




ChEESE

First Periodic Technical Report Part B



Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Spain, Italy, Iceland, Germany, Norway, France, Finland and Croatia under grant agreement No 101093038.



ChEESE-2P First Periodic Technical Report

Part B

PROJECT	
Contract Number	101093038
Project Name	Center of Excellence for Exascale in solid Earth
Project Acronym	ChEESE-2P
Project Website	https://cheese2.eu
Dissemination Level	CONFIDENTIAL
Nature	REPORT
Authors	Project coordinator and WP leaders
Contributors	All partners
Reviewers	All partners

REPORTING PERIOD	
RP Number	1
Duration	From 01/01/2023 (M1) to 31/12/2023 (M12)

HISTORY OF CHANGES		
Version	Date	Change
1	15/01/2024	Initial draft
2	26/01/2024	Revised by WP leaders
3	01/02/2024	Revised by all partners
4	28/02/2024	Final version, including suggestions from the RP1 review meeting
5	18/06/2024	Updated version after periodic report

Index

Table of Acronyms.....	4
Changes after first report review.....	6
Introduction.....	7
ChEESE-2P in a nutshell.....	7
Implementation phases.....	7
1. Explanation of the work carried out and overview of the progress.....	9
1.1. Objectives.....	9
1.2. Explanation of the work carried out per WP.....	12
1.2.1. WP1: Project Management.....	13
Overview of WP1.....	13
T1.1. Administrative and financial management.....	13
T1.2. Operational decision making.....	14
T1.3. Equality Committee.....	15
T1.4. Management of computational resources.....	15
T1.5. Data management plan and automatic FAIR validation.....	20
WP1 partner contribution.....	22
1.2.2. WP2: Exascale technical challenges in flagship codes.....	24
Overview of WP2.....	24
Exascale Ambition.....	24
T2.1. Code audits.....	25
T2.2. Increasing the performance portability of flagship codes for different accelerators.....	29
T2.3. Optimisation for Exascale.....	30
T2.4. Continuous Performance Efficiency Monitoring.....	30
WP2 partner contribution.....	30
1.2.3. WP3: Co-design with European HPC vendors and technologies.....	31
Overview of WP3.....	31
T3.1. Preparation of mini-apps.....	31
T3.2. OpenSequana prototypes (EuPEX).....	32
T3.3. RISC-V prototypes, the European Processor Initiative (EPI) and EuPilot... 32	
WP3 partner contribution.....	32
1.2.4. WP4: Simulation ecosystem.....	33
Overview of WP4.....	33
T4.1. Workflow implementation and optimization.....	33
T4.2. Lightweight performance analysis and optimization services for simulation ecosystem components.....	34
T4.3. Development, deployment and quality assessment tools.....	35
WP4 partner contribution.....	37
1.2.5. WP5: Farming Pilot Demonstrators.....	38
Overview of WP5.....	38
T5.1. Novel features (physics) in flagship codes.....	38
T5.2. Developments in Pilot Demonstrators.....	42

WP5 partner contribution.....	43
1.2.6. WP6: SCs and service enabling using EuroHPC infrastructures.....	44
Overview of WP6.....	44
T6.1. Setup of Simulation Cases and trials.....	44
T6.2. Exascale capability simulation cases and Scientific Grand Challenge....	44
T6.3. Exascale capacity simulation cases.....	44
T6.4. EuroHPC urgent computing service enabling.....	44
1.2.7. WP7: Dissemination, outreach and capacity building.....	45
Overview of WP7.....	45
T7.1. Dissemination and communication planning and tools.....	46
T7.2. Liaison activities.....	47
T7.3. Training activities.....	48
T7.4. Community building and roadmap activities.....	49
T7.5. Sustainability plan.....	50
T7.6. Collaboration tasks with complementary projects.....	50
WP7 partner contribution.....	53
1.3. Impact.....	54
1.3.1. Addressing the HORIZON-JTI-EuroHPC-2021-COE-01 call impacts.....	54
1.3.2. Addressing scientific impacts.....	58
1.3.3. Addressing technological impacts.....	58
1.4 Update of the plan for exploitation and dissemination of results (if applicable)..	59
2. Follow-up of recommendations and comments from previous review(s).....	59
3. Exploitation primarily in non-associated third countries (if applicable).....	59
4. Open science.....	59
5. Deviations from Annex 1 and Annex 2 (if applicable).....	61
5.1. Tasks/Objectives.....	61
5.2. Use of resources.....	63
5.2.1. Unforeseen subcontracting (if applicable).....	65
5.2.2. Unforeseen use of in kind contributions.....	65
Appendix 1: Exascale Computational Challenges (ECCs).....	66
Appendix 2. Flagship codes and scalability levels.....	67
Appendix 3. Pilot Demonstrators.....	68
Appendix 4. Simulation Cases.....	69
Appendix 5. EuroHPC systems.....	71
Appendix 6. Composition of the IUB.....	73
Appendix 7. ChEESE Technology Readiness Levels (TRL).....	75
Appendix 8. Peer-reviewed publications (M14).....	76
Appendix 9. Conferences and events (M14).....	78
Appendix 10. Communication and Dissemination (M14).....	80

Table of Acronyms

Acronym	Meaning
CA	Consortium Agreement
CI/CD	Continuous Integration / Continuous Deployment
ChEESE-2P	Center of Excellence for Exascale in Solid Earth - Second Phase
CoE	Center of Excellence
COLA	Collaboration agreement
CSA	Coordination Support Action
C2ISS	CASTIEL2 Information Sharing System
DA	Digital Asset
DevOps	Development Operations
DMP	Data Management Plan
DoA	Description of Action
DOI	Digital Object Identifier
Dx.y	Deliverable number y from WP number x
EC	Equality Committee
ECC	Exascale Computational Challenge
EGU	European Geosciences Union
EOSC	European Open Science Cloud
EPI	European Processor Initiative
EPOS	European Plate Observing System
EUPEX	European Pilot for Exascale
EWS	Early Warning System
FAIR	Findable, Accessible, Interoperable, Reproducible
FAIR-EVA	FAIR data Evaluator Validator and Advisor
GA	Grant Agreement
GO	General Objective
HE	Horizon Europe
IPR	Intellectual Property Rights
IUB	Industry and Users Board
JU	Joint Undertaking (EuroHPC)
KPI	Key Performance Indicator
LO	Liaison Officer
Mx	Project month number x (count starts on Jan 2023).
MSx	Milestone number x
NCC	National Competence Center
NDA	Non-Disclosure Agreement

PC	Project Coordinator
PEC	Project Executive Committee
PD	Pilot Demonstrator
POP	Performance Optimisation and Productivity (CoE)
PM	Person Month
PSB	Project Supervisory Board
RI	Research Infrastructure
RIA	Research and Innovation Action
SC	Simulation Case
SE	Solid Earth
SGC	Scientific Grand Challenge
SO	Specific Objective
SQAaaS	Software Quality Assessment as a Service
TCS	Thematic Core Service (EPOS)
TRL	Technology Readiness Level
UC	Urgent Computing
WF	WorkFlow
WP	Work Package

Changes after first report review

As a result of the first period review report received on 29th May 2024, the following revisions have added to v4 of this periodic report:

Financial part:

- Beneficiaries INGV, IMO, CINECA, IPGP (AE beneficiary of CNRS) and SU have justified in Section 5.2 the deviations of more than 25% on average personnel costs.
- Beneficiary BULL (EVIDEN) has corrected the cost category in the financial statement.

Deliverables:

- Deliverable D2.1 (First code audit report) has been resubmitted including:
 - A table summarising the POP metrics of the initial code versions (at M1) has been added in the Executive Summary.
 - A paragraph for each code has been added in the section "Brief description of the code and use-case" discussing the predicted application profiles at (pre-)exascale targets, based on the audit results.
- Deliverable D7.4 (Training plan) has been resubmitted including in Table 1 an additional column stating which trainings are solely organised by ChEESE and which are planned to be conducted in collaboration with other projects.

Technical report (this document):

- In the WP3 contribution section, the work carried out by each beneficiary has been detailed for T3.2 and T3.3.
- In the task T2.1 section (code audits), a new Table on exascale extrapolation based on smaller-scale (audit) profiles has been added (as in D2.1).

Introduction

ChEESE-2P in a nutshell

The second phase of the Center of Excellence for Exascale in Solid Earth (ChEESE-2P) is scientifically underpinned by a number of Exascale Computational Challenges (ECCs, see Appendix 1) and related geohazards existing on different domains of Solid Earth (SE), namely: computational seismology, magnetohydrodynamics, physical volcanology, tsunamis, geodynamics, and modelling of glacier hazards. With these background scientific targets, ChEESE-2P articulates around 3 main pillars:

- **Flagship codes.** Up to **11 flagship community codes** (see Appendix 2) will be optimised in terms of performance on different types of accelerators, code scalability, deployment, availability on EuroHPC systems following a CI/CD approach, performance portability across current pre-exascale systems and by co-designing with **7 mini-apps**, also on hardware architectures emerging from EuroHPC Pilots. Code optimization will also include heterogeneous single-node and multi-node performance, continuous efficiency monitoring using the Performance Optimisation and Productivity (POP) metrics, and quality and FAIRness assessment tools.
- **Pilot Demonstrators** (PDs). The CoE will farm **9 PDs** underpinned by concepts like multi-scale, multi-source, and multi-physics (see Appendix 3). In essence, PDs in ChEESE are computational workflows that make use of the flagship codes in specific problems addressing the target ECCs.
- **Simulation Cases** (SCs). The PDs will materialise in **15 SCs** representing capability and capacity use cases of particular interest in terms of science, social relevance, or urgency (see Appendix 4). The capability-oriented use-cases entail also 4 Scientific Grand Challenges (SGC), whereas capacity-oriented use-cases will focus on ensembles of petascale-level simulations for uncertainty quantification, hazard assessment, Urgent Computing (UC), or Earth's subsurface resource prospection among others. The SCs will produce relevant EOSC-enabled datasets and prototype services on aspects of geohazards like UC, Early Warning Systems (EWS), forecasts, and hazard assessments. An additional objective is to foster an emergency access mode in EuroHPC tier-0/tier-1 systems when geohazardous events (earthquakes, tsunamis and volcanic eruptions) materialise, including access policy recommendations.

Finally, ChEESE-2P will liaise, align, and synergize with other domain-specific European projects and longer-term mission-like initiatives like Destination Earth (DestinE). Emphasis will be given on the uptake by science, public administration and industry, including training and capacity building in cooperation with NCCs from critical countries.

Implementation phases

The ChEESE-2P implementation roadmap contemplates 3 main phases.

Implementation phase 1 (spanning roughly up to M12) with the following highlights:

- Project setup and overall planning, including quality guidelines (D1.1), Data Management Plan (DMP; D1.2), communication and dissemination plan (D7.1, D7.2), and training plan (D7.4). Setting all these plans constitutes the project milestone MS1.

- Collaboration plan with complementary projects (D7.3) including other CoEs and NCCs, as well as third-party stakeholders (public, industry) and other international and EU-funded activities.
- Signature of the collaboration agreement (COLA) with the CASTIEL2 Coordination Support Action (CSA). This is project milestone MS9.
- First code audit campaign to quantitatively set the baseline level of codes using hybrid POP metrics (D2.1) and drive the subsequent WP2 code optimisation activities.
- Collection of requirements for the 7 mini-apps (D3.1) for subsequent co-design with the European Processor Initiative (EPI) and the European Pilot for Exascale (EUPEX) projects during the second phase.
- Collection of requirements for implementation of workflows and PDs.
- Early deployment of selected codes on EuroHPC systems following a CI/CD approach (MS10). This includes, as a previous step, the obtain of quality badges from the Software and Services Quality Assurance as a Service platform (SQaaS) that resulted from the EOSC-Synergy project (<https://sqaaas.eosc-synergy.eu>) and the badge certification (assertion) through Badgr (<https://badgr.com>).
- Provision of computational resources for transversal WP2-WP4 activities through benchmark and development access modes on EuroHPC systems.
- Organisation of the 11th European Geosciences Union (EGU) Galileo Conference entitled "Solid Earth and Geohazards in the Exascale Era".

Implementation phase 2 (spanning roughly up to M36) with the following highlights:

- Increase the performance portability of flagship codes across different NVIDIA/AMD accelerators (D2.2), continuous code optimisation and efficiency monitoring through a ticketing system.
- Quantify the improvements in the flagship codes by updating the POP performance metrics after the second code audit round (MS5 and D2.3).
- Co-design with EuroHPC Pilots using the 7 mini-apps (MS2).
- Implement the necessary new functionalities (physics, couplings) in the 11 flagship codes as required by the PDs and periodic new code releases pushed to open git-based repositories and EuroHPC systems (MS3, D5.1).
- Further implementation of workflows (D4.1, D4.3) and PDs (D5.2 to 5.10).
- Setup of the Simulation Cases and trials (D6.1) and provision of computational resources through EuroHPC regular and extreme access modes.
- Execute the first production runs for the capability and capacity SCs.
- Draft a preliminary exploitation plan (D7.5).

Implementation phase 3 (up to the end of the project) with the following highlights:

- Final deployment of the PDs (MS6).
- Complete the 15 SCs through production runs and generate EOSC-enabled datasets integrated in the European Plate Observing System (EPOS) and elsewhere as detailed in the DMP (MS7 and D6.2 to D6.11).
- Enable 4 Urgent Computing (UC) services on target tier-0 EuroHPC pre-exascale architectures (Leonardo, LUMI, MN5) for an emergency access mode (MS8, D6.12).
- Perform user-driven evaluation exercises in collaboration with members of the Industry and User Board (IUB).

This First Periodic Technical Report summarises the project results during the first year of the project (from M1 to M12), which roughly coincides with the first

project implementation phase. **Further details can be found in the 8 deliverables already submitted to up date.**

1. Explanation of the work carried out and overview of the progress

1.1. Objectives

According to the Description of Action (DoA), ChEESE-2P has 4 high-level General Objectives (GO) and 15 Specific Objectives (SO) addressing:

- Exascale software engineering and simulation ecosystem (SO1 to SO4)
- Scientific excellence (SO5)
- Scientific-technical challenges (SO6 and SO7)
- Geohazard services and products (SO8 and SO9)
- Transversal integration around exascale computing and EuroHPC infrastructures (SO10 to SO15).

The following work has been done during the reporting period (M1-M12) in order to achieve the 15 SO, with **no significant schedule deviations**:

SO1: "To prepare 11 open-source community flagship codes for current and upcoming EuroHPC systems in the areas of computational seismology, tsunami modelling, physical volcanology, magnetohydrodynamics, geodynamics, and modelling of glacier hazards"

The transversal WP2-WP3 have already outlined a tailored code optimisation plan. The badge certification process is almost finished.

SO2: "To perform audit-driven code optimizations based upon 2 code audit rounds, one for defining each code baseline and detailing the code optimization plan, and a second for monitoring the performance of the codes along the project using the POP metrics as Key Performance Indicators (KPIs)"

The first audit campaign (up to M15) is finalising, D2.1 produced. Some difficulties were encountered (see T2.1 and Section 5.1 for details).

SO3: "To co-design mini-apps on OpenSequana (EuPEX) and RISC-V (EuPilot) exascale hardware prototypes ensuring an effective integration of co-design hints back to the main applications"

This concerns implementation phase 2, hardware prototypes are not available yet. The 7 mini-apps are in place (D3.1) and in git-based repositories. The requirements have been collected.

SO4: "To address specific simulation ecosystem challenges at large scale that go beyond the individual codes"

WP4 has started the collection of WorkFlow (WF) requirements and the automation of code deployment is going on (in collaboration with CASTIEL2).

SO5: “To develop and implement 9 new Pilot Demonstrators (PD) on different areas of SE addressing the 12 Exascale Computational Challenges (ECCs) and underpinned by the concepts of multiphysics, multiscale, multisource, and multihazard modelling”

Task 5.1 has started adding the necessary functionalities in flagship codes.

SO6: “To run 15 Simulation Cases (SC) derived from running on EuroHPC systems the 9 PDs at sites of particular relevance in Europe and beyond. This includes 9 capacity and 6 capability cases, 4 of which can, in turn, be considered as full-scale Scientific Grand Challenges (SGCs)”

This belongs to phases 2 and 3.

SO7: “To define a flexible ecosystem to ensure EOSC-enabled data quality management through an automatic FAIR data validation”

The DMP is in place (D1.2) including the data/software FAIRness evaluation tools and the plan for metadata scheme(s) and storage in centralised catalogues.

SO8: “To raise the Technology Readiness Level (TRL) of the PDs towards enabling downstream service prototypes addressing aspects of geohazards that are socially and economically relevant in Europe: disaster response, early-warning, and physics-based probabilistic hazard assessments”

This belongs to implementation phase 3.

SO9: “To foster the urgent computing access mode in EuroHPC tier-0/tier-1 systems for emergency management during high-impact geohazardous events (earthquakes, tsunamis and volcanoes) including related technical challenges and access policy recommendations”

This belongs to implementation phase 3.

SO10: “To act as an effective hub to foster the use of HPC resources and to roadmap the transition to exascale across the SE Community and related stakeholders”

The 11th Galileo conference was organised (May 2023) and the consensual document produced and shared with EuroHPC.

SO11: “To provide specialised training to scientists and end-users on the deployment and use of flagship codes and related services in collaboration with the National Competence Centers (NCCs)”

The training plan is in place (D7.4) with courses (HPC for geoscience series) running during 2024-2025. Up to 5 online webinars have been organised, summing a total of 194 participants.

SO12: “To consolidate and expand the role of the Industry and Users Board (IUB) in terms of identifying exploitable outcomes and promoting up-to-date market analyses”

The IUB inherited from ChEESE-1P has been reduced (25 members currently), retaining only the most active actors and defining two types of membership (full member and observer). Activities will start during the second phase, including exploitation analyses.

SO13: “To align and liaise with other projects (see Sec.1.2.7), mission initiatives (Digital Agenda, Destination Earth), Research Infrastructures (EPOS, FENIX), and federated environments (EOSC)”

11 liaison officers have already been appointed and will report periodically to the Project Executive Committee (PEC).

SO14: “To disseminate the CoE and the role of under-represented scientists to a wider-public audience, e.g. through the ShEESE initiative”

The former ShEESE activity has been overtaken by the Equality Committee (EC) in T1.3, which meets periodically and organised several actions (see T1.3 report).

SO15: “To collaborate with the CASTIEL2 CSA”

Substantial effort has been invested to meet this SO; see T7.6 report.

In terms of **milestones**, the percentage of achievement at M12 is shown in the Table below. Milestones MS9 and MS10 are delayed, see section 5.1 for explanation.

MS	Name	M	KPI	M12 (%)
MS1	Project setup phase	8	Internal management and communication tools in place, website operative, communication, data management and training plans defined (D1.1, D1.2, D7.1, D7.2, D7.3)	100
MS2	Preparation of mini-apps	18	Effective co-design with EuroHPC Pilots can start (D3.1)	75
MS3	Flagship code functionalities	24	Code version releases with the new functionalities (physics) available repository (D5.1)	too early
MS4	Mid-term review meeting	25	Check of the project program development (D1.3)	too early
MS5	Exascale code preparation	36	Improvements in the flagship codes performance metrics after the second code audit round (D2.2, D2.3)	too early
MS6	Exascale pilot deployment	42	PDs tested successfully (D4.4, D5.2 to D5.10)	too early
MS7	Production runs	48	SCs executed, EOSC-enabled datasets produced (D6.x)	too early
MS8	Community building and sustainability	48	UC services enabled, evaluation exercises with IUB members done, effective liaison with ecosystem initiatives, training completed	too early
MS9	Complementary Grants	6	Collaboration agreement has been signed	85
MS10	Application deployment	12	Applications developed by ChEESE-2P made available to users of EuroHPC	50

1.2. Explanation of the work carried out per WP

ChEESE-2P articulates around 26 tasks grouped in 7 WPs (see Table below). The total person effort (976 PMs) can be grouped in 3 broad “categories”: (i) transversal activities (WP2 to WP4, 427 PMs, i.e. a 43.75% of the total person effort) for HPC software engineering and hardware co-design; (ii) vertical activities (WP5 and WP6, 401 PMs, i.e. a 41.08% of the total person effort) for SE domain-specific science, production runs, and related outcomes and services and; (iii) other non strictly scientific-technical activities (WP1 and WP7, 148 PMs, i.e. a 15.17% of the total person effort) grouping the rest of activities.

WP	WP Title	Lead	Person Months	Start Month	End month
1	Project Management	CSIC	78	1	48
2	Exascale technical challenges in flagship codes	CIN	247	1	48
3	Co-design with European HPC vendors and technologies	BA	86	6	42
4	Simulation ecosystem	HLRS	94	1	48
5	Farming Pilot Demonstrators	CNRS	235.5	1	42
6	SCs and service enabling using EuroHPC infrastructures	INGV	165	13	48
7	Dissemination, outreach and capacity building	BSC	70.5	1	48
	TOTAL person months		976		

1.2.1. WP1: Project Management

Overview of WP1

The objectives of WP1 are:

1. Implement an efficient operational project management and give support to administrative and financial reporting, project legal aspects including Intellectual Property Rights (IPR), and timely delivery on the scientific and technical objectives of the project within the budget constraints.
2. Establish and enforce effective management and quality procedures for high-quality deliverables.
3. Effective communication between partners and detect project risks promptly, bringing them to resolution.
4. Promote gender and country balance through the Equality Committee (EC) activities.
5. Manage the computational resources according to project needs and foster the granting of additional EuroHPC resources through the different competitive access modes.
6. Implement a Data Management Plan (DMP) and an EOSC-enabled automatic FAIR software and data validation integrated in the CI/CD pipelines.

WP1 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
1.1	Administrative and financial management	CSIC	1	48	30
1.2	Operational decision making	CSIC	1	48	30
1.3	Equality Committee	LMU	1	48	25
1.4	Management of computational resources	CSIC	1	48	25
1.5	Data management plan and automatic FAIR validation	CSIC	1	24	35

T1.1. Administrative and financial management

The goal of this task, led by the Project Coordinator (PC), is to establish the procedures for project management (quality guidelines), including administrative management, financial reporting, timely production of deliverables, and perform eventual audits requested by EuroHPC and/or the co-funding National agencies. During the first reporting period, the following actions have been done:

- Writing and submission of deliverable D1.1 (Project Management and Quality Guidelines). This document was elaborated as a Handbook in order to ensure that all participants in the project understand the structure of the project, their roles, and the different responsibilities related to each role in the management structure. It also defines the internal procedure for the submission of Deliverables and Periodic Reports. D1.1 is a living document and may be updated as the project progresses.
- Assigning two internal reviewers for all deliverables submitted within this first reporting period. Supervision of the internal deliverable revision procedures set in place by the Consortium members during the kick-off meeting in Barcelona, in order to guarantee high quality results.
- Tracking of the status of all project deliverables and submission into the Funding and Tenders portal. Internal dissemination of these deliverables through our wiki and

accounting in the SACO repository and document collaborative tool (hosted and managed by the coordination partner CSIC).

- Carrying out of an internal financial monitoring at M6 in order to detect any misuse of resources (e.g. partner over/underspending). The Annex 2 of the GA has served as the basis for this reporting. No significant deviations have been detected for which no measures had to be taken.
- A total of **8 deliverables** have been submitted up to the reporting period (see Table below): 2 in WP1, 1 in WP2, 1 in WP3, and 4 in WP7.
- **NOTE:** deliverable D2.1 has been delayed by 1M (i.e. submitted by M13, see section 5.1 for reasons).

Deliverables submitted up to M12					
WP	Number		Deliverable name	Month	Contents
1	1.1	1	Project management and quality guidelines	2	Describes the project's internal management procedures as defined in T1.1
	1.2	4	First Data Management Plan	6	Describes the management life-cycle for all the digital assets to be collected, processed, and generated
2	2.1	7	First code audit report	12 (13)	First code audit round results and the recommendations arising from the analyses
3	3.1	8	Selected set of mini-apps and its requirements	12	Provide an account of hardware and software requirements of the already existing mini-apps
7	7.1	2	Project website	2	Release of ChESEE-2P project website with information about the project
	7.2	3	Communication and Dissemination Plan	3	Details activities to be undertaken as part of an overall project communication strategy
	7.3	5	Collaboration plan with complementary projects	6	Details the plan to ensure efficient collaboration with other CoEs for HPC applications and NCCs
	7.4	6	Training Plan	8	Identifies the training needs of researchers and other target audiences

T1.2. Operational decision making

Following the management structure described in the GA, T1.2 has provided support to a dynamic and inclusive decision-making process by:

- Managing the negotiation and signature of the Consortium Agreement (CA) and the Collaboration Agreement (COLA) between the CoEs and the CASTIEL2 CSA.
- Organisation of a two-day kick-off meeting (General Assembly) in Barcelona (M1, 19-20 January 2023). The meeting was hybrid in order to achieve maximum participation and all partners were well represented. This kick-off meeting was the first General Assembly out of the 3 that are foreseen along the project.
- Organising 9 meetings of the Project Executive Committee (PEC) and 3 meetings of the Project Supervisory Board (PSB). The resulting minutes and presentations are shared with the consortium partners through the internal Wiki.

- Ensuring effective internal communication through a series of tools provided by CSIC such as a project's zoom account, a series of mailing lists, a Wiki and a shared repository for editing and storing documents (SACO).

T1.3. Equality Committee

A project Equality Committee (EC) was set up during the kick-off meeting to promote and monitor our commitment with diversity and equality and to continue with the ShEESE ("She is") initiative that was launched during the first phase of the CoE. The EC is led by the partner LMU and, at present, it is composed of 7 members (see Table below) and meets once per trimester with the following results during the reporting period:

- Design of a common training series dedicated to equality, outlined in deliverable D7.3 (Collaboration plan with complementary projects). These workshops will be open to other CoEs and the first one will be offered by an expert on gender dimension in research.
- Prepare activities in the scope of the International Day of Girls in Women and Science (11th February 2024).
- A survey of project participants to have a picture of the ChEESE consortium composition and, eventually, show disparities and create awareness. The survey revealed the following number for the ChEESE consortium:
 - Overall gender balance: Male 59.8%, Female 40.2% (139 persons in total).
 - WP lead balance: Male: 72%, Female 28% (7 persons in total).
 - PSB balance: Male 71%, Female 29% (17 persons in total).
- All the activities of the EC are published in social media and also in the website to give visibility to our commitment to ensure equal opportunities to all groups.

No	Delegate	Institution
1	Duo Li	LMU
2	Arnau Folch	CSIC
3	Maria Jose Garcia	CSIC
4	Nathalie Favretto-Cristini	CNRS
5	Marta Garcia	BSC
6	Maria Paz Baghetti	BSC
7	Varvara Vedia	CSIC

T1.4. Management of computational resources

The goal of T1.4 is to assist consortium partners in the planification and provision of computational resources from competitive EuroHPC and National-level calls, ensuring alignment with the project roadmap during its 3 implementation phases. During the reported period (first implementation phase), several meetings were conducted in order to map the different project tasks, codes, PDs and SCs with the different EuroHPC systems (see Appendix 5) and the 4 types of access modes currently available to users. The resulting roadmap is summarised below (for the particular task of CI/CD on EuroHPC systems see the T4.3 section).

Task-to-access-mode mapping. The schedule to apply for computational resources for the different project tasks is summarised in the Table below with the following nomenclature for the 4 EuroHPC access modes: (B) Benchmark, (D) Development, (R) Regular, (E) Extreme scale.

Type of activity	Project task	EuroHPC mode				Project year			
		B	D	R	E	1	2	3	4
Code audit rounds	T2.1	x	x						
Code optimisation	T2.2, T2.3		x						
Co-design with mini-apps	T3.1		x						
Deployment and CI/CD	T4.3, T7.6	x	x						
Optimisation workflows	T4.1		x						
Code and PD developments	T5.1, T5.2		x						
Simulation setup	T6.1		x						
Capability SCs	T6.2			x	x				
Capacity SCs	T6.3			x					
UC service enabling	T6.4		x						

Code-to-system mapping. During the project, it is expected to continuously apply to benchmark and development access modes to perform the 2 code audit campaigns (T2.1), code optimisation activities (T2.2, T2.3), preparation of mini-apps (T3.1), CI/CD on EuroHPC systems (T4.3), optimisation of workflows (T4.1), code/PD developments (T5.1, T5.2), setup of SCs (T6.1), and enabling of UC services (T6.4). In terms of flagship codes, the following target systems and CPU/GPU partitions were identified for each flagship code (see Appendix 5 for the numbering of the EuroHPC system partitions):

Code / Partition	LUMI		LEO	MN5		VEGA		MELUX		KAROL		DIS
	1.1	1.2	2.1	3.1	3.2	4.1	4.2	5.1	5.2	6.1	6.2	7.1
SeisSol		x	x		x							
SPECFEM3D		x	x		x				x			
ExaHyPE			x	x	x	x	x					
Tandem		x	x			x	x	x	x	x	x	
xSHELLS	x	x		x							x	
HySEA			x		x				x			
FALL3D	x		x	x	x							x
OpenPDAC	x		x									
LaMEM	x							x	x			

pTatin3D	x	x	x	x						x	x	x
Elmer/ICE	x	x										x

During the reported period, ChEESE has been granted with **2 benchmark and 12 development access projects, summing up about 110,000 node/hours** on Leonardo, LUMI and Vega systems.

Code	Project reference (Access mode)	System	node/h	Start	End
SeisSol	EHPC-DEV-2023D07-03 1	Leonardo	3,500	14/07/2023	14/07/2024
	EHPC-DEV-2023D07-03 3	LUMI-G	10,000	03/07/2023	03/07/2024
SPECFEM3D	EHPC-DEV-2023D08-03 9	Leonardo	3,500	07/08/2023	07/02/2024
ExaHyPE	EHPC-DEV-2023D03-04 3	Leonardo	3,500	15/08/2023	15/08/2024
Tandem	EHPC-DEV-2023D10-04 4	LUMI-C	15,000	25/10/2023	25/10/2024
		LUMI-G	10,000	12/10/2023	12/10/2024
		Leonardo	3,500	16/10/2023	16/10/2024
xSHELLS	EHPC-DEV-2023D05-04 9	VEGA	1,000	01/06/2023	01/06/2024
HySEA	EHPC-BEN-2023B06-047	VEGA	400	15/06/2023	14/09/2023
	EHPC-DEV-2023D09-03	Leonardo	3,500	25/09/2023	25/03/2024
FALL3D	EHPC-DEV-2023D07-00 8	Leonardo	3,500	01/09/2023	31/08/2024
	EHPC-DEV-2023D10-02 8	LUMI-G	10,000	15/10/2023	15/10/2024
OpenPDAC	EHPC-DEV-2023D08-04 4	LUMI-C	15,000	30/09/2023	31/03/2024
LAMEM	project_465000139	LUMI-C	4,000	13/10/2023	01/02/2024
	EHPC-DEV-2024D02-01 7	LUMI-G	10,000	01/02/2024	01/02/2025

Code	Project reference (Access mode)	System	node/h	Start	End
SeisSol	EHPC-DEV-2023D07-03 ₁	Leonardo	3,500	14/07/2023	14/07/202 ₄
	EHPC-DEV-2023D07-03 ₃	LUMI-G	10,000	03/07/2023	03/07/202 ₄
pTatin3D	EHPC-BEN-2023B10-018	LUMI-G	12,000	15/10/2023	29/02/202 ₄
TOTAL			108,400		

Additional resources came also from competitive National projects:

Code	Project reference (Access mode)	System	time	Start	End
HySEA	AECT-2022-3-0015	MN4	0.81M core.h	01/11/202 ₂	28/02/202 ₃
	AECT-2023-1-0028	MN4	0.92M core.h	01/03/202 ₃	30/06/202 ₃
	AECT-2023-2-0033	MN4	0.6M core.h	01/07/202 ₃	31/10/202 ₃
xSHELLS	GENCI-A0140407382	Adastra, Jean-Zay, Joliot-Curie	85k GPU.h 7.1 M core.h	1/5/2023	30/4/2024
Elmer/Ice	project_462000007	LUMI-C/G	45,000 node.h	15/10/202 ₁	01/01/203 ₂

Capacity simulation cases. Task T6.3 will run the 9 capacity SCs as ensemble workloads (i.e. not needing tight interconnection between runs) that integrate scalable multi-physics/multi-domain models for UC and natural hazard assessment (see Appendix 4). The capacity SCs will be run during the second and third phase of the project with the following foreseen mapping with EuroHPC regular access mode:

SC	Involved codes (*)	Partitions (regular access mode)											
		LUMI		LEO	MN5		VEGA		MELUX		KAROL		DIS
		1.1	1.2	2.1	3.1	3.2	4.1	4.2	5.1	5.2	6.1	6.2	7.1
1.2	1, 2, 3, 4, 10		x	x									
1.3	1,3			x	x	x							
3.1	6			x		x				x			
5.1	7			x	x	x							x
5.2	7			x	x	x							x

6.2	8	x		x									
7.1	5		(X)			x		x					
9.1	11	x	x										(X)
9.2	11	x	x										(X)

(*) Flagship code number: (1) SeisSol, (2) SPECfem3D, (3) ExaHyPE, (4) Tandem, (5) xSHELLS, (6) HySEA, (7) FALL3D, (8) OpenPDAC, (9) LaMEM, (10) pTatin3D, (11) Elmer/Ice.

The mapping Table above can also be expressed in terms of the codes involved in the capacity SCs as follows:

Code	SC	Partitions (regular access mode)											
		LUMI		LEO	MN5		VEGA		MELUX		KAROL		DIS
		1.1	1.2	2.1	3.1	3.2	4.1	4.2	5.1	5.2	6.1	6.2	7.1
SeisSol	1.2, 1.3		x	x	x	x							
SPECfem3D	1.2		x	x									
ExaHyPE	1.2, 1.3		x	x	x	x							
Tandem	1.2		x	x									
xSHELLS	7.1					x		x					
HySEA	3.1			x		x				x			
FALL3D	5.1, 5.2			x	x	x							x
OpenPDAC	6.1, 6.2	x		x									
LaMEM	-												
pTatin3D	1.2		x	x									
Elmer/ICE	9.1, 9.2	x	x										x

Capability simulation cases. Task T6.2 will run the 6 capability SCs, 4 of which can be considered as SGCs (see Appendix 4). The capability SCs will be run during the second and third phase of the project with the following foreseen mapping with EuroHPC regular and extreme scale access modes:

SC	Involved codes (*)	Partitions (regular and extreme scale access modes)											
		LUMI		LEO	MN5		VEGA		MELUX		KAROL		DIS
		1.1	1.2	2.1	3.1	3.2	4.1	4.2	5.1	5.2	6.1	6.2	7.1
1.1	2		x	x			Systems NOT considered for capability SCs						
1.2	1		x	x		x							
2.1	2		x	x									
4.1	1, 3, 6			x	x	x							
8.1	9, 10	x											

8.2	4, 9, 10	x					
-----	----------	---	--	--	--	--	--

(*) Flagship code number: (1) SeisSol, (2) SPECFEM3D, (3) ExaHyPE, (4) Tandem, (5) xSHELLS, (6) HySEA, (7) FALL3D, (8) OpenPDAC, (9) LaMEM, (10) pTatin3D, (11) Elmer/Ice.

The mapping Table above can also be expressed in terms of the codes involved in the capability SCs as follows:

Code	SC	Partitions (regular and extreme scale access modes)											
		LUMI		LEO	MN5		VEGA		MELUX		KAROL		DIS
		1.1	1.2	2.1	3.1	3.2	4.1	4.2	5.1	5.2	6.1	6.2	7.1
SeisSol	1.2		x	x		x	Systems NOT considered for capability SCs						
SPECFEM3D	1.1, 2.1		x	x									
ExaHyPE	4.1			x	x	x							
Tandem	8.2	x											
xSHELLS	-												
HySEA	3.1, 4.1			x		x							
FALL3D	-												
OpenPDAC	-												
LaMEM	8.1, 8.2	x											
pTatin3D	8.1, 8.2	x											
Elmer/ICE	-												

T1.5. Data management plan and automatic FAIR validation

The goal of T1.5 is to develop and continuously implement a Data Management Plan (DMP) describing the management life-cycle for all the Digital Assets (DA) that will be collected, processed, and generated during the lifetime of the project. The DMP will be progressively implemented during the second and third project phases and will be aligned with the open science best-practices by means of an EOSC-enabled data/software quality approach able to assess (quantify) the level of FAIRness of the DAs. To this purpose, T1.5 contemplates an automatic assessment of data/software FAIRness with quantitative indicators relying either on:

1. Tools developed during the H2020 EOSC-synergy project (GA No 857647) for both software and data. Dockerized and online (the so-called service platform) versions of the EOSC-synergy validation tools already exist for easier deployment and testing in integrated environments, and will be connected to the ChEESE-2P data/software repositories (and services) to perform the corresponding FAIRness assessments.
2. The F-UJI platform, an automatic FAIR validation (data only) developed by the H2020 FAIRsFAIR project (GA No 831558) which is also promoted by CASTIEL2 as a transversal tool for all CoEs.


During the reported period the first version of the DMP has been produced (see D1.2 for details) embracing the **4 categories of DAs** shown in the Table below.

DA category	DA ID	DA name	Comments
Software	1.1	Codes	Flagship applications (source codes)
	1.2	Workflows	Resulting from T4.1
	1.3	Services	Resulting from T6.4
Scientific datasets	2.1	Experimental	Data coming from observations, measurements, or produced by detectors/sensors or by any other experimental device
	2.2	Benchmark	Datasets (experimental or synthetic) used to validate models and/or their performance
	2.3	Synthetic	Data generated by a simulation, particularly by the SCs
Metadata	3.1	Catalogues	Any (meta)data describing software and scientific data properties
	3.2	maDMP	Machine-actionable DMP
Other	4.1	Internal reports	Meeting minutes, internal notes to document the evolution of the project, calendar, resources management, mailing lists, etc
	4.2	Documents	Any documentation, either public or private, such as code documentation, technical notes, etc, not directly mentioned in the project deliverables list
	4.3	Deliverables	Project output documents

Repositories and FAIRness of DAs. DAs will be placed on centralised or distributed open repositories and evaluated using the EOSC-Synergy SQAaaS and the FAIR Evaluator, Validator and Advisor (FAIR-EVA) assessment tool that implements the Research Data Alliance (RDA) FAIR indicators and checks the FAIRness level of DA. This is summarised in the Table and Figure below (see D1.2 for further details).

DA ID	DA name	Repository	FAIRness	Metadata
1.1	Codes	- Native git-based repositories for development and CI/CD (mirrored to CASTIEL2) - Zenodo collection for DOI	SQAaaS	- Catalogues in git-based repositories - Model: CERIF - Format: JSON
1.2	WF	git-based repositories		
1.3	Services			
2.2	Benchmark	git-based	FAIR-EVA	- Data catalogue - Format: JSON
2.3	Synthetic (from SCs)	EPOS and other		- Model: CERIF extension - EPOS-DCAT-AP in an ingestion pipeline

4.3	Deliverables	• Digital.CSIC (with DOI)	NA	NA
-----	--------------	---------------------------	----	----



DMP implementation calendar. The timeline below shows the steps foreseen to implement the different aspects of the final DMP (D1.4 due at M34):

	Year 1	Year 2	Year 3	Year 4
SQAaaS for git-based code repositories	x			
Code catalogue (metadata)		x		
SQAaaS for WF and service repositories		x	x	x
WF and service catalogue (metadata)			x	x
Synthetic data from runs (PDs, SCs)		x	x	x
FAIR-EVA for datasets			x	x
Data catalogue (metadata)		x	x	x
Other DAs (e.g. deliverables)	x	x	x	x
Final DMP (D1.4)			x	

According to this plan, **activities up to M12 have been focussed on the SQAaaS for native git-based code repositories.** Results are reported in T4.3.

WPI partner contribution

All partners contributed to the Consortium Agreement (CA) finalisation and signature, took part in the personal F2F meetings and monthly teleconferences to monitor and manage project activities, managed the project activities within their organisation and

internally reviewed deliverables from other Partners. The Table below shows the WP1 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1.1	Administrative and financial management	L	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
1.2	Operational decision making	L	x	x		x	x					x	x					
1.3	Equality Committee	x	x						L			x						
1.4	Management of computational resources	L	x	x		x	x	x	x	x	x	x		x		x	x	
1.5	Data management plan and FAIR validation	L																

1.2.2. WP2: Exascale technical challenges in flagship codes

Overview of WP2

The objectives of WP2 are:

1. Address technical challenges (potentially at exascale) in the 11 flagship codes.
2. Exploit actual accelerators, in particular those targeting EuroHPC pre-exascale machines.
3. Audit the codes and improve single node, multi-node, and I/O performances at scale.
4. Continuously monitor the code performance efficiency during the project.

WP2 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
2.1	Code audits	HLRS	3 (27)	15 (36)	40
2.2	Increasing the performance portability of flagship codes for different accelerators	BSC	1	48	20
2.3	Optimisation for Exascale	CIN	6	48	10
2.4	Continuous Performance Efficiency Monitoring	BSC	18	48	0

Exascale Ambition

One of the main goals of ChEESE-2P project is to guarantee efficient (pre-)exascale execution in terms of capability or capacity runs. From the WP2 perspective:

- **Capability execution:** refers to an execution able to demonstrate excellent strong scaling up to a relevant fraction of a (pre-)exascale machine.
- **Capacity execution:** refers to an execution able to demonstrate limited strong scaling but excellent weak scaling spread on a relevant fraction of a (pre-)exascale machine (approximately in the number of ensemble members in a workflow).

According to above, tasks T2.1 and T2.3 are perfectly inline with the exascale goal, while T2.2 is needed to guarantee that the selected flagship code can reach good/excellent scaling performances across a wide range of heterogeneous hardware, without severe constraints (e.g. CUDA implementation that limits the availability on other accelerated platforms). As the real "issue" towards exascale is the size, the analysis already done in T2.1 needs to be refined taking into account the expected scalability (according to capability or capacity execution). Infact, the usability of T2.1 techniques with a pre-exascale execution pose severe limitations: size of the traces, usability of Score-P/Extrac recipes at large scale, I/O, network, etc. In the end, the complete WP2 methodology towards exascale will contain:

1. The use of TALP (possibly with GPU extension) to efficiently monitor the actual performance at scale or to extrapolate performance results to a very large scale execution. **This is part of WP2.**
2. Code restructuring (from 1. and/or 2. above in the T2.3 framework). This is part of WP2.

3. The use of a "scaling" workflow (to be sure that the workflow is not the bottleneck of the entire execution). This is part of WP4.
4. The use of a suitable budget to a given EuroHPC (pre-)exascale machine. This is not part of WP2.
5. The use of a selected tool able to measure the energy efficiency of the execution (if needed). This is not part of WP2.

WP2 aims pragmatically to target ONE or maximum TWO flagship codes (one for capacity, one for capability execution) with 1-5 of the above characterization.

T2.1. Code audits

The aim of this task is to assess the performance of ChEESE codes on the EuroHPC systems by means of 2 audit campaigns and with emphasis on Leonardo, LUMI and MN-5 (when deployed). Performance assessments will serve three goals: (1) identify code bottlenecks, (2) use the first audit round to set a performance baseline to monitor code evolution through the project and, (3) use the second audit round to quantitatively document the progress of flagship applications using the POP metrics. The results of the performance assessment campaigns will be documented in two deliverables (D2.1 and D2.3). Contribution to the task has two distinct roles: code owners should prepare performance assessment use-cases for their codes, BSC, and HLRS should support with performance analysis as well as liaison between the teams in ChEESE and the analysts in POP. During the reported period, T2.1 encountered 3 significant issues explained in detail in Section 5.1, namely: (i) POP3 not being funded until January 2024, (ii) access to EuroHPC systems and, (iii) lack of performance analysis tools installed on EuroHPC systems. Each of these issues had to be addressed individually and led to some delays. Ultimately, deliverable D2.1, expected by M12 (Dec 2023), was postponed to M13 (Jan 2024).

- T2.1 was mainly managed through the regular WP2 teleconferences. The lead of the task organised dedicated training on performance analysis tools and assessment methodology.
- At the project start, it was agreed that each code should perform both audit campaigns on the same target system (architecture) and with the same use case in order to facilitate the evaluations of improvements over the course of the project (for codes not initially accelerated it is admissible to do the first audit with a CPU version and the second with a GPU version). BSC and HLRS deployed customised installations of the tools mpiP, Extrae and Talp/DLB on Leonardo. On LUMI, it was not possible to share disk space across different compute projects, and we provided detailed step-by-step instructions on installing mpiP, Extrae and Talp/DLB. Similar instructions were provided for Vega.
- All code owners applied for access to at least one EuroHPC system and prepared (instrumented) codes and performance audit use cases.
- The first performance audit campaign was organised and managed by T2.1 lead at HLRS. Flagship codes have started their performance audit campaign on EuroHPC systems (task runs up to M15). Code owners selected their target system based on architecture, size, software environment, but also their availability (see Table below). Performance analysis tools were selected based on prior expertise of the code owners, supported programming model, and the preference of the performance analyst. Roughly half of the performance analysis were done with Score-P and Extrae

based tool chains respectively. In addition light-weight tools like mpiP and Talp/DLB were used for some of the audits.

- All partners involved in the task have contributed to D2.1 (First code audit report). The main part of this document consist of a chapter for each performance audit which briefly describes the respective flagship code, the selected performance audit use case, discussion of scaling curves (weak or strong as required by the use case), and the finally the results of the audit in terms of POP metrics as well as any observations and recommendations.

Code	Programing model	Version	System	Tool
SeisSol	MPI + SYCL + CUDA	1.1.3	Leo-B	
SPECFEM3D	MPI + CUDA	devel 4.0.0	Leo-B	Scalasca/2.6.1; Score-P/8.3
ExaHyPE	MPI + OpenMP target		Leo-B	Score-P/8.3
Tandem	MPI	1.1.0	LUMI-C	Score-P/8.3
xSHELLS	MPI + CUDA		Vega-G	
HySEA	MPI + CUDA	4.3.0	Leo-B	Extrac
FALL3D	MPI + OpenACC	8.2.1	Leo-B	Extrac and mpiP
OpenPDAC	MPI	1.0.0	LUMI-C	DLB/TALP and mpiP
LaMEM	MPI		LUMI-C	Extrac
pTatin3D	MPI + HIP		LUMI-G	Extrac
ELMER/ICE	MPI	fdd6c58f1	LUMI-C	Extrac

- At the time of writing this report, **10 of the 11 codes have analysed the resulting audit traces** and contributed to D2.1 including:
 - Baseline scaling behaviour.
 - Baseline POP performance metrics for the Focus of Analysis (FoA) parts of the code.
 - Communication structure (MPI).
 - Observations and recommendations.
- **Exascale ambition.** The Table below summarises the relevance for exascale of the audit cases.

Code	Audit case relevance for exascale
SeisSol	As part of a collaboration with the Texas Advanced Computing Centre, we have performed four full-machine runs with SeisSol on the Frontera supercomputer (8,000 nodes) during the last 2 years, simulating various scenarios with up to 500 mio grids cells. MPI scalability of SeisSol is therefore already well established. Our prime focus for the audit was therefore on GPU performance. In particular, we wanted to examine the GPU peformance achieved for SeisSol's local time scheme when applied to MPI partition sizes to be expected for our ChEESE-2P Pilot Demonstrator (PD) scenarios. The TPV5 scenario for the audit used a mesh with 17 million elements on up to 64 nodes of Leonardo. We expect that production scenarios for the Pilot Demonstrator PD1 (probabilistic seismic hazard analysis) will have meshes with 100-200 million elements, requiring huge ensembles (thousands) of scenarios. Scenarios for PD4 (complex


	multi-source tsunami modeling) may include extreme-scale runs with more than 1 billion grid cells. Hence, target scenarios are larger by a factor of at least 10x, maybe even 50x-60x. Respective weak scaling of the TPV5 speedup plots translates the range of 16-64 nodes to a range from 160 (10x16) up to nearly 4,000 (60x64) nodes. Hence, the audit scenario covers the MPI partition sizes to be expected for the PD scenarios.
SPECFEM3D	SPECFEM3D will be used for capability computing, which involves single runs requiring the maximum number of nodes available in the cluster. For the audit, we designed a run with small to intermediate sizes, ranging from 1,441,792 to 23,068,672 elements. We conducted both weak and strong scaling tests. The strong scaling tests revealed a bottleneck in MPI point-to-point (P2P) communication, specifically due to a lack of overlap between communication and GPU computation kernels. Despite this, our weak scaling performance was excellent. This is crucial for meeting the requirements of exascale capability runs, as it indicates that we can efficiently increase the mesh size and computational resources simultaneously. Our strategy of performing both weak and strong scaling tests allows us to identify and address specific performance issues, ensuring that SPECFEM3D is well-prepared for exascale computing environments.
ExaHyPE	Porting the ExaSeis seismic wave propagation model (based on ADER-DG discretisation) was only started in ChEESE-2P. The performance audit for ExaHyPE was executed on the first working GPU ADER-DG implementation of seismic wave propagation, using the LOH.1 benchmark as scenario. Our emphasis was on identifying bottlenecks in the parallelisation approach (e.g. load balancing) or related to the programming model (OpenMP target offload). ExaHyPE will primarily run as part of large workflows for uncertainty quantification (UQ), where most of the forward model executions will run on comparably few nodes. The range of 1-32 GPUs chosen for the audit is therefore representative. The size of the audit scenario (531,441 cells, approx. 3.2 billion degrees of freedom) reflects that of low- to medium-confidence forward executions in the UQ workflow. High-resolution models will be executed on models with 10x-100x more degrees of freedom (and on larger node counts). The respective scalability of the GPU implementation of ExaHyPE will be tackled in the second and third reporting period.
Tandem	Predicting application profiles for pre-exascale and exascale targets is challenging due to the audit being conducted on CPUs, while most Euro-HPC machines are primarily based on GPUs. The choice of performing the audit on CPU-based machines was motivated by the recent development of the code, as its portability to GPUs had not been thoroughly tested at the time the audit commenced. If the code were to run on a CPU-based exascale target, it could run efficiently if the mesh is large enough, with more than 1000 elements per rank. However, our experience on Frontera suggests that the current mesh partitioning strategy could become problematic if the MPI communicator has more than 100k ranks. Mitigation strategies are under investigation, and the use of GPU-based machines would likely relax this constraint.
HySEA	The main objective of Tsunami-HySEA is capacity execution, that is, running an ensemble of many simulations in parallel. For this purpose, we already have developed a version of Tsunami-HySEA capable of performing multiple simulations in parallel, each one running on one GPU. This version has been extensively used and tested by several institutions on a handful of EuroHPC systems, including a live demo carried out in the framework of the ChEESE project which consisted in simulating 540 scenarios using 68 nodes of the

	Marconi-100 supercomputer (272 GPUs). In the first audit performed in ChEESE-2P we considered the Tsunami-HySEA version devoted to capability (which is capable of simulating one single scenario using many GPUs), and a use case consisting in the simulation of the tsunami caused by the Tohoku earthquake occurred on March 11, 2011, using a mesh of 336.8 millions of cells. This mesh size is bigger than the mesh sizes that are usually employed in operational tsunami warning systems. The conclusions obtained in the audit revealed a high GPU usage, good GPU efficiencies and good strong scaling using up to 64 GPUs. Since the implementation of the numerical scheme is the same in both Tsunami-HySEA versions for capacity and for capability, and the communication techniques implemented in the capability version are more complex, the conclusions achieved in the audit regarding the GPU efficiencies of the capability version are useful for both versions. Despite that capability is not the main goal for Tsunami-HySEA, a good strong scaling was obtained for the considered use case, achieving a speedup of 57.4 using 64 GPUs.
FALL3D	FALL3D is mostly intended for capacity computing; i.e. ensembles of many scenarios potentially reaching a pre-exascale workflow. In operational settings, a single scenario typically requires a few nodes (up to 64 GPUs), with the target on time-to solution (urgent computing) rather than on strong scalability. The audit case was designed to push a single node of the Leonardo cluster with its 4 GPUs to the memory limit and to detect performance deficiencies. This addressed, extrapolation to large ensemble runs should be quite straight forward.
OpenPDAC	The OpenPDAC use-case selected for the audit corresponds, in terms of computational effort, to the "unit run" foreseen for Capacity computing. The target size of the unit run is $o(10^3-10^4)$ CPU cores (up to 64 LUMI-C nodes, i.e. 8192 CPU cores) . The exascale ambition (and challenge) of OpenPDAC-based Simulation Case SC6.2 consists in the orchestration and parallel execution of a large number ($\sim 100-1000$) of independent instances of such unit runs. On the other hand, the target unit run, besides being at the frontier of computational volcanology science, will be optimized to achieve in Simulation Case SC6.1, a time-to-solution compatible with Urgent Computing services to be deployed for ChEESE-2p end users.
LaMEM	The main objective of LaMEM is to be able to scale up the code to run on a substantial part of exascale machines. In the past, we managed a full-machine run on JUQUEEN (using $>450'000$ MPI ranks), with excellent weak scalability when considering the time per multigrid cycle. Yet, the time per timestep did not scale well, because the coarse grid required an increasing number of solution steps. In ChEESE-2P the objective is to be able to scale up realistic simulations, which requires including nonlinear visco-elasto-plastic rheologies, which requires the use of more accurate coarse-grid solvers. The first audit performed for LaMEM in ChEESE-2P confirmed that the coarse grid solution starts to completely dominate the solution time for >1024 MPI ranks. Accordingly, we will have to work on that aspect, which can hopefully be achieved by re-distributing coarse grid solvers on a subset of MPI-ranks. In addition, we have started porting LaMEM to make use of GPU-accelerators. At this stage it is unclear how much speedup that will bring.
pTatin3D	pTatin3d solves the Stokes equation for non-linear viscosities. Solving the non-linear problem is done by performing non-linear iterations during which we solve a linear problem. Therefore efficiently solving the linear problem is the bottleneck of the whole solve process, particularly for large viscosity contrasts. The audit case was designed to especially monitor that linear solve phase by

	simulating a ball sinking in a fluid with 6 orders of magnitude difference we address (1) high viscosity contrast, (2) gravity driven problem, (3) linear solve efficiency. This audit case was run with the first not yet optimized version of our GPU matrix-free kernel.
ELMER/ICE	For Elmer, we deliberately chose an already for modern CPU optimized (OpenMP SIMD) general purpose solver that we expected to have excellent POP metrics on LUMI-C (CPUs, only). The point is not to optimize the current solver, but to augment it with GPU target offload directives and interface to ROCm GPU enabled libraries for sparse matrix solutions (LUMI-G) to work towards capacity computing driven exascale applications (which inevitably need utilization of GPUs). By the nature of the Finite Element method applied in Elmer, the matrix assembly should scale linearly in a weak sense (i.e. increasing problem size and adding computational resources). The issue lies within the iterative solution of the matrix system. For the solution of the sparse system, we will utilize special pre-conditioners that confine the number of iterations (e.g., block-preconditioner for Navier-problems) and solve the resulting sub-problems with dedicated GPU enabled libraries (in our case RCalution on LUMI-G) that should guarantee to solve adequately increased problem sizes.

T2.2. Increasing the performance portability of flagship codes for different accelerators

Up to M12, T2.2 has been organised around three pillars: (i) collaborative planning across different WPs, (ii) task planning on a per-code basis and, (iii) dissemination/training on performance portable parallel programming approaches.



- To optimise the workload across work packages, special care was given to establish a shared workflow between WP2 and WP3. Specifically, this approach focused on the preparation of mini applications, serving as a mutual starting point for tasks T2.2 and T2.3. The diagram above clearly illustrates this integrated workflow/management, showcasing the intersection between these work packages.
- While a common workflow has been conceptualised at a high level, its implementation needs to be tailored to each individual flagship code, taking into

account their unique stages of development and specific needs for performance portability. To address this, initial meetings with the development teams of each flagship code were organised. These meetings aimed to understand the current status of the code and to identify potential avenues for future development and culminated in the **creation of a versioned/draft work plan document tailored to each code**. This work plan will serve as a reference for T2.2 activities as well to monitor the progress in due course. Moreover, it will provide an estimate of the planned effort that should be allocated internally to the T2.2 by the CINECA team through a ticketing system.

- In collaboration with CASTIEL2, a series of webinars have been organised within the objectives of T2.2 to provide a comprehensive overview of the latest and well-established trends in performance portable heterogeneous computing. These workshops delved into the application of parallel programming paradigms like Kokkos [\[link\]](#), Fortran DO CONCURRENT [\[link\]](#) and OpenCL [\[link\]](#), as well as the advanced usage of accelerated libraries like PETSc [\[link\]](#). The goal of this dissemination activity was to impart a thorough understanding of the capabilities and challenges associated with each approach to performance portability, enabling ChEESE-2P development teams to make informed decisions about the possible directions of work within T2.2.
- In conclusion, so far ChEESE-T2.2 has successfully defined an effective work strategy that takes into account WP level collaboration and code-specific task requirements. Future efforts will be directed towards solidifying the planned activities for each code and closely monitoring their progress, thereby advancing the project's goals in a structured and efficient manner.

T2.3. Optimisation for Exascale

Joint effort together with task T2.2 (as explained in the previous section)

T2.4. Continuous Performance Efficiency Monitoring

This task has not started yet

WP2 partner contribution

The Table below shows the WP2 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
2.1	Code audits	x		x		L	x	x	x	x		x		x		x	x	
2.2	Increasing the performance portability	x		x			L	x	x	x		x		x		x	x	
2.3	Optimisation for Exascale	x		x			L	x	x	x		x		x		x	x	
2.4	Continuous Performance Efficiency Monitoring		L															

1.2.3. WP3: Co-design with European HPC vendors and technologies

Overview of WP3

The objectives of WP3 are:

1. Extract to mini-apps the most relevant parts of applications according to targeted prototype architecture.
2. Leverage European HPC technologies through the SE Community applications.
3. Foster co-design with HPC vendors and technology and associated ecosystem (EuroHPC Pilots).

WP3 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
3.1	Preparation of mini-apps	CIN	6	18	30
3.2	OpenSequana prototypes (EuPEX)	EVIDEN	12	42	0
3.3	RISC-V prototypes, the European Processor Initiative (EPI) and EuPilot	CIN	19	42	0

T3.1. Preparation of mini-apps

In this task, the WP3 partners have selected and prepared a set of representative application parts which consist of extracted kernels or configurations of the full application that are called "mini-apps". Then, the hardware and software requirements of these mini-apps have been documented to prepare the co-design tasks T3.2 and T3.3 for EUPEX and RISC-V prototypes respectively. This work has been reported in D3.1 "Selected set of mini-apps and its requirements", submitted by M12. A tight collaboration with WP2 regarding this task has been set up as the mini-apps will be also used in T2.2 "Increasing Performance Portability".

- The summary table below shows the readiness of mini-apps, their use of CPU and GPU, the type of GPU backend used and the built-in metric to assess their

Mini-app	CPU	GPU	GPU Backend	Built-in Metric
Seissol-proxy			SYCL + SYCL/CUDA/HIP	Kernel times, FLOPS, bandwidth
Mini-SPECFEM3D			CUDA	Runtime
Tandem-static			PETSc	Kernel times
Mini-xSHELLS			CUDA/HIP	Average time per step
Tsunami-HySEA			CUDA/SYCL	Runtime
Mini-FALL3D			OpenACC	Runtime, iterations per second
Mini-pTatin3D			PETSc	Runtime, FLOPS

Ready
 In Development
 Not Available

performance.

- To support the WP3 partners in this task, especially the application leaders, some resources have been provided: a webinar dedicated to mini-app extraction has been organised by Eviden (BULL) to share their knowledge and best practices; and a set of guidelines for mini-app preparation have been presented (details available in D3.1 Appendix A).
- This task will continue during the first 6 months of 2024 by testing the mini-apps in collaboration with application developers, fixing potential bugs and ensuring the mini-apps and their defined test cases fully cover the needs for co-design. This could lead to new mini-apps and test cases if need arises.

T3.2. OpenSequana prototypes (EuPEX)

This task has not started yet.

T3.3. RISC-V prototypes, the European Processor Initiative (EPI) and EuPilot

This task has not started yet.

WP3 partner contribution

The Table below shows the WP3 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
3.1	Preparation of mini-apps	x					L	x	x	x		x	x				x	
3.2	OpenSequana prototypes (EuPEX)												L					
3.3	RISC-V prototypes, the EPI and EuPilot		L															

Note that T3.2 and T3.3 will start on M13 and M19 respectively. Some preparatory activities were conducted by task leaders (EVIDEN and BSC), with the rest of the partners not involved yet at the stage (access to the prototypes not available yet).

1.2.4. WP4: Simulation ecosystem

Overview of WP4

The objectives of WP4 are:

1. Address challenges at simulation workflow level that go beyond the main application codes.
2. Decrease overall time to solution by adopting a holistic view on the necessary workflow.
3. Exploit new methods of simulation setup, deployment, and data analysis.
4. Improve general usability for codes and workflows in the SE Community.

WP4 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
4.1	Workflow implementation and optimization	HLRS	1	42	20%
4.2	Lightweight performance analysis and optimization services for simulation ecosystem components	HLRS	25	48	-
4.3	Development, deployment and quality assessment tools	HLRS CIN	1	48	20%

T4.1. Workflow implementation and optimization

This task is intended to implement and optimise simulation workflows as a whole, reducing user interaction to a minimum and allowing fast deployment of simulations as well as handling of large ensemble workloads. Within the reporting period, various PDs focused on collecting the requirements for the workflows of the respective SCs. The actual planning of the workflows is tightly coupled with T6.1 (Setup of Simulation Cases and trials), which starts in M13. Early access to the targeted EuroHPC systems was requested to start implementing the workflow components (see T1.4). Unfortunately, access was granted later than planned due to delays in system deployment, leading to some initial delay in the workflow implementation. However, some early workflow specifications were already detailed following a template (see Figure 4.1.1) to characterise all components and data transfers, including the various files involved, a database, manual actions, scripts and already known bottleneck issues (if any):

- PD1 (SPECFEM3D steps) and PD2 (seismic tomography based on SPECFEM3D_FW)
- PD5 and SC5.1 (PVHA for the volcanic ash dispersal and deposition)
- PD6 and SC6.1 (Urgent high-resolution, 3D multiphase flow simulation of phreatic eruptions at Vulcano). Already implemented and made ready to run on the LUMI system by INGV. HLRS and INGV started analysing the workflow in more detail. A potential bottleneck component was identified here in the postprocessing part. The workflow is now subject to more detailed performance analysis by HLRS.
- PD9 (anticipated GLOFs simulations in Elmer/Ice). For illustrative purposes, Figure 1.4.2 shows the template filled for the PD9 case.
- The WF requirements will be reported on D4.1 (Initial report on workflow setups and workflow auditing, M18).

<Example Workflow>

<Contact + email>

Description: Example workflow just showing how to present using different boxes and arrows. It includes various files, a database, manual actions, scripts and an already known bottleneck issue.




Figure 4.1.1. Template provided by T4.1 for the workflow characterisation.

Description: Example workflow for anticipated GLOFs simulations in Elmer/Ice




Figure 4.1.2. Example of PD9 early workflow characterisation.

T4.2. Lightweight performance analysis and optimization services for simulation ecosystem components

There were no activities in this task as the start is intended for M25.

T4.3. Development, deployment and quality assessment tools

The objective of T4.3 is to improve code development, deployment and overall quality assessment. The following activities were done during the reported period:

- Ensure that all ChEESE codes are easily accessible to developers, software packagers, and HPC system maintainers as well as external tools and the user communities. In alignment with the DMP (see D1.2), all the source codes are already in native git-based open repositories as detailed in the Table below.

Code	Repo	Link
SeisSol	GitHub	https://github.com/SeisSol/SeisSol
SPECFEM3D	GitHub	https://github.com/SPECFEM
ExaHyPE	GitLab	https://gitlab.lrz.de/hpcsoftware/Peano
Tandem	GitHub	https://github.com/TEAR-ERC/tandem
xSHELLS	GitLab	https://gricad-gitlab.univ-grenoble-alpes.fr/schaeffn/xshells
HySEA	GitHub	https://github.com/edanya-uma/TsunamiHySEA
FALL3D	GitLab	https://gitlab.com/fall3d-suite/fall3d/
OpenPDAC	GitHub	https://github.com/demichie/OpenPDAC-11
LAMEM	GitHub	https://github.com/UniMainzGeo/LaMEM
pTatin3D	BitBucket	https://bitbucket.org/ptatin/ptatin3d/src/master/
Elmer/Ice	GitHub	https://github.com/ElmerCSC/elmerfem

- As a second activity and prior to any CI/CD on EuroHPC systems, T4.3 focussed on the set of quality checks derived from the EOSC-Synergy SQAaaS platform. This service harnesses the potential of the open badges specification and issues 3 levels of certification (bronze, silver and gold) rewarding quality achievements and compliance to software quality baselines. As a self-imposed **seal of quality**, ChEESE requires that all code repositories must have a certified badge before starting its CI/CD on the EuroHPC systems.
- At present, 1 code has a certified gold badge, 7 codes a silver, and 3 codes a bronze:

Code	Badge	Assertion link / Comment
SeisSol	Silver	https://eu.badgr.com/public/assertions/V1IRTJ4DSZessZ71umM-YQ
SPECFEM3D	Bronze	https://eu.badgr.com/public/assertions/SYNzEU1kSnyn33zQGnkTGQ
ExaHyPE	Bronze	https://eu.badgr.com/public/assertions/vg8PP-0FQlafJ-L-cRs_Xw
Tandem	Silver	https://eu.badgr.com/public/assertions/nM7pO3T0Sq2f1klkQ6gAGA
xSHELLS	Silver ¹	https://eu.badgr.com/public/assertions/50zMEGa5Sje8-QbgW5nSCQ

¹ The CeCILL licence in english is not recognized as it should (whereas it works in french). This has been reported, see <https://github.com/github/choosealicense.com/issues/1126>

HySEA	Silver	https://eu.badgr.com/public/assertions/S4Gk7HpLQwCCdu8BZOtc-A
FALL3D	Gold	https://eu.badgr.com/public/assertions/fokGiAA3QxGB-EMjZi3YHq
OpenPDAC	Silver	https://eu.badgr.com/public/assertions/s7DDg9SvRX6_IN8mUnkbuw
LAMEM	Bronze	https://eu.badgr.com/public/assertions/BKth1aOsRCqdLNhSKJjSZQ
pTatin3D	Silver	https://eu.badgr.com/public/assertions/Iv6uo7vhReCcKrrEf10IdA
Elmer/Ice	Silver	https://eu.badgr.com/public/assertions/DMhrwLxzQI-DBp_D7n89OQ

- Finally, ChEESE-2P took part actively in discussions with CASTIEL2 regarding their goal of setting up a code hosting, CI/CD-testing, and deployment service for the EuroHPC systems for all CoEs. Accounts for ChEESE staff to access the future CASTIEL2 services were just set up at the end of 2023 and we started to mirror our native repositories. However, several ChEESE codes already blueprinted CI/CD on selected EuroHPC systems independently of the activities with CASTIEL2.
 - SeisSol features CI/CD for every pull request on the public GitHub repository. A CPU and a GPU version are run on dedicated runner hardware (hosted at TUM). After the runs are completed, the solutions are compared to a pre-computed reference file.
 - ExaHyPE runs a CI/CD pipeline on the supercomputer (DINE) located in Durham, England. The run is triggered by commits to different branches that are selected as protected development branches (including the main branch). The pipeline directly runs on a dedicated compute node where the GitLab runner is installed.
 - T-HySEA pioneered CI/CD pipelines on VEGA.
 - FALL3D uses a CI/CD testing branch. The code from this is compiled for both CPU and GPU versions after a push on the Leonardo system. Then, both generated binaries (CPU/GPU) are executed in a test and the resulting output files compared. The success of the test is determined by the results falling within a configurable tolerance range. This testing workflow showed some challenges with the access to the EuroHPC systems for CI/CD because of the current security set up on the system, which requires a time-based one-time-password (2FA).

WP4 partner contribution

The Table below shows the WP4 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
4.1	WF implementation and optimization	x	x	x		L			x		x				x			
4.2	Lightweight performance analysis and optimization																	
4.3	Development, deployment and QA	x	x	x		L	L	x	x	x		x		x		x	x	

1.2.5. WP5: Farming Pilot Demonstrators

Overview of WP5

The objectives of WP5 are:

1. Implement novel features and coupling mechanisms in flagship codes for different PD workflows.
2. Develop the new workflows of the PD farm.

WP5 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
5.1	Novel features (physics) in flagship codes	CNRS	1	24	25
5.2	Developments in Pilot Demonstrators	CNRS	7	42	10

T5.1. Novel features (physics) in flagship codes

This task implements novel features in flagship codes that are necessary to the new generation of PDs (e.g. couplings for multiphysics, multiscale, multisource, forcing terms, etc). This is conceptually different from the T2.2/T2.3 activities, focussed on "algorithmics, porting, and optimization". The new versions incorporating these features are expected by M24 (D5.1). The following activities have been performed on the 11 flagship codes during the reporting period:

T5.1.1. SeisSol. The goal is to couple SeisSol to codes for geodynamics (pTatin3D, LaMEM) and seismic cycling (TANDEM), which simulate Earth's lithosphere deformation crust from millions of years down to thousands of years of entire seismic cycles. Ability to mix material models (elastic/elastic-anisotropic, acoustic, viscoelastic, poroelastic) for multiphysics simulations. Also, enable Cybershake capabilities including implementation of reciprocity and optimal "interfacing" between the Cybershake workflow and SeisSol (parameter files, formats of files, and model description). The main work done was to couple different codes with SeisSol:

- Coupling done with pTatin3d (100%).
- Ongoing implementation for LaMEM and TANDEM.

T5.1.2. SPECfem3D. The goal is to add attenuation and gravity kernels for computing different kinds of gravity anomalies. Extract and separate the computational core of SPECfem3D, which must be efficient on different exascale platforms, from the various higher-level workflows. To this purpose, a library called SEM-DK (Spectral Element Method - Development Kernel) will be developed in 5 steps: (1) extract the computational kernels already optimised during ChEESE-1P from SPECfem3D in order to make a version 0.1 of SEM-DK, (2) integrate the new WP2/WP3 developments and optimizations in SEM-DK0.1 or upgraded versions, (3) write an API for SEM-DK considering Fortran and C++ support to produce the version SEM-DK 0.2, (4) write a new version of SPECfem3D based on SEM-DK0.2 and validate against the up-to-date version of the code and, (5) write a new workflow for inversion (PD2) based on this new version of SPECfem3D. This will include new algorithms for multi-objective and

multi-physics inversion and functionals (seismology and gravimetry) as well as a new inversion scheme based on the Hessian matrix like Truncated Newton method.

- Implementation of gravity kernels (100%):
 - Mesher for a part of WGS84 ellipsoid Earth with topography added in the ellipsoid surface.
 - Cuda kernel to compute gravity field perturbation (Bouger anomalies, free air) from density perturbation.
- Begin implementation of SEM-DK (40%):
 - Mesher in 2D and 3D, defined a mesh partitioner for preparing the MPI version. The mesher outputs are compatible with the current version of SPECSEM3D which allows validation of SEM-DK with SPECSEM3D on the same setup.
 - Developed a sequential version of SEM-DK in 2D and 3D working on both CPU and GPU.
- Begin implementation of Newton methods for inversion workflow (25%):
 - Added a workflow to apply Hessian on vectors using the second adjoint method. This requires managing 4 SPECSEM3D simulations.
 - Inversion of Hessian in a free-matrix way using the conjugate gradient method.

T5.1.3. ExaHyPE. The goal is to develop UQ workflows with MUQ (MIT Uncertainty Quantification Library) and ExaSeis towards integration of improved earthquake sourcing (including dynamic rupture) and handling of more complicated geometries, as well as extend shallow water models in the context of landslide tsunamis.

- The integration of the ExaSeis framework into the newly redesigned ExaHyPE-2 code has been initiated. The framework runs for Cartesian grids using the ADER-DG solver. Future work will include adding curvilinear grids and PML as well as UQ workflows (25%).
- A landslide model has started to be developed in collaboration with the NGI, which will eventually be coupled with tsunamis. The landslide model is implemented in ExaHyPE-2 (20%).

T5.1.4. TANDEM. The goal is to add new multiphysics coupling, two-way linking to earthquake dynamics and wave propagation modelling with SeisSol, one way linking to long-term geodynamic modelling (LaMEM, ptatin3D) using stress, temperature, viscosity, density geometry as initial conditions and velocity / stress as boundary conditions”

- Two-way linking to earthquake dynamics and wave propagation modelling with SeisSol (T5.1.4) for subduction curvature vs. seismicity, fault roughness vs. seismicity.

T5.1.5. xSHELLS. The goal is to implement a dynamic and versatile subgrid-scale model based on deep learning to efficiently simulate turbulent geodynamo simulations that have barely been scratched by previous studies. Before dealing with three-dimensional geodynamo simulations, one needs to consider a lighter two-dimensional model of quasi-geostrophic convection in the Earth's core that retains some of the salient properties of a full geodynamo simulation, while lending itself more easily to the development of subgrid-scale models.

- One strategy to learn a stable subgrid-scale model is to take into account the temporality of the corresponding PDE. In order to train using this strategy, a differential solver that computes the time advancement of the filtered PDE is

needed. T5.1.5 implemented a spectral code for QG Boussinesq convection in an annulus with implicit-explicit integrators using JAX, a Python library supporting automatic differentiation, just in time compilation and much more.

- The next step in a learning workshop is to generate a dataset. When the boundary conditions are periodic, filtering commutes with partial derivatives and it is possible to only retain the first m Fourier modes. However, our system has no-slip velocity and imposed temperatures at the inner and outer boundaries of the annulus. Filtering Chebyshev modes in the radial direction is not preserving boundary conditions and we had to transform our fields to proper Galerkin bases in order to filter. Moreover, filtering velocities independently leads to non-zero divergence, which violates the Boussinesq equations. We thus use the filtered stream function to compute derived quantities such as velocities and vorticity. Now that we have a dataset containing objective filtered states, we will be experimenting with online learning with different neural networks models. Progress on the task is estimated at 25%.

T5.1.6. HySEA. The goal is to develop an operational version of the one-layer dispersive and non-hydrostatic HySEA code for earthquake generated tsunamis (Tsunami-HySEA) including nested grids and effective coupling of MeteoTsunami-HySEA and real atmospheric data. The progress for this contributions during the project are as follows:

- For the dispersive code:
 - A prototype running in a single GPU is already available. (100%)
 - A multi-GPU version is expected for M16 at most. (50%)
 - Nested grids will be implemented by M20. (20%)
- For the code for meteo-tsunamis:
 - A first prototype is running in a single GPU (100%) but needs testing and benchmarking.
 - The model is currently being tested with simple test cases (20%), this is the main task now regarding this code.
- Next step consists in considering more complex test cases (10%)
 - A test case from NTHMP has already been done.
 - Testing for real atmospheric data (0%).
 - Multi-GPU version of the code.

T5.1.7. FALL3D. The goal is to implement a mass-consistent meteorological downscaling for multiscale simulations

- A pre-processor of meteorological data based on a mass-consistent model is being implemented in order to interpolate meteorological data from mesoscale simulations (e.g. WFR model) to local-scale grids (20%).
- The capability to model resuspension by wind has been implemented and tested in well-characterised resuspension outbreaks. This will be part of code version 8.3 release (this is also a subject of an on-going PhD) (100%).

T5.1.8. OpenPDAC. The objective is to add new rheological models, modelling of the interaction between volcanic eruption mixture and water, including a description of the particle grain size distribution and the treatment of complex topographies (use of adaptive meshes).

- Development of a pre-processing tool to integrate in the simulation workflow capable of generating, from an ESRI ascii DEM file, the input files (STLs and dictionary) for the OpenFOAM mesh generation tool SnappyHexMesh. The final

computational grid should be refined at the desired level close to the topography and in the explosion region (defined modifying the original DEM). This will allow us to simulate processes at different scales, from sub-meters scales for the initial phases of the explosion and the interaction with the topography to decameter-scale for the rising plume in the atmosphere.

- Improve the modelling of water and the interaction with the volcanic mixture, by modelling the gas phase as a multicomponent phase (air+water vapour), and by testing the phase change models implemented in OpenFOAM to study the possibility to model water flashing.
- Implementation of the kinetic theory for granular flows for multiple solid dispersed phases to better simulate the behaviour of polydisperse gas-particle volcanic flows.

T5.1.9. LaMEM. The objective is to implement a julia interface for LaMEM to rapidly generate model setups of any region of the Earth while using available geophysical data.

- The Julia package `GeophysicalModelGenerator.jl` has been developed which allows to (100%):
 - Import data : Seismic tomography, GPS Velocity, Topography, EQ locations, Screenshots from published papers, Project lon/lat to cartesian
 - Create paraview files (VTK)
 - Manage inputs/outputs for geodynamics codes (eg LaMeM)
- Automatic installation of LaMeM on most available system (100%):
 - Automatically compiles for 105 systems, (mac, windows, linux). Includes PETSc, MPI & dependencies. Download binaries through the Julia package manager.

T5.1.10. pTatin3D. The contribution to the task T5.1.10 involves:

- Add support for Navier-slip type boundary conditions for spherical cap regional model domains using a symmetric Nitsche formulation (100%):
 - Implementation of the weak form for assembled and matrix-free operators of the Stokes system. The boundary conditions rely on a slip constraint and a stress or strain-rate constraint provided by the user.
 - The Nitsche method uses penalisation to enforce the slip constraint and adds symmetry in the weak form by differentiating the test functions for the discrete strain rate operator in the stress terms associated with the Navier-slip boundary condition.
- Accurate projection of material properties (stress and density) into a field defined on a mesh (100%):
 - We output pTatin3d results on a structured grid. This grid is utilised to construct a Q1 isoparametric mesh. An algorithm of point location is utilised to locate any point cloud (structured or not) in the Q1 mesh.
 - Interpolation functions of the Q1 isoparametric elements are utilised to perform a (tri-)linear interpolation from the Q1 mesh to the point cloud.
- Integration of pTatin3D output data formats with ASAGI to enable efficient import of 3D volume data into SeisSol (100%):
 - Utilising the procedure described previously, we generate a structured mesh containing the stress tensor, the density and the cohesion (for off-fault plasticity models) that is then written in Netcdf format to be read by ASAGI and interpolated on the Seissol unstructured mesh. A particular attention is necessary for the vertical component of the vector or tensor

valued variables as pTatin3d vertical is defined on the y axis while Seissol vertical is defined on the z axis.

T5.1.11. ELMER/ICE. The goal is to implement a coupled visco-elastic Reynolds flow type of model able to represent Glacier Outburst Floods (GLOFs). The Development of a coupled model for subglacial floods is done with the IMO collaboration by developing a simplified flow-line model (2D), with a purely elastic thin-plate model of the glacier underlain by a thin sheet of pressurised flood water. This model made it clear, that the coupling between the water sheet and the glacier needs to be in terms of piezometric potential for water flow at the ice–water interface that simultaneously creates the appropriate boundary condition for stress at the bottom of the ice as well as a potential gradient that forces the corresponding water flow in the thin sheet to guarantee compatibility with a changing subglacial cavity geometry. This implies that a numerical scheme needs to solve for the relation between a locally changing cavity and the hydro-potential. In the simple case, this is easily achieved with the implicit numerical solution of the equations for an elastic beam floating on a thin film of water; in a general case it would mean the necessity of a large amount of additional solution steps per time-step. In the full 3D application of PD9, this direct approach is not practical. Thus, we are adapting the method of artificial compressibility, which is already implemented in Elmer for thin fluid film problems, for the water-sheet that achieves the adjustment of the hydro-potential with an iterative approach. This adaptation was implemented as follows :

- Implementation of a visco-elastic model for the deformation of glacier ice by combining a linear elastic model with a non-linear Glen's flow law for viscous ice rheology (75%).
- Adaptation of an existing model for confined, thin-film fluid flow so that friction in the flow is according to Manning's law for turbulent flow, and (3) coupled the visco-elastic glacier model with the thin-film fluid model utilising artificial compressibility (75%).
- Coupling of the visco-elastic glacier model with the thin-film fluid model utilising artificial compressibility (75%).

T5.2. Developments in Pilot Demonstrators

Task T5.2 began at M7, i.e. at the time of the first periodic report 6 months had been dedicated to this task. Therefore, for some PDs, more development in flagship codes is necessary to make progress in T5.2. However, for all PDs, a roadmap has been defined, developments have been planned and strategies defined up to the end of the year. The difference in target Technology Readiness Levels (TRLs, see Appendix 7), ranging from 4 to 8-9, between the different PDs implies varying degrees of progress in their initial developments. At this stage, we can report on the progress made by PD1, PD3 and PD4

PD1 (Extreme-scale modelling of seismic hazards)

- BSC team's focus was on establishing ML-based solutions as fast analogs and candidates to supplement seismic simulations. This approach is considered as fundamental for parameter exploration and early assessment of hazards. During this review period, PD1 worked specifically on SC1.3. In particular, developing ML-based solutions for the Southern Iceland Seismic Zone (SISZ). This included pre-processing a large SISZ-specific data set in collaboration with BSC-IMO. The data was generated by means of the CyberShake platform and was used to train Machine

Learning (ML) models able to generate shaking maps and Uncertainty Quantification (UQ).

- In addition, BSC began extending the initial ML strategy. This continued effort consists in learning further models that rely on physics-based simulations to generate new inferences about seismic event source parameters. This work and its results are expected to complement traditional estimates of fault parameters required for wave propagation simulations.

PD3 (Global tsunami hazard and uncertainty quantification)

- A version of the workflow is available and it runs end-to-end, limited to the Mediterranean Sea (calculates hi-res probabilistic inundation model; the user has to provide local DEMs in the form of HySea grids).
- 3D Meshes were generated for 46 subduction zones worldwide; convergence tests to address the computational setup have been conducted for both the Mediterranean Sea and Pacific Ocean; Setup of HySEA ("Computational Engine") on Leonardo (+ Mercalli@INGV, HPC5@ENI); some massive runs were already conducted on Marconi100@CINECA with the ChEESE-1P version of the workflow.
- Interaction with GEM regarding interoperability with OpenQuake and first set of 3M tsunami runs (on Mercalli@INGV and HPC5@ENI) based on OpenQuake Monte Carlo simulation catalogue in the Mediterranean Sea
- We have utilised the Geo-INQUIRE ANTI-FASc earthquake slip generator, also distributed by the TCS Tsunami in EPOS.

PD4 (Complex multi-source tsunami modelling)

- The developments of the individual flagship codes used in PD4 are well underway. A conceptual workflow for coupling the different codes have been established, focusing on the model interoperability.
- Data for carrying out the corresponding Simulation Cases for testing the PD have been compiled. A literature review has been carried out for the Tonga case for mapping the code requirements and validation basis. Idealised test simulations have been carried out for individual flagship codes.

WP5 partner contribution

The Table below shows the WP5 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
5.1	Novel features (physics) in flagship codes	x		x				x	x	x		L		x	x	x	x	
5.2	Developments in Pilot Demonstrators	x	x	x	x			x	x	x	x	L		x	x	x	x	

1.2.6. WP6: SCs and service enabling using EuroHPC infrastructures

Overview of WP6

The objectives of WP6 are:

1. Define the 15 Simulation Cases (SCs) leading to capability/capacity production runs generating open EOSC-enabled datasets and hazard maps.
2. Execute 4 Scientific Grand Challenges (if full-scale resources are allocatable).
3. Enable 2 services in EPOS for tsunami and volcano Thematic Core Services (TSU-TCS and VO-TCS).
4. Enable 4 urgent computing services in EuroHPC infrastructures, including policy recommendations.
5. Engage the end-user community in the service co-design activities ensuring that the resulting products are well-suited for the applications and their users, including live demo exercises.

WP2 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
6.1	Setup of Simulation Cases and trials	INGV	13	18	0
6.2	Exascale capability simulation cases and Scientific Grand Challenge	CNRS LMU	19	48	0
6.3	Exascale capacity simulation cases	INGV	19	48	0
6.4	EuroHPC urgent computing service enabling	BSC	25	48	0

T6.1. Setup of Simulation Cases and trials

T6.2. Exascale capability simulation cases and Scientific Grand Challenge

T6.3. Exascale capacity simulation cases

T6.4. EuroHPC urgent computing service enabling

This WP has not started yet

1.2.7. WP7: Dissemination, outreach and capacity building

Overview of WP7

The objectives of WP7 are:

1. Identify and perform communication and dissemination activities in order to maximise the impact of the project among scientific audiences and target end-user stakeholders.
2. Liaise and coordinate with Research Infrastructures (EPOS), federated environments (EOSC), and other EuroHPC and REA (Horizon Europe) research projects to align and synergize efforts.
3. Identify and perform training activities in cooperation with NCCs in order to engage interested parties in the usage of codes and use cases.
4. Carry out community building and roadmap activities, with the main output being the Galileo Conference already granted by the European Geosciences Union (EGU).
5. Identify the exploitable results of the project and define the potential commercial strategies and services for the ChEESE-2Results (market study, business models, distribution channels and promotional strategy) to reach the market.

WP7 tasks					
Task	Name	Lead partner	Start month	End month	Percentage at M12 (%)
7.1	Dissemination and communication	CSIC	1	18	40
7.2	Liaison activities	CSIC	1	48	40
7.3	Training activities	BSC	6	42	30
7.4	Community building and roadmap	CSIC	1	48	40
7.5	Sustainability plan	BSC	25	48	0
7.6	Collaboration with complementary projects	CSIC	1	48	40

WP7 consists of ChEESE's efforts in dissemination, training, outreach and exploitation. Overall, ChEESE partners have participated in various dissemination and exploitation activities, which have helped increase notably the visibility and impact of the project. WP7 has set a number of Key Performance Indicators in D7.1 (Communication and Dissemination Plan), the status of which is reported in the Table below. Besides making sure that the project reaches its targets, WP7 is also responsible for the preparation of 8 deliverables, 3 of which have already been submitted.

Key messages	Activities	KPI	Status at M12
Press releases	Official statements	3	1
Impact on Media	Presence of ChEESE in mainly scientific media and, if possible, in popular media	30 appearances	20

Website sessions	Number of times a user has visited the website (sessions measured on Google Analytics)	1000 sessions/year	2345 sessions/year
Twitter	Engagement and organic traffic performance will be measured with Twitter Analytics and Metricool	Engagement: 1500 followers; Post Reach Percentage: 3%; Average Performance Rate: 2%	Engagement: 1155 followers; Post Reach Percentage: 3%; Average Performance Rate: 2%
LinkedIn	Effectiveness of the brand on the social network, which will be measured with LinkedIn Tools	Engagement: 600 subscribers; Post Reach Percentage: 2%; CTR: 2%.	Engagement: 501 subscribers; Post Reach Percentage: 2%; CTR: 2%.
Project Flyer	Leaflet with general information of the project	1	2
Roll Up	Printed banner of the project	1	1
Poster	Informational and educational material about ChEESE	1	2
Project Videos	A teaser to explain the project and its objectives and another video to explain the results	2	1
Journal Articles	ChEESE publications	30 articles in Q1/Q2 journals	16
Participation at events	Presentation of ChEESE results and research at conferences.	60	10
Events organisation	Hosting and organising the 11th EGU Galileo Conference	1	1

T7.1. Dissemination and communication planning and tools

Effective communication is pivotal for the success of the ChEESE project, ensuring that our groundbreaking research is shared effectively with our target audiences, including key stakeholders and the broader public. Here are some of the key initiatives under WP7:

- **Website:** The ChEESE website acts as the central hub for those interested in our project and activities. It presents in-depth research details and tracks all communication and outreach efforts. While primarily catering to the scientific community and end-users, we also engage the general public. In its first year, the

website attracted 2345 sessions mainly through direct and organic search, aligning with our yearly goals.

- **Press Releases:** By M12, we have issued one official statement, aiming for at least 3 by the end of the project. Future releases will highlight significant project milestones:
 - First Release: Introduction of ChEESE and its goals, available on our website since February 2, 2023.
 - Second Release: Mid-project progress update.
 - Third Release: Final scientific findings and their implications.
- **Impact on Media:** As of M12, ChEESE has been featured 20 times in scientific and general audience media, with a goal of 30 appearances by the end of the project.
- **Social Media Presence:** We utilise social media to extend our reach. ChEESE actively engages on Twitter and LinkedIn, targeting audiences most relevant to our project. The WP7 team oversees content sharing, complemented by partner contributions.
 - Twitter: Currently at 1151 followers, with a goal of 1500 by M48.
 - LinkedIn: With 501 subscribers now, we target 600 by M48.
- **Promotional Materials:** To enhance visibility at events and conferences, we've created:
 - Project Flyers: 2 leaflets summarising ChEESE, surpassing our one-flyer target. Available for download on our website.
 - Roll-Up: 1 printed banner for showcasing ChEESE at various events, meeting our M48 goal.
 - Posters: 2 informative posters detailing key project information.
 - Videos: 1 video elucidating ChEESE's objectives and results.

These combined efforts are crucial in driving ChEESE towards its communication objectives and widening our audience engagement.

- **Peer-reviewed publications:** 16 by M14, 11 in quartile Q1 (see Appendix 8 for details). Among these, 1 in Science and 2 in Nature Communications.
- **Conference Participation:** 10 by M14 (see Appendix 9 for details).
- **Other communication activities:** 33 by M14 (see Appendix 10 for details).

T7.2. Liaison activities

Different Liaison Officers (LO) were appointed during the kick-off to facilitate effective coordination with some pivotal initiatives (see Table below). This includes the collaboration with the Destination Earth (DestinE) initiative under the Digital Europe Programme, particularly focusing on digital twins in areas like geophysical extremes, and atmospheric, and ocean studies. On the other hand, the alignment with the 33 National Competence Centres (NCCs) established by the RIA EuroCC project and supported by the CSA CASTIEL is also important. These NCCs are instrumental in fostering European-level collaborations, including mentoring, twinning, and various networking activities under the CASTIEL project. The ChEESE project is also closely linked with several EuroHPC and H2020/HE projects such as EuPILOT, EuPEX, eFlows4HPC, DT-GEO, Geo-INQUIRE, and EsiWACE3, as well as the EPOS distributed RI and the European Open Science Cloud (EOSC). These collaborations are central to our goal of promoting seamless integration and mutual support among Research Infrastructures (RIs), within diverse federated environments like the EOSC, ENVRI, ESFRI, FENIX, and EuroHPC-JU.

Type	Project	Officer	Partner
Destination Earth		Arnau Folch	CSIC
CASTIEL-2 CSA and other CoEs	CASTIEL-2	Jose Gracia	HLRS
CASTIEL-2 CSA and other CoEs	EsiWACE3	Erwan Raffin	EVIDEN
CASTIEL-2 CSA and other CoEs	HPC3	Joan Farnos	BSC
EuroHPC and REA/HE projects	EuPILOT	Marta Garcia	BSC
EuroHPC and REA/HE projects	EuPEX	Erwan Raffin	EVIDEN
EuroHPC and REA/HE projects	eFlows4PC	Josep de la Puente	BSC
EuroHPC and REA/HE projects	DT-GEO	Ramon Carbonell	CSIC
EuroHPC and REA/HE projects	GEO-Inquire	Stefano Lorito	INGV
EuroHPC and REA/HE projects	EOSC	Arnau Folch	CSIC
ERIC	EPOS	Stefano Lorito	INGV

Liaison officers are expected to come up with concrete actions to promote cooperation during the second project phase. In a sense they are responsible for community building activities which will federate with end-users across the NCCs and IUB members, including training and capacity building to widen the access to codes and fostering transfer of know-how to end-user Communities.

T7.3. Training activities

The objective of T7.3 is to identify the training needs of researchers and other target audiences, as well as the organisation and logistical support of the training courses. The goal, as outlined in Specific Objective 11 of the GA, is to collaborate with NCCs in hosting accessible, inclusive training events and seminars. These events, free and balanced in terms of gender and geographic representation (i.e. favouring the EU13 member states admitted to the EU since 2004), will focus on the domains of computational seismology, volcanology, geodynamics, tsunami and meteo-tsunami modelling, and artificial intelligence. Details about the training courses can be found in D7.4 (Training Plan) submitted by M8, and foresee 6 training courses of the so-called "HPC for Geosciences Series" plus 1 course on AI (i.e. 7 in total):

No	Date	Theme	Title/Topic	Location
1	9-12 May 2023	HPC for Geosciences series	High-Performance Seismology CyberTraining	On-line
2	11-14 March 2024		Computational volcanology	In-person (BSC)
3	Q2 2024		Computational geodynamics	In-person (UM)
4	Q2 2024		Tsunami and meteo-tsunami modelling	Online + UMA
5	Q3/Q4 2024		Computational Seismology I	Online
6	Q3/Q4 2025		Computational Seismology II	Online

7	Q2 2024	Artificial Intelligence	AI and ML for Geosciences	In person (BSC)
---	---------	-------------------------	---------------------------	-----------------

On the other hand and up to M12, ChEESE partners have also organised 5 online webinars summing a total of 194 participants. The feedback obtained has been instrumental in shaping a tailored training program, leveraging the expertise of ChEESE partners to meet the specific requirements of our stakeholders. With this approach, we make sure that our training offerings are not only cutting-edge but also deeply relevant and impactful.

No	Date	Theme	Title/Topic	Participants
1	24/03/23	Other	ChEESE training event on EOSC-Synergy tools for SQAaaS	29
2	16/06/23	Programming interfaces and languages	Porting to GPU a Fortran code using intrinsic parallelism	48
3	28/06/23		OpenCL GPU programming for HPC Applications	55
4	07/07/23		Kokkos: A portable programming library for the exascale era	25
5	20/12/23		Device accelerated solvers with PETSc	37

T7.4. Community building and roadmap activities

A key focus for ChEESE is to foster a sense of community and pave the way forward with impactful initiatives, like the Galileo Conference, endorsed by the European Geosciences Union (EGU). In its first year, ChEESE hosted the 11th EGU Galileo Conference, centred around "Solid Earth and Geohazards in the Exascale Era." This event took place in Barcelona, Spain, from May 23rd to 26th, 2023.

- The Galileo conferences of the EGU are small-scale and informal events to promote in-depth discussion among experts on cutting-edge topics at the frontier of geoscience research. The event gathered 74 participants (53 men and 21 women) from 17 different countries to address both the challenges and opportunities that Exascale computing presents for geosciences. Most attendees came from European-based research institutes and academia, although other institutions worldwide were also represented by individuals (e.g. from the USA, China, and Latin America). The conference was articulated around 4 sessions linked to corresponding trending topics: preparation and optimization of High-Performance Computing (HPC) applications (Session 1), edge-to-end data workflows (Session 2), state-of-the-art in computational geosciences (Session 3), research infrastructures and EuroHPC access modes and policies (Session 4).
- Following the conclusion of the Galileo Conference, participants collaboratively developed a Consensual Document. This document reflects the collective agreement forged during breakout sessions, which were specifically tailored to probe into the future trajectory of solid earth research. These sessions yielded pivotal recommendations and insights aimed at propelling the field forward. Central to these discussions were the challenges and opportunities associated with HPC in the realm of solid earth sciences. The **Consensual Document** is available for reference as follows:

- Folch, A., Bhihe, C., Caviedes-Vouillième, D., de la Puente, J., Esposti Ongaro, T., Garg, D.,... & Zwinger, T. (2023). 11th EGU Galileo Conference: Solid Earth and Geohazards in the Exascale Era Consensual Document. <https://doi.org/10.20350/digitalCSIC/15439>
- For additional details on the 11th EGU Galileo Conference, including the agenda, participant information, and information about student grants awarded to Master Class students, please refer [to the Conference Report](#). This comprehensive document, which was submitted to the EGU at the event's conclusion, serves as an extensive resource for all related information.

T7.5. Sustainability plan

The Task will start in M25.

T7.6. Collaboration tasks with complementary projects

The objective of T7.6 is to foster collaborations with complementary projects. This task is pivotal in creating synergies and leveraging the collective strengths of various initiatives to enhance the overall impact and efficiency of our efforts in computational geosciences. The following activities are reported up to M12:

Collaboration with the CASTIEL2 CSA: during the reporting period, ChEESE has participated in many regular meetings and activities to plan, implement and monitor collaborations and to synchronise research and development activities with the CASTIEL2 CSA on several fronts including:

- Discussion and revision by the CSIC central legal department of the different versions (currently on v8) of the **Collaboration Agreement** (COLA) with the CoEs. The milestone for signing the cooperation agreement could not be met in time due to a delay in finalising the cooperation agreement in the CASTIEL2 CSA.
- Contribution to CASTIEL2 deliverable D7.1 (Collaboration Plan with the CoEs), where several Key Performance Indicators (KPIs) are defined closely to ensure that the CASTIEL2 Project Management Team (PMT) can accurately monitor the progress and success of the collaboration. CASTIEL2 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. CASTIEL submitted the resulting summary Table to EuroHPC for the Special Access request on the JU systems.
- Contribution to CASTIEL2 deliverable D2.1 (First-year report on NCC/CoE Networking, Mapping of Competences, Codes and Services, M12).
- Contribution to CASTIEL2 deliverable D2.4 (Legacy software report, M12) by providing information about background general and technical information and baseline state of ChEESE codes, as well as our initial development plans. This will be updated at M24 (CASTIEL2 deliverable D2.5) and M36 (CASTIEL2 deliverable D2.6).
- Contribution to CASTIEL2 deliverable D5.2 (Design strategy of the C2ISS and the evolved EuroCC Access Portal, M9). ChEESE collaborated in the design phase of the CASTIEL2 **Information Sharing System** (C2ISS) and related marketplace to understand the needs of CoEs when designing these two platforms and contribute to the knowledge pool and available information. The C2ISS results from CASTIEL's

WP2 and WP4 activities and aims at providing an overview of entities from the European Ecosystem on a single entry point accessible via APIs hosting information on, e.g., contact information, competencies, metadata, or training and events.

- Contribution to CASTIEL2 deliverable D5.8 (CI/CD platform, M12), which introduced the methodology applied in CASTIEL2 to coordinate the implementation of the EuroHPC CI/CD Pilot Platform and presented the EuroHPC CI/CD Pilot Platform architecture and roadmap for 2024. This will be updated at M24 (CASTIEL2 deliverable D5.9). ChEESE started its contribution to a **common continuous integration of the flagship codes and application deployment platform with automated testing** (CI/CD on EuroHPC systems; see also Task T4.3 report) through CASTIEL2 CI/CD meetings. The 11 codes in ChEESE answered the questionnaires to guide a description of how software is built, packaged, and installed. CASTIEL2 engaged CoEs in a pairing activity to get hands-on experience using and deploying selected codes on at least one EuroHPC JU supercomputers and, as a proof of concept for pairing systems, CI/CD pipelines were implemented for T-HySEA code on the VEGA system. On the other hand, CASTIEL2 prepared a GitLab instance that will host the CoE codes in a repository as the central hub for storing and managing source code in a structured and organised manner. It provides a centralised location where developers can contribute, track changes, and manage versions of their codebase. ChEESE has started the Unicast Deployment Path to leverage our existing Git repositories and write required GitLab Runners configured to authenticate with user credentials at a given EuroHPC JU supercomputers.
- Participation in the "**code-of-the-month**" initiative dedicated to presenting to the whole CoEs and CoC, those numerical codes that are thought to be of potential broad interest (including industrial applications). The Elmer/ICE code, with several other applications than geoscience (e.g. electric-machines, structural mechanics), was selected.
- Nominating the 2 ChEESE candidates to participate in the **Competence Centre and CoE Advisory Board (CAB)**; Tomaso Esposti Ongaro (INGV, ChEESE WP6 Leader) and Piero Lanucara (CINECA, ChEESE WP2 Leader) on topic 1 (Development of exascale-ready applications) and topic 2 (supporting supercomputing applications for Science & Innovation) respectively. The former ChEESE candidate was selected as a deputy member for topic 2 and, since M9, attends the regular meetings reporting back to the ChEESE PEC.
- Collaboration with CASTIEL2 WP5 (Communication and Dissemination) on the preparation of the EuroHPC Summit Week regarding elevator pitches and CoE Panel on daily work within the CoEs as well as collaborations within the ecosystem.
- Participation in monthly CASTIEL2 Project Management Team (PMT) and CoE leader meetings (from M6 onwards).
- Participation in the **monthly Training Coffee break meetings**. We presented our Equality training series during these meetings as outlined in the D7.2 Collaboration plan with complementary projects.
- Participation in the information collected by CASTIEL2 WP4 in interaction with the people in charge of industry in the CoEs of the CoEs (Industrial Interaction and Business Development) that was disseminated within the EuroCC2 network.
- Participation in the **hyperconnectivity study** of the EuroHyPerCon project, aimed at shaping the future of HPC in Europe by defining a long-term hyperconnectivity specification and implementation roadmap to meet Europe's future ultra-high-speed network requirements. The purpose of the study was to provide analysis and specification of the HPC (AI, Quantum) connectivity requirements in Europe, in views

to have a clear long-term specification of the hyper-connectivity service to be provided to EuroHPC, and the implementation roadmap to develop a federated, secure, and hyper-connected European HPC and data infrastructure, accessible via the cloud. workshop, ChEESE representatives attended the “Hyperconnectivity for European HPC Supercomputers – Feedback from HPC users and providers”, on November 22, 2023, at the Municon Conference Centre in Munich.

- ChEESE-2P researchers attend the “Code of the month” webinar for the industry working group, where the main features of a code of a CoE, the opportunities for this code to interact with industry as well as the available services for industry are discussed. Additionally, ChEESE-2P is attending the “Industry Coffee Breaks” organised by Castiel-2. In these events, CoEs, NCC and industry discuss HPC reliability, interest, regulations, access to HPC infrastructure (e.g. hardware, software, datasets), etc. SMEs are also at the core centre of the discussion in order to promote the adoption of CoEs results in coordination with NCC.

Collaboration with Industry:

- **Objective:** To harness industry expertise for identifying commercially viable outcomes of ChEESE-2P.
- **Actions:** Formation of the **Industry and Users Board (IUB)**, which serves as an independent mechanism for industry feedback on interim project results. The IUB includes representatives from industry SMEs, academic networks, and public governance bodies. Two types of IUB membership are established: full-time and collaboration, each with distinct roles and levels of involvement. Full-time members of the IUB are regularly invited to meetings and general assemblies, where their feedback directly influences the project's development. To maintain confidentiality, these members are required to sign a Non-Disclosure Agreement (NDA). On the other hand, collaboration members (also called “observers”), though kept informed about the project's progress, do not participate actively in discussions. Their role is more observational, allowing them to stay informed without direct involvement in decision-making processes. The classification of organisations into either full-time or collaboration memberships was determined by the Project Supervisory Board (PSB). See Appendix 6 for the current composition of the IUB.

Collaboration with National Competence Centres (NCCs):

- **Objective:** To establish European-level collaborations for sharing expertise in geophysics and emergency management.
- **Actions:** A survey to identify interested NCCs, leading to twinning and mentoring activities, particularly focusing on countries like Italy, Spain, Iceland, Turkey, and Portugal.

Collaboration with other Centres of Excellence (CoEs):

- **Objective:** To lead initiatives for equality and inclusion within the HPC community.
- **Actions:** Planning a training series dedicated to equality and gender balance in HPC, with workshops covering topics like inclusiveness, overcoming impostor syndrome, and leadership tools for women in HPC. ChEESE acts as the organiser, inviting other CoEs to participate.

Future Outlook & Next Steps:

- **Focus:** Continuous collaboration with other projects and initiatives, with an emphasis on shared knowledge and preparedness for upcoming challenges in computational geosciences.
- **Key Initiative:** Close liaison with the EU's Destination Earth (DestinE) initiative, focusing on the development of a digital replica of the Earth and its alignment with the EuroHPC Joint Undertaking.

Through all these collaborations, ChEESE aims to not only advance its scientific and technological objectives but also to create a more inclusive and interconnected community in the field of computational geosciences.

WP7 partner contribution

The Table below shows the WP7 partner contributions using the following IDs: (1) CSIC, (2) BSC, (3) INGV, (4) IMO, (5) USTUTT, (6) CIN, (7) TUM, (8) LMU, (9) UMA, (10) NGI, (11) CNRS (+IPGP), (12) EVIDEN, (13) CSC, (14) RBI, (15) UM, (16) SU, (17) E4 Analytics. An (L) indicates activity lead, an (x) indicates contribution.

Task	Main activity	Partner ID (beneficiaries + associated)																
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
7.1	Dissemination and communication planning	L	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
7.2	Liaison activities	L	x	x		x							x					
7.3	Training activities	x	L	x		x	x	x	x	x		x		x		x	x	
7.4	Community building and roadmap activities	x	L	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
7.5	Sustainability plan	x	L															
7.6	Collaboration task with complementary projects	L	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

1.3. Impact

The target impacts in the DoA are still relevant and on track; with no requirements for modifications. The following tables give the current (M12) status including, when applicable, the KPIs associated with the scopes of the call.

1.3.1. Addressing the HORIZON-JTI-EuroHPC-2021-COE-01 call impacts

Scope 1. *“Developing, maintaining, porting, optimising and scaling HPC application codes, addressing the full scientific and industrial workflow, particularly covering data aspects; testing and validating codes and quality assurance”*

Target	KPI	Status at M12
Periodic code auditing in collaboration with the POP3 CoE	Perform 2 code audit campaigns	On track, first audit round almost done
Port to accelerators (NVIDIA/AMD) and fine-tune the codes	All codes accelerated	too early to evaluate
Increase the performance portability of codes across EuroHPC systems	Performance metrics on EuroHPC tier-0 systems	too early to evaluate
Audit-driven optimizations. For level 1 codes, address also code resilience and fault-tolerance aspects	Run-time performance metrics from TALP	too early to evaluate. Tailored optimisation plan in place

Scope 2. *“Enhancing HPC applications and ensembles of coupled applications towards highly scalable, optimised codes”*

Target	KPI	Status at M12
Reduce bottlenecks in exascale capacity workflow executions	Increase the workflow metrics	too early to evaluate
Automate the code deployment process including workflow components	N/A	On track, see WP4

Scope 3. *“Streamline development, collaboration, automated testing and deployment processes throughout the application development and maintenance cycle, for example, by provisioning and using state-of-the-art development tools, platforms and software management models”*

Target	KPI	Status at M12
Use of agile methodologies	All codes in git-based repositories	Done, see T4.3
Software quality measures	Obtain EOSC-synergy badges for all codes	Done, see T4.3

Scope 4. *“Fostering pre-exascale-oriented codes and innovative algorithms that address societal challenges and/or are important for key scientific and industrial applications”*

Target	KPI	Status at M12
Address 12 ECCs identified by several HPC-community roadmaps	N/A	too early to evaluate
Farm 9 PDs	N/A	too early to evaluate

Scope 5. *“Provision of codes aimed at capacity pre-exascale systems for ensemble workloads that might not need tight interconnection between runs”*

Target	KPI	Status at M12
Run the 9 capacity SCs	Production runs performed and EOSC-enabled results/datasets available	too early to evaluate

Scope 6. *“Collaboration with EuroHPC JU systems and infrastructure federation initiatives to implement a robust and reliable automated deployment process for applications, in order to make novel developments timely available to the EuroHPC JU user communities”*

Target	KPI	Status at M12
Collaborate with CASTIEL2 on the CI/CD activities	All codes on at least 1 system	On track (with some delays)
Foster the end-user usage and uptake of codes through a containerized application deployment	Zenodo collection and git-based repositories in place	Done

Scope 7. *“Involvement in co-design activities (hardware, software, codes), including the collaboration with HPC vendors and the identification of suitable applications relevant to the development of European HPC technologies towards exascale”*

Target	KPI	Status at M12
Engage with the EuroHPC EuPEX and EuPilot projects to test on the available prototypes	The 7 mini-apps tested and with feedback integrated to main codes	7 mini-apps in place, prototypes not yet available

Scope 8. *“Enlarging and expanding HPC applications development and use, in particular for new user communities in EU countries and countries associated to Horizon Europe that are members of the EuroHPC”*

Target	KPI	Status at M12
Enhance the role of the IUB members	20 members	25 at present, with 2 types of membership
Enable resilient UC services and related policies on EuroHPC	4 UC service workflows tested and enabled on tier-0 systems	too early to evaluate
Exploitation plan for services and Engagement of target end-users related to industry, academia, public bodies and	Update the market analyses for the new service-oriented use cases	too early to evaluate

policy makers and update the service strategy		
---	--	--

Scope 9. *“Ensure wider access to codes and foster their uptake by user communities, in particular industry and SMEs”*

Target	KPI	Status at M12
virtual and trans-national access mechanisms to different user communities through Geo-INQUIRE	N/A	On track. VA/TA access rounds using ChEESE codes starting in Q1 2024.
Ensure FAIR validation of datasets, including EPOS-enabled metadata descriptors	Datasets integrated and/or compatible with EPOS standards	On track, see D1.2
Deliver a novel, high-resolution and uncertainty-based global tsunami model hazard map and related services	D6.4 (EPOS-TCS) and datasets available	too early to evaluate
Deliver the first multi-source European tephra hazard map at scale		too early to evaluate

Scope 10. *“Federating capabilities and integrating communities around exascale computing in Europe”*

Target	KPI	Status at M12
Interact strongly with EPOS towards the construction of TCS-VO and TCS-TSU	N/A	too early to evaluate (this will be D6.4 on M46)
Destination Earth (DestinE)	N/A	On track, proposals to integrate SE in DestinE on-going (March 2024)
Organise the Galileo Conference “Solid Earth and Geohazards in the Exascale Era”	Publication of 1 consensual document	Done

Scope 11. *“Address the skills gap in the targeted domain by specialised training and capacity building measures to develop the human capital resources for increased adoption of advanced HPC in industry and academia”*

Target	KPI	Status at M12
Organization of open, free, and gender and country-balanced (EU13) training events and seminars on the specific domains of SE	8 training events, including at least 1 course for each flagship code	Training plan in place. Courses scheduled by 2024-2025

Scope 12. *“Complementarity and synergy with National HPC Competence Centres and EU projects”*

Target	KPI	Status at M12
Feed in other large initiatives like Geo-INQUIRE and DT-GEO	N/A	On track, LO activities in place

Coordinate with SE/HPC ecosystem in all phases of the definition, implementation, validation and evolution of ChEESE-2P activities	N/A	Successful activity of the 11 LOs
--	-----	-----------------------------------

Scope 13. *"Include clear KPIs on the optimal employment of current and/or emerging HPC technologies, allowing the assessment of the progress towards the objectives, both in terms of outputs and ultimate impact"*

Target	KPI	Status at M12
Obtain code performance metrics (and optimise on-demand)	Increase in the POP metrics after the second audit round	too early to evaluate. First metrics available (D2.1)
TRL of the 9 PDs	Increase TRL of each PD by at least 2 points	too early to evaluate
Objective overall impact metrics	D7.6	too early to evaluate

Scope 14. *"Addressing the exascale and post exascale related technical challenges, such as load balancing; resilience; heterogeneity programming models, in particular accelerator-based architecture programming; run-time systems; workflow management tools; development environments and production environments"*

Target	KPI	Status at M12
Increase the performance portability of all flagship codes for different accelerators	N/A	too early to evaluate
exascale features enabling (resiliency, fault tolerance) in level 1 codes	N/A	too early to evaluate
bottlenecks beside the main HPC simulation codes using workflow management tools	N/A	too early to evaluate

Scope 15. *"Demonstrable advances of the targeted HPC applications towards highly scalable, optimised flagship codes and exascale performance (both computing and extreme data)"*

Target	KPI	Status at M12
Run the 6 capability SCs and 4 SGC	Production runs performed and EOSC-enabled results/datasets available	too early to evaluate
Apply to EuroHPC for competitive project access resources	Obtain 3 regular and 1 extreme scale EuroHPC competitive access projects	Phase 1 focussed on development access (11 projects granted)

Scope 16. *"Activities to improve the energy efficiency of applications, algorithms, methods, libraries and/or tools"*

Target	KPI	Status at M12
Energy consumption analysis on the EuPEX platform	N/A	too early to evaluate (D3.2)

1.3.2. Addressing scientific impacts

Target	Indicator (Goal/KPI)	Status at M12
Enhance the extreme-scale modelling of seismic hazards	Deploy PD1 and run SC1.1, SC1.2 and SC1.3	Too early, expected in phases 2 and 3 (WP5 and WP6)
Joint seismic and gravimetric data inversion at unprecedented level of resolution (100s of m)	Deploy PD2 and run SC2.1 at the Adriatic-Ionian region	
Advance current PTHA, exemplified in Global Tsunami Model and its interoperability with seismic hazard models and tools.	Deploy PD3 and run SC3.1 for the Global Tsunami Hazard Map	
Complex multi-source tsunami modelling	Deploy PD4 and run SC4.1 for the 2022 Tonga case	
Multiscale volcanic dispersal modelling combined with a multisource hazard assessment	Deploy PD5 and run SC5.1 towards the first European scale tephra hazard map	
Modelling phreatic explosions that occur in magmatic and hydrothermal systems of volcanic islands	Deploy PD6 and run SC6.1 at Vulcano	
Modelling geodynamo process across scales and train dedicated neural networks for subgrid-scales	Deploy PD7 and run SC7.1	
In geodynamics, model a self-consistent state of stress and couple it with the rupture propagation codes	Deploy PD8 and run SC8.1 and SC8.2 at Etna and the Hellenic arc	
Scenarios for rapidly rising GLOFs at the Eastern Skafá Caudron, Vatnajökull, Iceland (SC9.1) and contributing to workflow for investigations of hazard imposed by capsizing icebergs at Bowdoin glacier, Greenland	Deploy PD9 and run SC9.1 and SC9.2 in Iceland and Greenland	16
Q1-Q2 peer-reviewed publications	25	
Presentations, international conferences and exhibitions	50 abstracts and keynotes in events	

1.3.3. Addressing technological impacts

Target	Indicator (Goal/KPI)	Status at M12
Enable UC services in EuroHPC systems	enabling a UC service for	too early to evaluate

	earthquakes and volcanoes	
Consolidate Early Warning Systems (EWS) for Geophysical extreme events based on reliable physical modelling	N/A	too early to evaluate
Enhance Tsunami Early Warning Systems (TEWS) by embedding a probabilistic description of a wide range of source mechanisms in numerical modelling	N/A	too early to evaluate
Perform operational forecasting of geohazards	N/A	too early to evaluate

1.4 Update of the plan for exploitation and dissemination of results (if applicable)

The plan for dissemination and exploitation of results did not vary with respect to what is written in the DoA.

2. Follow-up of recommendations and comments from previous review(s)

Not applicable.

3. Exploitation primarily in non-associated third countries (if applicable)

Not applicable.

4. Open science

Early and open sharing of research. Task 4.3 ensures the FAIRness of all flagship codes, making the repositories easily accessible to developers, software packagers, and HPC system maintainers as well as external tools and the user communities. It makes use of CI/ CD pipelines and SQAaaS services built on the Research Data Alliance (RDA) FAIR indicators and with quantitative indicators of the FAIRness level (badges). In

addition, ChEESE is collaborating with CASTIEL regarding code hosting, CI/CD-testing, and deployment on the EuroHPC systems.

Reproducibility of results. This aspect of open science is addressed by the DMP, describing the management life-cycle for all the Digital Assets (DA) in ChEESE and will include not only codes but also workflows and services with their corresponding metadata catalogues in git-based repositories that will be opened at the end of the project. On top of this, WP4 will explore the use of the Workflows Hub Registry for the PDs, which will enable publishing workflow executions (e.g. the production runs) and share them as results in scientific publications through the generation of DOIs. Finally, metadata on specific datasets derived from the SCs will be in a data catalogue component (see D1.2).

5. Deviations from Annex 1 and Annex 2 (if applicable)

5.1. Tasks/Objectives

There are no significant deviations from the original plan in terms of tasks and objectives. However, as occurs in the other CoEs, two milestones are delayed (MS9 and MS10). In addition, it is worth to report a number of issues and difficulties that had to be faced during the first year, mostly originated by external factors:

Delay in MS9

- The JU expected the Collaboration agreement (MS9) signed by M6 with the CASTIEL2 CSA. This timing has proven to be unrealistic due to the complexity of the process, involving all partners from all CoEs, with many casuistics, legal aspects and comments/concerns from lawyer partner's teams. At present CoEs are discussing version 8 of the COLA with CASTIEL2.

Delay in MS10

- MS10 (applications made available to users of EuroHPC) was planned by M12. However, the methodology proposed by CASTIEL2 has been clarified only after their deliverable D5.8 (at M12). At present, ChEESE has mirrored all codes in the central hub, but the required GitLab Runners are not in place yet (as for the rest of CoEs).

Overwork related to CASTIEL2 activities

- There have been many collaboration activities with CASTIEL2 (see T7.6 for details), requiring a very substantial effort in terms of management (WP1, meetings, etc.) and collaboration activities (WP7, contribute to CASTIEL2 deliverables, reports, CI/CD activities, joint dissemination and training, etc., etc.). Because this was unforeseen when writing the DoA, most partners in the ChEESE consortium only have 2 PMs in WP1 and WP7 (for a 4 year project), a clearly insufficient number.

POP3 not being funded in 2023

- Contrary to our expectations, the POP3 CoE was not initially funded to start in January 2023 but afterwards, in January 2024. The effort (in terms of PM) originally allocated to T2.1 (audits) by BSC and HLRS for liaison and support, assumed that most of the actual code analyses would be done by the POP consortium staff. This required changes in the allocation of resources and roles and the following actions were taken:
 1. The first performance assessment has been a bit lighter than a full performance plan (according to the POP2 terminology) but provides, nonetheless, the full set of POP performance metrics.
 2. The code developers provided not only the use cases to the POP analysts, but also had to instrument their codes as well as to run the use cases and to extract the traces for the performance analysts. This was a significant unforeseen effort for the code owners.

3. BSC and HLRS provided guidance and organised training on the performance analysis tools, including deployment on the specific EuroHPC systems, instrumentation of the codes, and extraction of traces.
 4. In order to meet the deadline for the first deliverable D2.1, the code performance analysis was broken up in two stages: (i) determining the POP metrics and first observations and, (ii) interpretation of the POP metrics in terms of the code behaviour and root cause analysis. The POP metrics and first observation are reported in D2.1. Later on during the course of the project, further analysis of the performance traces will be done to address root causes. Priority is given to codes where POP metrics showed performance issues.
- Around M10 and when some of the performance audits had already started, individuals from the POP consortium (from BSC, HLRS, JSC, RWTH Aachen) kindly offered to do the analyses of the traces on a best effort basis and their help and assistance was essential to deliver D2.1 with a minor delay. The situation with POP has been regularised at M12, once they kicked-off.

Access to EuroHPC systems

- At the start of the project, CASTIEL2 shared with all CoEs that the EuroHPC JU wanted to facilitate some kind of fast-track access to CoEs and allocate resources for it. Accordingly, each CoE had to detail how many resources were expected, on which systems or partitions, etc. Some time later, CASTIEL2 communicated that there were some delays and that it was uncertain if the CoE special access would ever be implemented. Up to the end of the reporting period, the JU has not provided any kind of special access mode for CoEs yet.
- The expectations of a fast-track access initially discouraged the individual team members to apply for competitive access to EuroHPC systems through regular channels (benchmark and development). As early as M4, the ChEESE T2.1 lead strongly encouraged code owners to apply for benchmark access to EuroHPC systems, but this turned out to be too limiting to conduct full code audits (in terms of compute resources and project duration). It was not until M7-M8 that the code lead developers applied to development access calls (11 projects granted, see T1.4 section for details).
- At the time of writing the DoA, ChEESE-2P had planned to focus on 3 systems: Leonardo (classical GPUs with CUDA), LUMI (AMD APUs), and MareNostrum-5 (classical GPUs). Leonardo was of particular interest, as CUDA was widely available and porting of the codes relatively well understood for all codes in ChEESE-1P. On the other hand, LUMI's architecture (AMD APUs) was not widely available and required significant porting effort. At the start of the project, only LUMI was operationally available. Leonardo became fully operational and accessible at around M7 and MN-5 is still not operational at the time of writing.
- Added to these, sustained maintenance disruptions occurred (e.g. on Leonardo for several weeks and in coincidence with the start of our audit campaign). Overall, limited or delayed access to EuroHPC systems caused additional effort and some delays for T2.1 and D2.1.

Availability of performance analysis tools on EuroHPC systems

- Performance analysis relies on the availability of a number of tools. Due to the expertise amongst the ChEESE partners, we decided to employ one or more of mpiP, Extrae+Paraver, Score-P+Scalasca, and Talp/DLB for the different WP2 activities. None of the EuroHPC systems offered a reasonable subset of these. On the other

hand, some of these tools are highly specific to the particular architecture, the available parallel programming libraries and compilers, etc. As a result, the ChEESE staff had to deploy customised installations of the tools on Leonardo, often without previous skills. On LUMI, it was not possible to share disk space across different compute projects, and we provided detailed step-by-step instructions to code developers for self-installing mpiP, Extrae and Talp/DLB.

Shortcomings of the performance audit tools


- Some codes in ChEESE (SeisSol, ExaHyPE) make use of CUDA mapped memory (zero-copy memory) in order to avoid explicit data transfers between host and device. This is not properly implemented in the audit tools at stack, making the audit process difficult.

5.2. Use of resources


During the reporting period there have been minor underspending deviations (mostly in staff cost category).

- Underspending more accentuated in WP3, WP4 and WP7
- Initial recruitment issue of IT-skilled staff (this is being corrected)


The plots below show the use of PMs per WP, and the use of resources by partners. Only a 67% of the budget justified at M12 (assuming a linear interpolation from M1 to M48)



USE OF BUDGET YEAR 1



STAFF EFFORT YEAR 1



Revisions for the financial part:

- For INGV: A deviation of more than 25% on average personnel costs:
This deviation is due to the fact that their national collective agreement (CCNL 2019-2021) was updated at the end of 2022 increasing the annual costs of staff compared to the estimates contained in the contract. In addition, "actual personnel costs" include items (such as allowances under internal regulations) that are not calculated in the tabular costs we use for the preparation of project budgets in the proposals. Despite this, we are committed to delivering the planned action and achieving our outlined objectives as set up in the Grant Agreement.
- For IMO:
 - Specify the nature of the travel cost.
The nature of travel cost with total EUR 1,462, is the trip of IMO's IP, Dr. Benedikt Halldorson to attend the kick off meeting that took place in Barcelona, Spain 18. - 21.1.2023.
 - A deviation of more than 25% on average personnel costs
The reason for the average personnel cost to be higher than planned in RP1 is that the postdoc was not hired until 2024 in RP2. Therefore, the average

personnel cost is higher than expected in this RP since more senior staff members worked in the project..

- For CINECA: A deviation of more than 25% on average personnel costs: *The CINECA average personnel cost is higher than expected due to a high involvement of senior researchers during this first reporting period. In addition to this, the cost associated with the hours worked by senior employees had a high impact on the average cost because of the total worked hours. In the future we plan a larger involvement of junior researchers, and the average personnel cost will be closer to the budgeted one. Therefore, we are committed to deliver the expected work and fulfil our obligations.*
- For IPGP (AE of beneficiary CNRS): A deviation of more than 25% on average personnel costs: *During the preparation of proposal, the average personnel cost was estimated at 6200€/month. The actual cost of IPGP personnel participating in the project (Alexandre Fournier) is twice as high in part because he got promoted to Professeur de classe exceptionnelle after the proposal was submitted. Please note that IPGP will not request more EC contributions than initially planned. Regardless, IPGP will meet the objectives as outlined in the Grant Agreement. Alexandre Fournier has a permanent position and he will make sure that the objectives are met, in particular those objectives connected with Deliverable D6.8 led by IPGP.*
- For BULL: Personnel costs for BULL are categorised under a2 (unit costs) which is not in line with the DoA Annex 2 a1 (actual costs): *Personnel costs for BULL were erroneously categorized under A2 (unit costs) during the proposal phase. BULL consistently utilizes category A1 (actual costs) for expense reporting.*
- For SU: A deviation of more than 25% on average personnel costs: *The deviation observed in the first reporting period can be attributed to the hiring of the postdoctoral researcher in the second year, as this is when the majority of the work was scheduled to be undertaken. The work carried out during the first reporting period was done by senior staff, which explains the higher personnel costs. In any case, our institution will comply with all our obligations as outlined in the Grant Agreement.*

5.2.1. Unforeseen subcontracting (if applicable)

Not applicable

5.2.2. Unforeseen use of in kind contributions

Not applicable

Appendix 1: Exascale Computational Challenges (ECCs)

Exascale Computational Challenges (ECCs) in Solid Earth (SE) span a wide range of scales and disciplines and relate to fundamental scientific problems in understanding the Earth system from its core to the surface environment. Up to 12 ECCs on different SE domains underpin the ChEESE-2P scientific ambition:

Number	Solid Earth domain	ECC
ECC1	Computational seismology	Full-waveform inversion (FWI) and high-resolution sub-surface imaging
ECC2		Near real-time seismic scenarios
ECC3		Physics-based Probabilistic Seismic Hazard Assessment (PSHA)
ECC4	Tsunami modelling	Near real-time tsunami source inversion and characterization
ECC5		Faster Than Real Time (FTRT) tsunami scenarios
ECC6		Probabilistic Tsunami Hazard Assessment (PTHA)
ECC7	Computational volcanology	High-resolution atmospheric volcanic dispersal
ECC8		Volcanic explosions and interactions
ECC9		Probabilistic Volcanic Hazard Assessment (PVHA)
ECC10	Magnetohydrodynamics	The dynamo model and the Earth's magnetic field evolution
ECC11	Geodynamic modelling	Landslide triggering
ECC12		Lithospheric stress states and induced seismicity

Appendix 2. Flagship codes and scalability levels

ChEESE-2P considers 11 flagship codes: 4 in computational seismology, 1 in tsunami modelling, 2 in volcanology, 1 in magnetohydrodynamics, 2 in geodynamics and 1 in modelling glacier hazards. The Table below lists the flagship codes including code level, area, existence of mini-apps, lead partner and involved PDs. Area codes are: CS computational seismology, T tsunami modelling, MHD magnetohydrodynamics, V volcanology, GD geodynamics, GL glacier hazards.

No	Code	Level	Area	Mini-app	Lead	PDs
1	SeisSol	1	CS, T	yes	LMU	PD1, PD4
2	SPECFEM3D	1	CS	yes	CNRS	PD1, PD2
3	ExaHyPE	2	CS, T	no	TUM	PD1, PD4
4	Tandem	2	CS	yes	LMU	PD1
5	xSHELLS	1	MHD	yes	CNRS	PD7
6	HySEA	2	T	yes	UMA	PD3, PD4
7	FALL3D	2	V	yes	CSIC	PD5
8	OpenPDAC	2	V	no	INGV	PD6
9	LaMEM	1	GD	no	UM	PD8
10	pTatin3D	1	GD	yes	CNRS	PD8
11	Elmer/Ice	2	GL	no	CSC	PD9

In total, 6 codes are legacy from ChEESE-1P (SeisSol, SPECFEM3D, ExaHyPE, xSHELLS, HySEA and FALL3D) and 5 are new ChEESE-2P applications (Tandem, OpenPDAC, LaMEM, pTatin3D and Elmer/Ice). In terms of (strong) scalability and parallel performance, we define the "code level" as follows:

- **Level 0 codes.** Codes able to perform full capability runs or capacity workloads on exascale machines ("future tier-0 systems"). Good scalability (>90% efficiency) proven in more than 10,000 nodes (>10⁵-10⁶ cores). No tests at this scale exist yet.
- **Level 1 codes.** Codes able to perform full capability runs or capacity workloads on current pre-exascale machines ("future tier-1 systems" of the upcoming exascale era). Good scalability (>90% efficiency) proven in more than 1000 nodes (>10⁴-10⁵ cores). Up to 5 ChEESE flagship codes are considered to be already at or close to level 1 (SeisSol, SPECFEM3D, xSHELLS, LaMEM and pTatin3D).
- **Level 2 codes.** Codes able to perform full capability runs or capacity workloads on current petascale machines ("future tier-2 systems" of the upcoming exascale era).

era). Good scalability (>90% efficiency) proven in more than 100 nodes (>10³-10⁴ cores). Up to 6 ChEESE flagship codes are considered to be already at or close to level 2 (ExaHyPE, Tandem, HySEA, FALL3D, OpenPDAC and Elmer/Ice).

Appendix 3. Pilot Demonstrators

Pilot Demonstrators (PDs) are a proof of concept to address the 12 target ECCs and can result in several capability, capacity, or urgent computing Simulation Cases (SCs) depending on the underlying workflow and model configuration. ChEESE-2P considers 9 PDs underpinned by concepts like multiscale, multiphysics and multisource. The Table below lists the PDs in ChEESE-2P showing the corresponding area, involved flagship codes, and associated Exascale Computational Challenge (ECC). Area codes are: CS computational seismology, MHD magnetohydrodynamics, V volcanology, T tsunami, GD geodynamics, GL glacier hazards.

No	PD name	Area	Related flagship codes	Initial TRL	Target TRL	Concept	ECC
1	Extreme-scale modelling of seismic hazards	CS	All CS codes	5	7	multiphysics multiscale	ECC2 ECC3
2	Joint seismic inversion and tomography	CS	SPECFEM3D	4	6	multisource	ECC1
3	Global tsunami hazard and uncertainty quantification	T	HySEA	6	8-9	multiscale	ECC6
4	Complex multi-source tsunami modelling	T	HySEA, ExaHyPE, SeisSol	3	5	multisource multiphysics	ECC4 ECC5
5	Ensemble-based volcanic dispersal at multiple scales	V	FALL3D	5	7-8	multiscale	ECC7 ECC9
6	Multiphase 3D volcanic explosion modelling	V	OpenPDAC	4	7	multiphysics multiscale	ECC8
7	The Earth's dynamo model	MHD	xSHELLS	3	5	multiscale	ECC10
8	Geodynamics to geohazards	GD	LaMEM, pTatin3D	3	5	multiscale	ECC11 ECC12
9	Glacial outburst floods	GL	Elmer/Ice	2	4	multiphysics	none

Appendix 4. Simulation Cases

The PDs materialise in 15 Simulation Cases (SCs) that target 6 capability (T6.2) and 9 capacity (T6.3) use cases. SCs will run in production on EuroHPC tier-0/tier-1 using ChEESE-2P in-kind resources and making use of different EuroHPC access modes though competitive call resources. The ambition is to run up to 4 of the capability SCs as Scientific Grand Challenges (SGCs), i.e. as “full-scale” executions (60-90% of pre-exascale machines). SCs will produce FAIR and EPOS-compliant (i.e. identified by appropriate metadata) open datasets and, in higher-reaching TRL cases, will also serve as service prototypes for EPOS-TCS and to enable urgent computing services in EuroHPC systems (T6.4). This will be done in cooperation with the IUB members for the definition of functional requirements and through participation in live demo exercises. The Table below lists the SCs resulting from executing PDs in production runs and expected results. Nomenclature: SC1.2 means second SC derived from PD1, SC6.1 means first SC derived from PD6, and so on.

No	SC name	Type of run	SGC	Task(s)	Expected outcome
SC1.1	Anthropogenic noise and seismicity, wave propagation at higher frequencies in Rade d'Hyères	capability	yes	T6.2	open datasets
SC1.2	Next generation of earthquake simulators for the Hellenic arc and Iceland	capacity	-	T6.3	open datasets (earthquake catalogues)
SC1.3	Physics-based PSHA and urgent UQ for Iceland	capacity	-	T6.2,T6.3	UC service enabling
SC2.1	Seismic tomography of the Adriatic-Ionian region	capability	yes	T6.2	open datasets
SC3.1	The Global Tsunami Hazard Map	capacity	-	T6.3	TCS-TSU service in EPOS
SC4.1	Multi-source local high-resolution scenario with emphasis on the 2022 Tonga case	capability	yes	T6.2,T6.3	UC service enabling
SC5.1	Towards the European tephra hazard map	capacity	-	T6.3	TCS-VO service in EPOS
SC5.2	Volcanic dispersal at local scale: a reanalysis of the 2021 La Palma case	capacity	-	T6.2,T6.3	UC service enabling
SC6.1	Urgent high-resolution, 3D multiphase flow simulation	capability	-	T6.2,T6.3	UC service enabling

	of phreatic eruptions at Vulcano				
SC6.2	Long-term probabilistic hazard maps for phreatic eruptions at Vulcano island	capacity	-	T6.3	open datasets
SC7.1	MHD turbulence closure models in planetary cores using AI	capacity	-	T6.3	open datasets
SC8.1	Magmatic systems and induced landslides: the Etna case	capability	yes	T6.2	open datasets
SC8.2	3D evolution of tectonic deformation as driver of seismic hazard in the Hellenic arc	capability	-	T6.2	open datasets
SC9.1	Eastern Skaftá cauldron GLOF simulation, Iceland	capacity	-	T6.3	open datasets
SC9.2	Calving glacier hazard in Greenland	capacity	-	T6.3	open datasets

Appendix 5. EuroHPC systems

At the time of writing (Jan 2024), EuroHPC has 8 systems summing 15 different partitions in total and with the following characteristics:²

System	PFI ops	Partition	CPU	GPU	GFlop / watt	gCO ₂ / PFlop
LUMI @CSC (Finland)	386	1.1	LUMI-C 2,048 nodes 64 AMD EPYC	-	5,1	8,1
		1.2	LUMI-G 2,978 nodes 64 AMD Trento	4 AMD MI250X	51,3	0,8
LEONARDO @CINECA (Italy)	246	2.1	Booster (hybrid) 3,456 nodes 32 Ice Lake	4 NVIDIA A100	31,1	3,5
		2.2	Data centric 1,536 nodes 2 Sapphire	-	-	-
MN5 @BSC (Spain)	178	3.1	GPP Intel Sapphire	-	?	?
		3.2	ACC Intel Sapphire	NVIDIA H100	?	?
		3.3	NGT-GPP NVIDIA Grace (ARM)	-	?	?
		3.4	NGT-ACC Intel Emerald	Intel Rialto	?	?
VEGA @IZUM (Slovenia)	7	4.1	CPU 960 nodes 64 AMD EPYC	-	-	-
		4.2	GPU 60 nodes 64 AMD EPYC	4 NVIDIA A100	27	3,1
MELUXINA @LuxProv (Luxemb)	12	5.1	CPU 573 nodes 64 AMD EPYC	-	6,0	16,6
		5.2	GPU 200 nodes 32 AMD EPYC	4 NVIDIA A100	26,9	6,0
KAROLINA @NSC	9	6.1	CPU 720 nodes 64 AMD EPYC	-	5,6	26,4

² More info: https://eurohpc-ju.europa.eu/about/our-supercomputers_en

(Czech)		6.2	GPU	72 nodes 64 AMD EPYC	8 NVIDIA A100	27,2	5,4
DISCOV @STP (Bulgaria)	4	7.1	CPU	1,128 nodes 64 AMD EPYC	-	?	?
DEUCALION @MAAC (Portugal)	7	8.1	hybrid	1,632 nodes ARM A64FX	33 nodes NVIDIA A100	15	3,3

Appendix 6. Composition of the IUB

No	Name	Full IUB member	Signed NDA	Observers
1	Italian Civil Protection Department (DCP)	X	NOT YET	
2	Spanish Geographic National Institute (IGN)	X	YES	
3	European Plate Observing System (EPOS)	X	YES	
4	ARISTOTLE	X	YES	
5	AuScope	X	YES	
6	National Computational Infrastructure Australia (NCI)	X	YES	
7	Mexican National Seismic Service (SSN)	X	YES	
8	Buenos Aires Volcanic Ash Advisory Center	X	YES	
9	CAT-INGV	X	YES	
10	Centro PLINIVS	X	YES	
11	Icelandic Civil Protection			X
12	International Association of Seismology and Phys. Earth's Interior (IASPEI)			X
13	International Association of Volcanology and Chemistry of the Earth's Interior (IAVCEI)			X
14	Geohazard Supersites and Natural Laboratory (GEO-GSNL)			X
15	REPSOL			X
16	Schlumberger			X
17	Global Parametrics			X
18	ARM			X
19	FAULT2SHA			X
20	Modeling Collaboratory for Subduction (MCS)			X
21	Norwegian Water Resources and Energy Directorate (NVE)			X
22	International Center for Earth Simulation (ICES)			X
23	NVIDIA			X

24	Global Earthquake Modeling Foundation (GEMF)			X
25	Norwegian Meteorological Institute			X

Appendix 7. ChEESE Technology Readiness Levels (TRL)

Level	TRL Definition (H2020)	TRL for ChEESE PDs	Target
1	Basic principles observed	Equations and Quantities of Interest (QoI) defined. HPC parallel numerical code deployed.	Developers
2	Technology concept formulated	Efficiency metrics defined and strategy towards exascale preliminarily assessed.	Developers
3	Experimental proof of concept	Proof of scalability and efficiency of individual components. Validation performed on benchmark problems.	Developers
4	Technology validated in lab	Proof of models/data interoperability in workflows. Validation performed on benchmark problems.	Developers
5	Technology validated in relevant environment (industrially relevant environment in the case of key enabling technologies)	Individual components of the prototype service demonstrated in relevant environments. Error checking.	Field specialists
6	Technology demonstrated in relevant environment (industrially relevant environment in the case of key enabling technologies)	Component integration and interoperability. Use-case tests demonstrated in relevant environments. Assessment of required resources.	Field specialists
7	System prototype demonstration in operational environment	Individual components of the prototype service demonstrated in operational environments. Training of users.	End-users
8	System complete and qualified	Component integration and interoperability. Use-case tests demonstrated in operational environments. Assessment of Ethical aspects	End-users
9	Actual system proven in operational environment (competitive manufacturing in the case of key enabling technologies; or in space)	The service is deployed and operation-ready for different scenarios. Agreement on Ethical and Liability issues.	End-users

Appendix 8. Peer-reviewed publications (M14)

No	Title	DOI	Journal	Q
1	The EU Center of Excellence for Exascale in Solid Earth (ChEESE): Implementation, results, and roadmap for the second phase	10.1016/j.future.2023.04.006	Future Generation Computer Systems	Q1
2	Supplementary Material: Numerical Simulations of Seismoacoustic Nuisance Patterns from an Induced M 1.8 Earthquake in the Helsinki, Southern Finland, Metropolitan Area	10.5281/zenodo.8056327	Bulletin of the Seismological Society of America	Q1
3	Reconstructing tephra fall deposits via ensemble-based data assimilation techniques	10.5194/gmd-16-3459-2023	Geoscientific Model Development	Q1
4	11th EGU Galileo Conference: Solid Earth and Geohazards in the Exascale Era Consensual Document	10.20350/digitalCSIC/15439	Digital CSIC	NA
5	Modeling and Quantifying Parameter Uncertainty of Co-Seismic Non-Classical Nonlinearity in Rocks	10.1029/2023JB027149	Journal of Geophysical Research: Solid Earth	Q1
6	The complex dynamics of the 2023 Kahramanmaraş, Turkey, Mw 7.8-7.7 earthquake doublet	10.1126/science.ad0685	Science	Q1
7	How Does Thermal Pressurization of Pore Fluids Affect 3D Strike-Slip Earthquake Dynamics and Ground Motions?	10.1785/0120220205	Bulletin of the Seismological Society of America	Q1
8	Dynamic rupture models, fault interaction and ground motion simulations for the segmented Húsavík-Flatey Fault Zone, Northern Iceland	10.1029/2022JB025886	Journal of Geophysical Research: Solid Earth	Q1
9	Dueling dynamics of low-angle normal fault rupture with splay faulting and off-fault damage	10.1038/s41467-023-37063-1	Nature Communications	Q1
10	Broadband dynamic rupture modeling with fractal fault roughness, frictional heterogeneity, viscoelasticity and topography: The 2016 Mw 6.2 Amatrice	10.1029/2022GL098872	Geophysical Research Letters	Q1
11	Uncertainties, complexities and possible forecasting of Volcán de Colima energy emissions (Mexico,	10.5194/npg-30-571-2023	Non linear Processes in Geophys.	Q1

	years 2013–2015) based on a fractal reconstruction theorem			
12	Eruption plumes extended more than 30 km in altitude in both phases of the Millennium eruption of Paektu (Changbaishan) volcano	10.1038/s43247-023-01162-0	Nature Communications	Q1
13	A diffuse interface method for earthquake rupture dynamics based on a phase-field model	10.1029/2023JB027143	Journal of Geophysical Research: Solid Earth	Q1
14	Dynamics, interactions and delays of the 2019 Ridgecrest rupture sequence	10.1038/s41586-023-05985-x	Nature	Q1
15	Equivalent Near-Field Corner-Frequency Analysis of 3D Dynamic Rupture Simulations Reveals Dynamic Source Effects	10.1785/0220230225	Seismological Research Letters	Q1
16	Instantaneous physics-based ground motion maps using reduced- order modeling	10.1029/2023JB026975	Journal of Geophysical Research: Solid Earth	Q1

Appendix 9. Conferences and events (M14)

No	Partner	Details	Date
Participation in a conference			
1	CSIC	Arnau Folch presents the second phase of ChEESE at the EGU General Assembly 2023 with the short talk "The euroHPC Center of Excellence for Exascale in Solid Earth"	2023-04-23
2	CSC	Tómas Jóhannesson presents "Numerical modelling of rapidly-rising glacier outburst floods" as a short talk at the EGU General Assembly 2023	2023-04-25
3	UMA	Marc de la Asunción talked about the development of Tsunami-HySEA in the talk "Towards exascale-prepared codes for tsunami simulation" at the 11th EGU Galileo Conference: "Solid Earth and Geohazards in the Exascale Era"	2023-05-23
4	CSIC	Arnau Folch presents the second phase of ChEESE in the short talk "HPC projects in the Solid Earth ecosystem" at the 11th EGU Galileo Conference.	2023-05-23
5	CINECA	Piero Lanucara explains ChEESE-2P with the short talk "Leonardo: A Simulator4Earth" at the 11th EGU Galileo Conference.	2023-05-23
6	CSIC	Alejandra Guerrero presents the poster "Volcanic ash dispersal and deposition workflow on HPC" at the 11th EGU Galileo Conference: "Solid Earth and Geohazards in the Exascale Era"	2023-05-24
7	IMO	Farnaz Bayat presented the poster "A first look at the calibration of near-fault motion models to synthetic big data from CyberShake's application to the Southwest Iceland transform zone" at the 11th EGU Galileo Conference.	2023-05-25
8	CNRS	Vadim Monteiller presented "The challenges of full-wave HPC modelling of seismo-acoustic propagation generated by the detonation of UXOs of large charge weights in a shallow water environment" at the international conference UACE 2023 (Underwater Acoustics Conference & Exhibition) in Greece	2023-06-26
9	BSC	Marisol Monterrubio speaks in an invited talk about "Machine-Learning based Estimator for ground motion Shaking maps" for the Seismology and Artificial Intelligence congress.	2023-09-01
10	CSC	Thomas Zwinger explains ChEESE work with the presentation "Pilots with workflows issues Glacier hazards" at the workshop "Next-generation HPC workflows for natural hazards"	2023-09-13
11	BSC	Rut Blanco presents the poster "Machine-Learning based Estimator for ground motion Shaking maps" on the AGU	2023-12-13
Organisation of workshops and other events			

1	BSC, CSIC, CINECA	Webinar: ChEESE training event on EOSCs synergy tools for SQaaS	2023-03-24
2	ALL	Organisation of the 11th EGU Galileo Conference: Galileo Conference: Solid Earth and Geohazards in the Exascale Era.	2023-05-23
3	EVIDEN	Eviden presented ChEESE at TERATEC Forum 2023 - Europa Village	2023-06-01
4	BSC, CSIC, CINECA	Webinar: Fortran intrinsic parallelism and its application to GPU programming	2023-06-13
5	BSC, CSIC, CINECA	Webinar: OpenCL for GPU computing.	2023-06-28
6	BSC, CSIC, CINECA	Webinar: Kokkos - A portable programming library for the exascale era	2023-07-07
7	ALL	ChEESE supported and participated in the Castiel initiative: Code of the month	2023-08-15
8	BSC	Co-organisation with eFlowsHPC and DT-GEO of the workshop "Next-generation HPC workflows for natural hazards"	2023-09-13
9	ALL	Castiel Social Media Strategy	2023-09-14
10	BSC, CSIC, CINECA	Webinar: Device accelerated solvers with PETSc	2023-12-20

Appendix 10. Communication and Dissemination (M14)

No	Partners	Type of activity	Details	Starting Date
1	CSIC	Print Material	Print material to promote ChEESE: roll up	2023-01-05
2	CSIC	Press release	Initial press release about ChEESE-2P.	2023-01-07
3	CSIC	Communication Campaign	ChEESE mentioned on EuroHPC JU. Link .	2023-01-26
4	CSIC	Communication Campaign	ChEESE mentioned on Inside HPC. Link	2023-02-03
5	CSC	Communication Campaign	CSC work on ChEESE covered on the CSC website. Link .	2023-02-09
6	CSIC	Communication Campaign	Article about ChEESE-2P: "Building on success: continuation of ChEESE". Link	2023-02-10
7	All	Communication Campaign	Social Media Campaign to promote International Day of Women and Girls in Science	2023-02-11
8	ALL	Website	The post "Second phase of ChEESE, the European HPC Centre of Excellence in Solid Earth takes off" is published on ChEESE website. Link .	2023-02-11
9	CSIC	Website	ChEESE mentioned on ANSA. Link .	2023-03-02
10	CSIC	Print Material	Print material to promote ChEESE: brochures	2023-03-02
11	INGV	Communication Campaign	ChEESE mentioned on Il Giornale della Protezione Civile. Link .	2023-03-02
12	INGV	Communication Campaign	Story about ChEESE on the official INGV website. Link .	2023-03-04
13	BSC	Communication Campaign	News regarding the establishment of the associated Computational Geophysics Unit on the BSC website. Link .	2023-03-09
14	CSIC, INGV, CINECA	Website	ChEESE mentioned on Ingenio IT. Link .	2023-03-14
15	ALL	Communication Campaign	Booth at the ISC High Performance 2022 with EuroHPC JU	2023-05-29
16	ALL	Website	The post "Explore the Galileo Conference Presentations: A Gateway to Exascale Geosciences" is published on ChEESE website. Link :	2023-06-03

17	ALL	Website	The ChEESE endorses a successful Galileo Conference on Solid Earth and Geohazards in the Exascale Era, details of which can be found on the ChEESE website.	2023-06-05
18	ALL	Website	"Delve into the 'Galileo Conference Presentations: A Gateway to Exascale Geosciences' through the latest post on the ChEESE website. Link ."	2023-06-07
19	ALL	Print Material	Print material to promote ChEESE: brochures	2023-07-17
20	BSC	Communication Campaign	Interview with Marisol Monterrubio about "Supercomputing and AI to Minimize the Impact of Earthquakes" and the impact of ChEESE.	2023-09-15
21	UMA	Communication Campaign	ChEESE work is mentioned on the TV show "Los reporteros" from Canal Sur, discussing the work of EDANYA Research Group and the University of Málaga.	2023-09-16
22	ALL	Website	"Marking the 'International Day for Disaster Risk Reduction,' the ChEESE website features a dedicated post to this significant day. Link ."	2023-10-13
23	ALL	Website	"A pivotal task in 'Performance Portability within WP2 of ChEESE-2P has been detailed in a new post on the ChEESE website. Link ."	2023-11-01
24	ALL	Communication Campaign	A video outlining the objectives of ChEESE was showcased at the EuroHPC JU stand during SC 2023 in Denver.	2023-11-12
25	UMA	Communication Campaign	Science Communicator Emilio Doménech delves into Spain's tsunami risk in the feature 'Crisis Cabinet: Tsunami.' Notable experts from UMA such as Jorge Macías, who are key contributors to ChEESE and participate in the discourse.	2023-11-22
26	UMA	Communication Campaign	Article from Spanish newspaper CadizDirecto featuring a tsunami simulation for the ChEESE project by the UMA partners.	2023-12-01
27	UMA	Communication Campaign	Premiere of an engaging film featuring ChEESE simulations offering an in-depth look at the earthquake and tsunami threats facing the Iberian Peninsula: "¿Estamos preparados para el tsunami?"	2023-12-11

28	BSC	Communication Campaign	In the article titled 'BSC Inaugurates MareNostrum 5, Marking a Historic Investment in Spanish and European Science,' HPCWire mentions the project of ChEESE CoE. Link .	2023-12-21
29	BSC	Communication Campaign	Spanish news outlet SINC features ChEESE CoE's groundbreaking project: 'The Barcelona Supercomputing Center's quest to craft a digital twin of Earth for climate change assessment.' Link .	2023-12-21
30	BSC	Communication Campaign	20minutos, a Spanish news platform, discusses ChEESE CoE in an article detailing 'The pivotal elements of the Marenostrum 5 supercomputer, its objectives, and its prospective advantages'" Link	2023-12-21
31	All	Communication Campaign	Video about ChEESE first year as a project and the differences between the first and the second phase. Posted on social media: linkedin & twitter.	2024-01-01
32	UMA	Communication Campaign	Science Communicator Rob Mueller-Albrecht talks about ChEESE on his article "HySEA Tsunami Simulation Aided by SYCL" in his blog platform Medium. Link .	2024-01-11
33	UMA	Communication Campaign	ChEESE appears on the documentary on Spanish TV show "Comando Actualidad" about "La España Inundable" ("Floodable Spain").	2024-01-20