms and parallel programming models", and "programming languages and methods for HPC". These HPC topics are well covered by the EXCELLERAT partners in their regular training schedule, while EXCELLERAT specialises in HPC for exascale.

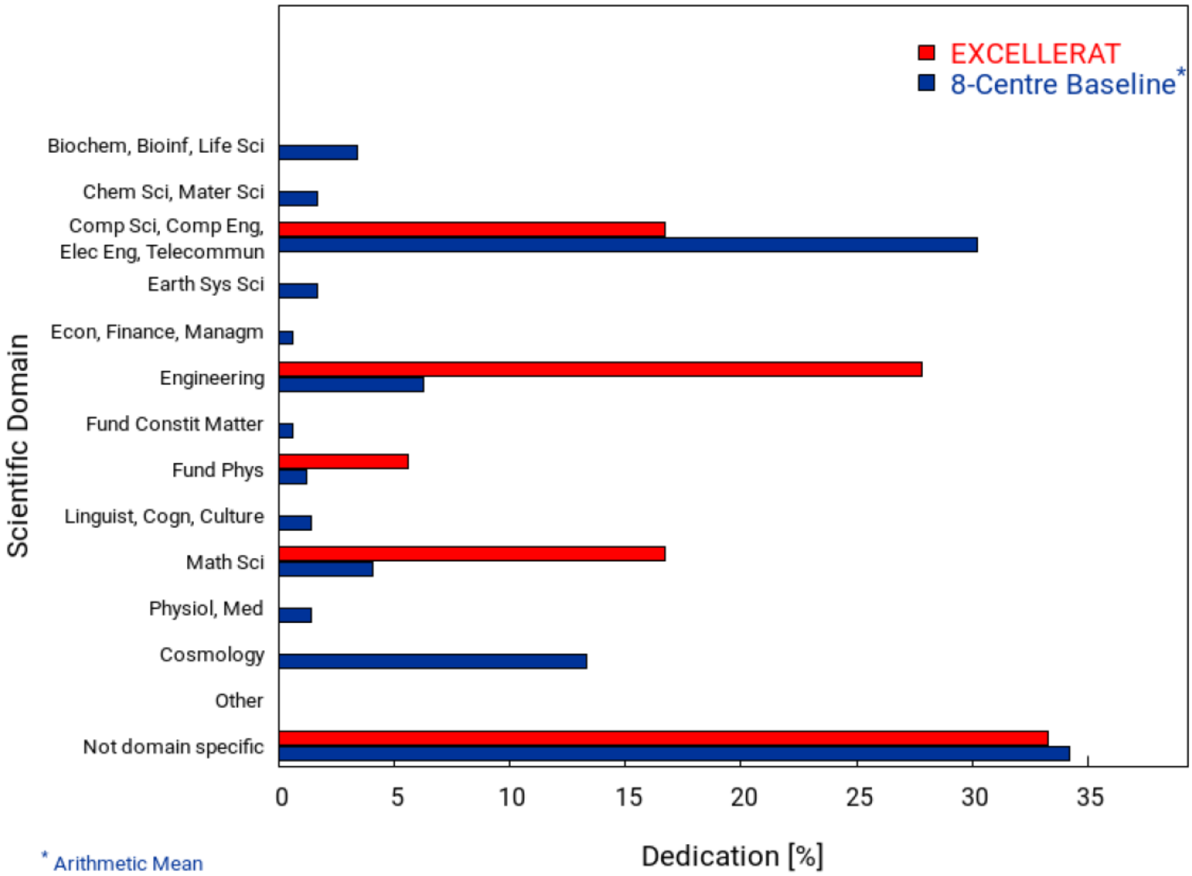

Figure 13: Profile representing the Scientific Domain of EXCELLERAT training and the eight-centre baseline from 09/2019 to 08/2021.

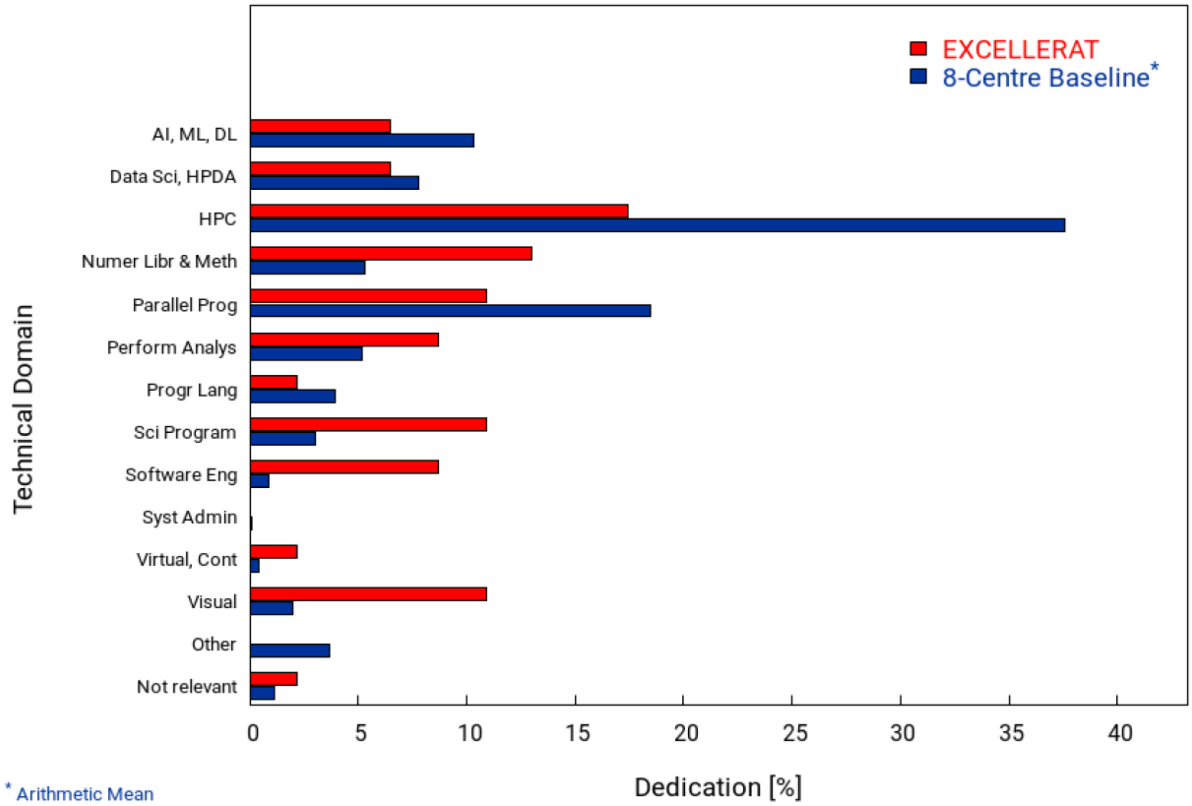

Figure 14: Profile representing the Technical Domain of EXCELLERAT training and the eight-centre baseline from 09/2019 to 08/2021.

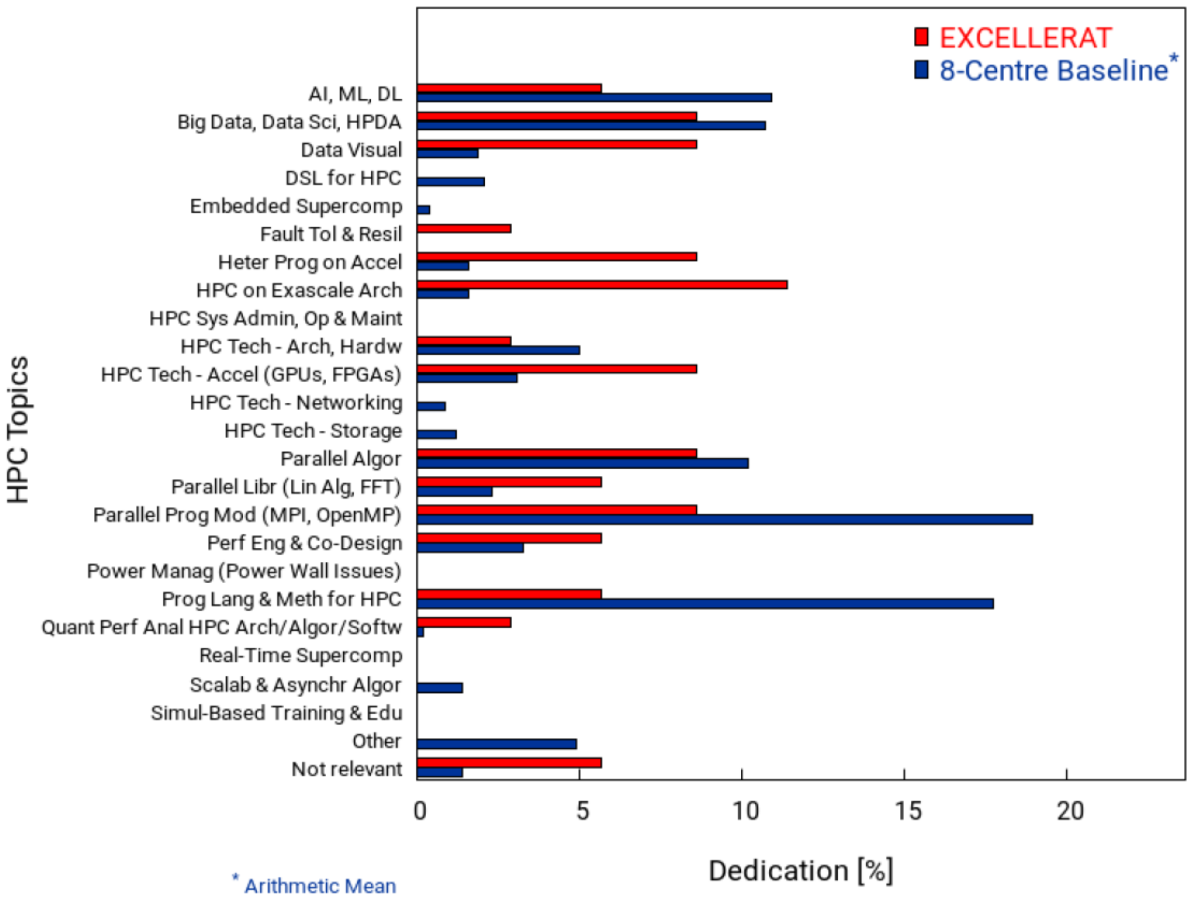

Figure 15: Profile representing the HPC Topics of EXCELLERAT training and the eight-centre baseline from 09/2019 to 08/2021.

4.2 Performed EXCELLERAT Training Activities

As described in D5.5 and in the previous reports, training activities within EXCELLERAT aim at filling specific gaps in the existing offer of the involved partners. On the one hand, when developing a new activity, the audience demand and the results of the training assessments are considered, for instance by integrating EXCELLERAT-specific interdisciplinary aspects into established activities of the consortium members. On the other hand, training related to the flagship codes or affiliated tools is provided, either by EXCELLERAT experts or by the Interest Groups. Compared to the few activities in M14-M23, due to the sudden switch to online teaching, the online program since M24 has been larger and enriched by some face-to-face events. A short summary about the challenges in transitioning to online courses can be found in D5.5, Section 4.3.

In D5.5, it was described how the Service Portal can act as the Knowledge Hub of the CoE, allowing to browse for training and education events, to deploy learning material for logged-in users, and to submit requests for dedicated training. Even though the portal has undergone a restyling since the submission of D5.5 (see Section 2), this did not affect the training section[11] with any major changes. A new section "*External training events*" has been added to cluster both training activities offered by EXCELLERAT partners and those by other CoEs or PRACE. Since 2022, EXCELLERAT Training events are additionally visible in the HPC in Europe Portal [12] managed by EuroCC and CASTIEL.

[11] https://services.excellerat.eu/searchevents/training

When listing the training activities, the respective training category is indicated as follows, according to the definition introduced in D5.1 (Section 5.2):

- Parallel Programming [PAR]
- Computational Fluid Dynamics [CFD]
- Scientific Visualization [VIS]
- Compute Cluster: Usage and administration [CLU]
- Performance Optimization and Debugging [PRF]
- Data in HPC [DAT]
- Programming Languages for Scientific Computing [LNG]
- Scientific Conferences and Workshops [C+W]
- Training for special communities [COM]
- Others [OTH]

In M24-40, the training effort below has been conducted. For each activity, these characteristics are specified: format, date, training category as above ([PAR] etc.), short description, resulting collaborations and material.

- "CINECA Virtual School on Numerical Methods for Parallel CFD" [CFD] (**online** from 2020-11-30 to 2020-12-11): Day 5 of the School, dedicated to Spectral Methods, included three presentations by EXCELLERAT experts at KTH. Related **material** is available on the Service Portal.
- "AMD Machine Learning workshop" [DAT] (**online** on 2021-01-20 and 2021-01-21): In the framework of a **collaboration** between AMD and USTUTT, the EXCELLERAT Interest Group AMD provided their first training of this kind in Europe. This remote one-day training was offered twice and allowed participants to actively use the AMD GPUs through the open-source development platform ROCm for HPC GPU computing [13]. The training also covered the configuration of an AMD-based machine learning environment.
- "First Joint CoE Technical Workshop" [C+W] (**online** from 2021-01-27 to 2021-01-29): The event originated as a **collaboration** between EXCELLERAT and the CoEs ChEESE and HiDALGO and has been organised by members of USTUTT, TERATEC and CINECA. It consisted of four technical sessions (load balancing, in-situ and remote visualisation, co-design, GPU porting), each with presentations and discussions lead by the involved CoEs and a few guests. **Material**: Links to the recording and documentation are available on the Service Portal.
- "Data analytics for engineering data using machine learning (ML4SIM)" [DAT] (**online** on 2021-03-19, from 2021-12-13 to 2021-12-14, and from 2022-01-20 to 2022-01-21): As part of the EXCELLERAT training program, Fraunhofer SCAI in cooperation with HLRS offered three instances of a workshop addressing the preparation, analysis and interpretation of numerical simulation data by machine learning methods. Besides the introduction of the concepts such as clustering, dimensionality reduction, visualisation and prediction, this course provided hands-on tutorials using the libraries numpy, scikit-learn and pytorch as well as the Fraunhofer SCAI tool SimExplore. Since the second edition, the workshop has been extended with a second day dedicated to prediction with deep learning, and to the interpretability of machine and deep learning methods. **Material**: The Jupyter Notebooks used for the hands-on sessions can be downloaded from the Service Portal and executed locally or on a cluster.
- "Numerical methods for Large Eddy Simulation" [CFD] (**online** from 2021-04-12 to 2021-04-16): This is a CERFACS training on the EXCELLERAT flagship code AVBP to solve compressible Navier-Stokes equations for laminar and turbulent reactive flows,

in 2D and 3D, on unstructured and hybrid meshes, with third-order Taylor-Galerkin schemes.

- "Fortran for High Performance Computing" [LNG] (**online** on 2021-05-04): During this Fortran web-seminar conducted by Wadud Miah (Southampton University) and organised by TERATEC, language constructs that are pertinent to HPC were discussed, as well as recent attempts to revive Fortran and encourage the community to be more involved in the future of the language. Following the web-seminar, a blog article [14] has been released to deepen the topics which arose during the Q&A. **Material**: Links to the recording and documentation are available on the Service Portal (Figure 16).
- "AMD GPU Training" [PRF] (**online** from 2021-06-21 to 2021-06-22): The EXCELLERAT Interest Group AMD provided an USTUTT-internal training as a deep dive into the CDNA (AMD GPU) architecture and ecosystem. Participants also got an overview of the HIP software.
- "Data compression of numerical data sets with the BigWhoop library" in "From Machine Learning to Deep Learning: a concise introduction" [DAT] (**online** from 2021-06-28 to 2021-06-30): The introduction to BigWhoop with hands-on sessions was part of an USTUTT machine and deep learning three-day course. The BigWhoop library for data compression is part of the EXCELLERAT Data Exchange and Workflow Portal[12]: As an efficient data reduction tool, BigWhoop can be applied to generic numerical datasets to minimize I/O bottlenecks and to optimise data storage. **Material** about BigWhoop is available on the Service Portal. Upon request, material about the whole course can be made available.
- "Training on Nek5000 at ENCCS" [CFD] (**online** on 2021-08-17): This course has been a collaboration between the Swedish National Competence Centre ENCCS and EXCELLERAT. KTH experts gave an overview of the capabilities and features of the EXCELLERAT flagship code Nek5000[13] with hands-on sessions to install, run and post-process some simple flow cases. **Material**: Links to the recording and documentation, including the exercises, are available on the Service Portal.
- "Scientific Visualization with COVISE and Vistle" [VIS] (**on-site at USTUTT** from 2021-09-30 to 2021-10-01): This two-day course was about the tools Vistle[14] and COVISE [15] to visualise simulation results on desktop and in Augmented Reality and Virtual Environments. The COVISE part of the course is part of the regular training portfolio of HLRS at USTUTT. Within EXCELLERAT, instead, in-situ capabilities depending on commonly used in-situ interfaces have been developed for Vistle. The EXCELLERAT extension of the course on Vistle is therefore new since 2021 (see also D5.5). **Material**: A Vistle tutorial, slides and further material of the course can be accessed from the Service Portal.
- "OpenFOAM training workshop" [CFD] (**online** from 2021-12-07 to 2021-12-08): Within another **collaboration** between ENCCS and EXCELLERAT, KTH experts gave an introductory training in OpenFOAM, diving also into the topic of CFD in HPC. The goal has been to pave the way in using the more specialized software that is the focus of EXCELLERAT (namely Nek5000, see e.g. the "Training on Nek5000"), also for a less specialised audience. Teaching **material** used during the course is publicly available in github, or can be downloaded directly from the Service Portal.

[12] https://services.excellerat.eu/viewcode/9
[13] https://services.excellerat.eu/viewcode/7
[14] https://services.excellerat.eu/viewcode/2

Further scheduled events from M41 on are finally listed below:

- “Numerical methods for Large Eddy Simulation” [CFD] (**online** from 2022-04-11 to 2022-04-15): This is a second edition of the 2021 CERFACS training on AVBP mentioned above.
- “Data analytics for engineering data using machine learning ML4SIM” [DAT] (**online** from 2022-05-09 to 2022-05-11 and from 2022-05-23 to 2022-05-25): After evaluating the participants’ feedback of the previous course instances in order to optimise the learning outcome, the online Fraunhofer SCAI workshop will be offered two more times as a three-day course. The main target group of these further instances are industrial end-users, e.g. belonging to the EXCELLERAT Interest Groups or to the EXCELLERAT partners’ stakeholders.
- “Scientific Visualization with COVISE and Vistle” [VIS] (**on-site at USTUTT** from 2022-09-29 to 2022-09-30): This will be the 2022 edition of the analogous visualisation course described above.

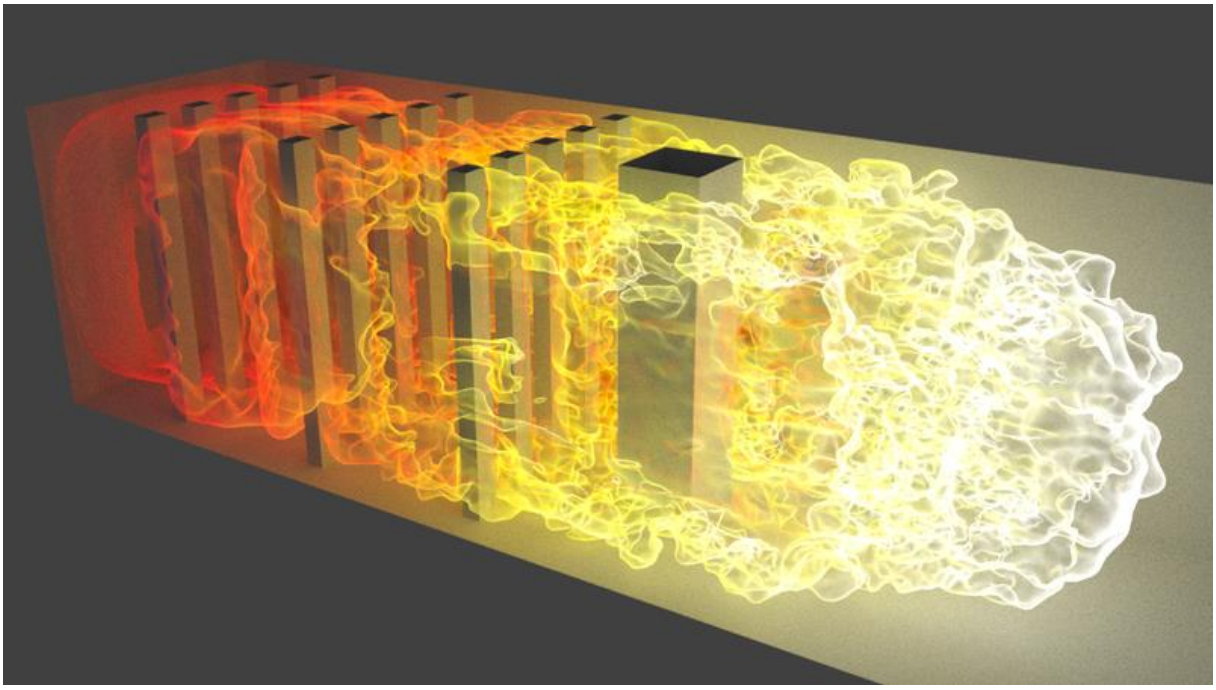

Figure 16: Explosion simulation using AVBP (courtesy CERFACS).

4.3 KPI Analysis

In D5.1 (Section 2.3), five functional and quality-related KPIs have been proposed, which can be seen in Table 5. The quality-related KPIs (5.4.4 and 5.4.5) refer to the ACSI Score (American Customer Satisfaction Index) also defined in D5.1. In the table, both the targets as defined in D5.1 and the current score until M40 are reported. According to the definition, we computed the final target at M42 for KPIs 5.4.1 and 5.4.2, while, for KPIs 5.4.3, 5.4.4 and 5.4.5, the average score over the whole period.

KPI number	Description	Definition	Target	Target at M42	Score at M40
5.4.1	Number of participants in EXCELLERAT training (per year)	*N_participants_per_year*	**≥20** for the first year after MS4 completion, +20% for the final period	52 [in total]	652
5.4.2	Size of developed material in number of lecture units	*Size_training_material*	**≥30** at the end of the first year after MS4 completion, +50% at the project conclusion	45 [in total]	40
5.4.3	Gender ratio of the participants	*R_female_male_quotient*	50%	50%	12%
5.4.4	ACSI score: Satisfaction about EXCELLERAT training outcome (via survey/feedback sheet) as average per year	N_{ACSI}*_training_per_year*	≥80	≥80 [total average]	81,38
5.4.5	ACSI score: Quality and uptake of developed EXCELLERAT material also by non-EXCELLERAT stakeholders as average per year (via survey/feedback sheet)	N_{ACSI}*_G_quality_material*	≥80	≥80 [total average]	78,81

Table 5: KPIs for Task 5.4

KPI 5.4.1 ("Number of participants") is largely satisfactory. As mentioned in D5.5, online events could reach a larger audience than the on-site events which were expected in D5.1 before the pandemic. Moreover, several channels to successfully promote training courses have been established throughout the project, also depending on the target group of each activity: the FocusCoE network of CoEs, the EuroCC and CASTIEL networks, the stakeholders' circle of each EXCELLERAT partner.

The Service Portal has been acting as the Knowledge Hub of the CoE since the release in M14, allowing to obtain a comprehensive and structured overview of the developed training material for each course. The number of material units (KPI 5.4.2) steeply increased year by year (7 in Y1, 3 in Y2, 33 in the last period) almost reaching at M40 the ambitious target of 45 material units.

The goal of ≥50% female participants could not be reached, halting at 12% in total at M40 and reaching a peak of 20% in Y2. The 50% goal is still an arguably ambitious target for training in engineering and technical domains. Some events with large participation (e.g., the AMD courses, the Fortran web-seminar) had a largely male audience, while some events with a smaller, more selective audience (e.g., the Joint CoE Workshop) had almost one third of female participants. Strategies to improve on gender equality adopted by other CoEs or by EuroCC and CASTIEL (e.g., encourage non-male presenters, allocate places for non-male participants in advance) will be more rigorously put into practice in the next funding period (if granted).

As it was already the case in the previous report D5.5, the satisfaction targets (KPIs 5.4.4 and 5.4.5) were overall reached (or very narrowly missed). An EXCELLERAT-specific online survey has been elaborated to evaluate the project’s training events and implemented directly in the EXCELLERAT internal service BSCW [16]. However, EXCELLERAT contributions to external (or co-organised) activities relied on different tools, for example the customisable PRACE-questionnaires via Indico [17].

5 HPC Provisioning

Task 5.5 "HPC Provisioning" implements the CoE hub for internal HPC resources provisioning, to support other WPs' activities. The task compiles and updates the list of available services, provides a point of contact for all CoE members to access specific services, support and documentation, and makes sure that the requests coming from the CoE are addressed by the service provisioning functions of the HPC centres. In other words, this task implements a link between the CoE and the European HPC infrastructure, including both PRACE members and EuroHPC pre-exascale prototype owners.

5.1 Provisioned Resources and Allocation

Computational resources have been provided to the Consortium, specifically by PRACE [18]. 0.5% of the total resources available for each PRACE call are reserved for the CoEs as selected by the European Commission under the E-INFRA-5-2015 and INFRAEDI-02-2018 calls for proposals. EXCELLERAT asked for the computational resources available for the Calls 18 to 24. The assigned resources are reported in Table 6 below. For details regarding the architecture's specifications of the involved HPC systems, please see Table 8 in Annex 3: Internal Resources.

Awards (cores/hours)	PRACE 18	PRACE 19	PRACE 20	PRACE 21	PRACE 22	PRACE 23	PRACE 24	Total
Marconi BDW	80,000	45,000						125,000
Marconi KNL	1,000,000	750,000						1,750,000
Marconi100			1,620,000	875,000	275,000	300,000		3,070,000
HAWK		1,150,000	1,000,000	550,000	280,000	224,500		3,204,500
JUWELS Cluster	100,000	175,000	70,000	45,000	40,000	55,000		485,000
JUWELS Booster				55,400	38,000	16,000		109,400
Joliot Curie Rome		1,715,000	950,000	350,000	150,000	122,000	650,000	3,937,000
Joliot Curie KNL		150,000	1,120,000	240,000	93,750	176,500	150,000	1,930,250
Joliot Curie SKL	350,000	180,000	170,000	170,000	87,000	107,250	244,600	1,308,850
MareNostrum4	700,000	240,000	150,000	240,000	100,000	92,500		1,522,500
SuperMUC-NG			302,500	86,000	65,000			453,5000
Piz Daint	450,000	850,000	935,000	250,000	510,000	50,000		3,045,000

Table 6: HPC Resources allocated to Call PRACE 18-24 (cores-hours)

The accounts are open on whole the different clusters and the access is granted on request to the various partners.

The wiki page of EXCELLERAT has been updated with all the up-to-date information regarding the HPC access[15]. Specifically, a new section has been added, *HPC Provisioning*, where HPC resources available to the Consortium provided by PRACE are updated, see screenshot below:

HPC Provisioning [edit]

- HPC resources on the various cluster for the Consortium provided by PRACE (Task 5.5)
- Last update: Apr. 2022
- The following page summarizes the different HPC Access procedures

Awards (Allocated)	Galileo Broadwell (CINECA): core/hours	Marconi100 (CINECA): local core/hours-(node/hours)	Hawk (HLRS) core/hours	JUWELS (Julich) core/hours	Joliot Curie AMD (Rome)-(core/hours)
PRACE Call 18	80.000	50.000 - (1.562)	--	--	--
resource used	148.860	0	--	--	--
resource used (Percentange %)	**186**%	0%	--	--	--
Start / End allocation	19-08-2019 / 18-08-2020	10-07-2019 / 08-08-2020	--	--	--
Account	Pra18_Excell_0	Pra18_Excelera_0	--	--	--
PRACE Call 19	45.000	37.500 - (1.171)	1.150.000	--	150.000
resource used	213.546	5.660	--		--
resource used (Percentange %)	**475**%	15%	--	--	--
Start / End allocation	30-10-2019 / 06-01-2021	30-10-2019 / 06-01-2021	--	--	--
Account	Pra19_Exceller_0	Pra19_Exceller_1	--	--	RACOE006
PRACE Call 20	--	147.272 (4602)	1.000.000	275.000	Cumulative with Call 21
resource used	--	0	0	0	--
resource used (Percentange %)	--	0%	0%	0%	--
Start / End allocation	--	24-04-2020 / 26-04-2021	/ 31/10/2020	01-07-2019 / 31-03-2021	--
Account	--	Pra20_Exceller	--	prcoe05	--
PRACE Call 21	--	79.546	550.000	100.000 (JUWELS Cluster) 55.400 (JUWELS Booster)	1.255.138
resource used	--	1478	0	2.8 (JUWELS Cluster) 6798.7 (JUWELS Booster)	272943
resource used (Percentange %)	--	1.9%	0%	0% (JUWELS Cluster) 12.3% (JUWELS Booster)	22%
Start / End allocation	--	08-10-2020 / 31-03-2021	/ 31/10/2020	01-07-19/30-09-2021	16-10-2020/01-04-2021
Account	--	Pra21_EXCELLERAT	--	prcoe05	RACOE006
PRACE Call 22	--	25.000 (781)	280.000	40.000 (JUWELS Cluster) 38.000 (JUWELS Booster)	--
resource used	--	416	172137.4	11.243 (JUWELS Cluster) 13818.1 (JUWELS Booster)	16305
resource used (Percentange %)	--	1.7%	61.5%	20.29% (JUWELS Cluster) 24.9% (JUWELS Booster)	10.87%
Start / End allocation	--	01-04-2021/30-09-2021	--	01-10-2021/31-03-2022 (JUWELS Cluster) 23-11-2020/31-03-2022 (JUWELS Booster)	/01-10-2021
Account	--	Pra22_Excellerat	PrcCoEExce	prcoe05	RACOE006
PRACE Call 23	--	27274 (852)	224500	275000/Cumulative (JUWELS Cluster) 55400/Cumulative (JUWELS Booster)	122000
resource used	--	0	527097	15019/C (JUWELS Cluster) 33262/C (JUWELS Booster)	19966
resource used (Percentange %)	--	0%	**235**%	5%/C (JUWELS Cluster) 60%/C (JUWELS Booster)	16%
Start / End allocation	--	01-10-2021/30-09-2022	/31-03-2022	01-10-2021/31-03-2022 (JUWELS Cluster) 23-11-2020/31-03-2022 (JUWELS Booster)	/31-03-2022
Account	--	Pra23_Exceller	PrcCoEExce	prcoe05	RACOE006
PRACE Call 24	--	--	--	--	650000
resource used	--	--	--	--	0
resource used (Percentange %)	--	--	--	--	0%
Start / End allocation	--	--	--	--	25-04-2022/01-04-2023
Account	--	--	--	--	RACOE006
Upon Request/Special Petition	--	--	--	--	--
resource used	--	--	--	--	--
resource used (Percentange)	--	--	--	--	--
Start / End allocation	--	--	--	--	--
Account	--	--	--	--	--

Figure 17: Screenshot of HPC Provisioning from the wiki

[15] https://kb.hlrs.de/excellerat/index.php/EXCELLERAT_Wiki

5.2 KPI Analysis

Table 7 presents the KPIs for this task, as defined in D5.1.

In particular, we can mention one project that has been approved. STRIKE Project, PI Cerfacs. Budget requested: 30 M core/hours. Moved to Irene KNL, with budget reduction.

KPI number	Description	Definition	Target	Score at M40
5.5.1	Percentage of approved proposals in competitive calls for resources for EXCELLERAT projects	*N_prj_appr / N_prj_submitted * 100*	>50%	100
5.5.2	Number of yearly allocated cpu_hours for EXCELLERAT projects	*N_cpu_hours*	$\geq 10^6$ for PY2, $\geq 10^7$ for PY3	ca. 65 M

Table 7: KPIs for Task 5.5

6 Conclusion

With a stable version of the service portal available, the successful submission of this deliverable D5.7 and its companion D5.6 at M42, last of a series of preparatory deliverable reports, the contribution of WP5 to Milestone MS6 "*Final Reports of all project outcomes and project close*" is achieved.

In the last reporting period, the service portfolio was expanded, implementing the directives obtained by the other work packages, and a significant effort was devoted in collaboration with WP7 to expand the external user base.

Most of the KPIs were met, in particular those related to usage and satisfaction. A most notable exception is the gender balance of trainees, but this reflects a well-known general situation, that the EXCELLERAT project alone has limited power of control over.

All partners expressed their commitment keeping the service portal active after the end of the Project, and to exploit and expand the offered services.

7 References

[1] "EXCELLERAT Service Portal," [Online]. Available: https://services.excellerat.eu.

[2] "PAinG-Flow website," [Online]. Available: http://flumacs.site.uniroma1.it/people/battista.

[3] "HiDALGO CoE website," [Online]. Available: https://hidalgo-project.eu/. [Accessed 17 12 2019].

[4] "HiDALGO success story," December 2021. [Online]. Available: https://www.excellerat.eu/success-story-improving-the-hidalgo-coes-urban-air-pollution-pilot/?utm_source=mailpoet&utm_medium=email&utm_campaign=excellerat-newsletter-10.

[5] "STREAmS website," [Online]. Available: https://github.com/matteobernardini/STREAmS.

[6] "The finite volume method for solid mechanics (project)," [Online]. Available: https://www.researchgate.net/project/The-finite-volume-method-for-solid-mechanics.

[7] "exaFOAM project website," [Online]. Available: https://exafoam.eu/.

[8] "DrivAer website," [Online]. Available: https://www.mw.tum.de/en/aer/research-groups/automotive/drivaer/.

[9] "SZE website," [Online]. Available: https://math.sze.hu/en_GB/home.

[10] "AIA website," [Online]. Available: http://www.aia.rwth-aachen.de/en/forschung.

[11] "CASTIEL project website," [Online]. Available: https://www.castiel-project.eu/.

[12] "HPC in Europe website events page," [Online]. Available: https://hpc-portal.eu/upcoming-events-courses.

[13] "The AMD ROCm Platform website," [Online]. Available: https://rocmdocs.amd.com/en/latest/index.html.

[14] "LNG event blog article," [Online]. Available: https://cerfacs.fr/coop/webinar-modern-fortran.

[15] "COVISE website," [Online]. Available: https://www.hlrs.de/covise.

[16] "BSCW website," [Online]. Available: https://www.bscw.de/en/. [Accessed 17 12 2019].

[17] "Indico website," [Online]. Available: https://getindico.io.

[18] "PRACE RI website," [Online]. Available: http://www.prace-ri.eu/. [Accessed 17 12 2019].

[19] P. Gualtieri, F. Picano, G. Sardina and C. Casciola, "Exact regularized point particle method for multiphase flows in the two-way coupling regime," *Journal of Fluid Mechanics,* no. 773, pp. 520-561.

[20] X. Zhu, E. Phillips, V. Spandan, J. Donners, G. Ruetsch, J. Romero, R. Ostilla-Mónico, Y. Yang, D. Lohse, R. Verzicco, M. Fatica and R. J. A. M. Stevens, "AFiD-GPU: A versatile Navier–Stokes solver for wall-bounded turbulent flows on GPU," *Computer Physics Communications,*, no. 229, p. 199–210, 2018.

[21] P. Costa, E. Phillips, L. Brandt and M. Fatica, "GPU acceleration of CaNS for massively-parallel direct numerical simulations of canonical fluid flows.," *Computers & Mathematics with Applications,* no. 81, p. 502–511, 2021.

[22] M. Bernardini, D. Modesti, F. Salvadore and S. Pirozzoli, "STREAmS: A high-fidelity accelerated solver for direct numerical simulation of compressible turbulent flows.," *Computer Physics Communications,* vol. 263, no. 107906, 2021.

8 Annex 1: EXCELLERAT Onboarding Final report: PAinG-Flow

Organisation (team)
FLUMACS

Members

Prof. C. M. Casciola - University of Rome Sapienza, Mechanical and Aerospace Department

Prof. P. Gualtieri - University of Rome Sapienza, Mechanical and Aerospace Department

Prof. F. Battista - University of Rome Sapienza, Mechanical and Aerospace Department

Dr. Francesco Salvadore, CINECA

Abstract

The project purpose is the implementation of the Exact Regularized Point Particle (ERPP) method to exert the particle/fluid momentum coupling in multiphase turbulent flows on a code able to run on supercomputers with GPUs.

The engineering challenge

The solid phase is modeled by evolving each single particle in a Lagrangian way with the same temporal scheme of the Eulerian phase. The method aims at regularizing the force that a small particle exerts on the fluid on a physical ground. As a particle moves along its trajectory, it generates a small-scale vorticity field that can be evaluated in a closed form. Indeed, the disturbance flow is shown to obey the unsteady Stokes problem with appropriate boundary conditions to account for the wall effects. The disturbance field produced on the scale of the particle is then regularized by the fluid viscosity. Hence, the process of vorticity generation and viscous diffusion allows to model on a physical ground without any “ad hoc” numerical artifact the inter-phase momentum coupling.

Application software code details

The code will solve the incompressible Navier-Stokes equation in wall-bounded and free-shear conditions. The fluid Eulerian phase is solved by discretizing the Navier-Stokes equation on staggered grid with a centered second-order finite difference scheme in space and with a third-order four-steps Low-Storage Runge-Kutta scheme in time. The code is written in Fortran 90 language and uses MPI library for parallelization. The GPU version is expected to be based on the CUDA Fortran paradigm.

Activity summary and outcomes

The goal of the project is to build a code for an incompressible turbulent flow in a channel with two periodic directions (streamwise and spanwise) and a non-homogeneous direction normal to the walls where the particles backreaction is modeled using the ERPP method (see [19]). The code must be able to take advantage of modern NVIDIA GPUs to simulate particularly challenging physical cases. We focused on physical cases characterized by a very large number of particles. In fact, the turbulence alteration due to the particles is determined by two dimensionless groups, namely the particle mass loading and the particle Stokes number. It’s worth observing that the Stokes number represent the response time of the particles when the background flow changes in time as in a turbulent flow. In the limit of small Stokes number, the particles behave like tracers, i.e. they follow the fluid particle trajectories. However, in two-way coupled simulations where a momentum exchange between the particles and the fluid occurs, it is not trivial the effect of the other control parameter, i.e., the mass loading. For instance, it is interesting to understand the behavior of turbulence when particles with progressively smaller and smaller Stokes numbers are considered keeping constant the mass loading. In fact, when the Stokes number is reduced the particles tend to behave like tracers (hence one would expect a negligible interaction with the carrier flow) but on the other hand the mass loading is fixed calling for a considerable momentum exchange between the two phases. This limit is achieved by considering particles with a smaller and smaller diameter (limit of the Stokes number to zero) [19]and by progressively increasing the number of particles to compensate the reduction of the mass of a single particle. It turns out that the study of this particular, though relevant behavior of a particle laden flow, requires the simulation to handle of an increasing number of particles hence calling for highly efficient and hybrid parallelization strategies.

In the first phase of the project, we started considering a pre-existing code that already implements the desired physics (configuration and algorithms). We then performed a requirements and feasibility analysis in view of the project objectives. The initial code, written in Fortran and parallelized with MPI, is working but it has some significant efficiency limits that make practically impossible the simulation of very demanding cases. In addition, some features, although valid, make the code particularly difficult for a GPU porting. We summarize the critical features of the initial code:

- The original code uses FFT-2D transforms both to solve the Poisson solver and to perform derivatives along the two periodicity directions, while using finite differences only to compute derivatives along the wall-normal direction. This strategy is particularly inefficient when considering GPU-enabled code since FFTs require onerous data transpositions between MPI processes.
- The original code uses an MPI decomposition in physical space consisting of a homogeneity component and the wall-normal component. When considering particles in turbulence, due to turbophoresis, they tend to accumulate at the wall. Consequently, the adopted MPI decomposition leads to marked load unbalancing and therefore performance degradation.

We summarize then how the two problems exposed can be resolved in an updated code:

- The updated code needs to use FFT-2D only for the solution of the Poisson problem while all derivatives can be computed using finite-difference methods. This will result in the need to use staggered grids also in x and y, in addition to z as it was the case in the initial code
- The updated code must use an MPI decomposition in physical space only in the directions of homogeneity. By leaving the wall-normal direction undivided, turbophoresis will not produce significant unbalancing effects.

In addition to the two fundamental points now presented, it arises the fact that the original code has a weakly modular structure – e.g., for particle handling – while, in view of GPU porting, a more robust structure of the various parts of the code is preferable. Therefore, considering the various factors and the significant changes required, it was decided to proceed with a complete rewrite of the code, instead of updating the initial code. However, some parts of the initial code were still used at least at the algorithmic level.

For the choice of the GPU paradigm, several options were evaluated (e.g., OpenACC, OpenCL) and CUDA-Fortran was chosen. This constitutes a good compromise between fast code porting – thanks to the help of automatic kernels – and achieving good performance where required thanks to the full functionality of CUDA features. The use of CUDA-Fortran is particularly convenient also for the use of some libraries already ported to GPUs for codes in the same research area (see [20], [21]).

Another choice was to produce a single code capable of running on both CPUs and GPUs. The strategy of maintaining a single source is an important advantage in view of the maintainability of the code. The different compilation modes allow to manage the different versions of the code. In particular, we chose an approach similar to the one used in [22].

The original code uses the *p3dfft* library for parallelization of FFTs. In the new code, we chose instead to use the *2decomp* library since GPU porting widely used in other community codes ([20], [21]) is available for this library. GPU porting of parallel FFT transforms is particularly critical as it requires to properly implement the overlap between communication, CPU-GPU copies, and computation to decrease the impact of transpositions on global elapsed times. The new code therefore considers, as anticipated, an MPI decomposition in physical space along the x and y directions, while the wall-normal z direction is not decomposed.

As for the implementation of the particles and the ERPP method, this is instead, to our knowledge, the first existing GPU implementation. In a good number of cases (about 30), the loops were adapted by simply using the automatic cuf kernels. In some cases, it was necessary to explicitly write CUDA kernels due to inherent limitations of the automatic kernels. In particular, the use of the CUDA function *atomicAdd*, which is needed to accumulate particle feedback in the Eulerian field, required writing a kernel explicitly. Explicit kernels had to be written in other cases as well for other reasons. In writing the new code, importance was given to minimize branches in the code between CPU and GPU versions. In this regard, several functions were written so that they could be called from both the CPU version of the code and the GPU version. Although this required some effort in design and implementation, the final result pays off in terms of readability and fewer lines of code.

Another significant aspect of the new code is that, unlike the original code, all memory is dynamically allocated. This means that a single executable can handle all possible physical configurations by simply modifying an input file. Through this input it is for example possible to choose the domain size, the number of particles, the type of particles (one-way or two-way) or even disable particle computation altogether. An automatic mesh management along the wall-normal direction has also been implemented starting from the Reynolds number specified in the input. This allows the code to be executed without any preparatory step but only by setting the input file. Finally, the post-processor that extracts statistics from the saved results has also been heavily rewritten.

The new code is working and complete with all its components, even if only the GPU compilation has been adequately tested. As a validation some runs already performed in the past with the original code were executed, and it was verified that the resulting statistics are completely equivalent. The code was then used to simulate a certain range of Stokes numbers. Two example figures are shown below.

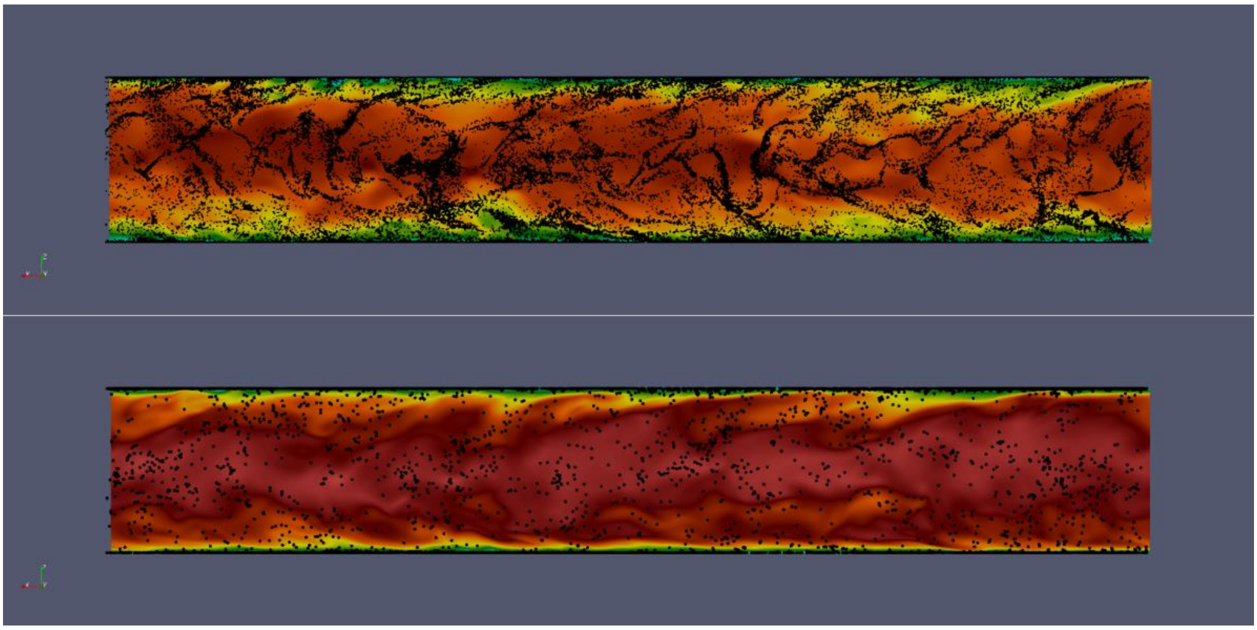

Figure 18: PAinG-Flow simulations examples

Two slices are shown displaying streamwise (horizontal) and wall-normal (vertical) direction. The top figure presents the Stokes=10 case (4 million particles) while the bottom figure has Stokes=10 (half a million particles). For Eulerian field the friction Reynolds number is 180. In the figures is clearly visible the turbophoresis phenomenon.

The most computationally challenging case that has been performed is the one at Stokes=1.6 which requires the simulation of 61 million particles, a high number even considering the two-way ERPP modeling. The case was run at 16 GPUs and required an elapsed time of 0.44s per time-step. The case at Stokes=5 (11 million particles) was instead shot at 8 GPUs requiring an elapsed time of 0.40s for each time-step. This shows that the code has reasonable performance in this range of use cases, even if it is not a true scalability analysis. In fact, in addition to the particle simulation, the evolution of the Eulerian part should also be considered. In the physical cases currently of most interest, the Eulerian part is characterized by a small number of points. However, the main goal of the project was to obtain a code capable of simulating ranges like those simulated, and this, also thanks to the use of modern accelerated architectures, is now fully possible. In the future, the code is however ready to handle even more massive Eulerian computational grids.

As far as the analysis of the physical results is concerned, the work is still in progress, and in particular a paper is in preparation to be submitted to a high-level fluid dynamics journal.

9 Annex 2: Assessment of Training Activities 2022

The description of the methodology used for the assessment has been adapted from CASTIEL D3.2 "*Training, Twinning and Mentoring Plans and Achievements*", Section 2.2. Authors of this specific section are Siegfried Höfinger, TU Wien, and Martina Blazkova, BSC, who contributed to develop the procedure for the CASTIEL analysis. The data for the EXCELLERAT assessment have been collected both by EXCELLERAT WP5 and by CASTIEL.

Training by eight EXCELLERAT partners and training by EXCELLERAT have been compared and analysed. Purpose of such an analysis is to assess how EXCELLERAT training counterbalances the offer already provided by the EXCELLERAT partners, how such an offer is enhanced through EXCELLERAT and which gaps need to be filled. The analysis encompasses training events in two academic years, as well as the joint data during the two periods:

- 09/2019-08/2020
- 09/2020-08/2021

Data on the training events by 6 EXCELLERAT partners had been collected through the CASTIEL training mapping survey among the National Competence Centres (NCCs) in Q4/2020 (see CASTIEL D3.2):

- BSC (within NCC Spain)
- CERFACS (within NCC France)
- CINECA (within NCC Italy)
- EPCC (within NCC UK)
- HLRS (within NCC Germany)
- KTH (as PDC Center for High Performance Computing within NCC Sweden)

The following 2 EXCELLERAT partners were not involved in the CASTIEL mapping and therefore provided their data directly to EXCELLERAT WP5:

- FRAUNHOFER
- RWTH - AIA

Data on EXCELLERAT training events (5 events between 09/2019-08/2020 and 7 events between 09/2020-08/2021) were collected by EXCELLERAT T5.4.

Unlike the CASTIEL assessment, only three categories were selected for the EXCELLERAT analysis: Scientific Domain, Technical Domain, HPC Topic. The list of the options for each category follows:

- Scientific Domain:
    - Biochemistry, bioinformatics and life sciences
    - Chemical sciences and materials
    - Computer sciences, Computer Engineering, Electrical Engineering, Telecommunications
    - Earth system sciences
    - Economics, finance and management
    - Engineering
    - Fundamental constituents of matter
    - Fundamental physics
    - Linguistics, cognition and culture
    - Mathematics

  - Physiology and medicine
  - Universe sciences / cosmology
  - Other
  - Non scientific domain-specific
- Technical Domain:
  - Artificial intelligence (AI), machine and deep learning
  - Data science and high-performance data analytics (HPDA)
  - High-performance computing (HPC)
  - Numerical libraries and methods
  - Parallel programming
  - Performance engineering
  - Programming Languages
  - Scientific programming
  - Software engineering
  - System administration
  - Virtualisation, containers
  - Visualisation
  - Other
  - Non technical domain-specific
- HPC Topic:
  - Artificial Intelligence, Machine and Deep Learning
  - Big Data, Data Science and High-Performance Data Analytics
  - Data Visualisation
  - Domain Specific Languages for HPC
  - Embedded Supercomputing
  - Fault Tolerance and Resilience
  - Heterogeneous Programming on Accelerators
  - HPC on Exascale Architectures
  - HPC Systems Administration, Operations and Maintenance
  - HPC Technology - Architecture, Hardware
  - HPC Technology - Accelerators (GPUs, FPGAs...)
  - HPC Technology - Networking
  - HPC Technology - Storage
  - Parallel Algorithms
  - Parallel Libraries (Linear Algebra, FFT...)
  - Parallel Programming Models (MPI, OpenMP...)
  - Performance Engineering and Co-design
  - Power Management (Power Wall Issues)
  - Programming Languages and Methods for HPC
  - Quantitative Performance Analysis of HPC Architectures, Algorithms and Software
  - Real-time Supercomputing
  - Scalability and Asynchronous Algorithms
  - Simulation-based Training and Education
  - Other
  - Not relevant

Further disclaimers for the analysis:

- The analysis was conducted before the EXCELLERAT "Training on Nek5000 at ENCCS" in 08/2021, which could therefore not be included.
- Data were collected from CASTIEL throughout Q4/2020, therefore events in 2021 were still being planned.
- Two centres (CERFACS and EPCC) did not distinguish between the two time periods. Therefore, the same data were used for the analyses in the two time periods.

Profiles – the Central Tool Used in this Analysis

A graphical representation has been developed in the form of profiles, characterising the individual categories listed above for EXCELLERAT training (red bars) as opposed to the rest of the training carried out by the other 8 partner institutions (blue bars). An example is given in Figure 19.

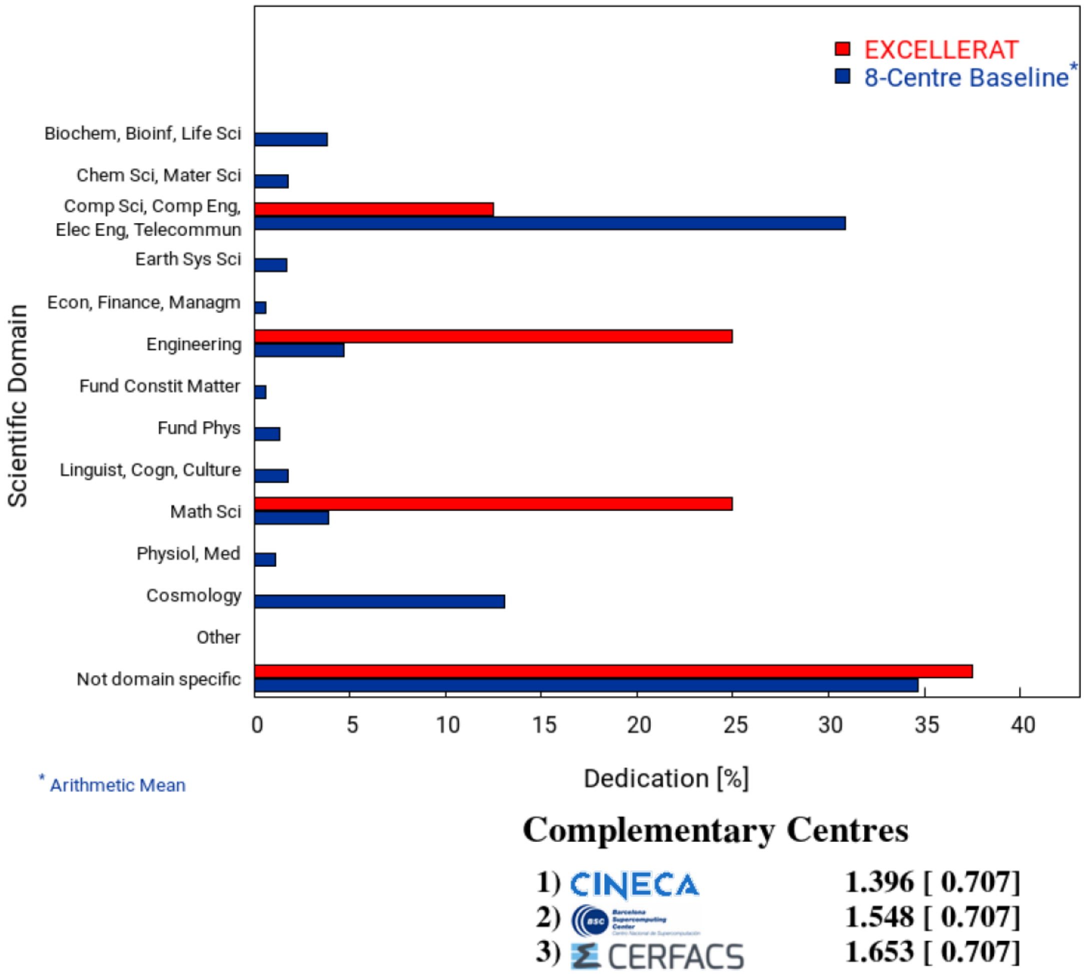

Figure 19: Profile representing the Scientific Domain of EXCELLERAT training and the eight-centre baseline in 09/2019-08/2020.

All available options for this particular category are drawn on the y-axis. The x-axis is given in units of percentages of the total number of mentions of criteria particular option by all various training events. Each training was assigned a label designating specific associations. For example, if a particular training event was tailored for "Biochemistry", then a label of 'a' was assigned. If, on the other hand, the training event had "Chemical Sciences" as Scientific Domain, then a label of 'b' was assigned. Further, if both scientific domains were considered, then both labels, ‘a’ and ‘b’, were assigned. Labels do just refer to individual options (hereafter also termed classes) mentioned on the y-axis.

The first step in computing a profile as shown in Figure 19 is to sum up all mentions of all labels, ‘a’, ‘b’, ‘c’ … altogether which formally represents 100% of a particular training programme. Next, partial sums of all mentions of individual labels are computed separately (i.e., just for labels ‘a’, just for labels ‘b’, …) and put into perspective by relating them to the overall sum. This latter ratio can certainly be expressed in percentages and the results are precisely the meaning of the **red bars** shown in Figure 19.

Since an individual “set of red bars” is obtained for every centre involved, an average score can also be determined representing the **mean of the 8 centres other than EXCELLERAT itself**, which is also included in Figure 19 in the form of blue bars. This average of all the **8 centres** will constitute a reference or baseline, against which the **assessment of the EXCELLERAT training activities** can be carried out. Taken together, all the red and blue bars will form a complete profile and differences between red and blue bars in a particular class will hint at a particular “strength” or a certain degree of “less than average dedication” of **EXCELLERAT training** (or a centre’s training) in that specific class. For example, what can directly be inferred from Figure 19 is that EXCELLERAT training is more focused on Engineering and Mathematics but less focused on Computer Sciences compared to the partners’ average.

Dealing with percentages implies that the sum of all red (or blue) bars will have to result in 100%. So, whenever we detect some overcapacity in a class (red bars are longer than blue bars), it follows that there must also be some “less than average” capacity in another class, otherwise the normalisation could not be upheld (both sets of bars have to sum up precisely to 100%).

Arithmetic Mean

For this analysis, a simple arithmetic mean has been used to derive the “blue bars”: Each centre contributed with a factor of 1/8 to the average baseline. In the CASTIEL analysis however, a weighted arithmetic mean has been used based on the ratio #training events / million citizens of each NCC (see CASTIEL D3.2 for more details).

Identifying Complementary Profiles

The score below the diagram in Figure 19 (“Complementary Centres”) indicates the matching capacity of the centre analysed in the picture, in terms of complementing another centre’s offer for the analysed category. E.g., looking at Figure 19, EXCELLERAT training, in terms of Scientific Domain, would be ideally complemented by the offer of CINECA, FRAUNHOFER, and RWTH. Such a matching was performed by CASTIEL for the NCCs in EuroCC, but it is outside the scope of this EXCELLERAT assessment.

Results

We summarise here the results of the analysis, considering the three categories mentioned above (Scientific Domain, Technical Domain, HPC Topics). The EXCELLERAT offer is compared to the baseline formed by the eight EXCELLERAT centres for the **joint periods** 09/2019-08/2020 and 09/2020-08/2021:

- **Scientific Domain (**Figure 20**)**: The EXCELLERAT offer shows a great amount of non-domain-specific training (comparable to the baseline) and is significantly above the baseline in the domains of engineering and mathematics. This was expected, based on the engineering vocation of EXCELLERAT. Keeping in mind that all applying domains of a training course had to be selected during the evaluation, it does not surprise that mathematical aspects dominate, since they are also a prerequisite or a necessary introductory step in most EXCELLERAT courses (e.g., in a CFD course). Moreover, it is expected that the option "non-domain specific" would be selected whenever the specific topic is lacking from the list. In fact, topics of EXCELLERAT training are often specialised e.g. on a particular code or library, and cannot be found in this taxonomy.

- **Technical Domain (**Figure 21**)**: The EXCELLERAT offer positions itself above the baseline with respect to numerical libraries and methods, performance analysis, scientific programming, software engineering and visualisation. These topics are all pertinent to EXCELLERAT applications, and EXCELLERAT is expected to fill in training gaps in these fields. On the other hand, several EXCELLERAT partners offer training on generic AI, Data Science and HPC topics, which justifies the EXCELLERAT ranking below the baseline in these domains.
- **HPC Topic (**Figure 22**)**: EXCELLERAT outperforms the baseline in the topics of data visualisation, accelerators, HPC on exascale architectures, parallel libraries, performance engineering and analysis: These are all core-topics of the EXCELLERAT mission. Among the categories where EXCELLERAT is below the average, we can mention "parallel algorithms and parallel programming models", and "programming languages and methods for HPC". These HPC topics are well covered by the EXCELLERAT partners in their regular training schedule, while EXCELLERAT specialises in HPC for exascale.

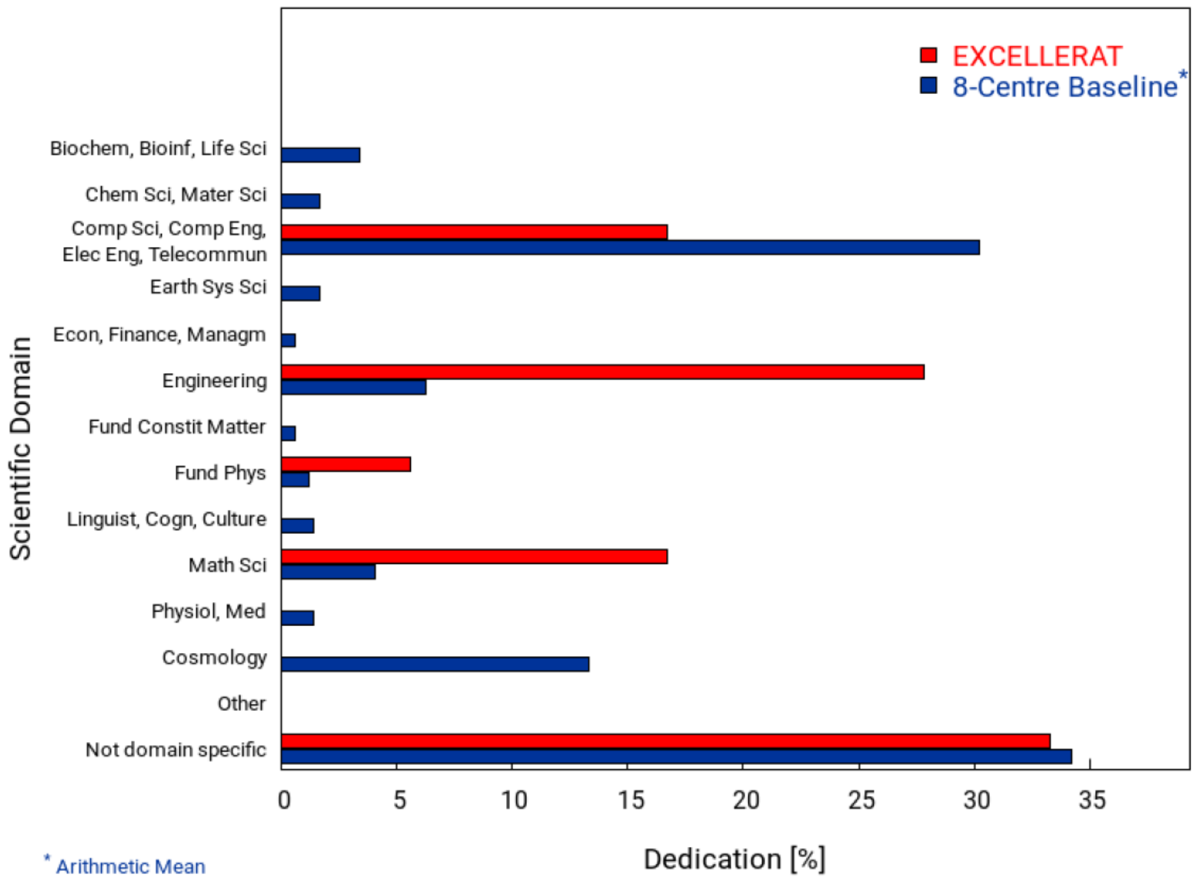

Figure 20: Profile representing the Scientific Domain of EXCELLERAT training and the eight-centre baseline in 09/2019-08/2021.

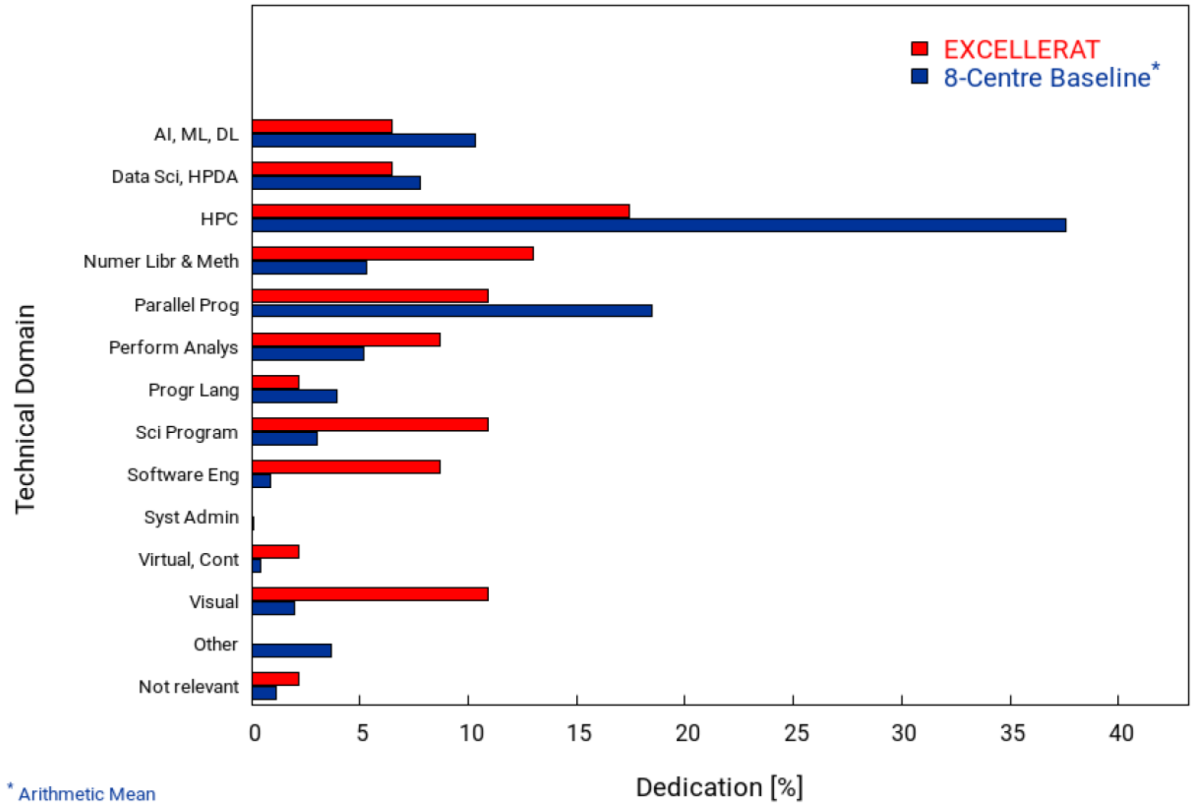

Figure 21: Profile representing the Technical Domain of EXCELLERAT training and the eight-centre baseline in 09/2019-08/2021.

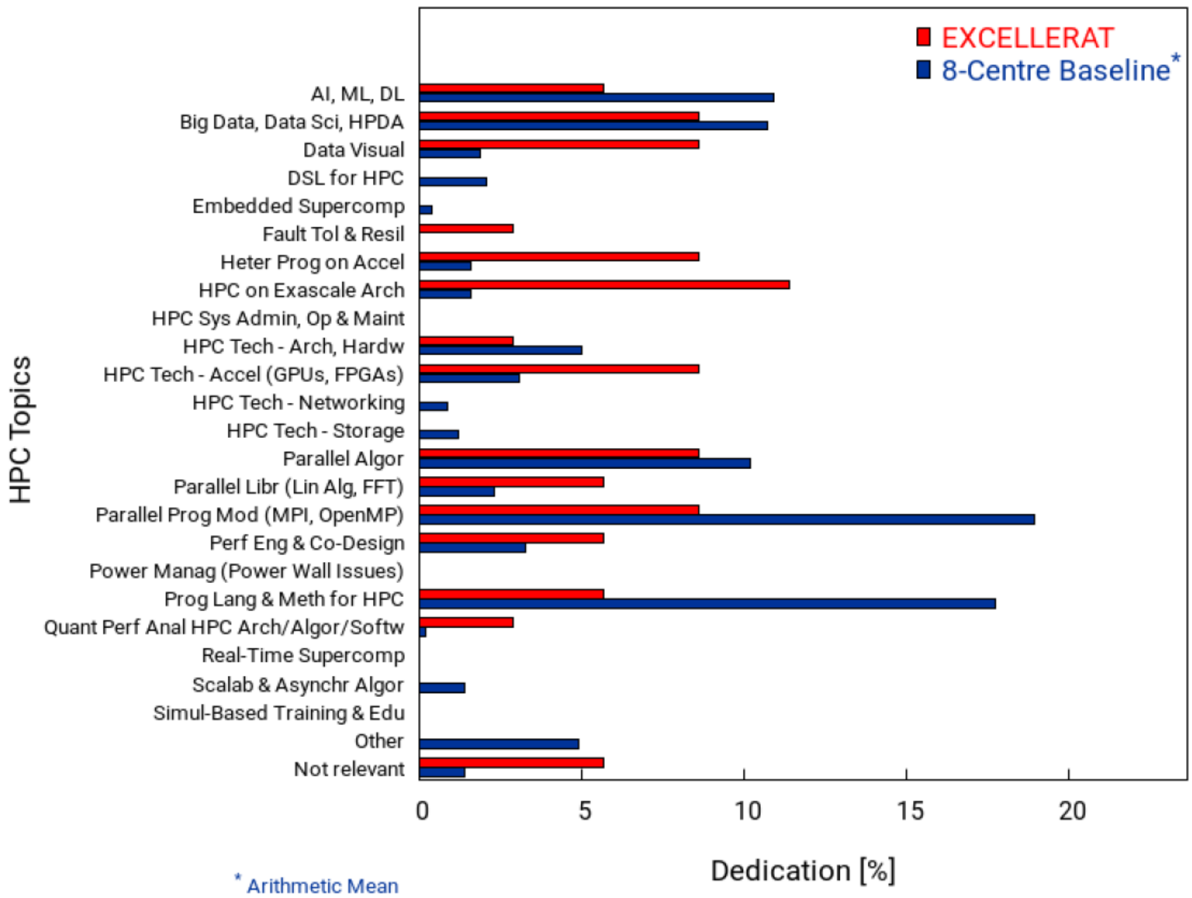

Figure 22: Profile representing the HPC Topics of EXCELLERAT training and the eight-centre baseline in 09/2019-08/2021.

10 Annex 3: Internal Resources

The resources available to the Consortium are continuously updated. The up-date list of Tier-0 resources is provided by the most recent Prace Call. The Table reported below is extracted from the Technical Guidelines of the Prace Call 22[16].

		HAWK	***Joliot-Curie Rome***	***Joliot-Curie SKL***	***Joliot-Curie KNL***	***JUWELS Cluster***	***JUWELS Booster***	***Marconi100***	***Mare Nostrum 4***	***Piz Daint***	***SuperMUC-NG***
System Type		HPE	Bull Sequana	Bull Sequana	Bull Sequana	Bull Sequana	Bull Sequana	IBM Power 9 AC922 Whiterspoon	Lenovo	Hybrid Cray xC50	Lenovo ThinkSystem
Compute	Processor type	AMD Epyc Rome	AMD Epyc Rome	Intel Xeon Platinum 8168 2.7 GHz	Intel Knights Landing	Intel Xeon Skylake Platinum 8168	AMD EPYC Rome	2 * IBM POWER9 AC922 at 3.1 GHz per node	Intel Xeon Platinum 8160 2.1 GHz	Intel® Xeon® E5-2690 v3 @ 2.60GHz (12 cores)	Intel Skylake Xeon Platinum 8174
	Total nb of nodes	5 632	2 292	1 656	828	2 511	936	980	3 456	5 704	6480
	Total nb of cores	720 896	293 376	79 488	52 992	120 528	44928	31 360	165 888	68 448	31 1040
	Nb of accelerators /node	n.a.	n.a	n.a.	n.a.	n.a.	4	4 GPU per node	n.a.	1 GPU per node	n.a.
	Type of accelerator	n.a.	n.a	n.a.	n.a.	n.a.	NVIDIA next Gen.	NVIDIA® Volta® V100, Nvlink 2.0, 16GB	n.a.	NVIDIA® Tesla® P100 16GB	n.a.
Memory	Memory / Node	256 GB	256 GB	192 GB	96 GB DDR4 + 16 GB	96 GB	512 GB DDR4-3200	256 GB DDR4 + up to 1.6 TB Optane Memory per node	96 GB (200 nodes with 384GB)	64 GB	96 GB
Network	Network Type	Infiniband HDR	Infiniband HDR 100	Infiniband EDR	BULL BXI	InfiniBand EDR	InfiniBand HDR	Mellanox Infiniband EDR	Intel Omni-Path Architecture	Cray Aries	Intel Omni-Path Architecture
	Connectivity	9D enhanced Hypercube	Dragonfly +	Fat Tree	Fat Tree	Fat Tree	Dragonfly +	DragonFly+	Fat Tree	Dragonfly	Fat tree within island (786 nodes) pruned tree between islands

[16] https://prace-ri.eu/wp-content/uploads/Technical_Guidelines_Call_22.pdf

		HAWK	***Joliot-Curie***	***JUWELS***	***Marconi100***	***MareNostrum 4***	***Piz Daint***	***SuperMUC-***
Home file system	type	NFS	NFS	GPFS	GPFS	GPFS	GPFS	GPFS
	capacity	100 TB	0.5 TB	2.8 TB	200 TB	32 TB	160 TB	256 TB
Work file system	type	Lustre	Lustre	GPFS	GPFS	GPFS	GPFS	GPFS
	capacity	25 PB	9.2 PB	2.3 PB	3 PB	4.3 PB	6.3 PB	33 PB
Scratch file system	type	n.a.	Lustre	GPFS	GPFS	GPFS	Lustre	GPFS
	capacity	n.a.	5.2 PB	9.1 PB	2 PB	8.7 PB	8.8 PB	17 PB
Archive	capacity	On demand	On demand	On demand	On demand	n.a.	n.a.	On demand
Minimum required job size	Nb of cores	8192	1 024		2 nodes	1024	6 nodes	960

Table 8: Updated list of internal available HPC resources