



SCALABLE PARALLEL ASTROPHYSICAL CODES FOR EXASCALE

Collaboration Plan with Complementary Grants

Deliverable number: D6.1

Version 1.0/1.0



Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441

Project Information

Project Acronym: SPACE
Project Full Title: Scalable Parallel Astrophysical Codes for Exascale
Call: Horizon-EuroHPC-JU-2021-COE-01
Grant Number: 101093441
Project URL: <https://space-coe.eu>

Document Information

Editor:	Andrea Mignone - UNITO
Deliverable nature:	Report (R)
Dissemination level:	Public (P)
Contractual Delivery Date:	30.06.2023
Actual Delivery Date	30.06.2023
Number of pages:	24
Keywords:	Collaboration Plan, Cooperation, CASTIEL2
Authors:	Marisa Zanotti – ENGINSOFT Andrea Mignone – UNITO Lubomir Riha – IT4I@VSB Eva Sciacca – INAF
Peer review:	Samuel Tootle (Goethe-Universität Frankfurt - GUF), Klaus Dolag (Ludwig Maximilian University Munich - LMU)

History of Changes

Release	Date	Author, Organization	Description of changes
0.1	11.05.2023	Gino Perna, ENGIN-SOFT	Started the document
0.2	29.05.2023	Marisa Zanotti, ENGIN-SOFT	added Relevant parts
0.3	08.06.2023	Andrea Mignone	Added chap. 2
0.4	14.06.2023	Eva Sciacca, INAF	Contribute to Chapter 2
0.5	16.06.2023	Lubomir Riha	Contribution to Chapter 2
0.6	21.06.2023	Marisa Zanotti	Final adjustments
0.7	28.06.2023	Eva Sciacca	Overall minor adjustments
1.0	29.06.2023	Andrea Mignone	Acronym list and final format.

Scalable Parallel Astrophysical Codes for Exascale

DISCLAIMER

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain. Neither the European Union nor the granting authority can be held responsible for them.



The space above and below the message intentionally is left blank.

Executive Summary

The Deliverable 6.1 encloses the details of the collaboration with other Centre of Excellence (CoE)s (as appointed by Euro-HPC and coordinated by CASTIEL-2) as well as a number of dissemination and training activities.

From code development and profiling activities, in fact, we identified a number of dissemination topics built from technical know-how which are likely to be relevant in other research area, such as: i) Parallelization, data communication and integration with the Message Passing Interface (MPI) , ii) memory layout of multi-dimensional arrays, iii) Energy efficiency, iv) Artificial Intelligence and Machine Learning, v) High Performance visualization. These topics will be addressed in training workshop and other events.

In this framework, Scalable Parallel Astrophysical Codes for Exascale (SPACE) will also collaborate with National Competence Centres for High Performance Computing (NCC) and hosting entities in order to foster the network and get support for event organization.

Contents

1	Introduction	5
1.1	Context and purpose	5
2	Collaboration with other Centres of Excellence	6
2.1	Communication & Dissemination	6
2.2	Training	7
2.3	Workshop & Events	7
2.4	Further collaborations with CoEs fostered by CASTIEL2	8
3	Collaboration with NCC	9
3.1	Communication & Dissemination	9
3.2	Training	9
3.3	Workshop & Events	9
4	Collaboration with Hosting Entities	10
4.1	Communication & Dissemination	10
4.2	Training and Events	10
5	Conclusions	11
A	Castiel2 D1.2 Collaboration Plan with the CoEs	13

List of Figures

List of Tables

List of Acronyms

CASTIEL	Coordination and Support for National Competence Centres on a European Level
CD	Continuous Development
CI	Continuous Integration
CoE	Centre of Excellence
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
DL	Deep Learning
GPU	Graphic Processing Unit
HPC	High Performance Computing
KPI	Key Performance indicators
ML	Machine Learning
MPI	Message Passing Interface
NCC	National Competence Centres for High Performance Computing
NCCL	NVIDIA Collective Communications Library
NoC	network-on-chip
SPACE	Scalable Parallel Astrophysical Codes for Exascale

1 Introduction

1.1 Context and purpose

This deliverable is the reference document for Task 6.1, under WP6, that EuroHPC has requested all [CoE](#) to include during the Grant Agreement phase.

The purpose of this document is to provide an overview of the collaboration activities with complementary grants (other CoEs, [NCCs](#) and Hosting Entities) that [SPACE](#) is planning to engage during the project life. This document is therefore linked to the Collaboration Agreement, a collaboration contract, that [SPACE](#), as well as the other CoEs, will formalise with Coordination and Support for National Competence Centres on a European Level ([CASTIEL](#)) at due time.

2 Collaboration with other Centres of Excellence

2.1 Communication & Dissemination

Code development and profiling activities foreseen during the [SPACE](#) project lifetime will lead to technical know-how and skill acquisition in the field of High Performance Computing ([HPC](#)) that may be relevant in several other research areas which are potentially connected with other [CoE](#)'s [1] activities. This will provide a first set of knowledge, skills and qualities useful for communication and dissemination activities. We briefly recap, more specifically, which technical details that are most likely to be covered during these actions.

- **Parallelization, data communication and integration with [MPI](#).**

Possible connections with other [CoEs](#): Plasma PEPSC / MAX CoE / CHEESE

As one of the project tasks focuses on performance profiling & benchmarking of parallel applications, we expect valuable know-how transfer while these activities will be accomplished on different EU-HPC parallel platforms. In particular, this action will focus on (i) node-level performance effects related to the memory system and the network-on-chip ([NoC](#)) (interconnect for the modules comprising a system-on-chip), and (ii) cluster-level effects related to the coordinated messaging protocols that implement [MPI](#) collectives (eg. barrier and reductions). Actions will be taken to improve vectorization, reduce latency, and parallelization at fine and coarse levels to reduce bottlenecks. This will be carried out by working at the node level (improving/developing vectorization, workflow, multithreading and task-based approaches where advisable) and multi-node level (e.g., working on [MPI](#) communication and on remote direct memory access).

Another important aspect that deserves investigation is to what extent [MPI](#) can be fully deployed in new application, or used in conjunction with other Compute Unified Device Architecture ([CUDA](#))-aware parallel libraries, such as The NVIDIA Collective Communications Library NVIDIA Collective Communications Library ([NCCL](#)), for inter-GPU communication.

- **Memory Layout of multi-dimensional arrays.**

Possible connections with other [CoEs](#): Plasma PEPSC / MAX / CHEESE

Efficient exploitation of accelerated hardware requires substantial re-factoring of the application workflow and memory organization of multi-dimensional arrays which is priority issue we aim to overcome. Depending on the target hardware (Graphic Processing Unit ([GPU](#)) or Central Processing Unit ([CPU](#))), in fact, array indexing may work differently leading to unwanted performance losses and extra latency.

- **Energy efficiency.**

Energy efficiency is currently a hot topic in [HPC](#) and is also recognized by the EuroHPC. Within the [SPACE CoE](#) we will improve the energy efficiency of [SPACE](#) applications using dynamic tuning methodology originally developed within the H2020 READDEX project. This methodology identifies significant regions in a parallel application and defines optimal hardware and runtime system settings for each region. These settings could be the number of active threads, CPU core frequency, or the CPU uncore frequency. The final list always depends on the underlying hardware.

The essential part of the dynamic tuning approach is the runtime system, which is able to analyze the parallel application behavior for different configurations and detect the optimal ones for each region. The second key role of the runtime system is to perform the dynamic tuning itself. This means that as an application goes from one significant region to another, the runtime applies the optimal hardware/runtime settings. In the [SPACE CoE](#) we use MERIC runtime developed at IT4I@VSB.

The energy savings in this approach are gained by the optimal usage of the hardware resources and by reducing the performance of the CPU units that are not needed to execute at full speed. For instance, if a code is memory bound, the CPU cores can run on a lower frequency. In this case, the code performance is not impacted and at the same time, the CPU power consumption is reduced. The same could be done to compute bound regions. We can reduce the performance of L3 caches and memory subsystems (i.e. the uncore units of a processor) and save energy. At the same time CPU cores, registers, L1 and L2 caches run at full speed.

These are the main principles that our methodology relies on. However, within complex HPC applications, the situation is difficult and optimal settings need to be found by the runtime system.

- **Artificial Intelligence and Machine Learning.**

Possible connections with other CoEs: RAISE / TREX / COEC

Within the [SPACE CoE](#) we aim at designing, implementing and evaluating Machine Learning (ML) components in astrophysics pipelines. Use Cases include off-line processes for the transformations and enrichment of data, such as dimensionality reduction, segmentation or classification assessing its performance to produce a final evaluation on the advantages and disadvantages of the produced ML solutions. Furthermore, [SPACE](#) is exploring ML techniques to be coupled with run-time astrophysical simulations. The main objective of these models will be the exploration of large-scale simulations, eventually comparing with large observational datasets and the application of ML techniques to mine them in order to create surrogate models that can combine information from both numerical and observational data and to explore the use of Deep Learning (DL) surrogate models. This approach will then be used to reconstruct part of the simulation outputs at run-time while the simulation is progressing to replace the original higher time-consuming modules. ML-based post processing of data and the use of surrogate models are widely leveraged in numerical simulations in the broader scientific and mathematics communities to mimic the original expensive simulation, therefore, we expect this theme would be of interest in other CoEs.

- **High Performance Visualization.**

Possible connections with other CoEs: NOMAD / HIDALGO2

The [SPACE WP3](#) focuses on investigating fast and interactive visualisation tools and techniques, using services at the [HPC](#) facilities near the data, but with high speed delivery of the results to the users. The ultimate goal is to provide visualisation capabilities that are fully interactive and can access and visualise large data sets at a remote repository. We will tackle both simulated and observed 2D and 3D datasets. There is significant expertise in the [SPACE](#) consortium in visualisation of multivariate datasets developed around a few long running projects on the following visualisation modules: VisIVO ¹; Blender-based rendering²; OpenVDB, Houdini, and Arnold data pipelines for visualization. This activity may be of interest to other CoEs when Exascale codes produce large scale datasets to be visualized and analyzed in a fast and efficient way near the data which will be valuable to different fields such as biology, fluid dynamics, or engineering simulations.

2.2 Training

A fundamental aspect in the CoE's activities is the training of newly recruited personnel. Since GPU programming has emerged in the past years as a relatively new discipline in computer science, a significant endeavour is necessary in order to endow graduate and Ph.D. students with the necessary programming skills and technical knowledge. CUDA, OpenACC, OpenMP and OpenCL - among others - are only some of the programming language paradigms that the CoE may decide to adopt to port codes to efficiently exploit exa-scale or near exa-scale hardware. It appears clear, therefore, that this step is essential in determining the overall CoE's success while pursuing the established goals during the project lifetime. At the same time this action reflects the CoE's commitment and intent to convey and make accessible information and content that is not easily available. To this end, hired staff will be strongly encouraged to actively participate - in particular during the first 36 months - in a number of training events, such as hackathons and bootcamps, typically consisting of coding events in which teams of code developers wish to port their applications to run on GPUs or optimize their applications that already run on GPUs.

Considering the collaborative approach of [SPACE](#) towards other CoEs, partners' representatives, it would be available to participate as invited speakers and attendees in training initiatives of other Centre of Excellence, as well as to jointly organize educational initiatives focused on the topic highlighted in the previous section.

2.3 Workshop & Events

Among the activities fostered by CASTIEL, the effective collaborations among CoEs, to jointly address cross-cutting issues, is something that [SPACE](#) is planning to carry out by organizing joint events and workshops with other Centres of Excellence and other European Projects.

The topics that should be the focus of these events are already presented in the previous sections. The idea is to find a common basis of interests either in a common scientific area (cosmology and astrophysics for

¹<https://visivo.readthedocs.io/>

²<https://www.blender.org/>

instance), or in the technologies and tools to be used (machine learning techniques, visualisation tools, artificial intelligence, etc).

Such workshops should be either organized as "stand-alone" events, or included in broader events (as HiPEAC or ISC for example), where a wider public audience can participate and join the discussion. Preliminary contacts in this perspective have already been taken with representatives of other CoEs and with other EuroHPC and Horizon projects.

2.4 Further collaborations with CoEs fostered by CASTIEL2

Thanks to the support of CASTIEL2, SPACE CoE is participating in the Special Access Scheme initiative (better detailed in the annex) and representatives of SPACE CoE are involved on monthly activities, coordinated by CASTIEL2 itself. The starting point was the NCCs-CoEs online meeting in April, that offered an overview on CoEs objectives and NCCs activities, and that provided SPACE with the possibility to SPACE CoE to present itself to the other complementary grant and to start widening its network. In particular, SPACE CoE champions and deputies are taking part in the different working groups (Training and Communication, for instance), to share ideas and best practice, as well to make information circulate efficiently. On a monthly bases SPACE CoE is also committed to contribute in the Continuous Integration (CI) and Continuous Development (CD) task first with its technical experts.

3 Collaboration with NCC

One of the specific actions supported by CASTIEL2 is to facilitate the connection between CoEs and NCCs [2] in order to foster their collaboration. Thanks to CASTIEL, SPACE CoE has the possibility to have regular networking meetings with NCCs where common topics and cooperation opportunities are discussed. In particular, the initiatives that SPACE is considering to carry out, with the assistance and support of NCCs, are detailed in the following sections.

3.1 Communication & Dissemination

Since 32 NCCs are institutionally engaged in the coordination activities in all HPC-related fields at the national level and serve as a contact point for customers from industry, science, (future) HPC experts, and the general public alike, SPACE is considering to take advantage of their presence in all European countries to widen its network, in terms of communication and dissemination actions.

Each of the 33 national competence centres will act locally to map available HPC competencies and identify existing knowledge gaps. The competence centres will coordinate HPC expertise at national level and ease access to European HPC opportunities for research and scientific users, public administration but also in different industrial sectors, delivering tailored solutions for a wide variety of users.

In such perspective, NCCs can offer as local reference point also for CoEs, becoming "communication partners" towards their stakeholders. And vice versa; SPACE is committed to ensure regular communication concerning NCCs initiatives to the consortium.

Information about such opportunities are shared during the Communication coffees, organized by CASTIEL2 on a monthly basis and participated by the workpackage leader for communication and dissemination.

3.2 Training

NCCs are highly involved in training and educational activities. The mutual collaboration in this case would consist of making the partners aware of the training opportunities provided by the NCCs, and possibly being involved with different levels of engagement:

- audience, participating either physically or remotely to webinars and courses on project-specific topics;
- (invited) speakers, in accordance to the relevance of the topics,
- co-organizers, providing content and speakers to the local organizers. NCCs can ensure their support "hosting" a CoE training initiative in the "training portal", accessible through the EuroCC, and promoting it through its communication channels.

Information about such opportunities are shared during the Training coffees, organized by CASTIEL2 on a monthly basis and participated by the workpackage leader for training and engagement.

3.3 Workshop & Events

This collaboration can be extended with a wider approach, where CoEs, NCCs and hosting entities could collaborate in workshops or hands-on organization thereby creating strong synergies for a broader audience.

Preliminary contacts and considerations, on such regard, have been already discussed after the NCCs-CoEs meeting in April, with:

- NCC Luxembourg - LuxInnovation - LuxProvide (for MeluXina)
- NNC Czechia - IT4I (for Karolina)
- NCC Italy - CINECA (for Leonardo)

In particular, according to the pairing operation, SPACE has been assigned to Karolina and Meluxina.

4 Collaboration with Hosting Entities

The main objective of the [SPACE CoE](#) is to deploy its applications onto EuroHPC supercomputers. The redesign activity, performed internally by code developers, has to be necessarily integrated in an implementation phase, when the resources of supercomputers are made available.

CASTIEL 2 is responsible for the coordination of the common continuous integration and the deployment platform across all EuroHPC supercomputers. [SPACE](#) therefore relies on CASTIEL 2 and its special access scheme, whose first result has been the pairing the the hosting entities and related machines.

As such, [SPACE](#) will primarily take advantage of the following HPC infrastructures:

- **KAROLINA** - Karolina is a petascale EuroHPC supercomputer located in Ostrava, Czech Republic. It is supplied by Hewlett Packard Enterprise (HPE), based on an HPE Apollo 2000Gen10 Plus and HPE Apollo 6500 supercomputers. Karolina is hosted by IT4Innovations National Supercomputing Center.
- **MELUXINA** - MeluXina is a petascale EuroHPC supercomputer located in Bissen, Luxembourg. It is supplied by Atos, based on the BullSequana XH2000 supercomputer platform and hosted by LuxProvide.

Since CINECA is partner of [SPACE CoE](#), and also acting as NNC Italy, we expect to foster the engagement of third hosting entities, in particular as far as the dissemination and training activities are concerned.

4.1 Communication & Dissemination

As per NCCs, hosting entities can benefit from mutual support with [CoE](#), being part of the wider communication network.

4.2 Training and Events

Training activities could be hosted at Supercomputing Centre premises, selecting both "hardware-related topics", equally interesting for all [CoEs](#) and "scientifically-relevant topics", considering either the fields of interest or the tools and technologies involved. Some options that [SPACE](#) is considering to address are included in section 2 of this document and such ideas could be further investigated in the next months to become subjects for dedicated workshops.

5 Conclusions

This deliverable is due at M6, but the activities [SPACE](#) is planning to perform in collaboration with the complementary grants will be organized and distributed throughout the whole duration of the project. In this perspective, [SPACE](#) is considering to enhance its connection [CoEs](#), [NCCs](#) and hosting entities by the end the year in order to set the Key Performance indicators ([KPI](#)) that will allow an adequate evaluation of the common activities. Both the list of [KPIs](#) and the monitoring of the collaboration plan will be included, reported and updated in the deliverables of WP5 (due at M18, M30, M42 and M48).

From the administrative point of view, [SPACE](#) is expected to sign a Collaboration Agreement with [CASTIEL](#) and the other [CoEs](#) in order to formalize their common goals and to have the possibility to share information and data useful to meet the related [KPIs](#).

References

- [1] “EuroHPC-JU Centres of Excellence for HPC applications ,” <https://eurohpc-ju.europa.eu/centres-excellence-hpc-applications.en>.
- [2] “EuroCC National Competence Centers ,” <https://www.eurocc-access.eu/about-us/meet-the-nccs/>.

A Castiel2 D1.2 Collaboration Plan with the CoEs


DIGITAL-EUROHPC-JU-2022-NCC-01



**Coordination & Support
for National Competence Centres on a European Level Phase 2**

Project Number: 101102047

D1.7 Collaboration Plan with the CoEs



EuroHPC
 Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101102047. The JU receives support from the Digital Europe Programme and Germany, Italy, Spain, France, Belgium, Austria.

Work package:	1	Management
Author(s):	Miriam Koch Jisika Yono Dennis Hoppe Guy Lonsdale Guntram Berti	HLRS scapos
Approved by	Project Management Team	07.06.2023
Reviewer 1	Maria Sancho	BSC
Reviewer 2	Jisika Yono	HLRS
Dissemination Level	Sensitive	

Date	Author	Comments	Version	Status
2023-05-05	Miriam Koch, Dennis Hoppe Guy Lonsdale Guntram Berti	Initial draft	V0.1	Draft
2023-06-01	Miriam Koch	Consolidated Version for Review	V1.0	Final

Sensitive
© 2023 Members of the CASTIEL 2 Consortium



List of abbreviations

C2ISS	CASTIEL 2 Information Sharing System (Web Platform developed in CASTIEL 2)
CD	Common Development
CI	Continuous Integration
CoE	Centre of Excellence
DoA	Description of Action
GA	Grant Agreement
GDPR	General Data Protection Regulation (EU Regulation for Data Handling)
IPR	Intellectual Property Rights
NCC	National Competence Centre
PMT	Project Management Team
WP	Work Package

Executive Summary

The Grant Agreement of CASTIEL 2 sets a number of goals for collaboration with the CoEs:

- Actively contribute to the coordination activities of CASTIEL 2
- Contribution to the knowledge pool and available information to the common portal C2ISS
- Establishing effective collaborations and jointly addressing cross-cutting issues
- Participating in regular meetings to plan, implement and monitor collaborations and to synchronize research and development activities
- Participating in benchmarking exercises
- Sharing results and best practices as relevant
- Joint publication and dissemination of results
- Joint events
- Actively contributing supported codes to a common continuous integration and application deployment platform with automated testing (e.g., using Special Access scheme in collaboration with CASTIEL 2) at least on all EuroHPC JU systems. That means that a system for code testing and deployment on EuroHPC JU Systems is developed.
- Establishing common best practices for IP management and development including effective measures to ensure code quality, reviews, testing, management and development cycles
- Actively advancing modularisation: implementation of concrete measures for identification of common routines/algorithms/modules, creation and extension of software libraries used by multiple codes across disciplines

This deliverable shows how the goals were divided into specific actions, objectives by CASTIEL 2 that the CoEs need to contribute to and goals that are targeted and controlled by mechanisms in the project.

Sensitive
© 2023 Members of the CASTIEL 2 Consortium



Table of Contents

1	Introduction	6
2	Specific Actions from the GA	7
2.1	Legacy Code	7
2.2	CI/CD Platform.....	7
2.3	Special Access Scheme	8
3	Mechanisms in the CASTIEL 2 project	8
3.1	Supporting Activities	8
3.2	CoE contribution to CASTIEL 2 objectives.....	9
4	Formal Framework for Coordination and Collaboration	10
5	Outlook & Next Steps	11
6	References and Applicable Documents.....	12

List of Figures

Figure 1: Overview about C2ISS.	10
--------------------------------------	----

List of Tables

Table 1: Overview of specific goals.....	7
Table 2: Overview of Follow-Up Deliverables.....	12

1 Introduction

The support of the Centre of Excellences (CoEs) is an integral part of the CASTIEL 2 project. This applies to all CoEs funded under the EuroHPC JU call HORIZON-EUROHPC-JU-2021-COE-01. The overall goals for the collaboration are defined in the GA:

- Actively contribute to the coordination activities of CASTIEL 2
- Contribution to the knowledge pool and available information to the common portal C2ISS
- Establishing effective collaborations and jointly addressing cross-cutting issues
- Participating in regular meetings to plan, implement and monitor collaborations and to synchronize research and development activities
- Participating in benchmarking exercises
- Sharing results and best practices as relevant
- Joint publication and dissemination of results
- Joint events
- Actively contributing supported codes to a common continuous integration and application deployment platform with automated testing (e.g., using Special Access scheme in collaboration with CASTIEL 2) at least on all EuroHPC JU systems
- Establishing common best practices for IP management and development including effective measures to ensure code quality, reviews, testing, management and development cycles
- Actively advancing modularisation: implementation of concrete measures for identification of common routines/algorithms/modules, creation and extension of software libraries used by multiple codes across disciplines

While drafting the proposal and Grant Agreement of CASTIEL 2, the consortium has taken this list into account and divided it into:

1. specific **actions**
2. topics to be covered by **mechanisms** that enable organic, individual collaborations between CoEs.
3. **Objectives** specific to CASTIEL 2 that affect the CoEs. An example is implementing the unified Web platform *C2ISS* affect the CoEs, as well, since the centres will be contributing valuable content to it.

The collaboration between CASTIEL 2 and the CoEs is formalised and described in two ways:

- 1) the Grant Agreements of both CASTIEL 2 and the CoEs describe mechanisms and goals for collaboration, and
- 2) the Collaboration Agreement (a collaboration contract), due in project month 6, regulates legal aspects, including confidentiality, access rights for use or GDPR issues.

Several key performance indicators (KPIs) will be defined in close cooperation with the CoEs to ensure that the CASTIEL 2 project management team (PMT) can accurately monitor the progress and success of the collaboration.

This deliverable gives an overview of all aspects of the planned collaboration activities that will be coordinated and initiated through CASTIEL 2. The corresponding Collaboration Deliverables of the CoEs, which are also supposed to be submitted by project month 6, will show their perspective and provide project-specific contributions.



2 Specific Actions from the GA

The CASTIEL 2 Grant Agreement defines specific collaboration actions, either formalised as a milestone or a deliverable, sometimes a combination of both. This section will describe these actions, provided in Table 1, in more detail.

Topic	Milestone(s)	Deliverable(s)
Find mechanisms for Legacy Code	-	D2.4, 2.5, 2.6 – Reports on Legacy Code
Develop a common CI/CD Platform	MS10, MS14	D5.8, 5.9, 5.10 – Reports on CI/CD Platform
Develop a Special Access Scheme	MS 6	-

Table 1: Overview of specific goals.

2.1 Legacy Code

Legacy Codes are software packages with a long lifetime of prior development. CASTIEL 2 and the CoEs will:

- Identify legacy codes included within the CoEs application code portfolios
- Identify critical legacy codes in terms of those requiring software re-engineering in order to facilitate migration and maintenance on JU systems (both current and future)
- Identify legacy code re-engineering requirements in terms of personnel resources and time-scales and suggestions for actions to meet those requirements (potentially outside the scope of current CoEs).

CASTIEL 2 will produce the Deliverables on this, that contain an overview over all codes of the CoEs and track their deployment to JU systems, specific problems and comments, uptake, portability, scalability as well as progress over the runtime of CASTIEL2

2.2 CI/CD Platform

The ultimate objective is to ease the deployment of CoE applications onto EuroHPC supercomputers. To lower this hurdle and enable code owners and end users to deploy with minimal overhead on multiple supercomputers, CASTIEL 2 coordinates the implementation of a common continuous integration and deployment platform across all EuroHPC supercomputers. In the first phase, CoEs can conveniently deploy their codes on several EuroHPC supercomputers, ideally with minimal or no code changes. Such a CI/CD platform also allows for automatic testing and code quality checks that will help improve the maturity of the CoE codes.

To foster the implementation of such an integration platform, CASTIEL 2 coordinates this activity while bringing together representatives from all current and future EuroHPC supercomputing sites and CoEs. CASTIEL 2 has set up early on a mailing list to facilitate discussions between all stakeholders and organizes monthly conference calls to guide the integration process and provide an open space for discussion.

A concept for a possible CI/CD environment based on the Jacamar¹ CI project is defined, ensuring a minimally invasive set-up for both the hosting sites and the CoEs. Since most

¹ <https://gitlab.com/ecp-ci/jacamar-ci>, accesses on 12 May 2023.



EuroHPC supercomputing sites operate some CI/CD environment (not necessarily Jacamar CI), CASTIEL 2 matched up to three CoEs with a hosting site to gain the first hands-on experience. Based on the lessons learned and additional requirements that CoEs might impose, CASTIEL 2 will develop the final CI/CD concept to be implemented across all EuroHPC supercomputing sites in close collaboration with all stakeholders.

2.3 Special Access Scheme

CASTIEL 2 will support the CoEs in achieving appropriate resource access to the EuroHPC supercomputers to enable the CoEs to reach their objectives for deployment and delivery of efficient, highly-performant application software on those systems. For this purpose, CASTIEL 2 prepared several surveys to gather initial requirements for access to and use of the EuroHPC supercomputers (for development, benchmarking and performance optimisation) from all CoEs. Based on a thorough analysis of the survey data, CASTIEL 2 compiled a detailed proposal submitted in project month 3 to the EuroHPC JU to justify a special access scheme specifically for the CoEs as a strategic part of the EuroHPC JU's programme.

The input provided by the CoEs has unveiled their plans to deploy about 60 different codes and pilots to the EuroHPC supercomputers. In general, their requests align with the allocated computer resource limits of the individual EuroHPC JU systems. However, the overall input reveals significant uncertainty regarding the required computing resources and their compatibility with different systems. To ensure the timely initiation of the CoE activities on the EuroHPC JU systems, CASTIEL 2 proposed two initial access models:

1. A "flat-rate" or "base-rate" approach for benchmark and development-type access. This approach assigns a fixed amount of resources to each code, enabling quick access to gather the necessary technical experience and conduct benchmarking runs on different machines. It includes access to all systems for benchmarking purposes and selected systems for development.
2. An individualized yet accelerated approach for regular and high-priority access.

Lastly, CASTIEL 2 recommends implementing a process that includes regular updates on resource estimates by CoEs and periodic checks on the usage of granted resources, allowing for interim adjustments when necessary.

3 Mechanisms in the CASTIEL 2 project

The specific actions are one side of the project. The second part is to support the CoEs in several aspects: to connect CoEs to each other and to the NCCs, foster industry interaction and coordinate communication activities. Furthermore, CoEs play a relevant role to some CASTIEL 2 objectives as the C2ISS or the Training Framework. An overview is given in this section.

3.1 Supporting Activities

Collaboration between CoEs and NCCs

This mechanism is situated in WP2, which is responsible for facilitating networking & collaboration across NCCs and CoEs, and fostering NCC/CoE collaboration. The first step is to support the operation of a communication channel between all CoEs (the HPC CoE Council). These communication channels will then facilitate regular discussions between NCCs and CoEs on topics of mutual interest to enable initiatives and working groups.

In parallel, WP2 will define and run workshops on non-technical, relevant topics and competences (soft skills, complementary competencies) based on the needs expressed by NCCs

Sensitive
© 2023 Members of the CASTIEL 2 Consortium



and CoEs – and make the materials available. Topics will either be sourced from the NCC/CoE discussions or from the knowledge CASTIEL 2 has from its management perspective.

Enabling Exchange of Expertise

Situated in WP3, activities are in place to assist the development of CoEs through an effective mentoring and twinning programme. Furthermore, the mobility of HPC specialists is enabled between communities, academia, and public and private sectors through the mentoring and twinning programme.

Also, WP3 will support the exchange of training-related best practices, available resources, knowledge and information within and between NCCs and CoEs. Moreover, CASTIEL 2 will connect NCCs and CoEs to the network of different European hardware and software suppliers to enable information exchange between them. This will also ease any co-design activities planned by CoEs.

Fostering Industrial Interaction

WP4 will support the European HPC CoEs and the NCCs to increase and extend engagements with established and potential industrial end-users (in particular SMEs). Objectives are to identify and implement actions and to provide specific support for CoEs to increase their outreach to industry through sectorial events, industry bodies or industry-specific communication channels.

Also, this WP will create and operate an “industry relations” working group with NCCs and CoEs as members, which will discuss and frame the challenges and requirements of establishing effective and impactful interaction with the industry.

Communication support

To give the CoEs more visibility on a European level, WP5 will revive and further develop the HPC CoE brand and represent it in all communication channels. Furthermore, this task will provide an overview of the communication skill levels within the CoEs and support them with seminars (e.g. communication of broad service portfolios, identifying and targeting new target groups) and individual counselling.

On more details about the different topics, please see the WP-specific deliverables of WP 2 to 4.

3.2 CoE contribution to CASTIEL 2 objectives

Training Framework

This task aims to facilitate access to training offered at the national and European levels to interested NCCs and other potential users (from industry, academia or the public sector). This includes improved coordination and increased availability of training activities on HPC across NCCs, CoEs and within the European HPC ecosystem through the centralised training portal of training opportunities in Europe.

This task will also establish effective cooperation and leverage synergies with other European initiatives in the area of training in HPC, including EUMaster4HPC and ETP4HPC, among others. A common certification baseline of training and modular training approaches will also be addressed. This task will therefore establish a task force inviting on the one hand NCCs, CoEs, and the previously mentioned European initiatives, and on the other hand certification experts such as the HPC Certification Forum and European certification agencies to jointly map out the situation and work on the training framework (which then is handed over to the CoEs and NCCs for prototypical implementation). Moreover, this task will coordinate the definition

and implementation of the European HPC training baseline for industry, SMEs and academia/research/the public sector.

C2ISS

This project starts off with a set of websites (www.euroccaccess.eu, www.hpccoe.eu), which will need to be integrated into one unified platform solution. The new Web platform (working title: C2ISS) will not only include information about training offers, but also about actors from the HPC system, their services & offerings and more. Figure 1 shows an overview about the individual components and services the portal will support.

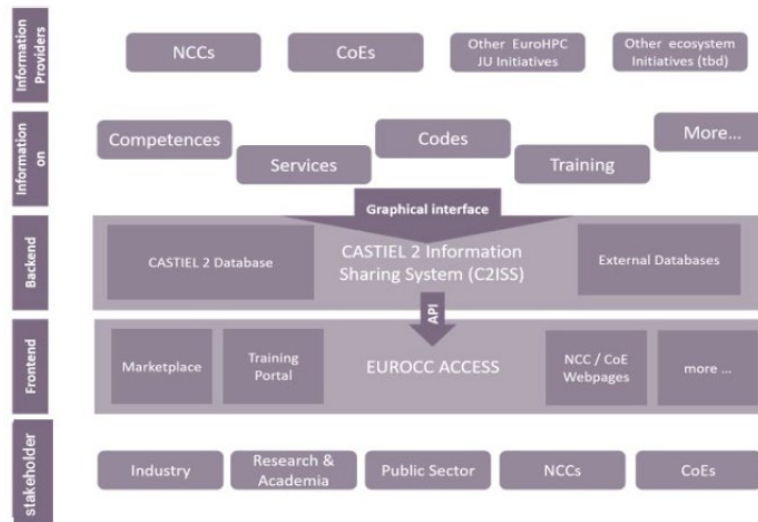


Figure 1: Overview about C2ISS.

To ensure that CoEs needs are met with the new Web platform, they will be integrated in all steps of the iterative design process of the C2ISS.

4 Formal Framework for Coordination and Collaboration

The previous sections described the actions and mechanisms for collaboration between CASTIEL2 and the CoEs. To efficiently achieve this, the exchange of data, information and knowledge is necessary. This is not a trivial task, since legal issues of confidentiality, IPR and GDPR are touched. WP1 is drafting a Collaboration Agreement, which regulates these legal issues. This Milestone (M11) is set in month six of the project.

Furthermore, WP1 has developed a set of KPIs to assess the effectiveness of collaboration. For example, the number of codes deployed via the CI/CD mechanism will enable CASTIEL2 to assess how well the mechanism is uptaken. To get an overview about all KPIs, please see D1.1 CASTIEL 2 and the European Dimension as well as D1.5: List of KPIs. These KPIs complement the other control mechanisms (reports of Work Packages on mechanisms and deliverables on the progress of specific goals).



5 Outlook & Next Steps

To recap, the collaboration between CASTIEL 2 and the CoEs is based on a legal framework consisting of the Grant Agreements and the Collaboration Agreement, quality control takes place through the KPIs developed in accordance with the CoE, reports of the work packages and deliverables to the specific goals. To follow the progress on collaboration, please see the deliverables in table 2:

Field of Activity	Topic	Follow-Up Deliverables (Due Date)
Specific Goals	Legacy Code	D2.4 Legacy Code Report (M12) D2.5 First Update of the Legacy Code Report (M24) D2.6 Second Update of the Legacy Code Report (M36)
	CI/CD	D5.8 CI/CD Platform (M12) D5.9 First Update of the CI/CD Platform (M24) D5.10 Second Update of the CI/CD Platform (M36)
	Special Access Scheme	D2.1 First year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M12) D2.2 Second year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M24) D2.3 Final year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M36)
Supporting Activities	Collaboration between CoEs & NCCs	D2.1 First year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M12) D2.2 Second year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M24) D2.3 Final year report on NCC/CoE Networking, Mapping of Competences, Codes and Services (M36)
	Exchange of Expertise	D3.1 Report on Training, Twinning, Mentoring and Mobility Activities in the First Year (M12) D3.2 Report on Training, Twinning, Mentoring and Mobility Activities in the Second Year (M24) D3.3 Report on Training, Twinning, Mentoring and Mobility Activities in the Third Year (M36)
	Support of Industry Interaction	D4.1 Initial Report on the Actions performed to support the NCCs and the COEs' Interactions with Industry (M12) D4.2 Intermediary Report on the Actions performed to support the NCCs and the COEs' Interactions with Industry (M24) D4.3 Final Report on the Actions performed to support the NCCs and the COEs' Interactions with Industry (M36)



	Support of Communication	D5.3 First report on awareness, impact and outreach (M12) D5.5 Second report on awareness, impact and outreach (M24) D5.7 Final report on awareness, impact and outreach (M36)
Contribution to common goals	Training Framework	D3.4 Implementation and adoption of the Training Framework (M6) D3.5 First Update of the Implementation and adoption of the Training Framework (M18) D3.5 Second Update of the Implementation and adoption of the Training Framework (M30)
	C2ISS	D5.2 Design strategy of the C2ISS and the evolved EuroCC Access Portal (M9) D5.4 Explanatory report on initial launch of the C2ISS and the evolved EuroCC Access Portal (M15) D5.6 Explanatory report on final version of the C2ISS and the evolved EuroCC Access Portal (M30)
General	KPIs	First Management Report (1.2) Second Management Report (1.3) Final Management Report (1.4)

Table 2: Overview of Follow-Up Deliverables.

As it can be seen in this deliverable, collaboration has started and is already producing first result. Once the Collaboration Plan in month 6 is in place, we expect collaboration to become even more fruitful and beneficial to all partners.

6 References and Applicable Documents

- [1] CASTIEL 2 project, <https://www.eurocc-access.eu>
- [1] CHEESE-2P, <https://cheese-coe.eu/>
- [2] BioExcel-3, <https://bioexcel.eu/>
- [3] MultiXscale, <https://www.multixscale.eu>
- [4] HiDALGO2, <https://hidalgo-project.eu/>
- [5] EXCELLERAT P2, <https://www.excellerat.eu/>
- [6] ESIWACE3, <https://www.esiwace.eu>
- [7] MaX, <http://www.max-centre.eu/>
- [8] SPACE, <https://www.space-coe.eu/>
- [9] Plasma-PEPSC, <https://plasma-pepsc.eu/>
- [10] CEEC, <https://ceec-coe.eu/>