EUROPEAN COMMISSION
Communications Networks, Content and Technology

Digital Excellence & Science infrastructure
High Performance Computing & Quantum Technology

AMENDMENT Reference No AMD-675191-12

Grant Agreement: 675191 — Excellence in SImulation of Weather and Climate in Europe (ESiWACE)

The parties agree to amend the Grant Agreement as follows ('**Amendment**'):

1. Change of Annex 1 (description of the action)

Annex 1 is changed and replaced by the Annex 1 attached to this Amendment.

All other provisions of the Grant Agreement and its Annexes remain unchanged.

This Amendment **enters into force** on the day of the last signature.

This Amendment **takes effect** on the date on which the amendment enters into force, except where a different date has been agreed by the parties (for one or more changes).

Please inform the other members of the consortium of the Amendment.

SIGNATURES

For the coordinator

For the Commission

Thomas LUDWIG with ECAS id nludtoma signed in the Participant Portal on 28/10/2016 at 11:44:42 (transaction id SigId-6737-lavx2OhylUzQkxE0rzizK2ae1MOvpQ7EklYfcbzxawMFWOJAMvYmzy3b8mMDyvFQzsFCGYgu7VzIc90qQhjkGL3K-PHslUMVSXYCSODP6fhoRgS-GjtSQ2n7taBHfbagIFK8pPWfzkmbN2RkKNfTgnit0W0). Timestamp by third party at
Fri Oct 28 12:45:01 CEST 2016

Enclosures:

Annex 1

EUROPEAN COMMISSION
Communications Networks, Content and Technology
High Performance Computing & Quantum Technology

ANNEX 1 (part A)

Research and Innovation action

NUMBER — 675191 — ESiWACE

Table of Contents

1.1. The project summary....3
1.2. The list of beneficiaries....4
1.3. Workplan Tables - Detailed implementation....5
1.3.1. WT1 List of work packages....5
1.3.2. WT2 List of deliverables....6
1.3.3. WT3 Work package descriptions....9
Work package 1....9
Work package 2....13
Work package 3....19
Work package 4....25
Work package 5....29
1.3.4. WT4 List of milestones....33
1.3.5. WT5 Critical Implementation risks and mitigation actions....34
1.3.6 WT6 Summary of project effort in person-months....36
1.3.7. WT7 Tentative schedule of project reviews....37

1.1. The project summary

Project Number [1]	675191	Project Acronym [2]	ESiWACE

One form per project

General information	
Project title [3]	Excellence in SImulation of Weather and Climate in Europe
Starting date [4]	01/09/2015
Duration in months [5]	48
Call (part) identifier [6]	H2020-EINFRA-2015-1
Topic	EINFRA-5-2015 Centres of Excellence for computing applications
Fixed EC Keywords	Computer sciences, information science and bioinformatics
Free keywords	weather, climate, model optimizations, exa-scale, HPC, scalability, usability, exploitability, workflow, software governance

Abstract [7]

ESiWACE will substantially improve efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms by supporting the end-to-end workflow of global Earth system modelling in HPC environment. This will be obtained by improving and supporting (1) scalability of models, tools and data management on state-of-the-art supercomputer systems (2) Usability of models and tools throughout the European HPC eco-system, and (3) the Exploitability of the huge amount of resulting data.We will develop solutions for cross-cutting HPC challenges particular to the weather and climate domain. This will range from the development of specific software products to the deployment of user facing services for both, computing and storage.ESiWACE leverages two established European networks, namely (1) the European Network for Earth System modelling, representing the European climate modelling community and (2) the world leading European Centre for Medium-Range Weather Forecasts. The governance structure that defines the services to be provided will be driven by the European weather and climate science community.Weather and climate computing have always been one of the key drivers for HPC development, with domain specific scientific and technical requirements that stretch the capability and capacity of existing software and hardware to the limits. By developing solutions for Europe and at European scale, ESiWACE will directly impact on the competitiveness of the European HPC industry by engendering new products, providing opportunities for exploitation beyond the project itself, and by enhancing the skills base of staff in both industry and academia.ESiWACE will be at once thematic, as it focuses on the HPC application domain of climate and weather modeling, transversal, as it covers several aspects of computational science, and challenge-driven, as climate and weather predictability represents a major societal issue.

1.2. List of Beneficiaries

Project Number [1]	675191	Project Acronym [2]	ESiWACE

List of Beneficiaries

No	Name	Short name	Country	Project entry date[8]	Project exit date
1	DEUTSCHES KLIMARECHENZENTRUM GMBH	DKRZ	Germany		
2	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	ECMWF	United Kingdom		
3	CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS	CNRS-IPSL	France		
4	MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV	MPG	Germany		
5	CENTRE EUROPEEN DE RECHERCHE ET DE FORMATION AVANCEE EN CALCUL SCIENTIFIQUE	CERFACS	France		
6	BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION	BSC	Spain		
7	SCIENCE AND TECHNOLOGY FACILITIES COUNCIL	STFC	United Kingdom		
8	MET OFFICE	METO	United Kingdom		
9	THE UNIVERSITY OF READING	UREAD	United Kingdom		
10	SVERIGES METEOROLOGISKA OCH HYDROLOGISKA INSTITUT	SMHI	Sweden		
11	NATIONAL UNIVERSITY OF IRELAND, GALWAY	ICHEC	Ireland		
12	FONDAZIONE CENTRO EURO-MEDITERRANEO SUI CAMBIAMENTI CLIMATICI	CMCC	Italy		
13	DEUTSCHER WETTERDIENST	DWD	Germany		
14	SEAGATE SYSTEMS UK LIMITED	SEAGATE	United Kingdom		
15	BULL SAS	Bull	France		
16	ALLINEA SOFTWARE LIMITED	ALLINEA	United Kingdom		

1.3. Workplan Tables - Detailed implementation

1.3.1. WT1 List of work packages

WP Number[9]	WP Title	Lead beneficiary[10]	Person-months[11]	Start month[12]	End month[13]
WP1	Governance, Engagement & long-term Sustainability	3 - CNRS-IPSL	47.00	1	48
WP2	Scalability	2 - ECMWF	202.00	1	48
WP3	Usability	4 - MPG	123.00	1	48
WP4	Exploitability	7 - STFC	106.00	1	48
WP5	Management & Dissemination	1 - DKRZ	36.00	1	48
		Total	514.00		

1.3.2. WT2 list of deliverables

Deliverable Number[14]	Deliverable Title	WP number[9]	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D1.1	Agreed portfolio of community tools	WP1	3 - CNRS-IPSL	Report	Public	6
D1.2	Roadmap to the implementation of 1km ESM ensembles	WP1	1 - DKRZ	Report	Public	46
D1.3	Business plan	WP1	1 - DKRZ	Report	Public	38
D1.4	First International HPC workshop	WP1	5 - CERFACS	Other	Public	7
D1.5	Second International HPC workshop	WP1	12 - CMCC	Other	Public	31
D2.1	Final workshop report on results and evaluation of demonstrators	WP2	1 - DKRZ	Report	Public	46
D2.2	Optimised community model code options tested on selected cases	WP2	6 - BSC	Other	Public	36
D2.3	Multithreaded or thread safe OASIS version including performance optimizations to adapt to many-core architectures	WP2	5 - CERFACS	Other	Public	24
D2.4	Multithreaded XIOS version adapted to many-core architectures, and supporting GRIB2 format	WP2	3 - CNRS-IPSL	Other	Public	36
D2.5	White paper on a strategy for full convergence of I/O and coupling tools	WP2	3 - CNRS-IPSL	Report	Public	36
D2.6	Report outlining a strategic approach for efficiency savings based	WP2	2 - ECMWF	Report	Public	36
D2.7	Report from study of data compression assessment	WP2	2 - ECMWF	Report	Public	36
D2.8	Implementation of IFS global 1 km atmosphere-only demonstrator and performance analysis	WP2	2 - ECMWF	Report	Public	24

Deliverable Number[14]	Deliverable Title	WP number[9]	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D2.9	Implementation of ICON global 1km atmosphere-only demonstrator and performance analysis	WP2	1 - DKRZ	Report	Public	24
D2.10	Implementation of NEMO global 1km ocean-only demonstrator and performance analysis	WP2	3 - CNRS-IPSL	Report	Public	42
D2.11	Implementation of EC-Earth 10km global coupled demonstrator and performance analysis	WP2	6 - BSC	Report	Public	42
D2.12	Implementation of ICON 10km global coupled demonstrator and performance analysis	WP2	1 - DKRZ	Report	Public	42
D3.1	ESiWACE Application Software Framework	WP3	4 - MPG	Report	Public	4
D3.2	Update of ESiWACE Application Software Framework (D3.1) prepared for the demonstrators, Version 2	WP3	4 - MPG	Report	Public	30
D3.3	Software specification for the third E2sCMS	WP3	4 - MPG	Report	Public	6
D3.4	Experiences with ESM Multi-model Ensembles for Educational Purposes	WP3	4 - MPG	Report	Public	14
D3.5	How to select, configure and install ESM software stacks	WP3	6 - BSC	Report	Public	14
D3.6	Update Handbook for system administrators	WP3	6 - BSC	Report	Public	36
D3.7	Software Stack for ESM– A Specification:Narrowing down D3.5 for the demonstrators	WP3	6 - BSC	Report	Public	18
D3.8	A Report from the use of D3.5 for the demonstrators	WP3	1 - DKRZ	Report	Public	30

Deliverable Number[14]	Deliverable Title	WP number[9]	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D3.9	SiWACE Scheduler development and support activities	WP3	8 - METO	Report	Public	18
D3.10	ESiWACE Scheduler development and support activities, v2	WP3	8 - METO	Report	Public	34
D4.1	Business model with alternative scenarios	WP4	1 - DKRZ	Report	Public	18
D4.2	ESD middleware design	WP4	1 - DKRZ	Report	Public	18
D4.3	Final implementation of the ESD middleware	WP4	1 - DKRZ	Other	Public	46
D4.4	Final report on alternative tape usage	WP4	7 - STFC	Report	Public	42
D5.1	Design and implementation of the intranet	WP5	1 - DKRZ	Websites, patents filling, etc.	Confidential, only for members of the consortium (including the Commission Services)	2
D5.2	Project public website	WP5	1 - DKRZ	Websites, patents filling, etc.	Public	2
D5.3	Media and Communication Plan	WP5	1 - DKRZ	Report	Public	2
D5.4	Data Management Plan (DMP)	WP5	1 - DKRZ	ORDP: Open Research Data Pilot	Public	6
D5.5	Dissemination and Exploitation Plan (EP)	WP5	1 - DKRZ	Report	Public	18
D5.6	Strategy for the Intellectual Property exploitation	WP5	1 - DKRZ	Report	Public	48

1.3.3. WT3 Work package descriptions

Work package number [9]	WP1	**Lead beneficiary** [10]	3 - CNRS-IPSL
Work package title	Governance, Engagement & long-term Sustainability		
Start month	1	**End month**	48

Objectives

Objectives
WP1 focuses on the governance of the ESiWACE. It is a networking activity in support to ESiWACE. More specifically WP1 aims at:
• Community engagement and governance: WP1 aims at better engaging the climate and weather modelling communities using Earth system models in the ESiWACE and ensure the ESiWACE serves the communities.
• WP1 will ensure:
o User representation in the governance of common software and in ESiWACE new developments
o Consistency and exploitation of possible synergies with other relevant activities in which the community is engaged (infrastructure projects, FETs, environment projects…).
• Enhancing community capacity in HPC: WP1 will foster efficient use of HPC through regular exchanges with the communities. It will also organize relations with PRACE and will deal with the PRACE resource allocation dedicated to the ESiWACE.
• Networking with industry and ETP4HPC: WP1 will network with ETP4HPC and relevant hardware and software industry.
• Strengthening strategy and business planning: WP1 will develop:
o An HPC roadmap for HPC and ESM for the climate and weather communities
o Upgrade the business plan, for which a first version is included in this proposal, and prepare for future activities, taking benefit from all the ESiWACE activities

Description of work and role of partners

WP1 - Governance, Engagement & long-term Sustainability [Months: 1-48]
CNRS-IPSL, DKRZ, ECMWF, MPG, CERFACS, BSC, SMHI, CMCC, Bull
This work package is led by: Sylvie Joussaume CNRS-IPSL (lead) and Joachim Biercamp DKRZ (co-lead)

Task 1.1 Engagement and governance [Lead: CNRS-IPSL; Participants: DKRZ, CERFACS, MPG, ECMWF] Networking activity, service activity
Task 1.1.1 Liaising with the community [Lead: CNRS-IPSL; Participant: ECMWF, DKRZ]
Task 1.1.1 will:
• Address HPC issues relevant for the ESM communitiy, in particular cross-cutting issues, through discussions at the General Assembly and HPC workshops [D1.4 and D1.5] as well as through feedback from the HPC Task Force (see Task 1.2.1), Scientific Advisory Board and supporting partners.
• Whenever needed, organize dedicated task forces for specific areas (HPC, Data, Demonstrators, Networks) and special interest groups for specific software (Scalable IO, new dynamical cores, No- Posix, Software defined networks, Agile methods…), gathering ESiWACE members, supporting partners and other experts; this will offer an opportunity to link and coordinate engineering resources between different centres.
• Liaise with other European projects involving ESM and HPC and deploy this interaction as a source of user feedbacks.

Task 1.1.2 Governance on community software [Lead: CNRS-IPSL; Participants: CERFACS, MPG]
Task 1.1.2 mainly concerns heritage software further developed and supported by the ESiWACE. It will:
• Establish general governance rules for community software, identifying different levels according to the maturity of software; the governance associated with each software shared in ESiWACE (e.g. OASIS, XIOS, CDOs, Cylc, NEMO) will be defined by identifying its stakeholder, steering and user groups [D1.1]; a stronger organisation will be tested with the coupler OASIS and lessons will be used for other shared software.
• Establish criterion and procedure for prioritizing user requirements and defining the ESiWACE activities on shared community software accordingly.

Task 1.1.3 Governance on innovative ESiWACE developments [Lead: DKRZ; Participants: CNRS-IPSL]

Task 1.1.3 mainly concerns new software that is developed by ESiWACE. It will ensure that this software will be usable by a large range of users. It will:

• Establish governance rules for new ESiWACE developments); it will define stages of development and review process by users.

• Apply a user-driven approach and collect users scientific and technical requests and problems regarding their applications / tools; this will involve partners, supporting partners and, when relevant, computing centres where applications are deployed.

Task 1.2. Enhancing community capacity in HPC [Lead: CNRS-IPSL; Participants: CERFACS, ECMWF, CMCC, BSC] Networking Activity

Task 1.2.1 Community building [Lead: CNRS-IPSL Participants: CERFACS, ECMWF, CMCC]

Task 1.2.1 will:

• Enlarge the ENES task force to better include the weather community; this task force will advise the ESiWACE on HPC issues for the community, such as relations with PRACE.

• Organise annual workshops, gathering the ESM community, to share experience on the state of the art in HPC and to discuss relevant cross-cutting issues.

Task 1.2.2 Strengthening the interface with PRACE [Lead: CNRS-IPSL; Participants: BSC]

Task 1.2.2 will:

• Manage relations with PRACE, in particular regarding the ESiWACE dedicated computing allocation, ensuring that this resource is used by project partners for developments in the framework of the ESiWACE (e.g. the performance benchmarking of WP2 T2) but also, through open calls, for external partners testing software developed with the support of the ESiWACE.

• Strengthen interactions with PRACE, through reporting on community needs, participating to PRACE foresights, interaction with PRACE projects.

Task 1.3. Strategic Interaction with HPC eco-system and HPC industry [Lead: DKRZ; Participants: ECMWF, BSC] Networking Activity

Task 1.3 will liaise with ETP4HPC and the European HPC eco-system. In this task, we will:

• Map the activities of ESiWACE to the Strategic Research Agenda (www.etp4hpc.eu/wp-content/uploads/2013/06/ETP4HPC_book_singlePage.pdf) and research topic timelines of ETP4HPC and contribute to their updating and extrapolation. Ensure close links with HPC related hard- and software industry and will foster their involvement in ESiWACE activities, in particular their participation to the annual HPC workshops.

• Establish links on the technical level with PRACE and other HPC centres. This will ensure that on one hand ESiWACE developments are informed by technical constraints and on the other hand that requirements of ESM applications are considered by the centres.

Task 1.4. Sustainability and business planning [Lead: DKRZ; Participants: CNRS-IPSL, Bull, ECMWF, SMHI] Networking Activity

Task 1.4.1 Develop a roadmap to implementation of an ensemble of ESMs with 1km resolution

The analysis of the demonstrator experiments (Task 2.2, D2.8 to D2.12) will provide valuable input for estimates of performance, scalability and limits of computability. This will be of great value for deriving the technical and organisational requirements for the operational production of predictions with coupled model ensembles at 1 km resolution in both atmosphere and ocean. The results of the analysis will build the basis for a roadmap to the implementation of an ensemble of ESMs resolving clouds and small scale eddies. The roadmap will detail the technical challenges to be addressed and technical requirements to be fulfilled in order to use such a system operationally for climate and weather simulations aiming at a better understanding of high-impact regional extreme events.

Tasks 1.4.2 Business plan

Task 4.2 will take care of the update of the business plan [D1.3]. Two alternative concepts will be developed: one for a second funded phase of ESiWACE and one for continued engagement without further EU funding. The concepts will include the findings and experiences from those tasks that deal with shared and newly developed software and will lay out how continued community support can be secured. This will include alternate support concepts and the potential for integrating the developed support schemes, tools and software into national and international funding schemes. A special focus will lie on evaluation of options and propositions for concrete steps towards an infrastructure of operational services using cloud- and eddy- resolving ESMs to support the science case.

Interactions with other Work Packages

This Work Package, in strong interaction with the WP5 management tasks, provides:

• Scientific guidance and leadership for all Work Packages and to the project as a whole, in particular to WP2 on common shared software and to WP2, WP3 and WP4 on new developments.

• Support, guidance and management assistance to all Work Packages as required and to the project as a whole.

This Work Package receives:
• Input from WP2 on user feedbacks regarding application of governance defined in T1.1.2
• Input from WP2, WP3 and WP4 on HPC issues and cross-cutting activities
• Input from WP2, WP3 and WP4 on needs for PRACE ESiWACE allocation and report from the use of allocation

Participation per Partner

Partner number and short name	WP1 effort
1 - DKRZ	13.00
2 - ECMWF	4.00
3 - CNRS-IPSL	17.00
4 - MPG	2.00
5 - CERFACS	3.00
6 - BSC	4.00
10 - SMHI	1.00
12 - CMCC	1.00
15 - Bull	2.00
Total	47.00

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D1.1	Agreed portfolio of community tools	3 - CNRS-IPSL	Report	Public	6
D1.2	Roadmap to the implementation of 1km ESM ensembles	1 - DKRZ	Report	Public	46
D1.3	Business plan	1 - DKRZ	Report	Public	38
D1.4	First International HPC workshop	5 - CERFACS	Other	Public	7
D1.5	Second International HPC workshop	12 - CMCC	Other	Public	31

Description of deliverables

D1.1 Agreed portfolio of community tools: D1.1 will report on shared software. It will define general governance rules for shared software and how they are applied for each software, in
particular their user group and how it interacts with ESiWACE (CNRS-IPSL, R, PU, PM6).
D1.2 Roadmap to the implementation of 1km ESM ensembles (DKRZ, R, PU, PM46).
D1.3 Business Plan (DKRZ, R, PU, M38).
D1.4 First International HPC workshop (CERFACS, OTHER, PU, PM7)
D1.5 Second International HPC workshop (CMCC, OTHER, PU, PM31)

D1.1 : Agreed portfolio of community tools [6]

Agreed portfolio of community tools

D1.2 : Roadmap to the implementation of 1km ESM ensembles [46]
Roadmap to the implementation of 1km ESM ensembles

D1.3 : Business plan [38]
Business plan: Updated business plan

D1.4 : First International HPC workshop [7]
First International HPC workshop : First International HPC workshop

D1.5 : Second International HPC workshop [31]
Second International HPC workshop: Second International HPC Workshop

Schedule of relevant Milestones

Milestone number[18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS1	Scheduler support and development plan (T3.3.1)	8 - METO	6	Allow the first deliverable D3.9 to meet the needs of the user community
MS3	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS5	Operational demonstrator of ESD middleware	1 - DKRZ	24	Demonstrator operative

Work package number [9]	WP2	Lead beneficiary [10]	2 - ECMWF
Work package title	Scalability		
Start month	1	**End month**	48

Objectives

• Demonstrate how to build and productively operate global cloud resolving models
• Improve efficiency and scalability of community models and tools and provide improved access to scientific and performance upgrades
• Provide support and training for community I/O and coupling tools and establish convergence strategies for the future
• In the exascale perspective, Investigate scientific and technical options for efficiency upgrades and methodologies to reduce volume of future high-resolution modelling output

Description of work and role of partners

WP2 - Scalability [Months: 1-48]
ECMWF, DKRZ, CNRS-IPSL, MPG, CERFACS, BSC, SMHI, ICHEC, CMCC, Bull
This work package is led by: Peter Bauer ECMWF (lead) and Sophie Valcke CERFACS (co-lead)

Task 2.1. Support, training and integration of state-of-the-art community models and tools
[Lead: SMHI; Participants: CERFACS, BSC, ECMWF, CNRS-IPSL], Networking Activity
This task fosters improved user support, enhanced training, and fastest integration of existing community models and tools, in response to the software governance defined in WP1, leading to enforced user-driven evolution of the community software. The metrics that will be used to monitor the two sub-task are the frequency of official releases, the number of downloads,
development tickets opened and closed, bug fixes, wiki page edits, and mails exchanged with
users.

Task 2.1.1 NEMO and EC-Earth [Lead: SMHI; Participants: BSC, ECMWF, CNRS-IPSL]
NEMO and EC-Earth are leading edge community applications both driven by consortium of European countries. These models, together with ICON, will constitute the heart of our high-resolution demonstrators (see T2.2).
NEMO is a modelling framework for oceanographic research, operational seasonal forecast and climate studies and EC-Earth represents a mature Earth System Model at the interface between weather and climate modelling. This task will facilitate community wide access and use of NEMO and EC-Earth latest versions, including version control and ticketing system, diffusion of best practice in using these models and will accelerate their update cycle. In particular, it will ensure a faster integration of OpenIFS, the open source version of the atmosphere component in EC-Earth (https://software.ecmwf.int/wiki/display/OIFS/OpenIFS+Home) resulting in more rapid benefit from on-going scientific developments for the community.

Task 2.1.2 OASIS coupler and, XIOS I/O server [Lead: CERFACS; Participants: CNRS- IPSL]
Use of the well-established OASIS coupler and the newer XIOS I/O server is constantly
increasing in the climate modelling community. Beyond traditional user support in the form of a mail help-desk, on-line forum, User Guide, tutorials, FAQs, trainings and hints for best practice that will be established for XIOS and maintained for OASIS, this task will be devoted to governed integration of specific user needs in the software. Using the governance structure defined in WP1 T1.1.2, user-specific problems and requirements will be analysed, bug fixes will be developed when needed, but new strategic functionalities will possibly also be added.

Task 2.2. Global high-resolution model demonstrators [Lead: ECMWF, Participants: DKRZ, BSC, CERFACS, CMCC, BULL/ATOS, CNRS-IPSL, MPG, SMHI], Joint research activity
The key science target is to reach global models with spatial resolutions that allow simulating convective clouds and small-scale ocean eddies. Hence, global cloud resolving models need spatial resolution of 1 km and similar resolution is required for the ocean. In operations, a throughput of 100 to 1000 simulated days per calendar day is required, ideally for ensemble simulations with 50 or more members. These numbers define the long-term strategic target for operational weather and climate prediction, respectively. The current task will build on the ESiWACE community models and demonstrate the computability of atmosphere-only, ocean-only and coupled global models at the highest possible resolution that is configurable today. The task will quantify how far away current capabilities are from the long-term strategic target.

The task will employ many-core architectures that are accessible to the ESiWACE partners. The computational and data handling performance of these configurations will be evaluated (including strong scaling) and extrapolated (through weak scaling) to the global 1 km resolution coupled model target.

The task will be carried out under advice from the scientific advisory board (SAB) and the science spokesperson. The task interacts with T1.4.1 "Develop a roadmap to implementation of an ensemble of ESMs with 1km resolution" and T1.4.2 "Business plan".

Task 2.2.1 Very high resolution atmosphere-only and ocean-only demonstrators [Lead: ECMWF, Participants: BSC, DKRZ, CNRS-IPSL, CMCC]
Based on the ECMWF IFS and the DKRZ ICON models, atmosphere-only simulations will be run at 1 km global resolution for limited forecast ranges to establish feasible configurations of memory and parallel task allocations on the Cray XC-40 Broadwell (ECMWF) and Bullx Haswell/Broadwell (DKRZ) HPC systems. At least two strong scaling experiments using the minimum required and the maximum accessible allocations of compute nodes will be run, respectively. This will establish a first-order scalability estimate for the two models at this resolution. The experiments will provide a quantitative estimate of the forecast throughput in "forecast days per day" for both models and two allocations per model. Standard tools (e.g. Extrae and Paraver developed at BSC) will be employed to characterize the cost distribution for each model in terms of time spent in model dynamics and physics, computation vs communication and their dependence on node allocation. Additionally a demonstrator of ocean-only simulations in a 1km global configuration will be tested on Bullx Curie at TGCC using the NEMO ocean platform.The aim is to point out the specific scalability issues of such an extreme configuration. The performance analysis will produce conclusions on specific performance bottlenecks, inter-model performance and scalability differences, and on the consistency of the performance across HPC systems – in relation to the performance target for forecast throughput (D2.8 ECMWF, D2.9 DKRZ, D2.10 CNRS-IPSL).

Task 2.2.2 High resolution coupled atmosphere-ocean demonstrators [Lead: BSC, Participants: ECMWF, DKRZ, STFC, SMHI, CERFACS, CMCC
The performance analysis in Task 2.2.1 will be extended to fully coupled atmosphere-ocean models, namely EC-Earth (OpenIFS, NEMO and OASIS coupler) and ICON (including its ocean component and YAC coupler). EC-Earth will employ the latest version of the IFS through OpenIFS (see T2.1.1) and NEMO-ORCA12 in the TCo1279-ORCA12 configuration, i.e. at 10 km atmosphere / ~10 km ocean resolution. ICON will use the same grid for ocean and atmosphere at 10 km resolution. The EC-Earth experimentation will be configured for the Mare Nostrum 4 (BSC) while the coupled ICON simulations will be run at DKRZ (see T2.2.1). Again, the minimum required and maximum accessible node allocations will be used to establish a strong scalability estimate. The performance analysis will be extended to the coupled system, the couplers (OASIS and YAC) and the I/O servers (XIOS and CDI-pio), and it will include an extrapolation towards resource requirements for a globally coupled model at 1 km given the available HPC architectures (D2.11 BSC, D2.12 DKRZ).

Task 2.3 Efficiency enhancement of models and tools [Lead: BSC. Participants: Bull, CMCC, DKRZ, ECMWF, CNRS-IPSL, SMHI, CERFACS, ICHEC] Joint research activity
Radical refactoring of ESM codes and tools will most probably be required to take full advantage of future exascale computing platforms. Until then however, existing codes still need to be incrementally revised to fully exploit massively parallel supercomputers accessible today, in particular in order to implement the global high-resolution model demonstrator described inT2.2

Task 2.3.1 Model optimisation [Lead: BSC; Participants: Bull, CMCC, DKRZ, ECMWF, CNRS-IPSL, SMHI]
This task will assess compute/communication performance of different codes, i.e. OpenIFS and NEMO (being part of EC-Earth) as well as ICON, with analysis tools that employ adequate tracing protocols, for example Extrae/Paraver, developed by BSC. Strategies for performance enhancement will be developed and optimizations such as overlap of communication and computation (e.g. YAXT for ICON), fewer but bigger communications, two-sided communication, asynchronous communications through partitioned global address space concepts (PGAS), improved code vectorization and use of OpenMP will be implemented. This task relates to Task 3.2 of WP 3 about software stack since Fortran-PGAS is currently not supported by all compilers. [D2.2, BSC].

Task 2.3.2 Tool optimization and enhancement [Lead: CNRS-IPSL; Participants: Bull, CERFACS, ICHEC]
Even before reaching the exascale, it is clear today that key developments are needed in our community tools to fit user needs in terms of performance and parallel functionalities. In this task, the following priority issues will be addressed:
• Currently, both the OASIS coupler and the XIOS I/O server support only MPI-task parallelism and it is known that multithreading is becoming essential in climate codes. For example, 240 threads are needed to fully exploit the 60-core Intel Xeon Phi processor. The client part of OASIS and XIOS will evolve to become multithreaded and/or thread safe in

order to fully fit the calling program multithreaded structure. Other developments improving the performance of OASIS, such as accelerating the performance of the mapping by better overlapping the source and destinations decomposition, will also be realized. [D2.3, CERFACS].
• XIOS, originally developed for the climate modelling community, naturally supports the netCDF file format. To extend its use to the weather community, support of the GRIB2 format is imperatively needed. It is planned here to enhance XIOS so to support GRIB2 format in parallel mode with a multi-process server and for all types of grids, i.e. regular but also curvilinear, reduced Gaussian, and unstructured grids. This work will be based on the semantic mapping between netCDF and GRIB developed in WP4 T4.2. [D2.4, CNRS-IPSL].
• It is now clearly recognized that the main functions of an I/O server and of a coupler, i.e. communication and interpolation of data (also envisaged in I/O systems for data output reduction) are extremely alike. In this task, we will work toward a convergence of these tools by proposing a unified API (Application Programming Interface) for both I/O and coupling in component models. In particular, XIOS will be upgraded so to include full coupling functionalities: fully parallel 2nd order conservative on-the-sphere regridding functions developed in the framework of the G8 ICOMEX project, communication of data between two parallel component processes, etc. The coupling performance of XIOS will be evaluated, and if judged appropriate, a strategy for full convergence of I/O and coupling tools will be defined (D2.5, CNRS-IPSL).

Task 2.4 Preparing for exascale [Lead: ECMWF. Participants: CERFACS, DKRZ, CNRS-IPSL, SMHI, BSC], Joint research activity
This task targets few precise developments to prepare for the next-generation model components that will be assembled for exascale systems by projects like CHANCE (for NEMO) and ESCAPE (for IFS), which have been recently submitted to the call FETHPC-1-2014 of H2020: Research & Innovation Actions (RIA).

Task 2.4.1 Concurrency and accuracy [Lead: ECMWF; Participants: CERFACS, DKRZ, CNRS-IPSL]
Weather and climate models integrate equations describing the equations of motion, mass and heat and add non-resolved process contributions through physical parameterizations. These scientific modules are usually executed sequentially on the same set of computing resources and the scientific performance of the overall model strongly depends on this sequence. This task will investigate the possibility in NEMO, ICON and OpenIFS of de-sequencing and thus enhancing concurrency between e.g. expensive parts of atmospheric physical parameterizations (such as radiation) and other calculations, or the sea-ice and biogeochemistry and the ocean dynamics.
Further, while the representation of models variables in double precision is common practice to ensure accuracy, recent studies suggest that significant efficiency gains can be achieved from single-precision representation. This option will be investigated for EC-Earth, in particular running long ensemble integrations.
Linking with the performance analysis in T2.1, this task will produce a report (D2.6, ECMWF) outlining a strategic approach based on concurrency and accuracy for efficiency gains in ESMs without jeopardizing the scientific performance of the models.

Task 2.4.2 Knowledge compression [Lead: ECMWF. Participants: BSC, SMHI]
Significant growth of ensemble weather and climate model output requires new approaches to output management without deteriorating the information contained in the model integrations. The dimensions that drive output volume are saved time steps, horizontal and vertical resolution as well as ensemble size. This task will investigate avenues for mostly reducing the ensemble dimension using operational NWP model and EC-Earth ensemble output. The task will deliver a report [D2.7, ECMWF] making recommendations for knowledge compression for ensemble weather and climate model output management.

Interactions with other Work Packages
This Work Package provides:
• Through the support and training on state-of-the-art community tools (see T2.1) WP2 will establish privileged contacts with users and is therefore in a strong position to provide feedbacks to WP1 T.1.1.2 for defining the rules of a governance for community tools.
• Input to WP1 in cross cutting issues.
• Input to WP1 on the need for resource.

This Work Package receives:
• The governance established in WP1 T1.1.2 will in turn help identify key strategic functionalities to implement in these tools as requested by the users. The allocations on tier-0 platforms dedicated to ESiWACE negotiated in WP1 T1.2.2 will provide an essential resource to realize the coupling and I/O benchmarks described in T2.2.2.
• Work done in T3.2 of WP3 about software stack will be greatly facilitate realization of T2.3.1 "Model optimisation".
• Supporting GRIB2 format in XIOS (see T2.3.2) will be based on the semantic mapping between netCDF and GRIB developed in WP4 T4.2.

Participation per Partner

Partner number and short name	WP2 effort
1 - DKRZ	10.00
2 - ECMWF	49.00
3 - CNRS-IPSL	36.00
4 - MPG	3.00
5 - CERFACS	27.00
6 - BSC	22.00
10 - SMHI	20.00
11 - ICHEC	6.00
12 - CMCC	9.00
15 - Bull	20.00
Total	202.00

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D2.1	Final workshop report on results and evaluation of demonstrators	1 - DKRZ	Report	Public	46
D2.2	Optimised community model code options tested on selected cases	6 - BSC	Other	Public	36
D2.3	Multithreaded or thread safe OASIS version including performance optimizations to adapt to many-core architectures	5 - CERFACS	Other	Public	24
D2.4	Multithreaded XIOS version adapted to many-core architectures, and supporting GRIB2 format	3 - CNRS-IPSL	Other	Public	36
D2.5	White paper on a strategy for full convergence of I/O and coupling tools	3 - CNRS-IPSL	Report	Public	36
D2.6	Report outlining a strategic approach for efficiency savings based	2 - ECMWF	Report	Public	36
D2.7	Report from study of data compression assessment	2 - ECMWF	Report	Public	36

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D2.8	Implementation of IFS global 1 km atmosphere-only demonstrator and performance analysis	2 - ECMWF	Report	Public	24
D2.9	Implementation of ICON global 1km atmosphere-only demonstrator and performance analysis	1 - DKRZ	Report	Public	24
D2.10	Implementation of NEMO global 1km ocean-only demonstrator and performance analysis	3 - CNRS-IPSL	Report	Public	42
D2.11	Implementation of EC-Earth 10km global coupled demonstrator and performance analysis	6 - BSC	Report	Public	42
D2.12	Implementation of ICON 10km global coupled demonstrator and performance analysis	1 - DKRZ	Report	Public	42

Description of deliverables

D2.1: Final workshop report on results and evaluation of demonstrators (DWD,R, PU, PM49).
D2.2: Optimised community model code options tested on selected cases (BSC, OTHER, PU, M36).
D2.3: Multithreaded or thread safe OASIS version including performance optimizations to adapt to many-core architectures (CERFACS, OTHER, PU, PM24).
D2.4: Multithreaded XIOS version adapted to many-core architectures, and supporting GRIB2 format (CNRS-IPSL, OTHER, PU, M36).
D2.5: White paper on a strategy for full convergence of I/O and coupling tools (CNRS-IPSL, R, PU, PM36).
D2.6: Report outlining a strategic approach for efficiency savings based on concurrency and accuracy (ECMWF, R, PU, PM36).
D2.7: Report from study of data compression assessment of dimensions required for Earth system model output archiving retaining information but reducing volume (ECMWF, OTHER, PU, PM36).
D2.8 Implementation of IFS global 1 km atmosphere-only demonstrator and performance analysis (ECMWF, R, PU, PM24).
D2.9: Implementation of ICON global 1km atmosphere-only demonstrator and performance analysis (DKRZ, R, PU, PM24).
D2.10: Implementation of NEMO global 1km ocean-only demonstrator and performance analysis (CNRS-IPSL, R, PU, PM42).
D2.11: Implementation of EC-Earth 10km global coupled demonstrator and performance analysis (BSC, R, PU, PM42).
D2.12: Implementation of ICON 10km global coupled demonstrator and performance analysis (DKRZ, R, PU, PM42)

D2.1 : Final workshop report on results and evaluation of demonstrators [46]

Final workshop report on results and evaluation of demonstrators

D2.2 : Optimised community model code options tested on selected cases [36]

Optimised community model code options tested on selected cases

D2.3 : Multithreaded or thread safe OASIS version including performance optimizations to adapt to many-core architectures [24]

Multithreaded or thread safe OASIS version including performance optimizations to adapt to many-core architectures

D2.4 : Multithreaded XIOS version adapted to many-core architectures, and supporting GRIB2 format [36]

Multithreaded XIOS version adapted to many-core architectures, and supporting GRIB2 format

D2.5 : White paper on a strategy for full convergence of I/O and coupling tools [36]

White paper on a strategy for full convergence of I/O and coupling tools

D2.6 : Report outlining a strategic approach for efficiency savings based [36]

Report outlining a strategic approach for efficiency savings based

D2.7 : Report from study of data compression assessment [36]

Report from study of data compression assessment of dimensions required for Earth system model output archiving retaining information but reducing volume

D2.8 : Implementation of IFS global 1 km atmosphere-only demonstrator and performance analysis [24]

Implementation of IFS global 1 km atmosphere-only demonstrator and performance analysis

D2.9 : Implementation of ICON global 1km atmosphere-only demonstrator and performance analysis [24]

D2.9: Implementation of ICON global 1km atmosphere-only demonstrator and performance analysis

D2.10 : Implementation of NEMO global 1km ocean-only demonstrator and performance analysis [42]

D2.10: Implementation of NEMO global 1km ocean-only demonstrator and performance analysis (CNRS-IPSL, R, PU, PM42).

D2.11 : Implementation of EC-Earth 10km global coupled demonstrator and performance analysis [42]

Implementation of EC-Earth 10km global coupled demonstrator and performance analysis (BSC, R, PU, PM42).

D2.12 : Implementation of ICON 10km global coupled demonstrator and performance analysis [42]

Implementation of ICON 10km global coupled demonstrator and performance analysis (DKRZ, R, PU, PM42)

Schedule of relevant Milestones

Milestone number[18]	**Milestone title**	**Lead beneficiary**	**Due Date (in months)**	**Means of verification**
MS1	Scheduler support and development plan (T3.3.1)	8 - METO	6	Allow the first deliverable D3.9 to meet the needs of the user community
MS3	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS6	Reports on user support, training, and integration of NEMO and EC-Earth community models	3 - CNRS-IPSL	30	Reports published on the intranet

Work package number [9]	WP3	Lead beneficiary [10]	4 - MPG
Work package title	Usability		
Start month	1	**End month**	48

Objectives

• Support scientific excellence through provision of effective HPC and big data infrastructures by allowing scientists to more easily design and carry out simulation campaigns that seamlessly exploit the existing multi-model framework, including the inherent value of model diversity.
• Considerably improve the ease-of-use of the software, computing and data-handling infrastructure for ESM scientists from the applications through the software stack to the hardware.
• Support the uptake of scheduling engines within the community through user-driven development, training and support services. This activity will absorb 40% of the effort in this work package, since it will provide a step change in the community's ability to cope with increasing suite complexity for both climate and weather applications in production and research modes. MetO, who have existing skills in Cylc and are supporting its adoption, will lead this task
• Reduce the skills gaps at individual centres by sharing best practice through worked examples using use-cases derived from user-driven engagement, the need to prepare extreme scale demonstrators, and governance.
• Propose avenues for co-design of system software and architectures between industry and applications

Description of work and role of partners

WP3 - Usability [Months: 1-48]
MPG, DKRZ, ECMWF, BSC, METO, UREAD, CMCC, ALLINEA
This work package is led by: Reinhard Budich, MPG (lead) and and Kim Serradell (BSC) together with Mick Carter (MetO) as co-leads

Task 3.1 ESM End-to-end Workflows Recommendations [Lead: MPG. Participants: DRKZ, BSC, UREAD, Allinea]
In this task, an Application Software Environment necessary for multi-model simulations will be specified and assembled in a prototype ESM workflow framework, which will be necessary to exploit the exascale opportunities of the future.

Task 3.1.1. Requirements and Specification[Lead: MPG; Participants: BSC, UREAD] Networking Activity
MPG, with the aid of BSC and UREAD, will host an introductory workshop where end-user requirements for environments and workflows will be discussed as well as best practices and lessons learned in specifying and implementing them. The aim of the workshop is to initiate the specification of a standard recommendation for an ESM End-to-end Workflows and Application Software Environments framework. Informed by experienced users, key players from the user and technical provider communities will be invited both from Europe and the US (National Centre for Atmospheric Research NCAR, Boulder, USA; Geophysical Fluid Dynamics Laboratory GFDL
Princeton, USA). Allinea will provide assistance for ensuring the ability of the design to support development and performance tools. The product will be versions [D3.2] of a white paper, published and kept up-to-date on the ENES portal by ESiWACE. For the first version [D3.1] of the white paper, several writing sprints of the smaller writers team will be scheduled.

Task 3.1.2. Development of Use Case [Lead: BSC; Participants: MPG, UREAD, Allinea] Joint research activity
Once first drafts of the white paper are circulated, a small team of scientific programmers from BSC (team lead), MPG and UREAD will start to convert the recommendations into a real life environment, our use case [D3.3]. This use case will be the workflow necessary for the ENES summer schools, planned to be held by UREAD at CSC28, Finland, – a PRACE Tier1 centre - in 2016 (3rdE2SCMS, European Earth System and Climate Modelling School). At the previous two E2SCMS, three GCMs were used to teach the students, each operated in its own framework. The use case will make it possible for students to better co-design and exploit the simulation exercises, by providing a unified framework. This way the use case shows on a somewhat smaller scale, what the system specification will be able to deliver to projects. The group will
start with first framework sketches very early in the project, and iterate through further drafts. Allinea will ensure the readiness of the prototype for their methods and tools, and that the prototype is able to integrate into a modern scheduling environment. The prototype will then be handed over to T3.1.3 in time for the 3rdE2SCMS to be tested in a provisional environment at CSC.

Task 3.1.3. First Installations [Lead: UREAD; Participants: BSC, MPG, Allinea] Joint research activity, Development of service
For these tests UREAD (team lead, organizer of the 3rdE2SCMS), BSC and MPG will install the software collection provided by T3.1.2 at CSC, and test it, aided by Allinea. Success metrics shall be established in advance. The ultimate test will be if the environment is usable for the 3rdE2SCMS and gets good ratings by the participants [D3.4]. The next step will then be to use the experience from the 3rdE2SCMS to develop the environment for the demonstrators.

Task 3.2 ESM System Software Stack Recommendations [Lead: BSC. Participants: DKRZ, Allinea, MPG, UREAD]
In this task, a methodology for maintaining a portable HPC system software stack will be developed for ESM applications, with their specific needs in terms of e.g. Pre- and Post- Processing, Analysis, Compilation and data handling software. This will support the exploitation of the pan-European HPC infrastructure and prepare for the demonstrators.

Task 3.2.1 Requirements and Specification [Lead: BSC; Participants: MPG, DKRZ, UREAD, Allinea] Networking Activity
An introductory, user-driven workshop will discuss best practices and lessons learned in specifying and implementing software stacks, and requirements for them. BSC will, with the aid of DKRZ, invite experienced system administrators, users and software architects from Europe, the US (GFDL, NCAR) and New Zealand (National Institute of Water and Atmospheric Research NIWA)). The aim of the workshop is to initiate the specification of a standard recommendation for an ESM System Software Stack. Allinea will provide assistance for ensuring the ability of the design to support development and performance tools. The recommendation will have the form of versions [D3.6] of a white paper, published on the ENES portal. The white paper should be a handbook for system administrators [D3.5] on how to select, configure and install ESM software stacks in order to be prepared for applications from our community and could serve as a base for a quality label for compute centres ("This centre is enabled for ESM applications"). The handbook needs scheduled updates.

Task 3.2.2 Prototyping [Lead: BSC; Participants: DKRZ] Joint research activity
Once first drafts of the white paper are circulated, a small team of scientific programmers will convert the recommendations to be expected into a real software stack. This will be installed at DKRZ, and used and tested also by non-DKRZ-administrators and users (BSC, other partner volunteers). It will enable the workflow necessary for the demonstrators [D3.8]. The group will start with first prototypes of the stack as soon as possible, and iterate through the development cycles as further drafts become available [D3.7]. The stack will then be handed over to T3.1.3 in time for the 3rdE2SCMS to be tested in different environments.

Task 3.2.3 First Installations [Lead: DKRZ; Participants: BSC, UREAD, Allinea] Joint research activity, Development of service
For those tests it will be necessary to install the software collection and stack provided by T3.1.2 and T3.2.2 on different machines at BSC, CSC, and UREAD and test it against agreed upon success metrics and evaluation plans. The ultimate test will be, if the stack is easy to install and maintain by the system administrators, and easily usable for the demonstrators. In parallel to this activity, different computing centres will be approached by the tasks participants in order to negotiate the conditions under which they would consider hosting the ESM system SW stack on their machines. The next step will then be to transport the stack to other PRACE- and topical tier
1-machines, using the different (configurations of the) operating system installed in the different centres: This activity is probably beyond the end of this project.

Task 3.3 ESM Scheduling [Lead: MetO. Participants: MPG, BSC, CMCC]
This task will maximise the chances of building a supported, user-driven community around the Cylc meta-scheduler for complex climate and weather suites on HPC systems.

Task 3.3.1: Scoping of the Work [Lead: MetO; Participants: MPG] Networking Activity, Development of service
A user-centric, initial workshop will establish the development priorities and the most effective form of user support services. It will, with the aid of WP1, establish a governance process for continued support and development activities within and after the project (consistent with the ESiWACE business plan [D3.1]). MetO will organize and evaluate the workshop; the other partners will help prepare, participate, and will assist in the evaluation. We plan to involve in this task two supporting institutions, NIWA and GFDL (see the letters of support provided in Appendix 1)

Task 3.3.2 Cylc Development [Lead: MetO; Participants: MPG] Development of service Informed by the workshop and the governance activity, under the Lead of MetO and assisted by BSC and MPG, in cooperation with NIWA and inviting involvement from GFDL, Cylc will be developed to support a wider community and set of platforms. The development will consist of two activities: Firstly to refactor code to ensure continued supportability and to allow a wider community engagement in the code of Cylc; second, to develop any new features required by the community according to the results of the governance activity in particular considering the challenges of exploiting exascale computing in climate and weather applications in a pan- European context [D3.9].

Task 3.3.3 Cylc Support Services [Lead: MetO; Participants: BSC] Provisioning of service
Informed by the workshop and the governance activity, training and support services will be provided to the community as prioritised by the governance activity. MetO will be responsible for level 2 support with level 1 support coming from local help-desk services. They will also be responsible for training, including the training of the local level 1 support.

Task 3.3.4 Scheduling Systems Development [Lead: CMCC; Participant: BSC] Networking Activity, Joint research activity
The areas of data analytics and multi-site, multi-model experiments need special attention in terms of their scheduling. Ideas for system architectures and design sketches for the appropriate software solutions will be developed and discussed to address these topics. Where systems already exist like Autosubmit http://ic3.cat/wikicfu/index.php/Tools/Autosubmit (supported by IS-ENES2 and adhering to PRACE guidelines), Ophidia http://ophidia.cmcc.it/, Cylc, and others, implementation strategies, possibly on the interaction between them, will be suggested informed by the governance process to be established like in T3.3.1, and interacting with T3.1 and T3.2. BSC and CMCC will share the task to organize and run the workshop and its evaluation, in close cooperation with WP1 and informed by the other WP partners [D3.10].

Task 3.4 Co-Design for Usability [Lead: ECMWF. Participants: MPG, Allinea]
If the community is to design efficient, powerful and usable simulation systems, a multi- disciplinary approach linking scientific software engineers with industry experts is essential. This task is to help develop such relationships employing a two-stranded strategy.

Task 3.4.1 Requirements Capture, Recommendations for Activities, and Implementation Methods for an ESM Co-Design Approach [ECMWF] Networking Activity
The aim of the task is to develop an exchange platform for the ESiWACE community and industry supporting the co-design of compiler standards. Output is a white paper that will take into account that boundary conditions and requirements will change during and after the project. In an introductory workshop, experienced developers from NWP centres, the seven ENES modelling communities, international centres, contact points from open source compiler and software builders, and key representatives of the industry on the working level will be invited by ECMWF. They will reduce skill and knowledge gaps by sharing best practices and lessons learned in co-design approaches and activities. These experts should not only be interested in co-design, but also in representing ESiWACE in bodies governing software standards for, e.g., compilers. A key aim will be to engender a unified voice to influence the HPC suppliers on behalf of a coordinated and increasingly convergent weather and climate community. The workshop will result in a white paper specifying recommendations for this interaction. The white paper will provide the basis for a comprehensive requirements capture, recommend activities to liaise with industry, solution providers, and standards bodies, and describe implementation methods to achieve an efficient lobbying in the future. The initial workshop should be organised back-to-back with one of the HPC workshops of ECMWF, ENES (planned) or NCAR. The sustained business plan of ESIWACE will contain a series of follow-on workshops for continuing the discussion within this group, revise requirements and recommendations, and report on implementation status.

Task 3.4.2 Usability Concepts: Joint developments by Communities, Solution Providers, and Industry [Participants: Allinea, MPG] Networking Activity
Simulation workflows in Earth system modelling for climate and weather are very complex, and do not serve high numbers of users. Rather, they are perceived as single-site problems, resulting in a natural tendency to single-site solutions. These are often not very well usable. Usability and complexity are problems in the software development field in general. The ESiWACE community would benefit substantially from sharing all aspects of e.g. complexity hiding or solutions optimized for usability from the software industry. Allinea will, with the aid of WP1 and ETP4HPC, initiate a series of workshops – only an initial one taking place during the lifetime of the ESiWACE co-located with another larger event of interest – bringing together Communities, Solution Providers, and Industry. In such workshops concerns about site dependencies, support structures, and many other aspects like future developments can be discussed, but also (their) solutions be presented. These workshops can serve as a platform to foster community involvement e.g. in standardization committees, where future common developments are decided, by establishing direct contacts to solution providers and industry: They are involved in such bodies and will provide appropriate links for more co- design activities between communities and industry.

Interaction with other Work Packages
This Work Package provides:
• Input to WP1 in cross cutting issues.
• Input to WP1 on the need for resource.
• Information to WP2 in particular T2.2 about system software stack available in most centres.
• Information to WP1 about the system software stack available in most centres.

• Input to W2 and WP4 regarding the requirement of tools and methods arising from the end-to- end workflow of ESM.
• D3.4 (T3.1) and D3.8 (T3.2) will provide input to the governance processes established by WP1.
This Work Package receives:
• Input from WP1 regarding requirements on new developments.
• Input from WP2, in particular task 1 on maturity and popularity of common software.
• For D3.2, D3.6 and D3.10 input from both WP2 and WP4 on recent shared software development.

Participation per Partner

Partner number and short name	WP3 effort
1 - DKRZ	6.00
2 - ECMWF	3.00
4 - MPG	18.00
6 - BSC	18.00
8 - METO	57.00
9 - UREAD	6.00
12 - CMCC	9.00
16 - ALLINEA	6.00
Total	123.00

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D3.1	ESiWACE Application Software Framework	4 - MPG	Report	Public	4
D3.2	Update of ESiWACE Application Software Framework (D3.1) prepared for the demonstrators, Version 2	4 - MPG	Report	Public	30
D3.3	Software specification for the third E2sCMS	4 - MPG	Report	Public	6
D3.4	Experiences with ESM Multi-model Ensembles for Educational Purposes	4 - MPG	Report	Public	14
D3.5	How to select, configure and install ESM software stacks	6 - BSC	Report	Public	14
D3.6	Update Handbook for system administrators	6 - BSC	Report	Public	36
D3.7	Software Stack for ESM– A Specification:Narrowing	6 - BSC	Report	Public	18

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
	down D3.5 for the demonstrators				
D3.8	A Report from the use of D3.5 for the demonstrators	1 - DKRZ	Report	Public	30
D3.9	SiWACE Scheduler development and support activities	8 - METO	Report	Public	18
D3.10	ESiWACE Scheduler development and support activities, v2	8 - METO	Report	Public	34

Description of deliverables

D3.1: ESiWACE Application Software Framework: A White Paper. Version 1, Specification of a Standard Recommendation for an ESM End-to-end Workflows and Application Software Environments Framework (MPG, R, PU, PM4).
D3.2: Update of ESiWACE Application Software Framework (D3.1), prepared for the demonstrators (MPG, R, PU, PM30).
D3.3: Software specification for the 3rdE2SCMS, narrowing down D3.1 for the 3rdE2SCMS (MPG, R, PU, PM6)
D3.4: Experiences with ESM Multi-model Ensembles for Educational Purposes: A report from the use of D3.1 for the 3rdE2SCMS(MPG, R, PU, PM14)
D3.5: How to select, configure and install ESM software stacks: Handbook for system administrators Specification of a standard recommendation for an ESM System Software Stack in the form of a white paper (BSC, R, PU, PM14)
D3.6:Update Handbook for system administrators (D3.5) (BSC, R, PU, M36)
D3.7: Software Stack for ESM– A Specification narrowing down D3.5 for the demonstrators (BSC, R, PU, PM18)
D3.8: Experiences with the ENES System Software Stack: A Report from the use of D3.5 for the demonstrators (DKRZ, R, PU, PM30)
D3.9: ESiWACE Scheduler development and support activities: a first report of Scheduler development and support activities for T3.3.2 and T3.3.3. (MetO, R, PU, PM18)
D3.10: ESiWACE Scheduler development and support activities, v2: a second report of Scheduler development and support activities, updating D3.9 and reporting on T3.3.4 (MetO & CMCC, R, PU, PM 34)

D3.1 : ESiWACE Application Software Framework [4]

ESiWACE Application Software Framework: A White Paper. Version 1: Specification of a Standard Recommendation for an End-to-end Workflows and Application Software Environments Framework

D3.2 : Update of ESiWACE Application Software Framework (D3.1) prepared for the demonstrators, Version 2 [30]

Update of ESiWACE Application Software Framework (D3.1) prepared for the demonstrators, Version 2

D3.3 : Software specification for the third E2sCMS [6]

Software specification for the third E2sCMS, Narrowing down D3.1 for the 3rdE2SCMS

D3.4 : Experiences with ESM Multi-model Ensembles for Educational Purposes [14]

Experiences with ESM Multi-model Ensembles for Educational Purposes: A Report from the use of D3.1 for the 3rdE2SCMS

D3.5 : How to select, configure and install ESM software stacks [14]

How to select, configure and install ESM software stacks: A Handbook for system administrators Specification of a standard recommendation for an ESM System Software Stack in the form of a white paper

D3.6 : Update Handbook for system administrators [36]

Update Handbook for system administrators (D3.5)

D3.7 : Software Stack for ESM– A Specification:Narrowing down D3.5 for the demonstrators [18]

Software Stack for ESM– A Specification: Narrowing down D3.5 for the demonstrators

D3.8 : A Report from the use of D3.5 for the demonstrators [30]

Experiences with the ENES System Software Stack: A Report from the use of D3.5 for the demonstrators

D3.9 : SiWACE Scheduler development and support activities [18]

ESiWACE Scheduler development and support activities. A First report of Scheduler development and support activities for T3.3.2 and T3.3.3.

D3.10 : ESiWACE Scheduler development and support activities, v2 [34]

ESiWACE Scheduler development and support activities, v2: Second report of Scheduler development and support activities,

Schedule of relevant Milestones

Milestone number[18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS2	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	9 - UREAD	9	Summer School operational
MS3	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS4	System software stack D3.7 handed over to T3.1.3 team	6 - BSC	18	Report published on the website

Work package number [9]	WP4	Lead beneficiary [10]	7 - STFC
Work package title	Exploitability		
Start month	1	**End month**	48

Objectives

Making the best use of HPC in Earth simulation requires storing and manipulating vast quantities of data. Often such manipulation involves handling very high volumes of data (by contrast to many commercial applications which sift, or query into, high volumes of data).
Existing storage hardware and software is not well adapted to these use modes, particularly, for binary data. Hence the overall goal of this work package is to address the exploitability of storage systems in weather and climate HPC, first by modelling existing and possible future systems, then by developing improved software and exploitation strategies for both disk and tape systems.
The work is broken into four specific tasks:
1. Understanding the business of storing and exploiting high volume Earth system data by developing a model which can be used to both understand the cost and performance of a range of storage strategies in weather and climate HPC workflows (and compare those with typical commercial strategies).
2. Developing a New disk storage layout for Earth system data to overcome the performance limitations of existing fixed on-disk formats and APIs by providing a novel data-specific layout interface. The new interface will be capable of supporting both interoperability between higher-level file formats and of customisation to local heterogeneous disk storage environments. It will be deployed in test mode to support the demonstrator activity.
3. Developing New tape access strategies and software customised for Earth system data by first modelling and simulating possible strategies, and then developing a new
software library which provides both higher bandwidth to tape storage and increased storage redundancy.

Description of work and role of partners

WP4 - Exploitability [Months: 1-48]
STFC, DKRZ, CMCC, SEAGATE
This work package is led by: Bryan Lawrence, STFC (lead) and Thomas Ludwig, DKRZ (co- lead)

Task 4.1 The business of storing and exploiting high volume climate data [Lead: DKRZ. Participants: STFC] Joint research activity and partly networking activity
System and data centre design and procurement are complex tasks. We will parameterise the required activities into a business model that could, for example, given building and power capacity, predict storage capacity and power drawdown. In and of itself, such a model is very straightforward, but building in the ability to understand how choices as to ensemble-size, model resolution, and run length, will impact both on compute and storage requirements, is a much more demanding. This new model will expose the very direct trade-offs between
scientific aspirations and physical limits (power, machine room and tape library size).
Specifically, the very significant increases needed in the proportion of budget devoted both to storage capital costs, and storage recurrent costs will be contrasted with future compute availability. The model will also be designed to help direct future research in I/O towards the most promising activities. Addressing this will require a short period of sustained working in the two main academic partners to develop, test and publish suitable business models. In a first step, the developed model will be disseminated to participating data centres and checked. In a second step, it will be published. DKRZ will develop model components to cover hardware, software and data centre aspects where STFC will incorporate domain-specific aspects for Earth system models and operation. A report describing the model and a lightweight implementation that allows experimenting with model parameters will be released as [D4.1].

Task 4.2 New Storage Layout for Earth system data [Lead: DKRZ; Participants: Seagate, CMCC, STFC] Joint research activity
I/O is being addressed at a number of levels within ESiWACE, since both current experience and projections have highlighted the performance/volume challenges ahead. The main goal of this task is to address performance and physical capacity issues associiated with writing, accessing, and Earth system data in the disk subsystems themselves. One key problem that can be addressed is separation of file metadata operations from data ac cess – this is the strategy some hybrid disk systems use to improve performance by putting file system metadata on SSD and data itself on spinning platters. However, with netCDF or GRIB data, the scientific relevant metadata and much of the information about actual

data layout is invisible to the file system (whether parallel or serial), which cannot then exploit acceleration techniques. To address this, we propose to develop and test a Earth system data (ESD) middleware library which understands these key formats, and which is customisable for different hardware environments (as will occur in differing data centre architectures). It will support, for example, writing the scientific metadata and possibly coordinate axes to SSD systems, and the binary data fields to traditional systems potentially split into multiple files. Where in-memory storage back ends are available, fast online data analytics will also become possible. The architecture is illustrated in the Figure below. The layout components will decide the data placement and orchestrate data access on the storage characteristics. Both blocking and non-blocking APIs will be supported to address typical use-cases. Tools will accompany the library so that data can be reconstructed as compatible with netCDF/GRiB and can be archived and exchanged. The design of the architecture will be documented in [D4.2], an operational demonstrator will be released (MS9) and a performance optimized version for all data centres released in [D4.3]. The MS9 milestone will be used to put in place the systems necessary to take data from at least one simulation of the demonstrator workflow. DKRZ will implement the base architecture and modules, while Seagate who will provide an optimised object storage backend, and co-develop ESD to integrate best-practices in storage access for the needs of the scientific domain. The layout components will be adapted and evaluated by CMCC and STFC by exploiting the petabyte-scale storage environments of their data centres. Additionally, CMCC will support the investigation of suitable memory and storage back-ends able to support horizontally scalable management of multidimensional scientific data. CMCC will specifically develop a back-end for the object store technology WOS31 from DDN to prevent vendor lock-in.

Task 4.3 - New methods of exploiting tape [Lead: STFC; Participant: DKRZ] Joint research activity
High-end climate computations are likely to generate such large volumes of data that not only will disk performance and cost be limiting, but the physical capacity in machine rooms, too. Leaving aside completely new storage media such as holographic storage, one possible way forward is to make better use of tapes. Traditionally, the climate community has used tape primarily for archive, that is, for backup, and recovery of small amounts of data from very big volumes. Sites such as ECMWF and the Met Office have developed bespoke environments (MARS and MASS respectively) which by controlling data and scientific metadata structures, and introducing carefully configured servers and storage cache, have extended the “write once, read infrequently” mode to allow higher performance environments, but these require large teams, and are still limited by serial processes. Similarly, at STFC, the JASMIN32 interface to tape currently utilises the CASTOR33 tape system designed by CERN. By contrast, some applications of HPSS have used RAIT concept (Redundant Array of Independent Tapes) and noting the complexity of tape management at scale, proposed some sophisticated strategies of using tapes in parallel. We propose to first model and simulate, then build, a prototype open source tool that offers configurable flexibility between high performance and increase in capacity vs. resilience. In contrast to existing approaches such as in HPSS, our approach can be deployed in the typical scenario where a library is equipped with different tape generations and technology. We would expect to evaluate the tape performance simulations in comparison to existing strategies in the climate community, and deploy the prototype tape library at STFC (M4.5). The system would be designed for insertion into existing workflows at other centres, but we expect an operational setting demonstration is mandatory to demonstrate its benefit. STFC focuses on the implementation and deployment whereas DKRZ leads the core modelling and simulation. The outcome of this will be documented in [D4.4].

Interactions with other Work Packages
This Work Package provides:
• Input to WP1 on status of development.
• The governance and community liaison tasks from WP1 will receive and incorporate the knowledge from the best-practices established in T4.4.1 according to the prediction of the developed model and research activities that also target future supercomputers.
• Input to WP2, T2.3.2 on the semantic mapping between netCDF and GRIB developed in WP4.
This Work Package receives:
• Input from WP1 regarding governance of new developed software.

Participation per Partner

Partner number and short name	WP4 effort
1 - DKRZ	45.00
7 - STFC	33.00
12 - CMCC	10.00

Partner number and short name	WP4 effort
14 - SEAGATE	18.00
Total	106.00

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D4.1	Business model with alternative scenarios	1 - DKRZ	Report	Public	18
D4.2	ESD middleware design	1 - DKRZ	Report	Public	18
D4.3	Final implementation of the ESD middleware	1 - DKRZ	Other	Public	46
D4.4	Final report on alternative tape usage	7 - STFC	Report	Public	42

Description of deliverables

- D4.1 Business model with alternative scenarios (DKRZ, R, PU, PM18).
- D4.2 ESD middleware design (DKRZ, R,PU, PM18).
- D4.3 Final implementation of the ESD middleware (DKRZ, OTHER, PU, PM46).
- D4.4 Final report on alternative tape usage (STFC, R, PU, PM42).

D4.1 : Business model with alternative scenarios [18]

Business model with alternative scenarios

D4.2 : ESD middleware design [18]

ESD middleware design

D4.3 : Final implementation of the ESD middleware [46]

Final implementation of the ESD middleware

D4.4 : Final report on alternative tape usage [42]

Final report on alternative tape usage

Schedule of relevant Milestones

Milestone number[18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS3	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS7	Prototypes of alternative storage backends	7 - STFC	33	Prototypes available
MS8	Prototype tape library for advanced tape subsystems	7 - STFC	36	Prototypes available
MS9	Implementation of ESD middleware at SFTC and	1 - DKRZ	36	ESD middleware finalised and rolled out to support the demonstrator.

Schedule of relevant Milestones

Milestone number[18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
	CMCC (to support the demonstrator)			

Work package number [9]	WP5	Lead beneficiary [10]	1 - DKRZ
Work package title	Management & Dissemination		
Start month	1	**End month**	48

Objectives

Objectives:
The general objective of this Work package 5 (WP5) is to ensure effective and smooth high- quality implementation of the project and delivery of innovation actions, and impacts. This includes general administrative and management practices that can be expected for programs founded via EU programs, including financial and scientific reporting. In particular this means
• to set-up and maintain a structure, procedures and tools that will allow a coherent and efficient technical and administrative management of the project;
• to keep the project on time and within the assigned budget;
• to identify and manage risks and to solve problems;
• to identify opportunities for improved results and collaboration;
• to coordinate the communication and interactions between work packages and partners;
• to coordinate the communication process between ESiWACE and the Scientific Advisory board (SAB);
• to coordinate the communication process between ESiWACE and the EC;
• to manage quality assurance;
• to establish and coordinate the dissemination and exploitation processes for ESiWACE.

The management structure and procedures of the project are extensively described in Section 3.2. More on innovation management (Section 1.4.2, 2.1.2, 2.2.2, 3.2.3), dissemination and exploitation measures as well as communication activities is already described in Sections 2.2.1., 2.2.3. In this text: that information is not repeated.

Description of work and role of partners

WP5 - Management & Dissemination [Months: 1-48]
DKRZ, ECMWF, CNRS-IPSL, MPG, CERFACS, BSC
Lead: Joachim Biercamp (DKRZ) lead, Peter Bauer ECMWF co-lead. Participants: all WPL and co-WPL.

Task 5.1 Coordination and Management [Lead: DKRZ. Participants: ECMWF] Networking activity
Coordination:
• The Coordinator will carry out the coordination and monitoring of scientific excellence within the project. This includes regular discussion with the Executive Management Board (EMB), Management Steering Board (MSB), Scientific Advisory Board (SAB). The management processes are described in more detail in Section 3.2.1.
• The Coordinator will ensure that planed work is carried out in time and budget. For this he will be supported by the Project Office (PO).
• The Coordinator will monitor the scientific review of reports and deliverables to the EC.
and ensure that any necessary scientific aspects are incorporated into the project.
• The WP leaders will ensure that the progress of milestones and deliverables is actively monitored and that they are delivered on time.
Management:
• The setting up of a consortium agreement in the very early stage of the project and prior to the signature of the grant agreement. The consortium agreement will regulate the consortium, rules for participation, and ownership and access to key knowledge (IPR, results, etc.).
• The PO at DKRZ will be in charge of managing the project using effective management procedures based on the project cycle management formal methodology. For more information on the tasks and composition of the PO, please refer to Section. 3.2.1. The PO will provide administrative / financial / legal support to all partners involved during the implementation of the project, this can include administrative tasks involved in the preparation, executing of and post-processing of major project meetings of the committees and panels.
• Management of the gender dimension: with support from the PO, the Coordinator will ensure that gender aspects of the project are fully considered within the research that is being undertaken, and that ESiWACE acts to promote gender equality wherever possible. For more details refer to Section 1.3.4.
• Liaison with the European Commission (EC), this includes:

a) The preparation of Project Periodic Reports, the Final Report and the Final Report on the EU Financial Contribution Distribution.
b) To maintain regular and comprehensive contact with the EC.
c) To ensure the appropriate follow-up of project obligations from the Grant Agreement (formal reporting: of science results and finances, project reviews, communication, and management).
d) To ensure that the appropriate EC representative is invited to the General Assembly meetings and any other relevant project meetings.
e) If there are any major difficulties within the project that cannot be resolved using the appropriate management structure, the Coordinator will liaise with the EC to seek advice and a solution.

Task 5.2 Risk management [Lead: DKRZ. Participants: ECMWF] Networking Activity
• The General Assembly will be responsible for dealing with risks, issues and benefits realization of the project.
• The Coordinator will be responsible for management of the risks within the project, and day-to-day maintenance of the risk registers will be undertaken by the PO (for more details refer to Section 3.2.4).
• Critical risks to project implementation, which have the potential to impact the project objectives being achieved, have been identified and described in Table 3.2.a.

Task 5.3 Innovation and IPR management, exploitation of results [Lead: DKRZ. Participants: ECMWF] Networking activity
• Management of knowledge and innovation is an integral part of our project. We will focus on the role and synergies between partners' experiences, competences, capabilities, and on how partners will protect, share, manage IPR capital actual exploitation. Detailing of the exploitation plans and preparation for innovation activities will be continuously followed up throughout the project. The innovation management is well integrated in the management structure of the project and in the work plan (see Section 3.2.3).
• The consortium agreement (CA) will be set up for regulating the ownership and access to key knowledge (IPR, data etc.) and scientific results, among other things, after the communication of the approval of the project by the European Commission and before the signature of the Grant Agreement with the European Commission.
• A strategy for the knowledge management, protection and for the exploitation of results, the Plan for Dissemination and Exploitation [D5.5], will be defined for the consortium in the early stage of the project, based on the principles explained extensively in Section 2.2. The strategy will be regularly updated during the entire project. Updates will be submitted to the EC as an integral part of the Project Periodic Reports.
• A Strategy for the Intellectual Property exploitation [D5.6] will be drafted at the end of the project for providing best practices in capturing and assessing the Intellectual Property and providing measures for exploitation after the end of the project.
• ESiWACE will provide open access to peer-reviewed scientific publications through a combination of golden open access and green open access, and it is voluntarily taking part in the European Commission Open Access Data Pilot for Research Data (see Section 2.2.2): we have included a Data Management Plan as a deliverable for project- month 6 [D5.4] to be drafted in compliance with the guidelines given on data management in the Horizon 2020 Online Manual. This deliverable will evolve during the lifetime of the project and represent faithfully the status of the project reflections on data management. Updates of the data management plan are thus planned and will be submitted to the EC as an integral part of the Project Periodic Reports.

Task 5.4 Dissemination [Lead: DKRZ. Participants: ECMWF] Networking activity
We have identified three stages for the dissemination of project results:
• Dissemination measures during the lifetime of the project.
• Dissemination measures in the closing phase of the project.
• Dissemination measures after the closure of the project.
As indicated above, we will comply with EC open access guidelines. Additionally, wherever results are suitable (content and size) to be distributed or stored by using online repositories, we will use ZENODO for disseminating the results of the project. For more details refer to Section 2.2.2 Task 5.5 Communication [Lead: DKRZ. Participants: ECMWF] Networking activity The communication activities of the project will involve all consortium partners and their respective staff, including researchers. The project office (PO) with the Scientific Officer (SO) is central part of the activities. As indicated in the table in Section 2.2.3, we have already foreseen tools for the implementation of our communication strategy; a Media and Communication Plan [D5.3] focusing on the most suitable tools for the most suitable audiences will be set up in PM2
of the project and regularly updated by the Project Office. We have identified different levels of communication activities:
• Communication activities to promote the project and its findings (ref. to Section 2.2.3 for more details). The website will be set up in PM2 [D5.2].

• Coordination of internal communication within the consortium, with the supporters, the user group committee and the scientific advisory board.
• Communication with other EU funded projects and the European Commission.
For more details refer to Section 2.2.3.

Interactions with other Work Packages
This Work Package provides:
• Scientific guidance and leadership for all WPs and to the project as a whole.
• Support, guidance and management assistance to all WPs as required and to the project as a whole.
• Support to organisations invited to participate to ESiWACE workshops, General Assemblies, training activities.
• Support to the Advisory Board: the Advisory Board will be consulted on ESiWACE activities and communities' needs.
This Work Package receives:
• Regular summaries of the scientific progress from each work package leader, so that the Coordinator can review progress and provide scientific guidance.
• Financial and administrative updates from all other WPs and partners.

Participation per Partner

Partner number and short name	WP5 effort
1 - DKRZ	24.00
2 - ECMWF	6.00
3 - CNRS-IPSL	2.00
4 - MPG	2.00
5 - CERFACS	1.00
6 - BSC	1.00
Total	36.00

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D5.1	Design and implementation of the intranet	1 - DKRZ	Websites, patents filling, etc.	Confidential, only for members of the consortium (including the Commission Services)	2
D5.2	Project public website	1 - DKRZ	Websites, patents filling, etc.	Public	2
D5.3	Media and Communication Plan	1 - DKRZ	Report	Public	2
D5.4	Data Management Plan (DMP)	1 - DKRZ	ORDP: Open Research Data Pilot	Public	6
D5.5	Dissemination and Exploitation Plan (EP)	1 - DKRZ	Report	Public	18

List of deliverables

Deliverable Number[14]	Deliverable Title	Lead beneficiary	Type[15]	Dissemination level[16]	Due Date (in months)[17]
D5.6	Strategy for the Intellectual Property exploitation	1 - DKRZ	Report	Public	48

Description of deliverables

D5.1: Design and implementation of the intranet, portal, wiki: The intranet is a tool for the legal, financial and administrative management of the project. It contains contractual documents, consortium documents, templates for the legal, financial administrative management, and copies of reports to the EC. The intranet will be web based (DKRZ, DEC, CO, PM2)
D5.2 Project Public website (DKRZ, DEC, PU, PM2)
D5.3 Media and Communication Plan (DKRZ, R, PU, PM2)
D5.4 Data Management Plan (DMP): The DMP will be drafted in compliance with the guidelines given on data management in the Horizon 2020 Online Manual (DKRZ, R, PU, PM6)
D5.5 Dissemination and Exploitation Plan (EP): Plan identifying types of potential pathways of market- oriented exploitation, converting or transforming knowledge will be identified, together with key factors for a successful innovation management. DKRZ, R, PU, PM18)
D5.6 Strategy for the Intellectual Property exploitation: Strategy for defining measures for exploitation "after the project" phase, providing evidence of best practices in capturing and assessing (DKRZ, R, PU, PM48)

D5.1 : Design and implementation of the intranet [2]

Design and implementation of the intranet (web based)

D5.2 : Project public website [2]

Project public website

D5.3 : Media and Communication Plan [2]

Media and Communication Plan (MCP)

D5.4 : Data Management Plan (DMP) [6]

Data Management Plan (DMP)

D5.5 : Dissemination and Exploitation Plan (EP) [18]

Dissemination and Exploitation Plan (EP)

D5.6 : Strategy for the Intellectual Property exploitation [48]

Strategy for the Intellectual Property exploitation

Schedule of relevant Milestones

Milestone number[18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS1	Scheduler support and development plan (T3.3.1)	8 - METO	6	Allow the first deliverable D3.9 to meet the needs of the user community
MS3	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS5	Operational demonstrator of ESD middleware	1 - DKRZ	24	Demonstrator operative

1.3.4. WT4 List of milestones

Milestone number[18]	Milestone title	WP number[9]	Lead beneficiary	Due Date (in months)[17]	Means of verification
MS1	Scheduler support and development plan (T3.3.1)	WP1, WP2, WP5	8 - METO	6	Allow the first deliverable D3.9 to meet the needs of the user community
MS2	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	WP3	9 - UREAD	9	Summer School operational
MS3	Establish governance rules for new ESiWACE developments	WP1, WP2, WP3, WP4, WP5	1 - DKRZ	12	Report published on the intranet
MS4	System software stack D3.7 handed over to T3.1.3 team	WP3	6 - BSC	18	Report published on the website
MS5	Operational demonstrator of ESD middleware	WP1, WP5	1 - DKRZ	24	Demonstrator operative
MS6	Reports on user support, training, and integration of NEMO and EC-Earth community models	WP2	3 - CNRS-IPSL	30	Reports published on the intranet
MS7	Prototypes of alternative storage backends	WP4	7 - STFC	33	Prototypes available
MS8	Prototype tape library for advanced tape subsystems	WP4	7 - STFC	36	Prototypes available
MS9	Implementation of ESD middleware at SFTC and CMCC (to support the demonstrator)	WP4	1 - DKRZ	36	ESD middleware finalised and rolled out to support the demonstrator.

1.3.5. WT5 Critical Implementation risks and mitigation actions

Risk number	Description of risk	WP Number	Proposed risk-mitigation measures
1	Limited engagement with the broader scientific community not directly represented by the ESiWACE partners.		ESiWACE relies already on two well-established European networks, ENES and ECMWF. In order to extend even more the developing culture of excellence to the broader climate and weather modelling community ESiWACE will organise HPC workshops fostering ESiWACE visibility, which, even more than GAs, will play a key role in keeping the community engaged. Interactions with the supporters (see also 3.2.5 Quality management) will help to assess as soon as possible if the community is aware of and content with ESiWACE services.
2	Uneven engagement of climate and weather modelling groups.		ESiWACE organises and promotes workshops on aspects common to climate and weather modelling. For assessment of engagement the same methods hold as above, with ENES and ECMWF taking special responsibly to assess engagement of climate and weather scientists, respectively.
3	Difficulty in achieving long-term sustainability.		The Business plan (see also Section 2.2.4) will be developed and continuously updated.
4	Limited strength of the governance structure established in WP1 to help defining priorities of developments for community tools.		Through continuous feedback with the user community WP2 will contribute to the coordination of model/tool development. ESiWACE representatives will moreover participate in developer community meetings
5	Scarcity of computing resources allocated by PRACE to ESiWACE: impossibility to test benchmarks for I/O and couplers on tier-0 platforms on very high number of cores.		ESiWACE will actively engage with PRACE to obtain allocations. ESiWACE partner resources will be exploited for tests on O(10000) computational processor cores.
6	Lack of significant improvements from the implementation of strategies for performance enhancement of tools and models.		ESiWACE will: Perform an early cost-benefit assessment using existing experiments, Provide an early definition of efficiency gain metrics for objective evaluation. Carry on further work and interaction with computer scientists to increase the gains
7	Change in the scientific response of the models as a result of increasing the concurrency of the different scientific modules of the models.		Accuracy and reproducibility requirements will be defined early in the project. ESiWACE will define a minimum experiment set-up required for evaluation.

Risk number	Description of risk	WP Number	Proposed risk-mitigation measures
8	Applicability of the knowledge compression only limited to selected user groups.		ESiWACE will devise a Strategy [D2.7] for addressing requirements for a wide range of applications.
9	Difficulties in catching the requirements for the application (T3.1) and system (T3.2) software stack, scheduling issues, co-design approaches, and related benefits, especially for inexperienced users and leaders.		WP partners will need to make convincing statements and arguments at the workshops, emphasizing the benefits to be expected from shared software development. The participation in the workshops needs to be spread as wide as possible in terms of skills and experience sets to be convincing.
10	Insufficient fulfilment of the needs of the communities involved by the published recommendations.		The evolving versions of the specifications will be appropriately fine-tuned.
11	Too high complexity of the meta-scheduler Cylc for research applications.		Cylc will be applied initially in operational and production mode where the benefits are better recognised.
12	Limited combinability of different scheduling approaches.		The number of approaches will be minimized, without imposing rules
13	Higher than expected complexity of the Earth system specific middleware (T4.1). Possible occurrence of design flaws in architecture and implementation.		ESiWACE involves experts from computer science and industry from the early design phase on.
14	Lack of funding for related projects in H2020 (ExaIO from Exascale10). Limited effectiveness of the co- design approach. Impossibility to merge ESiWACE results into a worldwide domain independent middleware. Slowing- down in the potential standardization for an API independent layer.		ESiWACE will support the Exascale10 working group by supplying the developed ESD middleware as proof-of-concept to build upon, which secures the impact of the working group and the effort.
15	Undeployability of RAIT as a consequence of lack of redundancy (for incompatibility of tape archives).		ESiWACE will interact with vendors and communicate and discuss on the required interface.

1.3.6. WT6 Summary of project effort in person-months

	WP1	WP2	WP3	WP4	WP5	Total Person/Months per Participant
1 - DKRZ	13	10	6	45	24	98
2 - ECMWF	4	49	3	0	6	62
3 - CNRS-IPSL	17	36	0	0	2	55
· CEA	0	0	0	0	0	0
· UPMC	0	0	0	0	0	0
4 - MPG	2	3	18	0	2	25
5 - CERFACS	3	27	0	0	1	31
6 - BSC	4	22	18	0	1	45
7 - STFC	0	0	0	33	0	33
8 - METO	0	0	57	0	0	57
9 - UREAD	0	0	6	0	0	6
10 - SMHI	1	20	0	0	0	21
11 - ICHEC	0	6	0	0	0	6
12 - CMCC	1	9	9	10	0	29
13 - DWD	0	0	0	0	0	0
14 - SEAGATE	0	0	0	18	0	18
15 - Bull	2	20	0	0	0	22
16 - ALLINEA	0	0	6	0	0	6
Total Person/Months	47	202	123	106	36	514

1.3.7. WT7 Tentative schedule of project reviews

Review number [19]	**Tentative timing**	**Planned venue of review**	**Comments, if any**
RV1	6	Brussels	Interim review
RV2	18	Brussels	
RV3	36	Brussels	
RV4	48	Brussels	

1. Project number

The project number has been assigned by the Commission as the unique identifier for your project. It cannot be changed. The project number **should appear on each page of the grant agreement preparation documents (part A and part B)** to prevent errors during its handling.

2. Project acronym

Use the project acronym as given in the submitted proposal. It can generally not be changed. The same acronym **should appear on each page of the grant agreement preparation documents (part A and part B)** to prevent errors during its handling.

3. Project title

Use the title (preferably no longer than 200 characters) as indicated in the submitted proposal. Minor corrections are possible if agreed during the preparation of the grant agreement.

4. Starting date

Unless a specific (fixed) starting date is duly justified and agreed upon during the preparation of the Grant Agreement, the project will start on the first day of the month following the entry into force of the Grant Agreement (NB : entry into force = signature by the Commission). Please note that if a fixed starting date is used, you will be required to provide a written justification.

5. Duration

Insert the duration of the project in full months.

6. Call (part) identifier

The Call (part) identifier is the reference number given in the call or part of the call you were addressing, as indicated in the publication of the call in the Official Journal of the European Union. You have to use the identifier given by the Commission in the letter inviting to prepare the grant agreement.

7. Abstract

8. Project Entry Month

The month at which the participant joined the consortium, month 1 marking the start date of the project, and all other start dates being relative to this start date.

9. Work Package number

Work package number: WP1, WP2, WP3, ..., WPn

10. Lead beneficiary

This must be one of the beneficiaries in the grant (not a third party) - Number of the beneficiary leading the work in this work package

11. Person-months per work package

The total number of person-months allocated to each work package.

12. Start month

Relative start date for the work in the specific work packages, month 1 marking the start date of the project, and all other start dates being relative to this start date.

13. End month

Relative end date, month 1 marking the start date of the project, and all end dates being relative to this start date.

14. Deliverable number

Deliverable numbers: D1 - Dn

15. Type

Please indicate the type of the deliverable using one of the following codes:

R	Document, report
DEM	Demonstrator, pilot, prototype
DEC	Websites, patent fillings, videos, etc.
OTHER	
ETHICS	Ethics requirement

16. Dissemination level

Please indicate the dissemination level using one of the following codes:

PU	Public
CO	Confidential, only for members of the consortium (including the Commission Services)
EU-RES	Classified Information: RESTREINT UE (Commission Decision 2005/444/EC)
EU-CON	Classified Information: CONFIDENTIEL UE (Commission Decision 2005/444/EC)
EU-SEC	Classified Information: SECRET UE (Commission Decision 2005/444/EC)

17. Delivery date for Deliverable

Month in which the deliverables will be available, month 1 marking the start date of the project, and all delivery dates being relative to this start date.

18. Milestone number

Milestone number:MS1, MS2, ..., MSn

19. Review number

Review number: RV1, RV2, ..., RVn

20. Installation Number

Number progressively the installations of a same infrastructure. An installation is a part of an infrastructure that could be used independently from the rest.

21. Installation country

Code of the country where the installation is located or IO if the access provider (the beneficiary or linked third party) is an international organization, an ERIC or a similar legal entity.

22. Type of access

VA	if virtual access,
TA-uc	if trans-national access with access costs declared on the basis of unit cost,
TA-ac	if trans-national access with access costs declared as actual costs, and
TA-cb	if trans-national access with access costs declared as a combination of actual costs and costs on the basis of unit cost.

23. Access costs

Cost of the access provided under the project. For virtual access fill only the second column. For trans-national access fill one of the two columns or both according to the way access costs are declared. Trans-national access costs on the basis of unit cost will result from the unit cost by the quantity of access to be provided.

Centre of Excellence in Simulation of Weather and Climate in Europe

ESiWACE

History of changes

Section	**Details of changes made**	**Changes made by whom**	**Date**
1	Table of beneficiaries removed	Kerstin Fieg /DKRZ	18/05/2015
3.1	Workpackage description removed	Kerstin Fieg / DKRZ	18/05/2015
3.1	List of Deliverables removed	Kerstin Fieg / DKRZ	18/05/2015
3.1	List of Milestones removed	Kerstin Fieg / DKRZ	18/05/2015
3.2	List of critical risks removed	Kerstin Fieg / DKRZ	18/05/2015
3.4	Summary of Staff Effort	Kerstin Fieg / DKRZ	18/05/2015
	Footer adapted	Kerstin Fieg / DKRZ	19/05/2015
4.2.1 Subcontracting	We have no subcontracting in this project, thus we have deleted the empty table in this section and replaced it with the wording "Not applicable"	Chiara Bearzotti /DKRZ	19/05/2015
4.2.2.	The text reported in correspondence of CNR (beneficiary Nr. 2) in the original table about the work performed by CEA and UPMC has been now converted into text. A breakdown of the costs requested by CEA and UPMC is now more clearly represented and reports now the amount of indirect costs as well.	Chiara Bearzotti /DKRZ	19/05/2015
SIGMA, Budget of beneficiary Nr. 2 CNRS	Part of the budget of CNRS (beneficiary Nr. 2) originally foreseen for its linked third parties CEA and UPMC has been now split correctly in the online budget table in SIGMA.	Chiara Bearzotti / DKRZ	19/05/2015
4.2.3	We have converted the original table into a text. We have added here the profile of Dr. Francisco Doblas-Reyes to this section, under the heading "Key personnel".	Chiara Bearzotti /DKRZ	19/05/2015

4.1	The scientist Laure Coquart who was named as staff of CERFACS was moved to CNRS were she is currently employed. We have been informed about this fact by our partner CNRS and corrected this mistake.	Kerstin Fieg / DKRZ	20/05/2015

History of changes introduced with the Amendment AMD-675191-12

Section	**Details of changes made**	**Changes made by whom**	**Date**
Logo in Part B	Removed from the DoA	Chiara Bearzotti	20 October 2016
Section 1.4, Part B	Sentence added at the bottom of the WP 2 description in this section.	Chiara Bearzotti	20 October 2016
Section 2.2.1, Part B	In the table, the wording “User Group Committee” has been replaced with “User Groups” in the table (Cell: Target Audience/Dissemination and Exploitation Strategy)	Chiara Bearzotti	20 October 2016
Section 2.2.3, Part B	Wording “User Group Committee” replaced with “**User Groups of community models and tools”**	Chiara Bearzotti	20 October 2016
Table 3.1a, Part B	New column added to this able for clarity with the indication of the estimated person-months per WP	Chiara Bearzotti	20 October 2016
Table 3.1a, Part B	Estimated person months for WP2 are now 202	Chiara Bearzotti	20 October 2016
Table 3.1a, Part B	• Co-Leadership for WP3 Met Office now co-leading together with BSC • In BSC: WP co-leader changed to Kim Serradell, due to changes in staff at BSC • Total estimated person months for WP3 are now 123	Chiara Bearzotti	20 October 2016
Table 3.1a, Part B	Estimated person months for WP4 are now 106	Chiara Bearzotti	20 October 2016
Section 3.1.1., Part B	Reformulation of the description of the workplan for WP2	Chiara Bearzotti	20 October 2016
Fig. 3.1a, Part B	Renaming of the Task T2.2 with this wording “Global high-resolution model demonstrators”	Chiara Bearzotti	20 October 2016
Table 3.1b, Part B	Gantt chart amended as necessary to reflect changes in the deliverable and milestone timing	Chiara Bearzotti	20 October 2016
Figure 3.2, Part B	Wording “User Group Committee” replaced with “User Groups”	Chiara Bearzotti	20 October 2016
Section 3.2.1, Part B	The parts of the text related to the “User Group Committee” have been removed.	Chiara Bearzotti	20 October 2016

Section 3.2.2, Part B	• Reformulation of the text related to the role of the User Group Committee, due to its replacement with the "User Groups of community models and tools" • Slight reformulation on the role of the General Assembly	Chiara Bearzotti	20 October 2016
WP4 and Deliverables list, Portal	D4.1 new timing: PM18 D.4.2 new timing: PM18 D 4.3 new timing: PM 46 D4.4 new timing: PM42 D4.5 deleted	Chiara Bearzotti	20 October 2016
WP4, Portal	Task 4.4. ha been removed to make effort available elsewhere in the project to support the demonstrator project.	Chiara Bearzotti	20 October 2016
WP4, Portal	• Task 4.2: ECMWF will not take part in this task. Part of the text has been amended. • In the section related to the person-month effort of the partners involved: ECMWF deleted from the participants in WP4	Chiara Bearzotti	20 October 2016
WP4, Portal	Objective 3 integrated with an additional sentence. Objective 4 removed.	Chiara Bearzotti	20 October 2016
WP5 and Deliverable list, Portal	D5.5 new timing: PM18	Chiara Bearzotti	20 October 2016
WP1, Portal	Task 1.1.1 Wording added to the second bullet point "demonstrators"	Chiara Bearzotti	22 October 2016
WP1, Portal	Task 1.4.1 has a new title "Develop a roadmap to implementation of an ensemble of ESMs with 1km resolution", its description" has been revised and enhanced	Chiara Bearzotti	22 October 2016
WP1, Portal	Task 1.4.2 Sentence added at the end of the paragraph	Chiara Bearzotti	22 October 2016
WP1 description and Deliverables List, Portal	D1.2 reformulated in the title and revised time of delivery (now PM46, before was PM36)	Chiara Bearzotti	22 October 2016
WP2, Portal	In the section related to the person-month effort of the partners involved: • ECMWF estimated person-month effort now estimated at 49 t was (i37 before) • MET Office estimated person-month effort is reduced from 3 to zero, thus it has been removed in the Portal. Met Office will nevertheless contribute to the activities in this work package through active participation in workshops	Chiara Bearzotti	22 October 2016
WP2, Portal	Objectives: one bullet point added, the existing bullet points have been integrated and amended with more specifications	Chiara Bearzotti	22 October 2016
WP2, Portal	Task 2.1.1 new sentence added: These models, together with ICON, will constitute the heart of our high-resolution demonstrators (see T2.2).	Chiara Bearzotti	22 October 2016
WP2, Portal	Task 2.2, Task 2.2.1 and Task 2.2.2 have been revised to introduce the Demonstrator	Chiara Bearzotti	22 October 2016
WP2, Portal	Task 2.3: additional sentence at the end of the paragraph	Chiara Bearzotti	22 October 2016
WP2, Portal	Task 2.3.2: additional sentence at the end of the first bullet point and slight reformulation at the end of the last bullet point	Chiara Bearzotti	22 October 2016

WP2, Portal	Task 2.4.2: CNRS-IPSL does not take part in this task any longer	Chiara Bearzotti	22 October 2016
WP2 description and Deliverables list, Portal	• D2.1 new title, new delivery date, moved to DKRZ • D2.3, D2.4, title modified in the wording • Newly introduced deliverables: D2.8, D2.9, D2.10, D2.11, D2.12	Chiara Bearzotti	22 October 2016
WP3 description, Portal	In the section related to the person-month effort of the partners involved: • MPG now with 18 person-months • Met Office now with 57 person-months	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Objectives list: fourth bullet point integrated with reference to the demonstrator	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Co-Leadership changed. • Co-Leadership for WP3 Met Office now co-leading together with BSC • In BSC: WP co-leader changed to Kim Serradell, due to changes in staff at BSC	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Task 3.1.3 new sentence added for reference to the demonstrators	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Task 3.2 new sentence added for reference to the demonstrators	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Task 3.2.2 and Task 3.2.3 added reference to the demonstrators	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Task 3.3.1: reference to D1.3 added in the text	Chiara Bearzotti	22 October 2016
WP3 description, Portal	Task 3.3.4: new sentence added to the text and two links provided in the text.	Chiara Bearzotti	22 October 2016
WP3 description and Deliverables list, Portal	• D3.2, title revise and integrated. • D3.5: delivery time changed to PM14 • D3.7: title integrated and delivery date changed to PM18 • D3.8: title integrated and delivery date changed to PM30)	Chiara Bearzotti	22 October 2016
Milestone, Portal	Milestones revised in the system to adapt to the demonstrators: MS2, MS3, MS4, MS6, MS7, MS8, MS9. Refined in their links to the work packages and in the time of delivery.	Chiara Bearzotti	22 October 2016
Deliverable list, Portal	Deliverable D5.4 changed to type "ORDP: Open Research Data Pilot"	Chiara Bearzotti	26 October 2016

TABLE OF CONTENTS

1 EXCELLENCE **6**
1.1 Objectives 6
1.2 Relation to the Work Programme 8
1.3 Concept and approach 9
1.4 Ambitions 15
2 IMPACT **18**
2.1 Expected impacts 18
2.2 Measures to maximise impact 22
3 IMPLEMENTATION **27**
3.1 Work plan — Work packages, deliverables and milestones 27
3.2 Management structure and procedures 31
3.3 Consortium as a whole 35
3.4 Resources to be committed 39
4 Members of the consortium **47**
4.1 Participants (applicants) 47
4.2 Third parties involved in the project (including use of third party resources) 86
5 Ethics and Security **87**
5.1 Ethics 87
5.2 Security 87

1 EXCELLENCE

Numerical weather prediction and climate modelling are highly dependent on the available computing power in terms of the achievable spatial resolution, the number of members run in ensemble simulations as well as the completeness of physical processes that can be represented. Both domains are also highly dependent on the ability to produce, store and analyse large amounts of simulated data, often with time constraints from operational schedules or international coordinated experiments. The ever increasing complexity of both numerical models and high performance computing (HPC) systems has led to the situation that today, one major limiting factor is no longer the theoretical peak performance of available HPC systems, but the relatively low sustained efficiency that can be obtained with complex numerical models of the Earth system.

The differences in model complexity as well as the temporal and spatial scales that were historically characteristic for climate and weather modelling are vanishing since both applications ultimately require complex Earth system modelling capabilities which resolve the same physical process detail across atmosphere, ocean, cryosphere and biosphere. With increasing compute power and data handling needs, both communities must exploit synergies to tackle common scientific and technical challenges. A joint climate and weather community engagement in research and service provision is urgently required and timely given the substantial societal investment in European wide infrastructures through the Copernicus services that are coordinated and managed by the European Commission.

ESiWACE will deliver the required research and services through leveraging two established European networks, namely (1) the European Network for Earth System modelling (ENES), representing the European climate modelling community contributing to the internationally coordinated experiments of the World Climate Research Program (WCRP) and the Intergovernmental Panel on Climate Change (IPCC) assessments, and (2) the world leading European Centre for Medium-Range Weather Forecasts (ECMWF), which is an independent European organisation supported by 34 member and cooperating states, and also the operator of the Copernicus services for atmospheric monitoring (CAMS) and climate change (CCCS).

1.1 Objectives

ESiWACE pursues the following objectives

ESiWACE will substantially improve the efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms.

> Improvements in efficiency, and hence in productivity, will be delivered by developing solutions for cross-cutting HPC challenges particular to the weather and climate domain. Solutions will range from the development of specific software products to the deployment of user facing services – and will encompass both computing and storage.

ESiWACE will support the end-to-end workflow of global Earth system modelling for weather and climate simulation in high performance computing environments.

> We will improve and support Scalability of models, tools and data management on state-ofthe-art supercomputer systems, Usability of models and tools throughout the European HPC eco-system, and the Exploitability of the results.

The European weather and climate science community will drive the governance structure that defines the services to be provided by ESiWACE.

We have already identified major themes which will immediately be addressed in the project. Representatives from the European weather and climate science community will contribute to setting priorities and define operational governance. A network of supporters will be formed and external organisations will be invited to participate in workshops, general assemblies, and training activities to support this process.

ESiWACE will foster the interaction between industry and the weather and climate community on the exploitation of high-end computing systems, application codes and services.

Our challenges mandate a trans-disciplinary approach. ESiWACE will foster co-design and vendor engagement through specific and generic industrial and academic interactions across the software and hardware spectrum.

ESiWACE will increase competitiveness and growth of the European HPC industry. Weather and climate computing have always been one of the key drivers for HPC development, with domain specific scientific and technical requirements that stretch the capability and capacity of existing software and hardware to the limits. By developing solutions for Europe and at European scale, ESiWACE will directly impact on the competitiveness of the European HPC industry by engendering new products, providing opportunities for exploitation beyond the project itself, and by enhancing the skills base of staff in both industry and academia.

How ESiWACE's objectives translate to specific goals in the work plan (section 3.1).

ESiWACE will:

- Provide **services** to the user community that will impact beyond the lifetime of the project. This will be obtained by **engaging** with the user community to set up governance structures with long-term objectives and work strategies. (WP1 and WP5)
- Improve **scalability** and shorten the time-to-solution for climate and operational weather forecasts at increased resolution and complexity to be run on future extreme-scale HPC systems.
  This will be obtained by investigating and implementing scientific and technical options for compute efficiency upgrades and model output reduction (WP2).
- Foster **usability** of the available tools, software, computing and data handling infrastructures. This will be obtained by identifying, designing and supporting the end-to-end workflow for climate modelling and weather forecasting applications in both research and production mode (WP3).
- Pursue **exploitability** of climate and weather model results. This will be obtained by addressing the major barriers that hinder the efficient use of the huge amounts of data produced by weather prediction and climate simulation, in particular the performance and volume limits of key storage technologies on a range of timescales (WP4).
- Establish governance of common software management to avoid unnecessary and redundant development and to deliver the best available solutions to the user community. This includes also the provision of **support and training** for software tools and **documentation of best practices** for efficient climate and weather simulations using state-of-the-art HPC resources (WP1 and WP2).

- Provide **open access** to research results and **open source** software at international level. This will be obtained by sharing results, codes and documentation with a broad audience from diverse communities beyond our own, thus developing a science culture of excellence (all WPs).
- Exploit **synergies** with other relevant activities and projects and also with the global weather and climate community (all WPs, but driven by WP1).Liaise with the **European HPC ecosystem**, in particular with the European Technology Platform for High Performance Computing (ETP4HPC[1]) and the Partnership for Advanced Computing in Europe (PRACE[2]) (WP1).

1.2 Relation to the Work Programme

ESiWACE responds to the call EINFRA-5-2015 – Centres of Excellence for computing applications.

ESiWACE addresses the specific challenge of the call EINFRA-5-2015:

HPC applications are essential to perform the complex multi-dimensional and multi-scale calculations and to handle the large-scale datasets necessary for modern numerical weather and climate simulation. ESiWACE will develop a European culture of excellence field by focusing on the needs of the user community. ESiWACE activities will join weather and climate modellers, software developers, vendors, and HPC centres: scientific discovery will evolve in parallel and in synergy with cutting-edge HPC development. The achievements of ESiWACE will go beyond the scientific and technological challenges they directly aim at, by impacting European society as well.

Given the challenges addressed and the activities planned, ESiWACE will be at once:

- **Thematic**, as it focuses on the HPC domain of climate and weather modelling,
- **Transversal**, as it covers several aspects of computational science (algorithms, analytics, numerical methods), and
- **Challenge-driven**, as progress in climate and weather predictability represents a major societal issue with multi-disciplinary implications

ESiWACE addresses the expected scope of the call EINFRA-5-2015:

- ESiWACE is a **user-driven** Centre of Excellence: the partners of the ESiWACE consortium are direct application users and owners. The ESiWACE governance model is explicitly based on implementing a user committee to provide feedback on on-going development of model codes and software tools, together with an advisory board of external scientific experts. The users have a well-defined role in the governance and are involved directly in the implementation of the activities.
- ESiWACE is **integrated**, encompassing not only HPC software but also relevant aspects of hardware, data management, data storage, connectivity, and security. Public and private partners will concentrate on development and scaling of climate and weather models toward the exascale, on evaluation and optimisation of specific model components and on data management on different architectures with, among others, development of methods for exploiting storage supports. The usability of all resources will be enhanced with tasks on benchmarking, co-design and development of best practices.
- ESiWACE is **multidisciplinary**, as it covers both HPC and weather and climate modelling areas, and fosters the interaction between industry and academia on the exploitation of highend computing systems, application codes and services communities.
- ESiWACE is **distributed**, by federating capabilities around Europe (France, Germany, UK, Italy, Spain, Ireland, Sweden), and will **ensure synergies** with existing national programmes and with several joint European efforts already in place such as the development of the

[1] http://www.etp4hpc.eu/

[2] http://www.prace-ri.eu/

coupled climate model EC-Earth[3], the ocean modelling platform NEMO[4] or the coupler module OASIS[4].

- ESiWACE **exploits available competences**, as it brings together general-purpose HPC centres involved in PRACE (Partners BSC and ICHEC), dedicated centres (Partners ECMWF, DKRZ, MetO, CMCC, STFC), and research institutions and private partners developing hardware and software. Each partner of ESiWACE contributes its own network of users of climate and weather simulations. All these groups and institutions are connected in a larger network of supporters to ESiWACE. ESiWACE will strengthen the scientific and technological exchange and therefore the EU competitiveness in these fields.

ESiWACE will directly address:

- **Provision of services** such as: (1) Optimisation of model codes and associated software tools as well as development of new methods for data storage to improve efficiency for petascale computing and prepare for exascale computing. (2) Benchmarking activities and governance to test and validate shared community software and to contribute to their maintenance. (3) Investigation into the potential for co-design through involvement (as partners) of hardware-vendors and an SME providing tools for software development. Addressing the skills gap in computational science through support and training of weather and climate scientists.
- **Synergy with the pan-European HPC infrastructure:** We will organize relations with PRACE ensuring adequate dedicated resource allocation and will work with PRACE and other HPC centres towards provision of an efficient and adequate software environment for weather and climate simulations using complex Earth system models. We will network with ETP4HPC, relevant hardware and software industry and the weather and climate scientists to foster efficient use of HPC in this context. In particular strengthen the European HPC strategy by developing a specific roadmap for weather and climate simulations.
- Long-term **sustainability** of support and services by developing options for sustained funding in form of a business plan. [D1.3]
- **Creation of communities around specific codes** by establishing or strengthening (when already in progress) the user-driven evolution of the community software via improved user support, enhanced training, and fast integration of existing models and tools (NEMO and ECEarth models, OASIS coupler, XIOS[5] I/O server, Cylc[7] meta-scheduler), in response to a welldefined software governance strategy.
- A **governance structure driven by the needs of the users**: The governance of ESiWACE will be driven by the user community, made up of experts from the fields of Earth system modelling across the public and private sector. The continuous interaction with the user community sought through various instruments (workshops, user groups, direct feedback, and general assemblies) will allow a progressive definition and adjustment of the scope of all ESiWACE activities

1.3 Concept and approach

1.3.1 Overall concept underpinning the project: main ideas, models, assumptions, transdisciplinary considerations.

ESiWACE targets the convergent use of Earth System Modelling for weather and climate science (ESM)[6]. Global Earth System Models, the post-processing methods to handle the vast amounts of

---

[3] EC-Earth: http://ecearth.knmi.nl/ [4]
NEMO: http://www.nemo-ocean.eu/

[4] OASIS: https://verc.enes.org/models/software-tools/oasis

[5] XIOS:http://forge.ipsl.jussieu.fr/ioserver/wiki [7]
Cylc: http://cylc.github.io/cylc/

[6] **Throughout this application the term „Earth System Modelling" (ESM) is used as short for „Earth System Modelling for weather and climate science". Earth System Modelling in a broader sense would also incorporate the solid Earth.**

data they produce, and their complex end-to-end workflow, are the central tools for weather and climate science - in both research and operations. The productivity that can be achieved with these tools is significantly limited by technical and structural properties of the available software, hardware, computing and data infrastructure. Many of these bottlenecks are domain-specific since they originate from special requirements inherent to climate and weather simulation. However, there are also examples of shortcomings that appear domain-specific but are merely historically artefacts. Therefore, overcoming these bottlenecks does not only require meteorological or climatological expertise, but multi-disciplinary efforts.

To address this scientific and technical challenges ESiWACE will leverage multi-disciplinary, world leading expertise hosted in Europe in the fields of weather forecasting, climate modelling, computer science and the provision of HPC services. ESiWACE will expand well beyond individual existing efforts and activities (see section 1.3.2) by a concerted multi-disciplinary action. The leading partners of ESiWACE are centres dedicated to climate research and weather prediction that combine a wide spectrum of expertise in research as well as operational production. This will ensure synergies with available services and domain specific technical competence, and it will also ensure that the achievments of ESiWACE directly translates into user benefit – both within the consortium and for the wider community. Moreover, each project task will be supported by domain and technical experts from specialized institutions and partners from industry, to foster co-design of basic computer code components, community models and tools in close touch with available and future hardware.

The central integrating part of our approach is "**Governance and Engagement".** We will create a governance structure, which ensures that the users from the weather and climate community (inside and outside the consortium) define the scope and priorities of the services to be provided by ESiWACE. Long-term sustainability of reliable support and service will be sought and achieved through the development of a business plan. In the preparation of the proposal the major themes to be addressed in ESiWACE have been identified in interaction with members of the wider scientific community. The emerging themes, having been mapped onto three technical work packages in ESiWACE, are:

The "Scalability" theme will coordinate performance assessment and upgrade development of state-of-the-art community models and tools, for example, ESMs, coupling and I/O technologies. To ensure that the community will be able to optimally exploit these models and tools, support and training will be provided. Community wide coordinated performance intercomparisons, along with review, analysis and implementation of scientific and technical solutions for efficiency increase will be carried out. Further towards the exascale horizon, selected strategies will be investigated to prepare for the next-generation models and tools, and the workflows associated with challenging volumes of high-resolution model output.

The "Usability" theme will focus on the ease-of-use of available tools, software, computing and data handling infrastructures for ESM. Aspects of application software and data handling environment as well as system software and hardware stack will be covered for both research and operations. Further, the end-to-end workflow in ESM will be addressed and improved by providing recommendations, example use cases and best practice definition. Prototypical support and development will be put in place for a common workflow solution for weather and climate modelling building on an existing, successful product.

The "Exploitability" theme addresses major data handling barriers that hinder exploiting HPC for ESM. In a first stage, ESiWACE will address how to best exploit existing and future storage hardware and software at peta- and exascale. ESiWACE will aim at improving the performance and capabilities of both disk and tape storage systems in ESM workflows. Disk performance will be improved by developing "Earth system data aware" software libraries that optimize information layout in heterogeneous disk storage environments. Tape performance and capacity will be improved by developing new "Earth system data aware" strategies and software.

1.3.2 National or international research and innovation activities linked with the project ESiWACE will establish strong links with a broad range of projects:

(a) ESiWACE will be a significant **beneficiary of the FP7 project IS-ENES2** (Infrastructure for the European Network for Earth System modelling)[7]. IS-ENES2 is the second phase project of the distributed e-infrastructure of models, model data and metadata of ENES (European Network for Earth System modelling[8]). It integrates the European climate modelling community, stimulates common developments of software for models and their environments, fosters the execution and exploitation of high-end simulations and supports the dissemination of model results to the climate research and impact communities.

ESiWACE will **extend and widen the scope of IS-ENES2**, exploiting developments from ISENES2, issues identified by IS-ENES2, and growing the community. In particular:

- ESiWACE will **take on governance** of selected software tools relevant to the three ESiWACE themes, extending the scope of the support to include the weather community, as well as delivering developments beyond the end of ISENES2, resulting in a long-term support and service activity for the weather and climate community.
- Several **concepts developed** in ESiWACE originate from networking activities in IS-ENES2, such as the need to invest in a common meta-scheduler Cylc, the need to develop common workflows, to address code convergence issues the benefit of annual international HPC workshops and to define common protocols to compare performance of Earth system models. Since IS-ENES2 is funded until March 2017, both projects will run in parallel for a period of ca. 18 months. We will profit from this time overlap to ensure the full transfer of results relevant to longterm support of Earth system models in HPC environments from the I3 project IS-ENES2 into ESiWACE.

(b) ESiWACE will exploiting **existing partner activities and collaborations aiming at improvement of efficiency of Earth system models**. ESiWACE will leverage these from primarily bilateral (or at best trilateral) activities by fostering common developments, information exchange and co-design, always pursuing the core objective of improving the effectiveness of the weather and climate community in Europe as a whole. Examples include the cooperation of CNRSIPSL, CMCC and Bull to improve the NEMO model, the cooperation of DKRZ, MPG and Bull to improve the ICON[9] model and its I/O performance, and the ECMWF Scalability programme relying on partnerships with Cray and NVIDIA. The latter aims at developing the next-generation forecasting system addressing the challenges of the future exascale high-performance computing and data management architectures. Similar activities are pursued by the MetO (GungHo and LFRic with UK NERC) and by MPG, DWD and DKRZ (HD(CP)^2 project with German BMBF support)[10].

(c) Another foundation activity in ESiWACE will be exploiting **existing and new relationships with groups working on storage challenges.** CMCC, DKRZ and STFC all have existing storage vendor engagement programmes, and a range of existing and planned European collaborations. ESiWACE will allow the leveraging of these activities towards common goals, and the dissemination of best practice into the wider community. Key relationships will include the EUDAT community (whether or not an expected EUDAT2 project is funded) and with the exascale I/O community (especially the European Open File System – EOFS – working group within the E10consortium[11]). ESiWACE partners established relationships with netCDF and GRIB technical teams and launched an early relationship with the HDF technical team[12]. At the working level we also share problems and solutions with the high-energy physics community through close links with the LHC storage community). On data dissemination, close collaboration with the Earth System

[7] http://is.enes.org

[8] https://portal.enes.org/

[9] ICON: http://www.mpimet.mpg.de/en/science/models/icon.html

[10] HD(CP)2: http://hdcp2.zmaw.de/

[11] E10: http://www.exascale10.com

[12] The netCDF http://www.unidata.ucar.edu/software/netcdf/ and GRIB http://en.wikipedia.org/wiki/GRIB are two major file formats in use in the community, modern versions of netCDF are based on version 5 of the Hierarchical Data Format, HDF, http://www.hdfgroup.org/HDF5/.

Grid Federation (ESGF), as supported within IS-ENES2 comes naturally. Partner DKRZ and STFC are the leading European organisations in technical support of ESGF.

(d) ESiWACE envisage interacting with future H2020 **Future and Emerging Technologies** research projects that deal with new long-term developments of Earth system models and their components and system models for future HPC-architectures. Existing project proposals of this category include the ESCAPE proposal led by ECMWF and the CHANCE proposal led by CMCC.

(e) Alongside these technical collaborations ESiWACE will **also liaise with established user communities and scientific projects** using ESM as a tool. This includes the EC-Earth consortium (currently chaired by SMHI), consortia managing regional weather applications such as ALADIN[13], HIRLAM[14], LACE[15], and COSMO[16] as well as existing and planned weather and climate European environmental (SC5-2014-1 "Earth systems" call under final evaluation) projects such as
(e) CRESCENDO, IMPULSE, PRIMAVERA, and the European Training Network Marie SkłodowskaCurie submitted project HPC4CM. Members of these groups will be invited to workshops and to participate in user groups for specific foster code to ensure ESiWACE activities match their requirements.

(f) ESiWACE will also **closely work with WMO programs** in the field of climate (World Climate Research Program, WCRP), and weather (World Weather Research Program, WWRP) in order to ensure ESiWACE developments serve international requirements. ESiWACE will cooperate with the CMIP panel coordinating international experiments (e.g. CMIP6), related intercomparison projects (such as OBS4MIPS[17]) and other WCRP and WWRP international projects in which the ESiWACE community are engaged. An agreed activity will be the planning of a few **joint workshops in coordination with WCRP and WWRP** working groups and programs (such as S2S[18], WGNE[19], WGCM[22]) to demonstrate usability, exploitability, seamless access, data processing integration for services development.

Letters of support and commitment to ESiWACE from several of these networks have been provided in Appendix 1.

1.3.3 Overall approach and methodology

The approach of ESiWACE follows the general overall approach: (i) networking to establish common requirement and governance methodology; (ii) joint research activities to develop towards the requirements, and (iii) service activities to provide both training and support.

The type of **networking activities** depends on the community's maturity in a given area. In some areas we can build on previous activities around established shared software especially from ISENES2. This includes the evaluation of coupling technologies and the support of the OASIS coupler, the development of I/O strategies, the development and support of the XIOS I/O server and the Cylc workflow engine, common analyses of model codes such as NEMO and EC-Earth. In this case **networking for existing community software** will foster community co-operation based on an agreed set of benchmarks and will drive continued research towards improvement of these tool through engagement with the existing user community. Dissemination of community software will be enhanced by improved user support and training as well as fastest possible integration of new software versions.

In other areas, the community's maturity is lower and will focus on **networking for new community software** with activities supporting to share specialist knowledge that exists but is not yet available across the community. An example is the development of suitable software stacks for HPC systems.

Networking activities will feature introductory workshops on specific tasks to foster cooperation between the partners, and include external stakeholders and the wider community. The outcomes of the workshops will be white papers in the form of web documents, followed by updates as demanded by changing environments and evolving requirements. A central overarching objective of the **networking activities** will be the provision of a governance framework for the sustainable provision of community codes and the ability to identify and to prioritize user-driven requirements.

---

[13] ALADIN: http://www.cnrm.meteo.fr/aladin/
[14] HIRLAM: http://www.hirlam.org/
[15] LACE: http://www.rclace.eu/
[16] COSMO: http://www.cosmo-model.org/
[17] OBS4MIPS: https://www.earthsystemcog.org/projects/obs4mips/
[18] S2S: Subseasonal to Seasonal Prediction Project -http://www.wmo.int/pages/prog/arep/wwrp/new/S2S_project_main_page.html
[19] WGNE: Working Group on Numerical Experimentation - http://www.wmo.int/pages/about/sec/rescrosscut/resdept_wgne.html [22] WGCM: Working Group on Coupled Modeling - http://www.wcrp-climate.org/wgcm/

This is underpinned by having a dedicated work package (WP1 on governance, engagement and long-term sustainability), which will take charge of **networking across communities**: science (climate and weather science), HPC ecosystem and HPC industry.

Joint research towards enhancement of existing and new development of systems and services is different across tasks, depending upon ground laying work achieved in other contexts like ISENES2 or others. The mechanism of governance of shared software and of innovative ESiWACE developments, which is secured through WP1, will make sure that the development work follows the requirements of the users. In particular the specifications for the development will be reviewed and discussed during the project. Additionally we will set up dedicated interest groups on crosscutting issues which are of interest for all of our three themes (for example the IO problem or the impacts of evolutions in the HPC ecosystem). Through these mechanisms we hope to maintain the flexibility to adjust the path of the research activities and to set appropriate goals for the continuation of our ESiWACE after the initial funding period.

Joint research activities will effectively contribute to quantitative and qualitative improvements of the services provided to the weather and climate community. Joint research on model and tool optimization will ensure that existing codes, which represent an integral part of the ESiWACE infrastructure, will evolve to better exploit massively parallel supercomputers, thereby enhancing the effectiveness of climate and weather research based on these codes. In addition, undertaking first steps towards the preparation of the next-generation model components will also help securing these efficiency gains in the longer term on exascale platforms. The research into development and provision of a common software stack and common runtime environment for ESM will considerably ease the path for new users and also for the deployment of complex model systems on different HPC platforms. The research on knowledge compression of ensemble data, efficient storage of ESM data (on disk and tape) and on the mapping between netCDF and GRIB file formats will lay the basis for providing support and service for a novel and much more efficient way to handle and share data within the weather and climate community.

Prototypical and limited **services** for European community climate models and tools were for the first time offered in the framework of IS-ENES and IS-ENES2 and are exploitable on-line through the ENES portal. ESiWACE will ensure that these services are maintained, enlarged and improved. This includes user-driven improvements of central models and tools and the delivery of support and training for existing community tools. **New services** will mainly be developed within the usability theme around the provision of a common HPC software stack and run time environment for weather and climate models. Both will proto-typically be installed and tested in a quasi-operational environment. As a first real use case the usage of the environment within an ENES summer school is foreseen. The **definition of a broader portfolio of services** is an integral goal of the ESiWACE project itself and will be an integral part of two central deliverables of the project: The HPC-roadmap for the climate and user community and the business plan for the future of this Centre of Excellence in Simulation of Climate and Weather in Europe.

1.3.4 How sex and/or gender analysis is taken into account in the project's content

Issues regarding gender equality, equal opportunity and diversity are considered extremely important. After careful thought, we consider the gender dimension of ESiWACE as neutral. The Consortium agrees to undertake actions during the course of the project to promote and guarantee gender equality in the project, ensuring it will act upon the EC recommendations listed in the "Gendered Innovation". Additionally, the Consortium is ready to contribute to surveys and investigations fostered by the European Commission.

At a project level, the Consortium is aware of the importance of attracting more high quality female researchers into the sphere of research and innovation and management, stimulating and promoting the progress of women in scientific careers. Currently, two of eight work package leaders / co-leaders are women. Gender balance among the personnel primarily responsible for carrying out the research and innovation activities will be ensured. The Consortium will support equal participation between men and women in the implementation of the actions and will aim to the extent possible for a gender balance at all levels of personnel assigned to any action, including at supervisory and managerial levels (ref. Art. 33 of the Grant Agreement). The promotion and monitoring of gender equality throughout the project will be the responsibility of the coordinator.

1.4 Ambitions

As indicated in section 1.3.1 ESiWACE focuses on three major themes, namely **Scalability**, **Usability**, and **Exploitability** addressing technical challenges and on a **Governance** activity to engage the community and to align the services along the requirement of the users..

Governance (WP1)

Currently, the productivity of climate and weather simulation is significantly limited by the technical and structural properties of the available software, hardware, computing and data-handling infrastructure. While selected individual solutions exist, the community does not benefit from a common strategy on how to develop portable methodologies that benefit a wider user group and that can be maintained and extended with complementary resources. The **ESiWACE ambition** is to introduce such a strategy through a governance of science software, models and tools with more developed community engagement. It will transpose existing and newly generated knowledge from current to future technologies and train young scientists. Further, ESiWACE will foster the interaction between industry and academia on the exploitation of high-end computing systems, application codes and services.

In our research field central, crucial software is generally developed in single institutions. It is not directly generating revenue. ESiWACE will nevertheless strive to engage and involve industry and solutions providers in long-term commitment to development, support and maintenance of such software, and to generate benefits and revenues of it from downstream usage, as e.g. for improved PDE solvers in the oil and gas industry. Enhanced shared software developments will be a method offering a huge **potential** in terms of focusing institutional resources onto science rather than onto technology and method development:

Further, ESiWACE will leverage significant **innovation potential** by bridging the perceived gap between **"traditional"** numerical weather prediction and climate modelling, which were understood to be physical only for the former, and climate time scale only for the latter: Both WMO and WCRP sport the concept that "weather" and "climate" work on space and time scales and with processes, which will become less and less differentiable, and are conceptually both described by the term "Earth system". ESiWACE will support this approach in the recognition that the weather and climate science communities have much more in common than usually perceived, and that the respective modelling or simulation approaches face problems which are very similar, especially when exascale comes into play.

These goals are both **technically and organisationally ambitious**, particularly in dealing with the necessity to integrate activities and responsibilities between institutions and across domains such as academia and industry. By centralizing these activities and providing an interface to other similar endeavours, we expect a big step forward for our community but also will be able to significantly contribute to reinforcement and growth of the European HPC ecosystem as forseen in the Strategic Research Agenda of ETP4HPC[20].

Scalability (WP2):

Historically, code is developed with a science focus, and code architectures or numerical techniques are developed by single institutes with a well-defined application in mind. Code optimizations are only applied incrementally to these code bases and within the limitations of the computing expertise at hand. This approach has produced very large legacy codes, which are difficult to manage within a wider community, difficult to migrate between different HPC

[20] "European Technology Platform for High Performance Computing ETP4HPC Strategic Research Agenda Achieving HPC leadership in Europe". (http://www.etp4hpc.eu/strategy/strategic-research-agenda/)

architectures and nearly impossible to adapt to future exascale compute and data handling environments requiring, in parts, fundamentally new code design, work flow management and scientific algorithms.

The solution to this dilemma is a combination of radical and incremental change of direction, for which ESiWACE will **pave the way** in tight collaboration with the weather and climate modelling community.

More radical re-engineering is rather difficult to perform for code used across institutes or communities due to user and application specific requirements and hardware. However, there is often **potential** for unifying the framework in which user-specific code is embedded as well as for identifying common and generic code components that can be managed and optimized as common libraries. ESiWACE will foster this through the **development of strategies** for (i) new unifying coupling and I/O functionalities, (ii) concurrency in crucial science code workflows, and for examining the trade-off between (i) precision and memory/speed and (ii) information content and data volume.

At the same time, less intrusive code adaptation is preferable where substantial efficiency gains can already be obtained without intrusive code re-design and on shorter time scales. ESiWACE will **provide this** by supporting more stringent (i) overlap of communication and computation, (ii) twosided and asynchronous communication, or (iii) the use of OpenMP in community models and tools. **Fundamentally new** in this context is a dedicated effort to define efficiency metrics (i) allowing an objective benchmarking of community tools and (ii) guiding code optimization efforts.

This **level of ambition** has never been realized in this community as it breaks away from the existing legacy code development approach and because it introduces a view ranging from support and evolutionary changes for existing tools towards designing next-generation systems suitable for the exascale.

ESiWACE adds **substantial innovation potential** in the following areas:

- Through unified and community wide science and computing strategies, the pull-through of new scientific and computing methodologies into community models and tools will be greatly facilitated and provide benefit for a much wider user group than before.
- Those European Copernicus services relying on weather and climate modelling will immediately benefit from developments and services.
- The present European excellence in scientific research and operational applications in weather and climate will be extended to the area of computing and data handling.
- The interaction with hardware and software providers will be greatly facilitated since co-design of code components and libraries will be facilitated, and hardware dependent performance will be easier to gauge. Thus procurements for future hardware will be optimized.
- The European weather and climate community will be prepared for the challenges presented by exascale HPC systems.

Starting point will be the establishment of so-called demonstrators of atmosphere-only, ocean-only and coupled ocean-atmosphere simulations, which will be run at the highest affordable resolutions (target 1 km) to estimate the computability of configurations that will be sufficient to address key scientific challenges in weather and climate prediction.

Usability (WP3):

Today, it is realized that sophisticated and flexible workflow solutions are increasingly important in production environments. However the emerging solutions are still far from universal and currently rare in the research environment. IS-ENES2 has established a growing appetite for a step change in capability of workflow solutions in the research environment and this proposal is able to capitalize on recent investments at NIWA[21], the MetO, MPG and others aimed specifically at this user base.

[21] The NZ National Institute of Water and Atmosphere

ESiWACE has the ambition to significantly improve the interaction between those with deep computing knowledge and those with the best scientific ideas. This way we will drive research in workflows solutions which offer a much greater potential for performance optimization in the noncomputer-architecture-minded sense, as does the standard way of experiment design and execution. This will allow for considerable advances in a number of areas:

- Large difficulties exist to organize and carry out multi-model ensembles (see projects like PRIMAVERA, IMPULSE, CRESCENDO) ESIWACE will develop an environment to remedy this situation, including education of young researchers.
- The complete stack from the science application down across the complete system and data handling software to the hardware is much more heterogeneous than it is healthy for the communities involved. ESiWACE will provide some counterbalance against the commoditization trend currently observable in the computing industry by testing solutions, proposing and proliferating standards, and educating young scientists in their use.
- Information about best practice and working examples is often missing. ESiWACE will improve this with its dissemination methods.
- Involvement with solution providers is low. Providing a funded platform to engage and exchange with industry, also directly with ETP4HPC by involving an SME is a new approach. ESiWACE will gain the attention of the computing industry via greater and better co-ordinated engagement for the ESM community. With this activity ESiWACE will be very involved with the milestones of ETP4HPC, "Programming Environment", from 2016/17 **on**.
- Increased complexity of earth system model suites and the need to automate more data pre- and post-processing means that there is the urgent need to find tools to free the scientists from the increasing burden as HPC resources grow. ESiWACE' activity on meta-scheduling, like the provisioning of the Cylc workflow engine, suitable for research and production environments and specifically developed for the climate and weather communities, provides the opportunity to give step-change improvements in the management of complex workflows.
- Dissimilar and disparate working environments and software stacks are a hindrance not only for multi-site, multi-model high-resolution full complexity Earth system model experiments, but also for the individual researcher needing to be flexible in terms of the usage of his computational and storage resources across different (topical or PRACE) sites, and for the software engineer in need of benchmarking his model or tool across different platforms. A huge potential lies in the provisioning of recommendations for shared common environments and software stacks across sites and architectures in terms of usability and maintainability.
- Rational scheduling of simulations based upon concrete parameters of the according experiments has the **potential to exploit** machines and resources much more elegantly than possible currently, and will be supported by ESiWACE through training and services for provisioning of technical support.

Exploitability (WP4):

Data volumes produced by the weather and climate community have been growing with a doubling time just short of two years for decades. Ongoing growth at these rates will stress **existing infrastructure** – both in terms of storage and bandwidth to and from storage – and will stress timeto-solution for most weather and climate workflows (where solution is defined as the production of knowledge or information, not data). It is **already recognized** that analysis time is a growing proportion of time-to-solution, and that trend will only be exacerbated without intervention, since increased use of tape will introduce even more latency and bandwidth constraints in the system (and increased use of tape will be mandated by the storage costs – both in terms of energy and finances).

The bandwidth issue is further compromised by difficulties with the POSIX file system in massively parallel environments with large numbers of files or high volume files with internal complexity that is hidden from the file systems (such as occurs with GRIB and netCDF). The complexity of multiple file formats also leads to difficulties for humans too: while software can mediate between formats, it can only do so when the semantic relationships are well understood – in the case of GRIB and netCDF, **they are not**.

ESiWACE will **drive innovation** in the following areas::

- **Currently**, planning for storage growth is mostly carried out separately from planning for operational and scientific activities, and all institutions carry out such planning independently. Storage vendors are not always aware of the requirements. Developing a parametric model for storage within ESiWACE which is inclusive of the scientific workflow is an **ambitious** activity, that has the **potential to significantly improve** the ability to both plan realistic scientific experiments and deliver optimal storage configurations – not least since vendors will be able to consider the workflow of multiple institutions when designing solutions. PRACE, EUDAT and other e-infrastructure providers will be better able to understand the storage requirements and costs of the weather and climate communities.
- **Moving beyond** POSIX file systems within the ESiWACE project **is too ambitious** for the resources and timescale available, but the proposed innovative approach to providing middleware that can accelerate I/O to scientific formats, even in POSIX file systems, will provide an important interim milestone, with significant **potential benefit** to the community – not least because it may end up saving significant amounts of money in comparison to traditional approaches (thus allowing greater volumes of data to be stored and handled efficiently). This **potential** is in part because the use of object stores and other nontraditional storage backend middleware will be **innovative** in our community, although it's success will depend on ambitious expectations for the performance of the middleware.
- **Current** tape systems in use in weather and climate cannot effectively employ parallelism to extract data from the same file, and rely on multiple copies of the same data to effect reliability against corruption. Developing a software library which can be used to provide more parallelism and more efficient storage of extra copies of data utilizing the Redundant Array of Independent Tape (RAIT) concept **will be ambitious**, particularly where several generations of tape drives and tape media share the same library, but it is necessary. While RAIT itself cannot be thought of as an innovation, making such a library available to the open source community will be new, and allow generic e-infrastructure providers deploy services similar to those such as dedicated providers such as ECMWF and the Met Office.
- Finally, in terms of exploitability, the conversion of data to information involves extracting meaning from data. Despite a quarter of a century of co-existence**, it has not been possible, yet** for the community to create a semantic mapping between GRIB and netCDF allowing users to reliably convert data (and information) between the two. Accordingly, developing such a mapping is both **ambitious and innovative**, and should lead to much greater interworking between the two communities – meaning that tools which improve performance for one type of data will be deployable in support of the other type of data. This will significantly support the objectives and infrastructure necessary to further the aims of the World Climate Research Programme.

2 IMPACT

2.1 Expected impacts

2.1.1 Expected impacts set out in the work programme, under the relevant topic

ESiWACE will <u>achieve all four impacts</u> indicated in the text of the call EINFRA-5-2015:

Impact 1) Improved access to computing applications and expertise that enables researchers and industry to be more productive, leading to scientific excellence

Simulations in weather and climate research always have been limited by the computational power available. Scientific progress would have been faster, if computers would have been more efficient, that is: more efficient per computational core, better scalable, easier to use, and better to exploit. Although growth in peak computing power came for free for a very long time, code optimization still was mandatory due to this limitation in efficiency. The exascale era will now bring the coercion not only to parallelize to much higher number of processes than so far, but also for heterogeneous processor architectures, as they are provided by commoditized products available today. Researchers both in ESM will suffer even more from the fact that their tool development needs to address performance and efficiency in a very technical sense, and not focus on ease of use or comprehensibleness. Further commoditisation of computing architectures will decrease usability for ESM. In many cases, suboptimal usability of the modelling environment leads to decreased productivity: Scientists need to concentrate on the tool instead of the scientific experiment they are forced to foster technical instead of the scientific skills. So, despite all technical efforts to improve scalability and exploitability, bad usability degrades time to solution, and, such, scientific excellence.

Scalability traditionally was at the heart of model and tool developers and is considered of major importance. But considering exploitability of data and models as well as their usability offers lower hanging fruits in terms of efficiency. Sharing software development will enable the ESiWACE user community to devote more concentration upon scientific topics, and methods. ESiWACE will also foster services for the use of modern software engineering methods and tools like version control systems, rapid coding environments, performance and debugging tools, and others.

With increased scalability, improved usability, enhanced exploitability, and a better balance between them, researchers will be empowered in keeping and extending their internationally leading role in the field of weather and climate research. The governance employed will improve community bind-in, enable more exchange within the weather and climate community, and ensure logical and traceable selection of development and service topics of interest for the community.

It is safe to assume that with the joint efforts of ECMWF and ENES, a large fraction of the EU community in ESM will be impacted by the innovations from ESIWACE, since all major European institutions are involved with or part of these two organisations. In particular, the developments within ESiWACE will benefit a large weather forecasting community in Europe since ECMWF represents the global medium-range weather forecasting interest of the majority of the European countries. Regional services organised through consortia such as ALADIN (16 countries), HIRLAM (11 countries), LACE (7 countries), and COSMO (7 countries) will immediately benefit from the outcomes of this project, with the first three consortia already directly sharing many computer code components with ECMWF, including the dynamical core. All consortia share a common research expertise with ECMWF in several areas handled by this project. Also, all European modelling groups participating to CMIP programs with their Earth system models are members of ENES.

Impact 2) Improved competitiveness for companies and SMEs through access to ESiWACE expertise and services

In terms of HPC capacity deployed, the EU lost 10 % of its high-end computing capacity from 2008 to 2010 , whereas other nations increased their efforts in this area during the same period. Fewer high-end computing resources available in the EU mean that scientific know-how which critically relies on HPC and influences the development of new HPC systems25, is weakening in Europe. The European Union has many successful scientific and engineering software firms and is strong in many important areas of parallel software development. In fact, the large majority of the principal parallel software applications in use at EU HPC sites has been created and is further developed in Europe. However, the mastering of advanced HPC hardware is closely linked to the associated software and losing out on one side inevitably leads to a loss on the other

By increasing efficiency in using super computers for one of the most challenging fields of applications for computer based simulations, namely the modeling of weather and climate. ESiWACE will contribute to increased competitiveness and growth of European HPC industry. ESiWACE will lay the foundation for new knowledge and skills. Scientific discovery will be boosted by HPC developments and by weather and climate modellers, software developers, computing scientists, vendors and target users working closely together striving towards a culture of excellence. The net gain in terms of scientific and technical innovation and efficiency will then enhance competitiveness for European science and industry and, by granting open access to research results, enable Europe- and world-wide transfer of the newly acquired state-of-the-art knowledge. In particular, ESiWACE will unite the weather and climate communities to provide a single, stronger interface to HPC suppliers and encourage them to engage (WP3). Those suppliers able to engage will not only strengthen their position with this user groups, but with a much wider weather and climate community globally.

Impact 3) European leadership in applications that address societal challenges or are important for industrial applications through better code performance and better code maintenance and availability

ESiWACE unites the European climate and weather communities in their struggle to keep abreast of technological progress on the way to exascale. Only with this merging of forces a level of competitiveness can be reached, which will enable both communities to maintain and even improve their high scientific performance standards in the area of climate and weather research, which are both of considerable scientific and societal relevance. Improved scalability of codes will lead to faster time-to-solution and decreased uncertainty in forecasts and projections. Enhanced maintainability and better usability of frameworks and re-usage of workflow elements will lower learning curves and improve applicability of complex workflows to new scientific problems. Standard environments on pan-European Infrastructures like PRACE will ameliorate their usability to new levels. Faster, less complex storage access will lead to speed-ups both in production and analysis modes, leading to faster production of scientific results, and enable exploitation of results with unprecedented speed.

Impact 4) Larger number of scientists and engineers will be trained in the use of computational methods and optimization of applications

There is only a small workforce available that has the adequate educational background and is well trained in HPC, especially in parallel programming. In addition, scientists that develop and maintain the computational tools and application codes often do not have an attractive career path. This hinders the exploitation of HPC in research and industry. By 2020 the computing power available in today's most performing HPC systems will be available on desktop systems. A well trained workforce capable of efficiently using this computing power is essential. Additionally, at Member States, there is still fragmentation of HPC in climate and weather domain and this leads to inefficient use of resources and only partial exchange of expertise.

In ESiWACE, the ESM scientist and scientific programmers will experience a new level of expertise available to them from the services provided by the ESiWACE workshop series, white papers, portals, improved support and training on community models (e.g. NEMO) and tools (e.g. OASIS, XIOS, Cylc), and sharing of best practices in trainings and mutual visits count to the measures employed.

2.1.2 Improving innovation capacity and the integration of new knowledge

The **user-driven, integrated, multidisciplinary, distributed approach** of ESiWACE will lay the foundation for new knowledge and skills.

ESiWACE will contribute to increased competitiveness and growth of European HPC industry by:

- Increasing efficiency in using super computers for one of the most challenging fields of applications for computer based simulations, namely modeling of weather and climate
- Bringing together weather and climate modellers, software developers, computing scientists, vendors and target users working closely together striving towards a culture of excellence, consequently boosting HPC developments and achieving a net gain in terms of scientific and technical innovation and efficiency;

- Uniting the weather and climate communities to provide a single, stronger interface to HPC suppliers and encourage them to engage, thus strengthening the position, growth and competitiveness of these user groups;
- Granting open access to research results, enabling transfer of the newly acquired state-of-theart knowledge to other actors on the European and global market.

2.1.3 Contribution to environmental and socially important impacts

The targeted improvements in efficiency and productivity will directly address important societal challenges around climate change and extreme weather prediction. Improved use of HPC will allow:

- More reliable weather and climate forecasting across scales.
- More accurate projections of future climate under various scenarios providing better guidance for climate change mitigation policies.
- Better understanding and prediction of weather and climate extremes impacting on economic and political decision making for emergency response and change adaptation.

As indicated in the EC Communication "High-Performance Computing: Europe's place in a Global Race", the race for leadership in high performance computing systems is driven by the need to address societal grand challenges more effectively.

Without HPC there would be no

- Projection of climate evolution;
- Forecasting of weather, which is necessary for planning our daily activities and dealing with severe weather conditions that can devastate lives and properties;
- Prevention and management of large-scale climate changes;
- Rational decision making in this area for policy makers, influencing public and private levels.

But HPC is also vital for the EU industrial capabilities:

- Returns on investment in HPC are extremely high: Companies and countries investing most in HPC lead in science and economics
- Advances in the area of HPC such as new computing technologies, software, or storage applications feed into the broader ICT industry and the consumer mass market, becoming available in households within five years of their introduction in high-end HPC**.**

2.1.4 Barriers/obstacles, and any framework conditions that may determine whether and to what extent the expected impacts will be achieved

Usability and complexity are problems in the software development field in general. Simulation workflows both in ESM are very complex, and are used by a comparatively low number of experts. Rather, they are perceived as single-site problems, resulting in a natural tendency to single-site solutions. These are often not very well usable, for other colleagues, but also across sites or communities. The communities might not be ready to take more concrete steps to turn shared software development into an every-day commodity way of working. The support by the individual centres might be too low since the usual way of working is the do-it-yourself method. They often trust their in-house developments more than quality-controlled software developed externally. Also, sharing of IPR and engaging in open-source projects are often perceived as contrary to earning appropriate merit and credentials.

The ESiWACE community needs to ensure and take all measures available so that it benefits from sharing all aspects of the software development and application processes not only between centres, but also between the communities. Language barriers, skill gaps and missing information links play crucial, obstructing roles. These have to be identified and remedied by the governance processes initiated by ESiWACE.

2.2 Measures to maximise impact

2.2.1 Dissemination and exploitation of results

Dissemination and exploitation activities are of high relevance in ESiWACE.

For maximising the impact of the project results, we have chosen a combination of measures for:

- **Dissemination and Exploitation**
- **Media and Communication.**

These activities will involve **all consortium partners** and their respective staff, including researchers from **climate, weather and HPC community**. These activities will be managed partly in WP1 “Governance, Engagement & long-term Sustainability” and partly in WP5 “Management and Dissemination”. More specifically under WP5 we have planned tasks (Task 5.3 and Task 5.4) dedicated to the implementation of a well-structured plan to support an effective sharing of the results within all relevant target stakeholders.

The work done in WP2-4 will provide the feed to WP1 and WP5. Dissemination of the concrete results of work completed in WP 2-4 will be done at different levels: Through the active engagement with potential users, the implementation of communication activities, the structuring of an exchange between the consortium and other countries. The partners in the project have professional communication and public engagement officers in their organizations, and we will take advantage of their network.

- <u>In the early stage</u> of ESiWACE, the **Dissemination and Exploitation Plan** [D5.5] and a **Media and Communication Plan** [D5.3] will be set up to steer the activities of all partners. The plans will be based on the elements reported in the Table 2.2a here below.
- <u>During the project implementation</u>, the contents of the plans will be updated on the basis on the development of the project, for maximizing the impact of ESiWACE in a consistent manner.
- ESiWACE focuses on the role and synergies between partners’ experiences, competences, capabilities and on spreading knowledge all over Europe.

The activities will include organising workshops, meetings and training activities with external stakeholders, such as industry and end-users; potentially interested groups.
Specific annual dissemination events will be organised by ESiWACE to be held at partner sites (General Assembly, Work packages workshops) and will help us publicize our work in such a way that the consortium will profit from publicity and generate interest in the project and its outcomes; encourage scientists to join our partner institutes, companies and activities, draw the attention of national and regional governments and other public and private funding sources to the needs of long-sustainability of our researches and innovation plans.

	Dissemination and Exploitation Strategy	**Media and Communication Strategy**
Objective	Enable and facilitate implementation of ESiWACE outputs for the climate abs weather community	Promote ESiWACE outputs and interact with the wider climate and weather community as well as HPC centres and industry

Target Audience	ESiWACE partners ESiWACE supporters User Groups European Commission	General public Climate research institutions and National weather services. Wider scientific community HPC industry ETP4HPC European Commission (as a multiplicator)
Instruments	ESiWACE Portal / wiki[D5.1]] ESiWACE workshops ESiWACE public website [D5.2]	ESiWACE website [D5.1] Progress reports to the EC Scientific publications
Access	Project partners General public	General public
Who is in charge?	WP1 / WP5	WP1 / WP5

Table 2.2a: Overview: the Pillars of the ESiWACE strategy for dissemination, exploitation and communication activities.

This comprehensive programme of offline activities, training, and dissemination activities will be complemented by a web presence to ensure the maximum engagement with **public and private sector** representatives of a full range of users and uses both within the scientific community and beyond this, in a wider European and international context. WP5 will establish a central webplatform based on the existing ENES portal and collaboration tools such as a wiki. The web presence will be based on:

- **Public Website** and its interactive facilities for the exchange of information and data. The website will include data (sample input and output), scientific outcome (papers, reports, documentation, and conference proceedings), and an online exchange function to track the communication between partners on global project topics. Exchange forums for discussion between partners will be linked to the respective repositories (D5.2).
- **Project Portal** (online platform): will provide limited access for project partners as well as the European Commission, and direct access to all material generated in the course of the project as well as quick looks at the project status. It will contain a repositories for project documentation (plans, progress tracking, reports, financial information) including **Wiki and project management tool as** the central collaboration platform for the partners and the supporters (D5.1).

Progress reports on the activities and the results will be submitted on regular basis to the European Commission. Summary reports in a language accessible for a lay audience will be made available on the project website.

Scientific and technical results of ESiWACE will be disseminated at European and international level through **scientific articles**, submitted for **peer-reviewed publication**, strategy papers, participation in conferences and workshops. ESiWACE results will be exploited at European and international level by weather and climate modelling groups (research institutions, weather forecast services) relying on HPC resources.

2.2.2 Data/software policy and management of intellectual property rights (IPR)

A key objective of publicly-funded research is that it should lead to the exploitation of results, which goes one step further than the mere production and dissemination of new scientific knowledge. Innovation is understood as any activity aiming to promote not only the dissemination, but crucially the subsequent exploitation of the results of the research and development projects. The strategic use and management of Intellectual Property (IP) in international research initiatives and in business is essential for strengthening the European scientific and technological base, boosting innovation and ensuring growth in the EU. In this context our consortium is aware that Horizon

2020 places much emphasis on systematic Intellectual Property exploitation strategies as a means to better protect innovation initiatives, and to reap commercial and economic benefits from EUfunded research.

The strategy for the knowledge management, protection, dissemination and for the exploitation of result, the will be defined in **Dissemination and Exploitation Plan** [D5.5]. The strategic document will be regularly updated during the entire project. Updates will be submitted to the European Commission as an integral part of the Project Periodic Reports. A final document, a Strategy for intellectual property exploitation [D5.6] will also be made available at the very end of the project. ESiWACE results will be exploited at European and international level by weather and climate modelling groups (research institutions, weather forecast services) relying on HPC resources.

- **Open access to peer-reviewed scientific publications**: Open access will be granted to all scientific publications resulting from ESiWACE with a combination of golden and green open access. We will make use of institutional and topic repositories for making our publications available. ESiWACE scientists, as EC grant recipients, will ensure that electronic copies of peer-reviewed scientific publications become freely available to anyone as soon as possible and in all cases no later than six months after publication. Additionally, wherever results are suitable (content and size) to be distributed or stored by using open access repositories, ESiWACE will use **ZENODO**[22] for disseminating the results of the project to larger audiences and larger networks, in full open access.
- **Open access to software and tools**: Software and tools developed in ESiWACE will be managed through an open-access license and the project will produce test suites under which the software can be operated for selected hardware options. External users will have free and full access to these tools and documentation through the project website. The software policy pursued in ESiWACE is fundamentally open. Data, documentation and training tools will be made freely available, and software produced by ESiWACE will be managed through an Apache-2 open source.

The procedures above will be monitored by the Project Office (see section 3.2.1 for more details).

2.2.3 Media and Communication activities for promoting the project and its findings

We understand communication as much more than a reporting duty: There is a major difference between a communication, which has been strategically planned with the expected societal impacts we would like to make happen, and ad hoc efforts made just for meeting contractual requirements. We are aware of the contractual obligations related to communication efforts requested by Horizon 2020, but our goal is to go beyond these and strive for **high quality outcomes**.

The communication element of the project will involve **all consortium partners** and their respective staff, including researchers. Already in the preparation stage of the proposal, we have set high priority in raising the awareness that communication is a continuous process, not a onetime effort when the project ends. The partners in the project have professional communication and public engagement officers in their organizations, and we will take advantage of their network.

Right from the beginning, the project will be guided by key partners for planning to achieve the desired outcomes, and on the basis of clearly identified objectives. Appropriate resources have been allocated to this task. As indicated in the section 2.2b table, we have already foreseen **a number of tools** for the implementation of our communication strategy.

[22] www.zenodo.org

Target Audience	Communication objective	Communication materials	Communication tools	Who is in charge
Project Partners and Supporters	Ensure an integrated project	Results of WP tasks, risks, benefits, queries, new scientific questions arising	Project wiki / collaboration Platform [D5.1], General Assembly meetings (GA), Work package Meetings, email, web and telephone conferencing (WebCo, Telco)	All Partners
Scientific Advisory Board (SAB)	Ensure the external SAB is fully informed	Progress reports, issues, deliverables, questions, problems arising	project wiki / collaboration platform, General Assembly meetings, WebCo, TelCos with MSB / EMB	Management Steering Board (MSB), Coordinator, Executive Management Board (EMB)
User Groups of community models and tools	Share relevant project progress and knowledge	Progress reports, deliverables, results, strategic documents	project webpage / portal, conferences, workshops, General Assembly meetings	Scientific Officer (SO), MSB, Coordinator, WP leaders (WPL)
Wider NWP and Climate Community	Share relevant project progress and knowledge	Progress reports, deliverables, Communication Plan, strategic documents	portal, conferences, workshops, peer reviewed articles	SO, WPL, Project Partners, Supporters, Project Office
Business Sector *(Hardware vendors, SMEs)*	Ensuring maximum benefit	Progress reports, deliverables, products documentation	user guides and training material, workshops	SO,WP1 and WP5
Other EU – Projects	Sharing the understanding of project and results	Progress reports, deliverables, and results	public project website / portal, project meetings with other EU projects, scientific conferences and presentations [D5.2]	SO, MSB, WP1 & WP5, All Partners
Public and Society	Share relevant project progress and knowledge, Ensure visibility of the project	Relevant results and implications, FAQ	public project webpage, press releases, media contacts, animations, flyers, public lectures, links to other relevant web contents, project progress and significant results, progress reports	SO, PO, WP1 & WP5, All Partners

*Table 2.2b: Overview of ESiWACE main communication and tools (**this list is not exhaustive**)*

2.2.4 The ESiWACE Legacy – Beyond the Project

A particularity of Earth System Modelling is the fact that many activities have much longer timescales than can be supported in a 4-year project. For example the typical lifetime of a weather, ocean or climate model is several decades during which the models evolve around some aging legacy code, corresponding to the order of 1000 man-years of development. Analysis of data from large endeavours such as the simulation for IPCC reports may also take many years. The computing hardware and most of the middleware layer of software change drastically in much shorter time scales. A central goal of ESiWACE is to identify and support components of the ESM workflow which have the potential to be sustainably supported on the longer time scale but be kept synchronized with evolving super computing environment in order to maintain the efficiency of the tools and consequently of the scientists.

The ESiWACE partners recognizes that there is a need to create a sustained funding for an infra structure to continue this work (e.g. Mitchell et al., 2012[26]). There is no direct market today for the products and services we will provide even if downstream users such as climate service centres or private weather services will benefit from it. Indeed, our commitment to open access would make it difficult to monetize products, which is consistent with most of our scientific institutional partners being non-profit organisations. To be successful in the long run, the centre will therefore to a large extent be dependent on its ability to raise national or international funds.

A central task and delivery of ESiWACE will therefore be the development of a business plan to address this need [D1.3]. The matter of long-term sustainability and commitment of support for Earth system models is of central importance for the climate and weather community as a whole and we will foster a strategic planning involving existing governance structures such as the ENES and ECMWF boards. This will be supplemented by community building achieved through ESiWACE workshops, task-forces and special interest groups and exploiting the ENES portal to provide for discussion and cooperation by leading partners.

Specific options to consider within the business plan will be developed and discussed, both in terms of on-going funding for ESiWACE itself, and in terms of the activities and products that might be engendered. Additional possibilities of partnering with the private sector will be addressed. For this we will build on input from our private partners and also from interactions with the ETP4HPC.

Specific actions we envision so far include:

- ESiWACE private partners will investigate the possibility of providing selected components of the ESM workflow as building blocks. The increased maintainability of the ESM applications and increased efficiency of the associated workflow through commercial products (e.g. adding to established tools such as debuggers) could be beneficial not only for ESM users, but at the same time applicable to other markets and thus be rendered profitable.
- A very concrete goal we have set in this respect is the long-term continuation of the series of annual high quality workshops on HPC. The partners will commit to organize these and funding through private partners will be secured.
- Commitment from the partners of ESiWACE and potentially some of the supporters to longterm support of individual software components from institutional funding, This is in their own interest, fulfilling their own needs, but also benefitting from other partners, through distributed and shared efforts (thus avoiding redundant re-development). ESiWACE, by strengthening the sharing of software development, is an important step in that direction.
- Investigation as to whether collaborative funding could be sought from existing and future downstream services and projects.
- Address “classical” funding schemes through the Joint Programming Initiative on Climate, national funding agencies or through subsequent calls of the European Commission.

[26] Mitchell J., Budich R., Joussaume S., Lawrence B. and Marotzke J. (2012), “Infrastructure strategy for the European Earth System Modelling community 2012-2022”, ENES Report Series 1, 33 pp.
URL https://verc.enes.org/community/about-enes/the-future-of-enes/ENES foresight.pdf

 The current ESiWACE proposal has some built-in features to guarantee that the work will be of benefit for the community whichever funding futures transpire in. Although we ask for funding for a three years period, we plan for project duration of 48 months. An initial phase of 4 months will be used to set up the governance structure of the project and establish a mood of cooperation. A closing phase of 9 months (i.e. project month 41 to 48) will see no more directly funded work in the technical work packages (WP2, WP3, WP4), but governance and management will be continued by the coordinating entities (ENES through CNRS-IPSL, DKRZ and ECMWF). This phase will be used to transfer the support and service into a next phase, be it directly funded as a continuation of ESiWACE or through continued commitments of the partners.

ESiWACE will close with a number of outputs. All project results will be made available in open access. We will present all its results and products in a coherent form at the large-scale by big events such as EGU and similar.

Dissemination measures *in the closing phase* of the project:

The final report of the project will include a plan for the use and dissemination of results, to demonstrate the added value and positive impact of the project on the European Union. A final publishable summary of the results will be made available to the Commission for dissemination in the public domain. This will include information on expected results, and their wider societal implications. The text will be drafted in a way to be understandable for a lay audience.

Dissemination measures *after the closure* of the project:

After the official end of the project, the results of the project will stay available on the project website hosted by DKRZ. The website archives all documentation related to the project, including publications, and will be accessible for 5 years after the end of the project

2.2.5 Achievement of Expected Impacts through Dissemination, Exploitation and Communication

The innovation of products, services and business is deeply linked to Europe's future economic growth. Target, audience and messages of the communication strategy will be clarified before deciding on the most appropriate media. With the appropriate choice of the tools mentioned above, ESiWACE aims at demonstrating how its research contributes to a **European Innovation Union** and how the public spending is accounted for, by providing a tangible proof that collaborative research and innovation actions add value by 1) showing how the European collaboration achieves more than would have been otherwise possible, contributing to competitiveness and solving societal issues; 2) showing how the outcome are relevant for our lives by introducing novel technologies and creating new jobs; 3) making better use of the results by making sure that they are taken up by the business sector and the scientific community to ensure follow up.

3 IMPLEMENTATION

3.1 Work plan — Work packages, deliverables and milestones

3.1.1 Brief presentation of the overall structure of the work plan

The work plan of ESiWACE is organized in five work packages (WP), three of which encompassing the bulk of the foreseen technical and scientific work (WP2, WP3, WP4) and two dealing with governance of ESiWACE products and services (WP1) and coordination of the project itself (WP5) respectively (Fig 3.1a).

Table 3.1a List of work packages, Leaders and Co-Leaders

Nr.	Work Package Title	Lead Institution short name	Co-Lead Institution, short name	Estimated person-months

WP1	Governance, Engagement & longterm sustainability	CNRS-IPSL, Sylvie Joussaume	DKRZ, Joachim Biercamp	47
WP2	Scalability	ECMWF, Peter Bauer	CERFACS, Sophie Valcke	202
WP3	Usability	MPG Reinhard Budich	BSC and MET OFFICE Kim Serradell & Mick Carter	123
WP4	Exploitability	STFC Bryan Lawrence	DKRZ, Thomas Ludwig	106
WP5	Management & Dissemination	DKRZ, Joachim Biercamp	ECMWF Peter Bauer	36
Total				514

WP1 "Governance, Engagement and long-term sustainability" addresses governance and strategic questions for ESiWACE, establishing the integration of the project into the existing scientific and HPC infrastructure and ensuring ESiWACE serves the community Europe-wide in a long-sighted, sustainable way.

WP 5 "Management and Dissemination" comprises coordination and management tasks, to monitor the work progress of the consortium, to identify factors of risk for ESiWACE, to establish efficient internal communication (mutual exchange with users and scientific advisors) as well as interaction with the European Commission, and to ensure punctual dissemination of the results and communication.

The activities of WP1 and WP5 are classified as **Networking Activities**.

The technical and scientific tasks have been mapped to the three overarching themes of ESiWACE based on the peculiar requirements of ESM simulations and on the various challenges this community faces, especially in view of the changing HPC landscape with new high-end multicore parallelized architectures and increasing data volumes. These scientific-technical WPs address short and long-term aspects within the following type of activities:

- Networking Activities (NA)
- Service Activities (SA)
- Joint Research Activities (JRA)

WP2 on **"Scalability"** demonstrates how to build and productively operate global cloud-resolving and eddy-resolving models, thanks to more efficient model codes and tools (model coupler, I/O libraries). User training and support as well as provision of improved codes are essential to ensure that the modellers take full benefit of the planned software optimization and development.

WP3 on **"Usability"** aims to considerably improve the ease-of-use of available tools, computing and data handling infrastructures. The workload focuses on structuring and supporting the end-toend workflow both in research and in production modes and on contributing in software development and in definition and spreading of best practices. WP3 will organize workshops and foster user exchange on the planned activities

WP4 on **"Exploitability"** tackles the major roadblocks that hinder efficient use of the considerable amounts of data produced by weather and climate simulations. WP4 plans to co-design with industrial partners' data access interfaces and storage layouts and to develop new methods of tape exploitation. WP4 aims furthermore at increasing the compatibility of the two most used data formats in climate and weather modelling.

The Leaders and Co-Leaders of the work packages are listed in table 3.1. **Project duration**

The overall project duration is of 48 months. During this period, we have foreseen three stages of the project (see Table 3.1.b):

- **Phase-in** from project month 1 to 4. This is the launching phase of all activities in the work packages, and for the setting up of the governance bodies of the project;
- **Core** starting in project month 5 and ending at the end of project month 40, in this phase all work packages are up and running and the groups are fully operational;
- **Phase-out** from project month 41 to end of project month 48: in this phase only the WP1 and WP 5 will be actively running, for finalising the management, dissemination, and exploitation activities and pursuing the sustainability of ESiWACE.

A more exhaustive, thorough description of the WPs and of the associated tasks and their required classification are to be found in the descriptions in the following pages

*Fig. 3.1a List of work packages and tasks (left) and graph showing the relation between work packages (right): The graph illustrates the spirit of creating a **centre** serving the community. The interaction of the three scientific/technical work packages which focus on the three ESiWACE themes is supported on the administrative level by WP5 and steered by WP1 to guarantee that the work is informed by and responds to community requirements.*

WP1 Governance and engagement

- T1.1 Engagement and governance
- T1.2 Enhancing community capacity in HPC
- T1.3 Strategic interaction with HPC ecosystem and HPC industry
- T1.4 Sustainability and business planning

WP2 Scalability

- T2.1 Support, training and integration of state of-the-art community models and tools
- T2.2 Global high-resolution model demonstrators
- T2.3 Efficiency enhancement of models and tools
- T2.4 Preparing for exascale

WP3 Usability

- T3.1 ESM end-to-end workflows Recommendations
- T3.2 ESM system software stack recommendations
- T3.3 ESM scheduling
- T3.4 Co-Design for Usability

WP4 Exploitability

- T4.1 The business of storing and exploiting high volume climate data
- T4.2 New storage layout for Earth system data
- T4.3 New methods of exploiting tape
- T4.4 Semantic mapping between netCDF and GRIB

WP5 Management and Dissemination

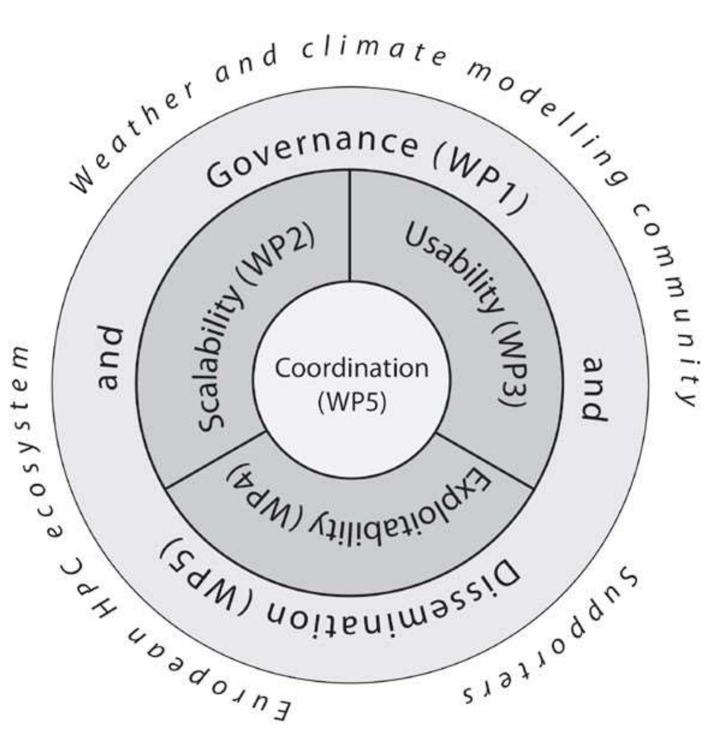

Table 3.1b Gantt chart: Timing of the different Work Packages. Deliverables, Milestones and Reports

	Phasing in				OPERATIONAL PHASE																																										Phasing out (*)	
	2015				2016												2017												2018												2019							
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48
	Sept	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	Ma	Jun	Jul	Aug
WP1						D1.1	D1.4																								D1.5							D1.3								D1.2		
WP2																								D2.3 D2.8 D2.9												D2.2 D2.4 D2.5 D2.6 D2.7						D2.10 D2.11 D2.12				D2.1		
WP3				D3.1		D3.3								D3.4 D3.5				D3.7 D3.9												D3.2 D3.8				D3.10		D3.6												
WP4																		D4.1 D4.2																								D4.4				D4.3		
WP5		D5.1 D5.2 D5.3				D5.4												D5.5																														D5.6
Milestones						MS1			MS2			MS3						MS4						MS5						MS6			MS7			MS8 MS9												
GA Meetings / Annual Plenum																																																
Reporting periods	Reporting period 1																		Reporting period 2																	Reporting period 3												
WP Progress reports																			PR																	PR											PR	
Financial reports																				FR																	FR											FR

(*) only for MGT & DISS activities

M = Milestone D= Deliverable WP = Work Package PR = Progress Report FR = Financial Report

Details of the Work Packages can be found in Part A of the Description of Action

List of Deliverables can be found in Part A of the Description of Action

List of Milestones can be found in Part A of the Description of Action

3.2 Management structure and procedures

Figure 3.2: Governance structure of ESiWACE

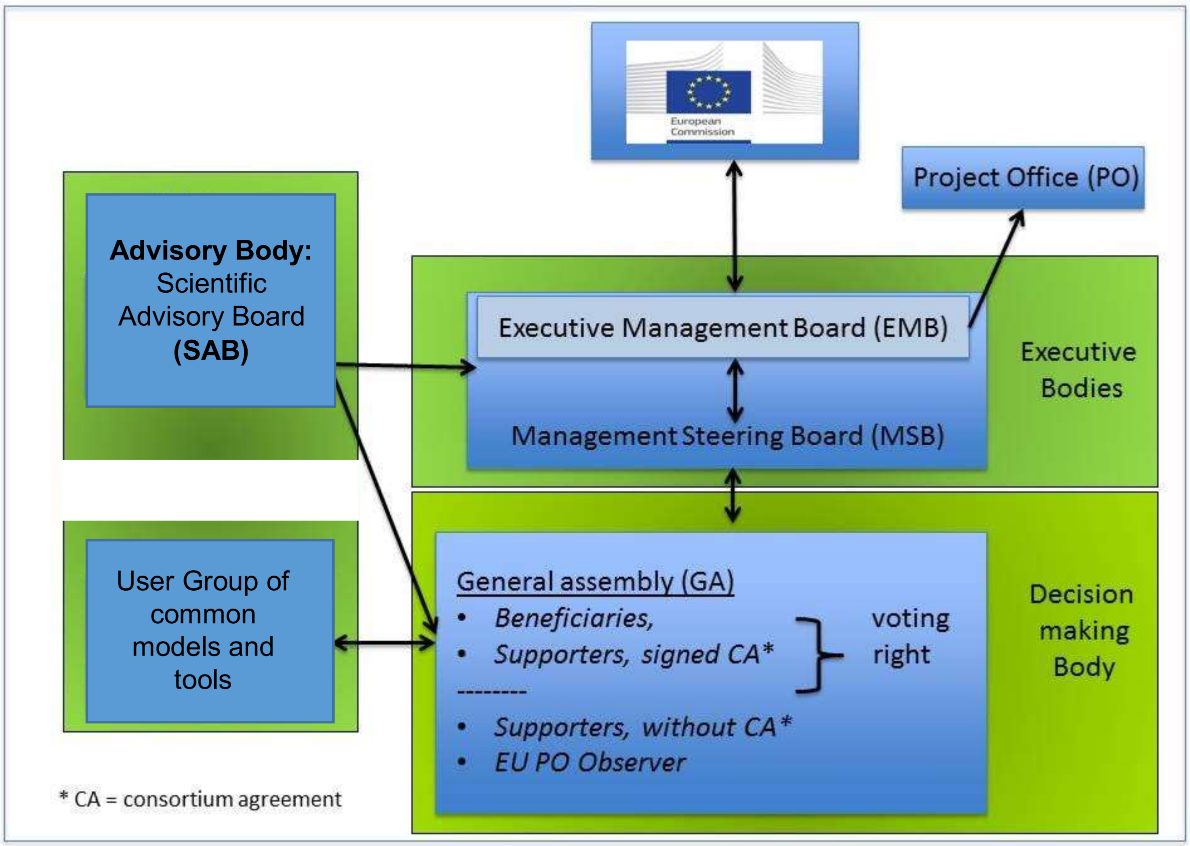

3.2.1 Description of the organisational structure and the decision-making

The main target of the governance structure described here is to enable fast flow of information between all parties involved in the ESiWACE. To minimize overhead and friction loss we keep the structure as simple as possible.

The consortium consists of 16 beneficiaries from science and industry from 7 countries.

Additionally, the consortium is supported by a number of institutions, called *supporters* (see Section 3.3.2), who have committed to support the project without being beneficiaries.

The picture below describes the organizational structure in the project.

In the governance structure we have foreseen roles for

- **Decision-making** bodies i.e. the General Assembly
- **Executive bodies** i.e. the Coordinator, the Management Steering Board and the Executive Management Board
- **Advisory** bodies i.e. the Scientific Advisory Board

In support to the above and to the beneficiaries, we will establish a **Project Office (PO).**

Decision-making body

The ultimate decision-making body of the consortium is the **General Assembly (GA)**. The GA is responsible for taking key decisions for the project as a whole, on the basis of issues raised by the Executive Bodies. It consists of one representative of each consortium partner (beneficiaries and supporters who signed the Consortium Agreement) and it is chaired by the coordinator. All consortium members have the right to vote. Decisions of the GA have to be made at least by simply majority. Supporting institutions (without having signed the Consortium Agreement) have the right to attend the GA and raise topics.

Formal meetings of the GA will be held during the annual project meetings.

Executive bodies (EB)

The Executive Bodies consist of the **Coordinator**, the **Management Steering Board (MBS)** and the **Executive Management Board (EMB)**.

- The **Coordinator** of the project is Joachim Biercamp (DKRZ) and Co–Coordinator is Peter Bauer, (ECMWF). They are responsible for the overall scientific coordination of the project and function as liaise with the European Commission on behalf of the consortium.
- The Management **Steering Board** (MSB) is responsible for the execution of the project. It proposes decisions to the General Assembly (GA) and is monitors their execution.

The MSB consists of

- Coordinator and Co-Coordinator
- Work package leader (WPL) and work package Co–Leaders
- Scientific Officer of the Project Office

MSB assists the coordinator in monitoring the project progress, proposing corrective and preventive actions, promoting the project and its findings, engaging with the public and media about the project, facilitating the communication within the consortium, as well as coordinating the reporting for the European Commission. Furthermore, MSB will implement, promote and monitor gender and diversity equality throughout the lifetime of the project. Awareness of best practice in gender and diversity equality can be shared and may instigate longer-term changes, taking into account regional and cultural differences.
The MSB will meet regularly every two months in a telephone/web conference or on person to secure the smooth information flow. Decisions in the MSB will have to be made at least by simply majority.

- The **Executive** Management **Board (EMB)** role is to execute the decisions taken by the GA, monitor the project progress and to propose corrective and proactive actions for the scientific coordination of the project. The EMB can react promptly when immediate questions have to be answered or intervention is required. The EMG consists of:

- Coordinator and the co-coordinator
- Up to 2 members of the **Management Steering Board** elected
- Scientific Officer of the Project Office

The EMB meets weekly in a telephone / web conference or in person.

Advisory body

The **Scientific Advisory Board (SAB)** is established for ensuring scientific evaluation of the project and links to other programs. Furthermore, it gives advice on the project's scientific approach and orientation by liaison with the EB. It is comprised of selected key international experts with a scientific high profile. For the SAB we have already contacted international experts (see Section 3.3.2), who have their interest in becoming SAB members.

Project office

The **Project Office (PO)** is in charge for the day-to-day management of the project, in a supporting role to the Coordinator, Co-Coordinator and the other governance bodies. It consists of:

- **Coordinator and Co-Coordinator**
- **Scientific officer (**SO**)**, she/he will play an important role in WP1 on the governance issues and. She/he will help CNRS-IPSL in liaising with the communities, in enhancing community capacity in HPC and liaising with PRACE. She/he will help DKRZ on the governance for new developments, on liaising with industrial partners and in the elaboration of the roadmap on HPC.
- **Administrative assistant (AA)** based at DKRZ, who is in charge of monitoring the financial and administrative implementation of the project. Moreover, the AA is supporting the partners in the implementation of the financial rules of Horizon 2020 and the requirements deriving from the Grant and the Consortium Agreement. The AA will be the contact person for the partners in preparing their interim financial payment requests, in dealing with the formalities of the financial reports and of the certificate of financial statements.
- **Web administrator** in charge of the project intranet, the website and web-based dissemination activities.

3.2.2 How the organisational structure and decision-making mechanisms are appropriate to the project

One main target of ESiWACE is to foster the inclusion and cooperation of within the weather and climate research community. A consequence is that networking activities play a major role in the project.

Special features of the project structure are the following:

- The rather limited size of core partners compared to the larger number of supporting partners contributing to selected tasks in single work packages;
- The large influence of the user community on thematic focus points of ESiWACE.

This requires special procedures of decision making and flexibility in the management / implementation of the activities at different levels (work package and tasks):

- The role of the **User Group of common models and tools** is central to ESiWACE is central: we consider it crucial for bringing in new ideas, requests for new research and tasks forces to be installed. The governance work package (WP1) will evaluate needs of the user community. Depending on the scientific relevance, the issues will be passed to WPs with the request to comment. This kind of interaction / feedback processes will ensure that the project responds to upcoming new themes and requirements in the most efficient way. In this way the project will be responsive to any internal and external opportunities that may arise.
- **Task Forces** can be installed by the Executive Bodies either when cross cutting issues will be identified (e.g. HPC task force) or when new and unforeseen themes arise during project run time, that require collaboration of project members and / or supporting partners.
- The **General Assembly** ensures that all partners and supporters involved in the project can have a voice in the decision-making of ESiWACE.
- The **Executive Bodies** collect inputs on specific topics from the advisory boards and user group committee, preparing the ground for the decisions of the General Assembly and being responsible for the execution of these decisions within the consortium.

3.2.3 Innovation management

The Consortium Agreement (CA) will be produced and signed before the signature of the Grant Agreement with the European Commission is an internal agreement establishing their rights and obligations of the consortium members with respect to the implementation and organisation of the action, in accordance with the Grant Agreement. Items regulated in the EC Grant Agreement will not be repeated in the CA. The CA will formalize project management procedures and regulate joint ownership, use of background and Intellectual Property Rights (IPR). Based on the agreement outlined in the project's CA, a structured strategy for the protection of intellectual property arising from the project will be updated and implemented with consensus of all parties (i.e. Dissemination and Exploitation Plan [D5.5]).

Effective innovation management within this project will require an overview of the project in its entirety and for this reason the Coordinator and the Scientific Officer will be in charge of innovation management. In practice, the Coordinator and the Scientific Officer 1) will ensure the development of a strategy and concrete actions of the consortium for the protection, exploitation and dissemination of the results of the project -including software licensing issues if applicable-; 2) will address and combine the technical, scientific and application aspects of innovation and benefit from the expertise regarding all these aspects represented in the consortium; 3) produce an overview of the WPs and the outcomes of ESiWACE to provide effective management and therefore exploitation of these both during and after the project; 4) track and propose commercialisation and exploitation strategies for the whole consortium; 5) give advice to the parties concerned about the ownership, access rights, legal and commercial implications, patents, publications, copyrights, etc.

In an early stage of the project the innovation management will be integrated into the Dissemination and Exploitation Plan [D5.5] and so becomes formal constituent of the project plan. **At a more mature stage** of the project, this plan will be adapted to take into account best practice methods of maximizing the value of intellectual property, for dealing with technology transfer/exploitation/protection, and with the assessment of IP and research results. Finally, **towards the end of the project**, the focus will be on defining a Strategy [D5.6] for the intellectual property exploitation after the project.

3.2.4 Risks and risk management strategy

Critical risks, which have the potential to impact the project objectives, are identified and collected in Part A, "WT5 critical implementation risks and mitigation actions". At the present stage, all these risks have been identified and analysed by the beneficiaries. The Scientific Officer (SO) will actively monitor the work progress in order to raise awareness on the occurrence of such risks throughout the duration of the project. As explained in Section 3.2.1, the Executive Management Board (EMB) is in charge of managing the risks on a day-to-day basis, and the GA will be responsible for deciding on proposed preventive and corrective actions. At each meeting of the EMB, the list of risks will be analysed and updated.

Critical risks for implementation can now be found in Part A of the Description of Action

3.2.5 Quality Management

The quality control management involves the product description and quality expectations of key deliverables, and an internal review and acceptance procedure. The procedures for the quality control process of deliverables will take into account Technology Readiness Levels (TRL). When a deliverable is ready for review, it will be forwarded to the Project Office (PO), which verifies its general compliance. Moreover internal quality control will be conducted by Management Steering Board (MSB) and PO to secure the highest possible scientific standard of the deliverable. If necessary, they can request a revision of the deliverable before approval.

On a more general level, the MSB will interact with the named representatives of the supporters to assure that the project results and services are well disseminated and to evaluate the perceived benefit that ESiWACE generates for the community.

3.2.6 Financial and Administrative Management

The Project Office is in charge of the financial and administrative management of the project. Regular reporting to the EC in the form of activity and financial progress reports is planned in the Gantt Chart (Table 3.1b). To ensure the transparency of the project management, all reports will be made available on the ESiWACE website. Only documents or part of documents concerning individuals will be kept confidential. The payments will be distributed to the beneficiaries according to the provisions of the Grant Agreement and the consortium agreement, and decisions taken by the General Assembly, if the case. Budget forecasts will be established as an input for MSB meetings.

3.3 Consortium as a whole

3.3.1 Consortium: the Beneficiaries

The ESiWACE consortium is built from two established European networks: the European Network for Earth System modelling (ENES), gathering the European climate modelling community contributing to the World Climate Research Program international coordinated experiments and Intergovernmental Panel on Climate Change assessments, and The European Centre for MediumRange Weather Forecasts (ECMWF) which is an independent intergovernmental organisation supported by 34 states. The ESiWACE partners are leading public and private institutions having long standing expertise in delivering products and services for HPC based climate research and numerical weather prediction. In particular we want to emphasize the scientific and technical expertise related to weather and climate research software as well as high performance computing. During the proposal preparation, partners showed high communication, collaboration, adaptation and team-working skills, which pave the way for a successful project. By this, we are prepared to translate the outcomes of ESiWACE into substantial societal benefit.

The consortium will be led by two topical super-computing centres each operating one of the largest computing facilities in Europe. The coordinator, the German climate computing centre (DKRZ) since nearly three decades is a provider for high performance computing and data management facilities and services tailored to the need of climate modelling. The coordinator will in its management and dissemination efforts be supported by ECMWF running an operational service, producing and disseminating numerical weather predictions.

The expertise of the scientific partners optimally matches our objective to develop solutions for cross-cutting HPC issues particular to the weather and climate domain. The consortium comprises world leading climate research institutions, operational weather services, super-computer centres including two PRACE members and also experts in computer science including university members. Table 3.3a provides an overview of the expertise and roles relevant for the objectives of ESiWACE.

The expertise of the three partners from industry matches the three main tools that our communities share with others but which need co-design and information exchange between weather and climate science and industry to efficiently address our needs on the long run: the super computer themselves, the storage hardware, the software environment. Bull SA, a leading European Computer vendor, has several climate and weather centres among its customers and accordingly is well acquainted with their applications. Seagate Technology is a world leader in storage solutions and through the acquisition of Xyratex Ltd, a leading provider of data storage technology has large customers from the climate and weather domain. ALLINEA, a UK-based SME, provides integrated software development tools, designed to handle the scale of today's systems. All industry partners (as well as ECMFW and BSC) are members of ETP4HPC.

Table 3.3a: Partners and principal expertise

	Climate Research	Operational numerical weather forecast	Computer Science	Development of models and tools	Operation of HPC-facilities
DKRZ	X		X	X	X
ECMWF		X	X	X	X
CNRS-IPSL	X			X	
MPG	X			X	
CERFACS			X	X	X
BSC			X	X	X
STFC			X	X	X
MetO	X	X		X	X
UREAD	X			X	
SMHI		X		X	
ICHEC			X	X	X
CMCC	X		X	X	X
DWD		X		X	X

3.3.2 The Supporters of ESiWACE

Major themes for ESiWACE have already been identified, but we expect representatives of the European weather and climate science community to carry out detailed priority setting, as well as operational governance which will continually adapt and improve our services. To foster and organize this process we started to form a group of supporters and were overwhelmed by the large number of reactions: up to now, we have received nearly 40 Letters of Commitment and Support (Table 3.3b) from consortia, institutions and individuals which qualify ESiWACE goals and the anticipated outcomes as very important for their own work and objectives.

The supporters of ESiWACE are legal national and international entities which are not beneficiaries but wish to support ESiWACE through complementary activities. Most of them commit explicitly to specific supporting activities such as

- implementation and evaluation of the ESiWACE software stack,
- active participation in ESiWACE workshops or meetings,
- contribution to software intercomparison or evaluation,
- to name contact persons to interact with ESiWACE on scientific and technical level,
- support of the dissemination of ESiWACE results and outcomes
- fostering cooperation between ESiWACE and projects or activities of supporter

All individual committed contributions can be found in the Letters of Commitment (Appendix 1).

The presence of the supporters will significantly strengthen ESiWACE and underline the importance attached to ESiWACE by the climate and weather community as well as HPC centers.
Moreover, the following leading scientists declared their availability to become members of the scientific advisory board (SAB):

- Isabelle Bey, Center for Climate System Modelling (C2SM), Switzerland
- Alison Kennedy, Edinburgh Parallel Computing Centre (EPCC), UK
- Rudolf Fischer, NEC Deutschland GmbH, Germany
- Peter Fox, Rensselaer Polytechnic Institute (RPI), USA
- Wilco Hazeleger, Netherlands eScience Center (NLeSC), The Netherlands
- Heikki Järvinen, University of Helsinki, Finland
- Thomas Schulthess, Swiss National Supercomputer Centre (CSC), Switzerland

The supporters will be listed on the project website with indication of their complementary activities. The application of new supporters is constantly evaluated by the Management Steering Board.

Table 3.3b: List of Letters of Commitment and Support (in alphabetical order of affiliation or consortium). For the various individual and institutional commitments, see Appendix 1

Signatories	Affiliated Institution / Federation / Network	Named contact person
W. Hiller Director of Computer Center	AWI, Alfred Wegener Institute, Helmholtz Centre for Polar and Marine Research, Germany	Bernadette Fritzsch
P. May; Assistant Director	Bureau of Meteorology, Australia	Kamal Puri
P. Pišoft, Senior Scientist	Charles University, Czech Republic	Thomas Halenka Michal Belda
I. Bey, C2SM Executive Director	C2SM, Center for Climate System Modelling, Switzerland	Isabell Bey (for SAB)
A.Will, Senior Scientist Member of Scientific Management Board	CLM-Community, Climate Limitedarea Modelling-Community, Germany	Klaus Keuler Andreas Will
V Eyring Chair of CMIP Panel	Coupled Model Intercomparison Project (CMIP) Panel	-
D. Ulmer, Director, EMEA Operations	CRAY U.K. Limited, UK	Philip Brown
T. C. Schulthess Director	CSC, Swiss National Supercomputing Centre, Switzerland	Will Sawyer
T. Christoudias, Research Scientist	The Cyprus Institute, Cyprus	NN
T. Beckers, Director HPC Sales EMEA	DataDirect Networks, Germany	-
R. Sausen, Head of ESM Department	DLR, Institut für Physik der Atmosphäre, Germany	NN
K. Krogh Andersen, Director of Research and Delvelopment	DMI, Danish Meteorological Institute, Denmark	Bent Hansen Sass
R. Döscher, Chair of Steering Committee	EC-Earth Consortium, Europe	Ralf Döscher
A.Bode, Director of LRZ	EnCompAS (proposal to H2020) Consortium	
A. Kennedy, Executive Director	EPCC, Edinburgh Parallel Computing Centre, UK	Alison Kennedy (for SAB)
E. Robinson, Executive Director P. Fox, President	ESIP, Foundation for Earth Science, USA	Erin Robinson
I. Pisso, Scientific Secretary	eSTICC, eScience tools for investigating Climate Change in Northern High Latitudes, Norway	-
P. Taalas, Director General	FMI, Finnish Meterological Institute, Finland	Hanelle Korhonen

U. Cubasch, Senior Scientist, Professor of Meteorology	Freie Universität Berlin, Germany	Ingo Kirchner
V. Balaji, Head, Modelling Systems Group	GFDL, Geophysical Fluid Dynamics Laboratory Princeton University, USA	V. Balaji
D. Turek, Vice President Technical Computing	IBM, USA	-
S. Negre, President	Intel Corporation SAS, France	-
M. Tsukakoshi, Director, Information Systems Department	JAMSTEC, Japan Agency for Marine-Earth Science and Technology, Japan	-
G. van der Steenhoven, Director General	KNMI, Royal Netherlands Meteorological Institute, The Netherlands	Jan Barkmeijer
Patrick Jöckel Speaker of the consortium	MESSy, Modular Earth Submodel System Consortium, Germany	NN
P. Bougeault, Director General	Météo-France, France	NN
R. Fischer, Senior Manager HPC	NEC Deutschland GmbH	Rudolf Fischer (for SAB)
W. Hazeleger, Director	NLeSC, Netherlands eScience Center	Wilco Hazeleger (for SAB)
M. Uddstrom, Principal Scientist	NIWA, The NZ National Institute of Water and Atmosphere, New Zealand	-
H. L. Tolman Dirctor Environmental Modeling Center	NOAA, US Dept of Commerce, National Oceanic and Atmosph. Administration, National Weather Service, National Center for Environmental Prediction ,USA	John Michalakes Hendrik Tolman
Ø. Hov, Director of Research	Norwegian Meteorological Institute, Norway	-
S. Kraemer, Director Business Development HPC-EMEA	NVIDIA, Germany	-
P. Fox. Director	RPI, Rensselaer Polytechnic Institute, USA	Peter Fox (also for SAB)
C. Heinze, Leader of project EVA	University of Bergen, Norway	-
H. Järvinen Senior Scientist, Professor of Meteorology	University of Helsinki, Finland	Heikki Järvinen (for SAB)
T. Fichefet, Senior Scientist, Professor of Physics	UCL-TECLIM, Universite Catholique de Louvain, Belgium	Pierre-Yves Barriat
T. Palmer, Director	University of Oxford, UK	Tim Palmer
D. Carlson, Director WCRP P.M. Ruti, Chief WWRC	World Meteorological Organisation (WMO): WCRP & WWRP	-

3.4 Resources to be committed

The total amount of EC contribution requested by ESiWACE amounts to 4,951,050 Euro distributed over 4 years. The total **direct costs** of the project amount to 3,960,839 Euro, the amount of the **indirect costs** amounts to 990,211 Euro.

The table below provides an overview of the planned budget and requested EC contribution per budget line *(figures below are indicative)*:

Budget lines	Planned budget (€)	EC contribution requested (€)	In % on the total direct+indirect costs
Personnel	3,420,439	3,420,439	69%
Other direct costs	540,400	540,400	11%
Subcontracting	-	-	0%
Total direct costs	3,960,839	3,960,839	80%
Total indirect costs	990,211	990,211	20%
Total direct + indirect costs	4,951,050	4,951,050	100%

Looking at the distribution of the direct costs only, and excluding the indirect costs, which are just a lump-sum attributed by the EC, the total direct costs budget is mostly devoted to staff (86%), corresponding to 514 person-months over 4 years, and the rest is assigned to “other direct costs” such as travels, organisation of workshops, summer schools, network meetings, and dissemination and communication activities (13.5 %) and audit costs (0.5%). We so not plan to subcontract tasks. The direct cost distribution across the work packages can be seen in the chart here below.

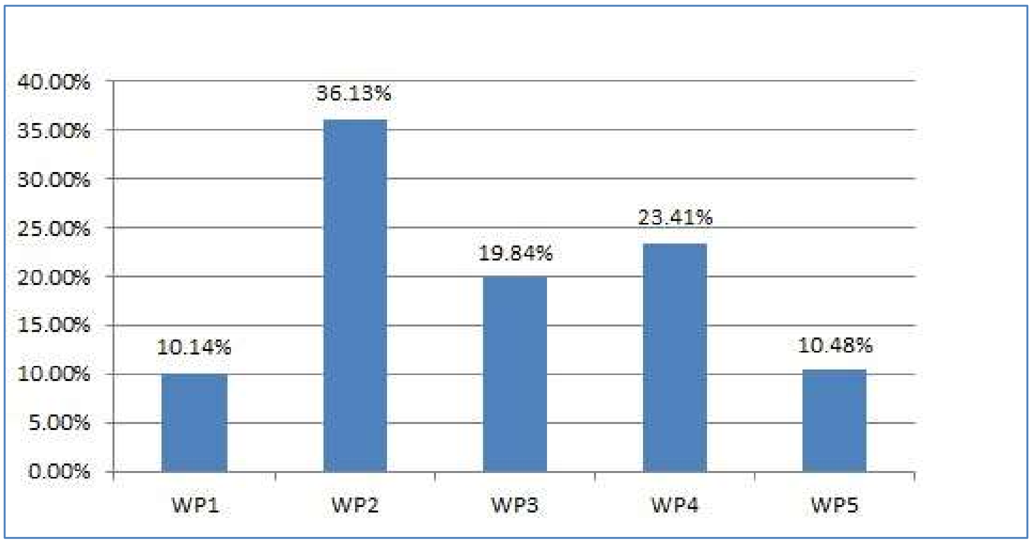

Chart 3.4 a: Distribution across the work packages of the direct costs. Figures are indicative.

The sum of the planned budget for the implementation of activities linked to the governance, engagement, planning of the long-term sustainability (WP1) and dissemination, communication and management tasks (WP5) correspond to 20% of the entire budget of the project.

Across the work packages, the distribution of the direct costs according to “personnel” and “other direct costs” highlights again how the WP5 in particular will take care of funding all the activities related to the funding of the work of the committees and panels mentioned in the governance, the dissemination activities, and will actually sustain financially the engagement and governance activities under WP1.

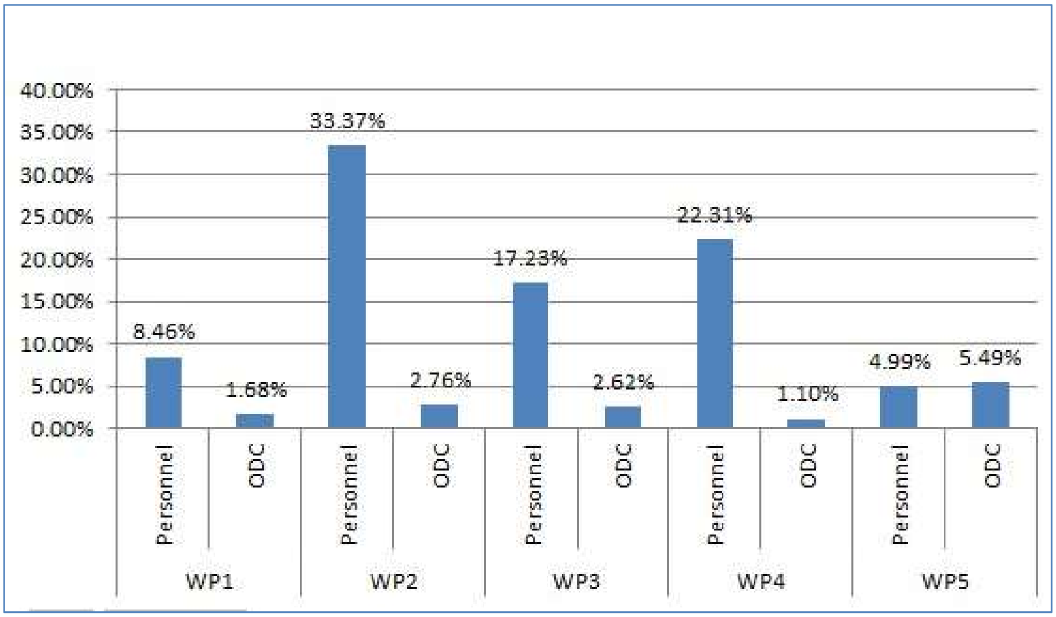

Chart 3.4 b: Distribution of the direct costs according to "personnel" and "other direct costs" (ODR) in the work packages. Figures are indicative.

Summary of staff effort can now be found in Part A of the Description of Action.

Details on the other direct costs

For participants where these planned budget exceed 15% of the personnel costs (according to the budget table reported in the administrative forms), we provide here a breakdown of the other direct costs and a justification for them.

Table 3.4b 'Other direct cost' item Figures below are indicative

	Cost (€)	Justification
1/ DKRZ		
Travel	37600	Travel budget for the project office, including the coordinator and the scientific officer, and for the DRKZ scientists for joining the annual meetings and project workshops, and for the coordination of the activities planned across all the other work packages.
Equipment	0	
Other goods and services	188000	• Organisation of the general assembly and project meetings, incl. invitation of the members of committees / panels / reviewers / external experts. • Organisation of meetings for tasks forces / special interest groups • Overall project publications / maintenance of web portal • Audit costs for the certificate of the financial statements (CFS) for approx.8000 euro
Total	225600	

	Cost (€)	Justification
4/ MPG		
Travel	10100	Travel to the workshops organised in the framework of WP3, travel for joining additional project meetings
Equipment	0	
Other goods and services	48000	Organization of workshops and networking meetings and the summer schools in the framework of WP3, including the invitation of external speakers and experts to project meetings
Total	58100	

	Cost (€)	Justification
5/ CERFACS		
Travel	12000	Travel to the workshops organised in the framework of WP1, travel for joining additional project meetings
Equipment	0	
Other goods and services	15000	Organization of workshops in the framework of WP1
Total		

	Cost (€)	Justification
9/ UREAD		

Travel	6000	Travel budget for joining workshops and general assembly meetings
Equipment	0	
Other goods and services	0	
Total	6000	

	Cost (€)	Justification
11/ ICHEC		
Travel	6000	Travel budget for joining workshops and general assembly meetings
Equipment	0	
Other goods and services	0	
Total	6000	

	Cost (€)	Justification
13/ DWD		
Travel	12000	Travel to the workshops organised in the framework of WP2, travel for joining additional project meetings
Equipment	0	
Other goods and services	4500	Organization of three workshops in the framework of WP2
Total	16500	

	Cost (€)	Justification
16/ ALLINEA		
Travel	6000	Travel budget for joining workshops and general assembly meetings
Equipment	0	
Other goods and services	0	
Total	6000	

Glossary

AA	Administrative Assistant
ALADIN	Aire Limitée, Adaptation Dynamique, Development International
API	Application Programming Interface
BMBF	Bundesministerium für Bildung und Forschung
CAMS	Copernicus Atmosphere Monitoring Service
CASTOR	CERN Advanced STORage manager
CCCS	Copernicus services and climate change
CDI-pio	Climate Data Interface – parallel i/o
CERN	http://home.web.cern.ch/
CHANCE	Co-design of High performance Algorithms and Numerics for oCean models at Exascale
CoE	Centre of Excellence
COSMO	Consortium for Small-Scale Modelling
CMIP	Coupled Model Intercomparison Project
CMIP6	Coupled Model Intercomparison Project Phase 6
CRESCENDO	Coordinated Research in Earth Systems and Climate: Experiments, Knowledge, Dissemination and Outreach, project application submitted for the call SC5-01-2014
DMP	Data Management Plan
EC	European Commission
EMB	Executive Management Board
EOFS	European Open File System
EP	dissemination and Exploitation Plan
ESD	Earth System Data
ENES	European Network for Earth System modelling
EOFS	European Open File System
ESCAPE	Energy-efficient Scalable Algorithms for Weather Prediction at Exascale, project application submitted for the call FETHPC-1
ESGF	Earth System Grid federation
ESM	Global Earth System Models [23]
ESMF	Earth System Modeling Framework
ETP4HPC	European Technology Platform for High Performance Computing
EUDAT	European Data Infrastructure
EUDAT2020	Project application submitted for the call H2020-EINFRA-2014-2
FET	Future and Emerging Technologies
GA	General Assembly
GCM	General Circulation Model
GFDL	Geophysical Fluid Dynamics Laboratory, Princeton
GRIB	GRIdded Binary
GungHo	Globally Uniform Next Generation Highly Optimized
G8 ICOMEX	ICOsahedral-grid Models for Exascale Earth System Simulations
HD(CP)2	High Definition Clouds and Precipitation for advancing Climate Prediction
HIRLAM	High Resolution Limited Area Model
HPC4CM	High performance computing for climate modelling, project application submitted for the call H2020-MSCA-ITN-2015
HPSS	High Performance Storage System

[23] **PLEASE NOTE: Throughout this application the term „Earth System Modelling“ (ESM) is used as short for „Earth System Modelling for weather and climate science”. Earth System Modelling in a broader sense would also incorporate the solid Earth.**

I3	Integrated Infrastructures Initiative
ICON	Icosahedral non-hydrostatic (general circulation model)
IFS	Integrated Forecast System
IMPULSE	Improved Modelling and Physical Understanding of Decadal to Centennial Scale Climate Changes, project application submitted for the call SC5-2014
I/O	Input/Output
IPCC	Intergovernmental Panel on Climate Change
IPR	Intellectual Property Rights
IS-ENES2	Infrastructure for the European Network or Earth System modeling – phase 2
JRA	Joint Research Activities
LHC	Large Hadron Collider
LACE	Regional Cooperation for Limited Area Modeling in Central Europe
LFRic	Lewis Fry Richardson
MARS	Meteorological Archival and Retrieval System
MASS	Mass Archive Storage System
MESSy	Modular Earth Submodel System
MSB	Management Steering Board
NA	Networking Activities
NCAR	National Centre for Atmospheric Research, Boulder
NEMO	Nucleus for European Modelling of the Ocean
netCDF	network Common Data Form
NIWA	The NZ National Institute of Water and Atmosphere
NWP	Numerical Weather Prediction
OA	Open Access
OASIS	Couper software see https://verc.enes.org/oasis
OBS4MIPS	Observations for Model Intercomparisons
OpenIFS	Open Integrated Forecast System
OpenMP	Open Multi-Processing
OpenPALM	Open Projet d'Assimilation par Logiciel Multiméthode
PAPI	Performance Application Programming Interface
PGAS	Partitioned Global Address Space
PO	Project Office
POSIX	Portable Operating System Interface
PRACE	Partnership for Advanced Computing in Europe
PRIMAVERA	PRocess-based climate sIMulation: AdVances in high resolution modelling and European climate Risk Assessment, project application submitted for the call SC5-01-2014
RAIT	Redundant Array of Independent Tapes
RIA	Research & Innovation Actions
SA	Service Activities
SAB	Scientific Advisory Board
SO	Scientific Officer
SSD	Solid-State Drive
S2S	Subseasonal to Seasonal Prediction Project
TRL	Technology Readiness Levels
WCRP	World Climate Research Programme
WGCM	Working Group on Coupled Modeling
WGNE	Working Group on Numerical Experimentation
WMO	World Meteorological Organization
WP	Work Package
WPL	Work Package Leader
WWRP	World Weather Research Programme
XIOS	eXtended Input/Output System
YAC	Yet Another Coupler
YAXT	Yet Another eXchange Tool

Bibliography of strategic documents

Council of the European Union, 2013, Conclusions on 'High Performance Computing, 3242nd COMPETITIVENESS (Internal Market, Industry, Research and Space) Council meeting Brussels, 29 and 30 May 2013

ENES foresight paper: Mitchell J., Budich R., Joussaume S., Lawrence B. and Marotzke J. (2012), "Infrastructure strategy for the European Earth System Modelling community 2012-2022", ENES Report Series 1, 33 pp. https://verc.enes.org/community/about-enes/the-future-of-enes/ENES_foresight.pdf

ETP4HPC, The European Technology Platform for High performance Computing, Vision Paper, http://www.etp4hpc.eu/publications/key-documents/

ETP4HPC, The Strategic Research Agenda: http://www.etp4hpc.eu/publications/key-documents/

European Commission Brussels, 15.2.2012, COM(2012) 45 final, Communication from the Commission to the European Parliament, the council, the European economic and social committee and the committee of the regions, High-Performance Computing: Europe's place in a Global Race

European Commission, 10th July 2013, Communication (COM(2013) 494 - "Public-private partnerships in Horizon 2020: a powerful tool to deliver on innovation and growth in Europe"

European Commission, Action 132: Invest in High Performance Computing, http://ec.europa.eu/digital-agenda/en/pillar-v-research-and-innovation/action-132-invest-highperformance-computing

European Commission, Press Release Brussels, 17 December 2013, EU industrial leadership gets boost through eight new research partnerships

European Commission, 17 December 2013, Factsheet High Performance Computing PPP: Mastering the next generation of computing technologies for innovative products and scientific discovery

European Commission, Guidelines on Open Access to Scientific Publications and Research Data in Horizon 2020

European Commission, Guidelines on Data Management in Horizon 2020

European Commission, Communication EU Research and innovation guidance for project participants

European Commission, Communicating EU Research and Innovation, A guide for project participants

European Commission, Guidance on the evaluation of innovation, social sciences and humanities and other aspects of H2020 proposals

European Commission, Guidance on Gender Equality in H2020

Mark Sawyer, Mark Parsons, 2011, A strategy for research and innovation through high performance computing, The University of Edinburgh

4 Members of the consortium

4.1 Participants (applicants)

Deutsches Klimarechenzentrum GmbH (DKRZ)

About the institute

DKRZ, the German Climate Computing Centre, is a national service provider which constitutes an outstanding research infrastructure for model-based simulations of global and regional climate and the investigation of the processes in the climate system. DKRZ's principal objectives are provision of adequate computer performance, data management, and service and support to use these tools efficiently. DKRZ operates one of the largest supercomputers in Germany and provides its more than 1000 scientific users with the technical infrastructure needed for the processing and analysis of huge amounts of data from climate simulations. This also includes training and support for related application software and data processing issues. DKRZ participates in many national and international projects aiming to improve the infrastructure for climate modeling. Through its research group on scientific computing DKRZ is linked to the Department of Informatics of the University of Hamburg.

DKRZ is a non-profit and non-commercial limited company with four shareholders. MPG (Partner 5) holds 55% of the shares of DKRZ (see http://www.dkrz.de/about-en/Organisation/gesellschafter for more references). The dependency relationship has been declared in the Part 2 – Administrative data of participating organisation of this application form

Contribution to the specific project

In ESiWACE, DKRZ will coordinate the entire project in WP5, co-lead WP1 and WP4. Moreover, DKRZ will contribute to WP2 and WP3.

Dr. Joachim Biercamp (male) holds a PhD in Physical Oceanography and has a long standing experience in supporting data intensive climate simulations. He is leading the Application department of DKRZ. His responsibilities include the organization of user support and the interaction with DKRZ's user group and scientific steering committee. He coordinated the procurement and benchmarking of the previous two generations of DKRZ super computers (a NEC SX6, and an IBM Power6 system). Both ranked within the TOP 30 of the TOP500 list. Joachim also was in charge of procuring the latest super computer, a BULLX system which is currently being installed and will be reach a peak performance of more than 3 PetaFLOPS and a file system of world leading 50 PetaBytes. He is involved in several national and international projects dealing with infrastructure for climate modeling. In particular he is member of the steering committee of the project HD(CP)2 aiming at development and operation of a cloud resolving version of the ICON model which is used for both, climate research and numerical weather prediction.

Dr. Kerstin Fieg (female) works in the Application department of DKRZ as work package leader in third-party funding projects. She holds a PhD in Meteorology has 20-year experience in supporting data intensive climate simulations and project management. She worked in the field of development, coupling and performance optimization of Earth System Models and has contributed to procurement and management of several HPC systems. Furthermore she worked as Project Manager for ORACLE Deutschland GmbH for 4 years. Currently Kerstin is involved in several national and international projects dealing with infrastructure to support climate modeling. Kerstin will support the WP5 and WP1

Prof. Thomas Ludwig (male) is the director of DRKZ and leader of the research group Scientific Computing of the DKRZ. His research activity is in the fields of high volume data storage, energy efficiency, and performance analysis concepts and tools for parallel systems. At DKRZ Thomas takes the responsibility for accomplishing its mission: to provide high performance computing platforms, sophisticated and high capacity data management, and superior service for premium climate science. Thomas will lead WP4.

Dr. Julian Kunkel (male), he is Principal Investigator in the group Scientific Computing at the DKRZ. Since 2006, Julian has been working on tracing environments and tools for client and server-side I/O. In 2013, he defended his thesis about the monitoring and simulation of parallel programs on application and system level. Julian is member in several program committees and chairing the Research Poster sessions of the International Supercomputing Conference (ISC) in 2014. Also, he is co-chair for ISC's Research Paper sessions since 2010. He was responsible for the University of Hamburg's contributions to the SIOX and ICOMEX project. Currently, with his role as member of the steering committee of the Exascale10 EOFS working group, he is focusing on system-wide monitoring, optimization of parallel I/O and on novel I/O interfaces for Exascale. Julian will support WP4.

Sonja Kempe (female) has more than 10-year experience in the management and administration of third-party funding projects. She is also in charge management of multi-partner consortia where DKRZ is involved. Sonja will support the WP5 in particular and the entire consortium in administrative and financial matters.

Irina Fast (female) graduated in Meteorology at the Free University Berlin, Germany. She has expertise in different areas of Earth System Modelling such as development of workflow infrastructure for performance of long-term climate simulations and data management, running of high-resolution simulations, processing and analysis of model output, as well as porting and benchmarking of climate models on various HPC systems, analysis and improvement of the computational performance, and implementation of asynchronous model output via CDI-PIO library. She will contribute to WP2 and WP3.

Dr. Kerstin Ronneberger (female) works in the Application Department of DKRZ. She holds a PhD in Earth System Research. Since 2005 she worked in several projects dealing with gridsoftware, data-access and workflow in Earth-System Modeling.

Dr. Panagiotis Adamidis (male) holds a PhD in Mechanical Engineering. He has 17 years of experience in the area of High Performance Computing with emphasis on parallel numerical algorithms and parallel programming models. Since 2006 he has been working for DKRZ, dealing with the development of parallel algorithms for earth system models and optimization issues at application level. He participates in the project HDCP2 (High Definition Clouds and Precipitation for Climate Prediction) where his focus is on enhancing scalability of the ICON model for future generation supercomputers. He will contribute to WP2.

Publications, and/or products or other achievements

1. Julian Kunkel, Michael Kuhn, and Thomas Ludwig: Exascale Storage Systems – An Analytical Study of Expenses. Supercomputing Frontiers and Innovations, pages 116–134, 06 2014.
2. Christopher Bartz: An in-depth analysis of parallel high level I/O interfaces using HDF5 and NetCDF-4. Master's thesis, University of Hamburg, 04 2014.
3. Dirk Meister, Jürgen Kaiser, Andre Brinkmann, Michael Kuhn, Julian Kunkel, and Toni Cortes: A Study on Data Deduplication in HPC Storage Systems. In Proceedings of the ACM/IEEE Conference on High Performance Computing (SC). IEEE Computer Society, 11 2012.
4. Konstantinos Chasapis, Manuel Dolz, Michael Kuhn, and Thomas Ludwig: Evaluating Power-Performace Benefits of Data Compression in HPC Storage Servers. In Steffen Fries and Petre Dini, editors, IARIA Conference, pages 29–34. IARIA XPS Press, 04 2014.
5. Stockhause, M., Höck, H., Toussaint, F., and Lautenschlager, M.: Quality assessment concept of the World Data Center for Climate and its application to CMIP5 data. Geosci. Model Dev., 5, 1023-1032, 2012, doi:10.5194/gmd-5-1023-2012.
6. Quadt, F., Düsterhus, A., Höck, H., Lautenschlager, M, Hense, A.V., Hense, A.N., Dames, M.: ATARRABI – A Workflow System for the Publication of Environmental Data. Data Science Journal, 11, 89-109, 2012, doi:10.2481/dsj.012-027.
7. Julian Kunkel, Alvaro Aguilera, Holger Mickler, Xuan Wang, Andrij Chut, Thomas Bönisch, Jakob Lüttgau, Roman Michel, Johann Weging. The SIOX Architecture – Coupling Automatic Monitoring and Optimization of Parallel I/O. In Supercomputing, Supercomputing, pp. 245–260, Springer International Publishing, ISBN: 978-3-319-075174, 2014
8. Andre Brinkmann, Toni Cortes, Hugo Falter, Julian Kunkel, Sai Narasimhamurthy Whitepaper: http://www.exascale10.com/ (not peer reviewed)
9. Panagiotis Adamidis, Irina Fast, Thomas Ludwig (DKRZ).Performance Characteristics of Global High-Resolution Ocean (MPIOM) and Atmosphere (ECHAM6) Models on LargeScale Multicore Cluster. Parallel Computing Technologies - 11th International Conference, PaCT 2011, Kazan, Russia, September 19-23, 2011. Proceedings; 01/2011

Projects, and/or activities, services

- EU FP7 IS-ENES Infrastructure for the European Network for Earth-SystemModelling, EU FP7 Project fostering simulations with global earth system models
- EU FP7 EUDAT: European Data Infrastructure; building and supporting a collaborative data infrastructure (http://www.eudat.eu)
- ScalES: Scalable Earth System Models, funded by BMBF (https://www.dkrz.de/Klimaforschung/dkrz-undklimaforschung/infraproj/scales/scales).
- C3-Grid & C3-INAD: Collaborative Climate Community Data and Processing Grid, funded by BMBF
- HD(CP)2: "High definition clouds and precipitation for advancing climate prediction" focusing on cloud resolving models, funded by BMBF
- ICOMEX: ICOsahedral-grid Models for EXascale earth system simulations (http://wr.informatik.uni-hamburg.de/research/projects/icomex/)

- Scalus: SCALing by means of Ubiquitous Storage (http://www.scalus.eu/) • SIOX: Scalable I/O For Extreme Performance (http://www.hpc-io.org/)
- A full list of projects led or participated by DKRZ is available here: http://www.dkrz.de/Klimaforschung-en/projects

DKRZ supports complex and compute intensive national and international collaborative projects: simulations are carried out and resulting data are managed and made available through the World Data Center Climate WDCC operated by DKRZ (http://www.dkrz.de/daten-en/wdcc) .

- IPCC AR4: Simulations for the Fourth Assessment Report of the Intergovernmental Panel on Climate Change
- IPCC AR5: Simulations within the frame of the international climate model comparison project CMIP5 and for the 5th climate assessment report of the United Nation
- Climate simulations for Europe with CLM: the regional climate model CLM, which is based on the weather forecast model "LM" by the German Weather Service, was used for ensemble simulations of regional climate changes in Europe.
- STORM - development of a high resolution climate model: STORM is a consortium project by several German climate research institutes with the common goal of developing a climate model with very high spatial resolution. The impact of small scale processes on the quality of climate simulations is planned to be evaluated on the basis of a 21st century simulation with the model
- Climate Simulations for Europe with CLM: the regional climate model CLM, which is based on the weather forecast model "LM" by the German Weather Service DWD, was used for ensemble simulations of regional climate changes in Europe.

Significant infrastructure, and/or major items of technical equipment

DKRZ is running one of Germany's most powerful high performance computers and world class data storage and archiving hardware. The current system, an IBM Power6 computer with 8500 processor cores and a peak performance of 160 TFLOPS will be replaced by a new BULL system in March 2015. The new system will be installed in two stages, starting with 36000 processor cores and a peak performance of 1.5 PetaFLOPS to be upgraded to ca 8000 cores and 3.2 PetaFLOPS in spring 2016.
The file system will be based on lustre and will initially have a net storage capacity of 20 PetaBytes in 2015 which will be upgraded to 50 PetaBytes in 2016.
For long term archiving of data DKRZ runs a tape library with the capacity to 75 PetaByte of data annually.
A data nodes and data portal of the Earth System Grid Federation (ESGF) are run by DKRZ (http://esgf-data.dkrz.de

European Centre for Medium-Range Weather Forecasts (ECMWF)

About the institute

The European Centre for Medium-Range Weather Forecasts (ECMWF) is an international organisation supported by 34 European and Mediterranean States. ECMWF's longstanding principal objectives are the development of numerical methods for medium-range weather forecasting, the operational delivery of medium-to-seasonal range weather forecasts for distribution to the meteorological services of the Member States, to lead scientific and technical research directed to the improvement of these forecasts, and the collection and storage of appropriate meteorological data. ECMWF has extensive competence in operating complex global forecasting suites on high-performance computers and in transitioning top-level science from research to operations exploiting innovative approaches in computing science to fulfil the tight runtime and delivery constraints required by Member States. ECMWF has signed the delegation agreement with the European Commission to operate the Copernicus Atmospheric Monitoring Service and the Copernicus Climate Change Service.

Contribution to the specific project

ECMWF's contribution to ESiWACE is dedicated to WP1, WP2, WP3, WP4 with coordinator in WP2 .Further, ECMWF will co-coordinate Work Package 2, through the support of the integration of the OpenIFS model in the climate community EC-Earth system, the contribution to detailed performance assessment and code optimization work enhancing the level of concurrency and the overlap of communication and computation. ECMWF will also investigate the information content of ensemble model output to propose ways for significantly reducing data volume produced by long model integrations.

.

Dr Peter Bauer (male) is co-coordinating Work Package 2 of this proposal. He joined ECWMF in January 2000 and heads the Model Division iin the research Dept. that comprises the physical and numerical aspects of numerical weather prediction. Before joining ECMWF, he was leading a DLR research team on satellite meteorology in Cologne, Germany. His background covers physical modelling, data assimilation and satellite remote sensing. He obtained his masters and PhD degrees in meteorology from the universities in Cologne and Hamburg, respectively. During his career, he was awarded research fellowships by NOAA and NASA, and a science award by DLR. He is the author and co-author of 100 peer-reviewed scientific journal papers, and his publications have an h-index of 34. He is a member of several scientific advisory committees at the international level (WMO, ESA, EUMETSAT) and has extensive experience with managing international research projects. At ECMWF, his current duties also include the management of the transition of new model cycles from research to operations and he is the manager of the recently launched Scalability Programme. He has contributed to the EU project EURAINSAT (FP5) and numerous other ESA/EUMETSAT/NASA/BMBF projects.

Dr Tiago Quintino (male) is the Team Leader for Data Handling within the Forecast Dept. at ECMWF. He and his team develop software for data encoding and decoding, pre and postprocessing of meteorological products, storage and perpetual archival of weather observations and forecast data. Tiago develops the meteorological archival software (MARS) and the high performance I/O middleware (FDB). Previously, he worked for the Von Karman Institute for Fluid Dynamics on high performance CFD software for aerospace applications and contributed to multiple EU funded projects. He is also the author of 25 journal publications and book chapters in the area of high performance scientific computing.

Dr Glenn Carver (male) will contribute to Work Package 2 of this proposal. He joined ECMWF in 2011 and leads the OpenIFS project to deliver academic versions of the IFS forecast model. Before joining ECMWF he was a senior research fellow at the University of Cambridge, UK where he led development of a community atmospheric chemistry model and was PI on two national awards and co-PI on two others totaling 680kEuro. He has over 40 peer-reviewed publications with 1300 citations and an h-index of 17.

Dr Filip Vana (male) will contribute to Work Package 2 of this proposal. He joined ECMWF in early 2012 to work on the OpenIFS project. Before joining ECMWF, he was working as the Area Leader for Dynamics for the LACE (Limited Area model for Central Europe) Consortium representing the Czech weather service (CHMI) there. His background covers model dynamics & numerics including variational methods, physics and parallelization aspects. He was a member of Scientific Advisory Committee of LACE and the Committee for Scientific and System/maintenance Issues

(CSSI) of Aladin consortium (Cooperation of over 15 National Meteorological Services of Europe and Northern Africa in the field of Numerical Weather Prediction (NWP). He has 8 per-reviewed publications.

Publications, and/or products, services or other achievements

1. Geer A.J., P. Bauer and C.W. O'Dell, 2009: A revised cloud overlap scheme for fast microwave radiative transfer in rain and cloud. J. Appl. Meteorol. Climat., **48**, 2257-2270.
2. Carver, G.D., Váňa, et al., The ECMWF OpenIFS model, EGU General Assembly Abs., 15, 2013.

Projects, and/or activities

1. EU FP7 project **CRESTA** (2011-2014): Collaborative Research into Exascale Systemware, Tools & Applications (CRESTA) on investigating the efficiency gains from a Partitioned Global Address Space (PGAS) implementation for numerical weather prediction.
2. EU FP7 ERC Advanced Grant project **PantaRhei** (2013-2018): Inter-disciplinary integrated forecasting system for fluid flow.
3. EU H2020 project **MACC-III** (2014-2015): Monitoring atmospheric composition and climate based on IFS as main global model component leading to **Copernicus Atmospheric Monitoring Service**.
4. EU H2020 project **ERA-Clim2** (2013-2016): Global coupled reanalysis covering the 20th century based on the IFS as main global model component to be integrated in **Copernicus Climate Change Service**.

Significant infrastructure, and/or major items of technical equipment

- ECMWF's computer facility includes supercomputers (current Cray XC-30 clusters are ranked 19th and 20th in top500 list), archiving systems and networks.
- ECMWF produces operational forecasts, archives and disseminates global model output to member states under tight schedules employing its computing and data handling infrastructure.
- Both Copernicus services ECMWF will be operating will be supported by the ECMWF computing and data handling infrastructure.

Centre National de la Recherche Scientifique (CNRS-IPSL)

About the institute

The Centre National de la Recherche Scientifique (CNRS) is the main French public research institution under the responsibility of the French Ministry of Education and Research. CNRS acts here in the name of the Institut Pierre Simon Laplace (IPSL), which is a federal institute located in Paris and composed of 9 research laboratories working on global environmental and climate studies. IPSL gathers about thousand scientists and represents more than a third of the French research potential in atmospheric and oceanic sciences. Main laboratories from IPSL involved in ESiWACE are Laboratoire des Sciences du Climat et de l'Environnement, and Laboratoire d'Océanographie et du Climat. One of the main objectives of IPSL is to understand climate variability, both natural and anthropogenic, and future evolution, at global and regional scales.

IPSL's work relies on the development of Earth system models of different complexity (e.g. IPSLESM). IPSL is strongly involved in IPCC Working Group 1. CNRS-IPSL coordinates IS-ENES and IS-ENES2, and has also been involved in several European projects such as ENSEMBLES, METAFOR, EMBRACE. CNRS-IPSL was a pioneer in developing since the 1980s a numerical model of the global physical ocean taking into account the HPC issues from the very beginning. This led to the NEMO Consortium involving different research/operational oceanography centres in Europe (CNRS through LOCEAN and LGGE, CMCC, INGV, Mercator Ocean, UK Met Office, NOC), which join efforts for the sustainable development of NEMO (www.nemo-ocean.eu). Today, CNRS is, among the consortium partners, the largest contributor in terms of number of experts to the NEMO System Team, including the Scientific Leader and the Project Manager, and leading the NEMO HPC working group. CNRS-IPSL with the Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA) have developed XIOS, a software library dedicated to efficient IO management for climate models.

Contribution to the specific project

CNRS-IPSL will lead ESiWACE's WP1 and substantially contribute to WP2 on NEMO and XIOS

Dr. Sylvie Joussaume (female) is a researcher at CNRS since 1983. She is an expert in climate modelling. She has been coordinating IS-ENES (EU FP7) phases 1 and 2 since 2009 and is Chair of the ENES Scientific Board. She is Review Editor in IPCC-AR5. She is involved in the Management Committee of JPI Climate (Joint Programming Initiative on Climate, collaboration between 14 European countries to coordinate jointly their climate research and fund new transnational research initiatives). She has been vice chair of the PRACE scientific committee since 2011 and will chair PRACE SSC in 2015.

Marie-Alice Foujols (female) is a research engineer at CNRS. She is in charge of coordinating technical aspects of the Earth system modelling activities at IPSL, and the head of the IPSL engineers group working across the IPSL laboratories. She has 20 years expertise in high performance computing.

Sébastien Denvil (male) is a research engineer at CNRS. An applied mathematician, he holds a Master's degree in "Information treatment and data processing". He joined the global climate modelling group at CNRS-IPSL where he is responsible for the long IPCC-type simulations and for the distribution of model outputs. He is strongly involved in international metadata standards and in ESGF.

Claire Lévy (female), software engineer at CNRS, is NEMO Project Manager since 2008. She has been working on NEMO (previously OPA) development since the 1990s. Claire will contribute to WP2.

Françoise Pinsard (female)**,** software engineer at CNRS, is a member of the NEMO R&D team at LOCEAN and will contribute to WP2.Françoise Pinsard started to worked on NEMO 10 years ago when the model was running on the Earth Simulator in Japan. She was working on the benchmarking and optimisation of the code in order to obtain the maximum of parallelism from the code. Since then, she is still working on the HPC aspect in NEMO, for example by developing specific tools to identify and quantify bottleneck in NEMO scalability.

Dr Sébastien Masson (male) is a researcher in oceanography and climatology at Université Pierre et Marie Curie (UPMC), linked third party to CNRS. Member of the NEMO Developing Team for 10 years, he has been running the model on high performance computers for more than 15 years, such as the Japan Earth Simulator, national facilities within "GENCI grand challenges", and on PRACE projects. He is the PI of the 4-year PULSATION project (French National Research Agency). He is responsible for the upscaling part of the WP2 of project EMBRACE (EU FP7) and participates to the projects IS-ENES and IS-ENES2 (EU FP7)

Dr Yann Meurdesoif (male) is an engineer at CEA, linked third party to CNRS, and has a background in theoretical physics. He has been working for several years as a consultant for high performance computing at the CEA supercomputing centre and has then acquired a strong background on parallelism, code porting and optimization on a large variety of supercomputer. He has developed the parallel versions of several components of the IPSL Earth system model and is responsible at IPSL for the development of XIOS, a software library dedicated to efficient IO management for climate models.

Arnaud Caubel (male) engineer at CEA, linked third party to CNRS, in charge of coordinating technical aspects of the IPSL Earth System climate model and its environment. He has expertise on coupling aspects and high performance computing and is involved in the use of XIOS (WP2).

Dr. Laure Coquart (female) has a PhD in fluid dynamics from Univ. of Paris VI. She is now a research engineer in the Climate Modelling and Global Change team where she tests and validates the development of the OASIS coupler on different computing platforms. Dr Coquart also provides user support by email and via the OASIS forum or by visiting research teams in France. Through this support, Dr Coquart developed a strong expertise in coupled climate modelling and established working contacts with many climate-modelling groups in Europe.

Publications, and/or products, services or other achievements

1. Dufresne J.L., **M.A. Foujols**, **S. Denvil**, and 57 other authors including **S. Joussaume**, **S. Masson**, Y**. Meurdesoif**, 2013: Climate change projections using the IPSL-CM5 Earth System Model: from CMIP3 to CMIP5. *Climate Dynamics*, 40, 2123–2165, doi:10.1007/s00382-012-1636-1.

2. Masson S. , P. Terray, G. Madec, J.J. Luo, T. Yamagata, K. Takahashi, impact of intradaily SST on ENSO characteristics in a coupled model variability, *Climate Dynamics*, 39, pp 681-707, 2012.

3. Masson, S., M. -A. Foujols, P. Klein, G. Madec, L. Hua, M. Levy, H. Sasaki, K. Takahashi and F. Svensson, 2007: OPA9 — French Experiments on the Earth Simulator and Teraflop Workbench Tunings, *High Performance Computing on Vector Systems*, p 25-34, Springer, 2007

4. Mitchell J., Budich R., Joussaume S., Lawrence B., Marotzke J. (2012), Infrastructure strategy for the European Earth System modeling community: 2012-2022, ENES Report Series 1, 2012, 33 pp https://is.enes.org/the-project/communication/ENES%20foresight.pdf

Projects, and/or activities

- EU FP7 Project **IS-ENES** phase 1 (2009-2013), phase 2 (2013-2017) is the distributed einfrastructure of models, model data and metadata of the European Network for Earth System Modelling (ENES). IS-ENES projects, coordinated by CNRS-IPSL, combine expertise in climate modelling, computational science, data management and climate impacts. They aim at fostering the integration of the European Climate and Earth system modelling community, enhancing the development of Earth system models for the understanding of climate variability and change, supporting high-end simulations, and at facilitating model applications to better predict and understand climate change impacts on society.

- EU FP7 Project **EMBRACE** (Earth system Model Bias Reduction and assessing Abrupt Climate change, 2011-2015): bringing together the leading Earth system models (ESMs) in Europe around a common set of objectives to improve our ability to simulate the Earth system and to make reliable projections of future global change. The project has a number of key goals: to reduce the main, known biases in existing European ESMs, to fully evaluate ESM simulation capabilities and improvements made in the project, to increase the realism of, and interactions between, the physical and biogeochemical components of ESMs, to assess the risks of abrupt or irreversible changes.

- ANR (French National Research Agency) Project **PULSATION** (Peta-scale mULti-gridS oceanATmosphere coupled simulatIONs) aims to explore new pathways toward a better representation of the multi-scale physics that drive climate variability.

- ANR (French National Research Agency) Project **CONVERGENCE (**Convergence facing Big Data era and Exascale challenges for Climate Sciences).aims to develop a platform capable of running large ensembles of simulations with a suite of models, to handle the complex and voluminous datasets generated, to facilitate the evaluation and validation of the models and the use of higher resolution models.

Significant infrastructure, and/or major items of technical equipment

- CNRS-IPSL is strongly involved in IPCC Working Group 1 through authors and review editors. CNRS-IPSL is one of the seven European modelling groups developing and performing international global model inter-comparison projects, such as CMIP5.

- CNRS-IPSL also contributes to the NEMO HPC Group, active in analysing the future possible developments of NEMO and of its components in view of the evolving computational landscape in the coming 5-10 years. These activities include the conception of new solvers and mathematical algorithms.

Max-Planck-Institut für Meteorologie (MPG)

About the institute

The Max Planck Institute for Meteorology (MPG) performs basic research in the interest of the general public. Its mission is to understand the Earth’s changing climate.

5. It comprises three departments (The Atmosphere in the Earth System, The Land in the Earth System, The Ocean in the Earth System) and hosts independent research groups focused on: Fire in the Earth System, Forest Management in the Earth System, Sea Ice in the Earth System, Stratosphere and Climate, Turbulent Mixing Processes in the Earth System.

Scientists at MPG investigate what determined the sensitivity of the Earth system to perturbations such as the changing composition of its atmosphere, and work towards establishing the sources and limits of predictability within the Earth system. The MPG develops and analyses sophisticated models of the Earth System which simulate the processes within atmosphere, land, and ocean. Such models have developed into important tools for understanding the behavior of our climate. Models form the basis for international assessments of the climate change. Targeted in-situ measurements and satellite observations complement the model simulations.

MPG is committed to informing public and private decision-makers and the general public on questions related to climate and global change.
Together with the University of Hamburg, MPG runs an international doctoral programme, the International Max Planck Research School on Earth System Modelling (IMPRS-ESM) to promote high-quality doctoral research into the Earth's climate system, hosting approximately 50 PhD students per year.
MPG is actively involved in the cluster of excellence "Integrated Climate System Analysis and Prediction" (CliSAP), a research and training network whose goal is to bridge the gap between natural sciences, economics and humanities, creating synergies for analysing natural and humancaused climate change and developing scenarios for the future.
The MPG is the major shareholder of German Climate Computing Centre (DKRZ GmbH), the coordinator of the ESiWACE project. DKRZ is an outstanding research infrastructure for modelbased simulations of global climate change and its regional effects. DRKZ provides tools and the associated services needed to investigate the processes in the climate system, computer power, data management, and guidance to use these tools efficiently.

Contribution to the specific project

MPG's contribution to ESiWaCE is dedicated to WP1, WP2, WP3, WP5
Reinhard Budich (male) is responsible for the strategic IT partnerships of MPG. As an Oceanographer he has been working in Earth system modelling for over 20 years. He has held responsibility for the IT at MPG in many roles, was the Technical Coordinator of the FP7 project ISENES, is member of the board of ENES, was running the COSMOS network as a project manager and was Director of the FP5 PRISM project. Since 2001 he has been running the ENES office in Hamburg. Reinhard was involved as technical coordinator in ISENES1, has initiated the International summer school series E2SCMS of ENES, is work package leader in ISENES2 and member of the Council of EUDAT. Reinhard also is a member of the IUAC of GEANT. Reinhard will lead the WP 3 "Usability" and supports WP1 and 2.

Dr. Luis Kornblueh (male) graduated at the Technical University of Darmstadt, Germany and received his PhD in Meteorology from University of Hamburg. His first task at MPG has been refactoring and parallelizing a comprehensive GCM (ECHAM). At the same time he has been participating in Science- and Mission-Advisory groups of ESA and EUMETSAT projects for radiooccultation instruments and working on the influence of those on numerical weather prediction. During the last couple of years he focused more on the performance of models in detail: Optimization on the loop level, correct and fast OpenMP implementations, MPI handling, and I/O strategies. Recently, the problem of optimizing workflows is a new subject he is looking into, because there is more potential to gain improvements compared to hard-core optimization. Luis will be involved in WP 2 and 3.

Karl-Hermann Wieners (male) is scientific software developer with the Computational Infrastructure and Model Development group at MPG. His primary interest is scientific workflow, data processing, and database applications. Before joining MPG, he has worked for 8 years as software engineer and technical consultant in an SME, with customers in retail and media industry. At MPG, he has contributed to the EU Project IS-ENES, and is currently contributing to IS-ENES2. Karl-Hermann will be involved in WP 2 and 3.

Chiara Bearzotti (female) is in charge of Unit for project development and project monitoring at the MPG. Her role is to support scientists granted third-party funding in managing their projects and implementing them according to the requirements of the funding agencies, not only on contractual and financial issues, but also on issues related to Open Science /Open Access, intellectual property rights management, communication, dissemination and public engagement. Chiara has degrees in economics, project design and project management, rounded up by sound legal knowledge of EU law; Chiara is an Open Access Ambassador at the MPG. In the last 16 years, Chiara has been working in the project development and management both on the side of those applying for funding and for those, like European Commission agencies, funding projects. She has been the project manager for the FP7 research projects THOR and NACLIM, ESFRI European XFEL project and she has been involved in several Horizon 2020 applications and grant agreement preparations. Given the institutional link between Max Planck Meteorology and DKRZ, Chiara will support DKRZ in the implementation of the activities foreseen under WP5.

Publications, and/or products, services or other achievements

1. Bonaventura, L., Redler, R., & **Budich, R. G.** (2012). Earth system modelling 2: Algorithms, code infrastructure and optimization. Heidelberg: Springer.
2. **Budich, R. G.** (2011). Network-based approaches to climate knowledge discovery: Climate Knowledge Discovery Workshop, Hamburg, Germany, 30 March to 1 April 2011. EOS, Transactions of the American Geophysical Union, 92(47), 425-425.
3. Ford, R., Riley, G., Redler, R., & **Budich, R. G.** (2012). Earth system modelling 5: Tools for configuring, building and running models. Heidelberg: Springer.
4. Giorgetta, M. A., Jungclaus, J. H., Reick, C. H., Legutke, S., Bader, J., Böttinger, M., Brovkin, V., Crueger, T., Esch, M., Fieg, K., Glushak, K., Gayler, V., Haak, H., Hollweg, H.D., Ilyina, T., Kinne, S., **Kornblueh, L.**, Matei, D., Mauritsen, T., Mikolajewicz, U., Mueller, W. A., Notz, D., Pithan, F., Raddatz, T., Rast, S., Redler, R., Roeckner, E., Schmidt, H., Schnur, R., Segschneider, J., Six, K., Stockhause, M., Timmreck, C., Wegner, J., Widmann, H., Wieners, K.-H., Claussen, M., Marotzke, J., & Stevens, B. (2013). Climate and carbon cycle changes from 1850 to 2100 in MPI-ESM simulations for the coupled model intercomparison project phase 5. Journal of Advances in Modeling Earth Systems, 5, 572-597
5. Balaji, V., Redler, R., & **Budich, R. G.** (2013). Earth system modelling 4: IO and postprocessing. Heidelberg: Springer.
6. Hiller, W., **Budich, R. G.,** & Redler, R. (2013). Earth system modelling 6: ESM data archives in the times of the grid. Heidelberg: Springer.
7. Puri, K., Redler, R., & **Budich, R. G.** (2013). Earth system modelling 1: Recent developments and projects. Heidelberg: Springer.
8. Riedel, M., Wittenburg, P., Reetz, J., van de Sanden, M., Rybicki, J., von St Vieth, B., Fiameni, G., Mariani, G., Michelini, A., Cacciari, C., Elbers, W., Broeder, D., Verkerk, R., Erastova, E., Lautenschlager, M., **Budich, R. G.**, Thiemann, H., Coveney, P., Zasada, S., Haidar, A., Buechner, O., Manzano, C., Memon, S., Memon, S., Helin, H., Suhonen, J., Lecarpentier, D., Koski, K., & Lippert, T. (2013). A data infrastructure reference model with applications: Towards realization of a ScienceTube vision with a data replication service. Journal of Internet Services and Applications, 4: 1.
9. Stevens, Bjorn, Marco Giorgetta, Monika Esch, Thorsten Mauritsen, Traute Crueger, Sebastian Rast, Marc Salzmann, Hauke Schmidt, Jürgen Bader, Karoline Block, Renate Brokopf, Irina Fast, Stefan Kinne, **Luis Kornblueh**, Ulrike Lohmann, Robert Pincus, Thomas Reichler, and Erich Roeckner (2013) Atmospheric component of the mpi-m Earth system model: Echam6. JOURNAL OF ADVANCES IN MODELING EARTH SYSTEMS, 2013
10. Valcke, S., Redler, R., & **Budich, R. G.** (2012). Earth system modelling 3: Coupling software and strategies. Heidelberg: Springer.
11. Wan, H., M. A. Giorgetta, G. Zängl, M. Restelli, D. Majewski,L. Bonaventura, K. Fröhlich, D. Reinert, P. Rípodas, and **L. Kornblueh**. (2013) The icon-1.2 hydrostatic atmospheric dynamical core on triangular grids - part 1: Formulation and performance of the baseline version. Geoscientific Model Development Discussions, 6:59–119

Projects, and/or activities

- EU FP7 Project **EUDAT** (http://eudat.eu/) is building a data e-infrastructure which supports collaborations for, among others, climate modellers (ENES). STFC is leading the work package on scalability in the (now extended till March 2015) project as well as the task force on authentication and accounting
- EU FP7 project **IS-ENES2**: Infrastructure for the European Network for Earth System modelling - Phase 2

Significant infrastructure, and/or major items of technical equipment

- MPG is the major shareholder (55%) of the supercomputer at the German Climate Computing Center (DKRZ)
- MPG is one of the key laboratories developing and performing large model intercomparison exercises, such as CMIP5.

Centre Européen de Recherche et de formation Avancée en Calcul Scientifique (CERFACS)

About the institute

CERFACS (http://www.cerfacs.fr), established in 1987 in Toulouse (France), is currently one of the world's leading research institutes working on efficient algorithms for solving large-scale scientific problems. The CERFACS Climate Modelling and Global Change team conducts basic scientific research and high-level technical developments in the field of climate studies. In particular, the team develops the OASIS coupler currently used by more than 40 climate-modelling groups in Europe and around the world. CERFACS activities in high performance computing encompass assembling high-resolution coupled climate based on state-of-art component models, porting and optimising them on a variety of platforms. CERFACS has participated in CMIP5 and one of its main scientific objectives is to make significant contributions to the understanding of the world climate variability on regional to global scales and to climate impact studies at seasonal-todecadal time scales. CERFACS is getting also involved in building new approaches to deal with large data volumes produced in climate science together with large data centres in Europe. Thanks to its strong expertise in code coupling and the central role played by the OASIS coupler in the European climate community, CERFACS was heavily involved in the set-up of the IS-ENES1 (2009-2012) and IS-ENES2 (2012-2016) projects and now actively participates in IS-ENES2 as leader and co-leader of 2 work packages and leader of the HPC task force. CERFACS is also involved in several other e-infrastructure and scientific FP7 European projects: PREFACE, 20142018 (WP leader); CLIPC, 2014-2018 (participant); SPECS, 2012-2016 (WP leader); EUDAT 2011-2015 (Task leader).

Contribution to the specific project

CERFACS is co-leading WP2 and is leading T2.2 on performance benchmarking for I/O and coupling technologies. Within this work package, CERFACS is of course also involved in the tasks addressing user support and training for the OASIS coupler (T2.1.2) and optimization of the coupler (T2.3.2). Thanks to its experience on governance for the OASIS software, CERFACS is participating to WP1 in T1.1.2 about the governance for community software. Finally, CERFACS also commits to continue leading the HPC task force and to organize an HPC workshop (WP1, T1.2.1).

Dr. Sophie Valcke (female) holds a "highly qualified" research engineer position at CERFACS where she is working on high-resolution atmosphere-ocean-ice coupled modelling and is leading a team of 4 engineers developing the OASIS coupler. Thanks to her expertise in HPC for climate, Dr Valcke currently sits on the PRACE Access Committee, participated to the PRACE Scientific Case Panel and in the related working group for the European Exascale Software Initiative. Dr Valcke is CERFACS Principal Investigator for the current IS-ENES2 project and was CERFACS PI in ISENES1 and METAFOR project. These projects favour Dr Valcke's interaction with many climatemodelling groups in Europe and with other groups internationally developing coupling frameworks, such as the USA-led ESMF or the NCAR Community Earth System Model (CESM). Dr Valcke also played a key role in the set-up of the International Working Committee on Coupling Technologies that organizes global efforts related to the characterization, comparison, and benchmarking of Earth system model coupling technologies (http://earthsystemcog.org/projects/iwcct/).

Publications, and/or products, services or other achievements

1. Dunlap, R., M. Vertenstein, **S. Valcke**, and T. Craig, 2014. Second Workshop on Coupling Technologies for Earth System Models, Bull. Amer. Meteor. Soc., 95, ES34–ES38, doi:10.1175/BAMS-D-13-00122.1.

2. **Valcke, S.** (2013): The OASIS3 coupler: a European climate modelling community software, Geosci. Model Dev., 6, 373-388, doi:10.5194/gmd-6-373-2013. http://www.geosci-model-dev.net/6/373/2013/gmd-6-373-2013.pdf

3. **Valcke, S**., T. Craig and **L. Coquart**, 2013. OASIS3-MCT User Guide, OASIS3-MCT_2.0, Technical Report TR/CMGC/13/17, Cerfacs, France.

4. **Valcke, S**., V. Balaji, A. Craig, C. Deluca, R. Dunlap, R. Ford, R. Jacob, J. Larson, R. O'Kuinghttons, G. Riley, M. Vertenstein (2012): Coupling technologies for Earth System Modelling, Geosci. Model Dev., 5, 1589-1596, doi:10.5194/gmd-5-1589-2012. http://www.geosci-model-dev.net/5/1589/2012/gmd-5-1589-2012.pdf

5. Voldoire, A., E. Sanchez-Gomez, D. Salas y Mélia, B. Decharme, C. Cassou, S.Sénési, **S.Valcke**, I. Beau, A. Alias, M. Chevallier, M. Déqué, J. Deshayes, H. Douville, E. Fernandez, G. Madec, E. Maisonnave, M.-P. Moine, S. Planton, D.Saint-Martin, S. Szopa, S. Tyteca, R. Alkama, S. Belamari, A. Braun, **L. Coquart**, F. Chauvin. 2011. The CNRMCM5.1 global climate model: Description and basic evaluation. Clim. Dyn. Special Issue, DOI: 10.1007/s00382-011-1259-y.

Projects and/or activities

CERFACS provides advanced HPC training in code coupling for PhD & post-doctoral students and engineers (see http://www.cerfacs.fr/19-25708-Home.php). CERFACS is involved in the current EU and national projects:

- EU FP7 project **PREFACE**, 2014-2018: Variability, previsibility and impacts of climate change on the Tropical Atlantic.
- EU FP7 project **CLIPC**, 2014-2018: Climate change indicators.

- EU FP7 project **SPECS**, 2012-2016: High-resolution decadal prediction.
- EU FP7 project **IS-ENES2**, 2012-2016: Development and implementation of OASIS3-MCT coupler, coupling technology performance benchmarking, multi-model multi-member highresolution ESM experiments, standardisation and metadata, development of a climate date portal for impact community.
- EU FP7 project **EUDAT,** 2011-2015: Infrastructure for scientific data in the context of Big Data
- ANR (French Agence Nationale de la Recherche) project **CONVERGENCE** 2013-2017: Infrastructure for the French climate modelling community.

Significant infrastructure, and/or major items of technical equipment

CERFACS has indoor HPC capacity of 75 peak Tflops, thanks to a BULL Intel SandyBridge cluster and HP AMD MagnyCours cluster.

Barcelona Supercomputing Center(BSC)

About the institute

The Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-CNS, BSC henceforth), created in 2005, has the mission to research, develop and manage information technology in order to facilitate scientific progress not only in computer science but also in a large range of applications. More than 350 people from 40 different countries perform and facilitate research into Computer Sciences, Life Sciences, Earth Sciences and Computational Applications in Science and Engineering at the BSC. The BSC is one of the eight Spanish “Severo Ochoa Centre of Excellence” institutions selected in the first round of this prestigious programme in Spain, as well as one of the four hosting members of the European PRACE Research Infrastructure. The BSC hosts MareNostrum III, a Tier-0 PRACE system currently ranked as the 24th most powerful supercomputer in Europe (57th in the world) with 1 Pflop/s capacity. In addition, the BSC hosts other High-Performance Computing (HPC) resources, among which it is worth mentioning MinoTauro, a hybrid system with GPUs incorporated.

The Earth Sciences Department of the BSC (ES-BSC) was established with the objective of carrying out research in Earth system modelling. The ES-BSC conducts research on air quality, mineral dust and climate modelling and strongly contributes to the scientific and technological advancement in atmospheric and mineral dust modelling. In this sense, the ES-BSC develops and maintains a state-of-the-art mineral dust model: NMMB/BSC-Dust. The excellent results of the group on this field have contributed to the recently creation of the first World Meteorological Organization (WMO) Regional Meteorological Center specialized on Atmospheric Sand and Dust Forecast, the “Barcelona Dust Forecast Center”. In which the NMMB/BSC-Dust model has been selected as the reference mineral dust model. Currently the model provides mineral dust forecasts to the World Meteorological Organization (WMO) Sand and Dust Storm Warning Advisory and Assessment System (SDS-WAS) Northern Africa- Middle East-Europe (NAMEE) Regional Centre that is managed by a consortium between the Spanish Weather Service (AEMET) and BSC. Furthermore, BSC and UNEP are collaborating in the development and implementation of the WMO SDS-WAS West Asia Regional Centre in which the NMMB/BSC-Dust is designed to perform mineral dust forecast simulations. ES-BSC also undertakes research on the development and assessment of dynamical and statistical methods for the prediction of global and regional climate on time scales ranging from a few weeks to several years. The EC-Earth model is used for this purpose. The formulation of the predictions includes the development and implementation of techniques to statistically downscale, calibrate and combine dynamical ensemble and empirical forecasts to satisfy specific user needs in the framework of the development of a climate service.

The high performance capabilities of MareNostrum III and the close collaboration with the Computer Sciences department allow efficiently increasing the spatial and temporal resolution of atmospheric modelling systems, in order to improve our knowledge on dynamic patterns of air pollutants in complex terrains and interactions and feedbacks of physico-chemical processes occurring in the atmosphere, as well as to push the boundaries of global climate prediction. To this, it should be added the increasing collaboration between the different BSC departments on the rapidly growing field of Big Data. Therefore, BSC offers a unique infrastructure to carry out the range of Earth system simulations on which the ES-BSC is a reference worldwide.

Contribution to the specific project

BSC's contribution to ESiWACE is concentrated in WPs 1, 2 and 3.
In WP3 the BSC leads task 3.1. ES-BSC has a vast experience in deploying climate models not only on the Marenostrum3 supercomputer, but also on several other HPC platforms like Archer (Edinburgh), CCA (ECWMF; Reading), etc. BSC has already organized a summer school on climate modelling and therefore has the knowledge required to successfully perform the task. Also, the department can take advantage of the strong links between climate modellers, technical staff of the department and system administrators.

In WP2 leads task 2.3. ES-BSC can offer a wide range of services (based on software and hardware) to ESiWACE to analyse and improve performance of Earth system models. Increasing climate model efficiency requires the application of a wide set of tools to analyse and understand the behaviour of these models running in a parallel environment. BSC, through the Tools Group, develops tools like Paraver or Dimemas that can easily help the user of these codes to understand the behaviour of the code and identify possible bottlenecks and hardware related problems of the application. Furthermore, BSC provides the OmpSs programming model and COMPSs Superscalar (both developed in the institution). OmpSs can exploit parallelism through data dependencies or use different devices (GPU's, accelerators) in a transparent way for the user and COMPSs Superscalar is a programming model that aims to ease the development of applications for distributed infrastructures, such as clusters, grids and clouds.

Prof. Francisco Doblas-Reyes (male) is an expert in the development of seasonal-to-decadal climate prediction systems and the head of the ES-BSC. He is involved in the development of the EC-Earth climate forecast system since its inception. He was an IPCC lead author in the Fifth Assessment Report, serves in WCRP and WWRP scientific panels, is a member of the ENES HPC Task Force, has participated in a number of FP4 to FP7 projects and is author of more than 100 peer-reviewed papers. He is shaping BSC's plans for the development of a weather and climate modelling service that brings the latest developments of HPC and Big Data research to the Earth science community, increasing at the same time the resilience of the European society to weather, air quality and near-term climate extremes.

Mr. Oriol Mula-Valls (male) is an engineer that acts as system administrator of the local network, maintains the environment where the software development takes place and coordinates the design of new solutions for the problems that continuously affect the modellers.

Mr. Kim Serradell (male) is an engineer that as research support engineer. He has a vast experience in the implementation of operational and research meteorological and air quality models on different supercomputing platforms. He is interested in analysing the behaviour of these models with different tools and techniques in order to improve the performance of such codes.

Mr. Domingo Manubens (male) is an engineer that acts as software engineer. He is involved in the design of new solutions in the framework for managing complex global climate coupled models workflows in supercomputing environments. He is the lead developer of Autosubmit. He is interested in facilitating the collaboration during develop and test processes and in the continuous integration of codes developed by the modellers.

Publications, and/or products, services or other achievements

1. Pérez, C., K. Haustein, Z. Janjic, O. Jorba, N. Huneeus, J.M. Baldasano, T. Black, S. Basart, S. Nickovic, R.L. Miller, J.P. Perlwitz, M. Schulz and M. Thomson (2011). Atmospheric dust modeling from meso to global scales with the online NMMB/BSC-Dust model-Part 1: Model description, annual simulations and evaluation. Atmos. Chem. Phys. 11, 13001-13027.

2. Doblas-Reyes, F.J., I. Andreu-Burillo, Y. Chikamoto, J. García-Serrano, V. Guemas, M. Kimoto, T. Mochizuki, L.R.L. Rodrigues and G.J. van Oldenborgh (2013). Initialized nearterm regional climate change prediction. Nature Communications, 4, 1715, doi:10.1038/ncomms2704.

3. Donners, J., C. Basu, A. McKinstry, M. Asif, A. Porter, E. Maisonnave, S. Valcke and U. Fladrich (2012). Performance Analysis of EC-EARTH 3.1, PRACE Whitepaper (http://www.prace-ri.eu/IMG/pdf/Performance_Analysis_of_EC-EARTH_3-1-2.pdf).

4. Johnson, C., A. Carter, I. Bethune, K. Statford, M. Alava, V. Cardoso, M. Asif, B.S.A. Schuberth and T. Weinzierl (2013). PRACE DECI (Distributed European Computing Initiative) Minisymposium. Applied Parallel and Scientific Computing-Lecture Notes in Computer Science, 7782, 43-60, doi:10.1007/978-3-642-36803-5_3.

5. Markomanolis, G.S., O. Jorba and J.M. Baldasano (2014). Performance analysis of an online atmospheric-chemistry global model with Paraver: Identification of scaling limitations. International Workshop on High Performance Computing for Weather, Climate, and solid Earth Sciences held in conjunction with 2014 International Conference on High Performance Computing & Simulation.

6. Asif, M., A. Cencerrado, O. Mula-Valls, D. Manubens, A. Cortés and F.J. Doblas-Reyes (2014). Case study in large scale climate simulations: Optimizing the speedup/efficiency balance in supercomputing environments. 14th International Conference on Computational Science and Its Applications, doi:10.1109/ICCSA.2014.57.

Projects, and/or activities

- EU FP7 project **IS-ENES**: Infrastructure for the European Network of Earth System Modelling.

- EU FP7 project **SPECS**: Seasonal-to-decadal climate Prediction for the improvement of European Climate Services.

- PRACE Tier-0 project **HiResClim**: High Resolution Climate Modelling. ○

Significant infrastructure, and/or major items of technical equipment

- BSC-CNS hosts MareNostrum, the most powerful supercomputer in Spain. MareNostrum has a peak performance of 1,1 Petaflops, with 48,896 Intel Sandy Bridge processors in 3,056 nodes, and 84 Xeon Phi 5110P in 42 nodes, with more than 104.6 TB of main memory and 2 PB of GPFS disk storage.

- BSC-CNS hosts MinoTauro, an NVIDIA GPU cluster with 128 nodes, each one carrying two Intel E5649 (6-Core) processor and two M2090 NVIDIA GPU cards.

- BSC-CNS has a dedicated storage for the Earth Sciences Department of more than 600 TB net space that will be grown to 2 PB in the following two years. This system also acts as ESG node for the Spanish climate modelling community.

Science and Technology Facilities Council (STFC)

About the institute

The Science and Technology Facilities Council (STFC) is one of seven research councils in the UK. Its facilities, instruments, and expertise support an extremely wide range of research at universities, in research councils, and industry. The research council is one of Europe's largest multi-disciplinary research organisations, and has considerable world-class research expertise in areas ranging from nanostructures to lasers, from particle physics to cosmology, from high performance computing and supercomputing to peta-scale data management.

STFC, in particular, provides services to NERC, the National Environment Research Council, so support climate modelling and climate data management. These services include data services for CEDA, the Centre for Environmental Data Archival, and the JASMIN high performance computing cluster, which is backed by the largest Panasas storage system in the world (or at least the largest which is not secret.). The services also include the tapestore, which for NERC, holds both backups of the data archives, and provides an elastic storage facility for research data. (The tapestore infrastructure also supports other communities, including the LHC Tier-1 data centre.)

As a partner in this project, STFC's work is split across the Scientific Computing Department (SCD) which provides the expertise in data storage and processing, and the RAL Space department which will provide the expertise in climate modellingmodelling, climate data formats, and environmental data management and the data formats used. Together they are currently responsible for in excess of 10 PB of environmental data, in millions of files, held primarily on spinning disk, supporting many national and international projects.

Contribution to the specific project

The STFC primary interest and contribution to ESiWACE is twofold: Firstly, projections for future data growth in the STFC environmental data archive have have already led to major concerns as to how the community ambitions can be met. STFC will both generate project requirements, and trial new solutions and technologies developed by ESiWACE. The challenge lies not just in the volume, but also in the ability to ingest data at the required rates, in providing both archives and working repositories, in making data available for analysis in high performance computing clusters, and in providing support for the data throughout its lifecycle, from initial capture to publications of the results.

The second objective lies in supporting data interoperability. While format interoperability is to some extent a solved problem, semantic interoperability of data held in differing formats is not. CEDA has considerable expertise in supporting metadata management, but has not been able to resource NetCDF/GRIB semantic interoperability, despite the enormous importance of such interoperability to the Weather and Climate community. The work proposed here will meet significant international programmatic objectives.

Professor Bryan Lawrence is the Director of CEDA. He also holds a joint position as professor of weather and climate computing at the University of Reading, where his role also includes leadership of the Models and Data Division of the National Centre for Atmospheric Science. With a PhD in atmospheric physics, he has collaborated widely in the field of climate research and the development of computational models, alongside his interest and leadership in data science. Professor Lawrence has in excess of one hundred relevant publications, and was recently awarded the AGU Leptoukh Lecture for "significant contributions to informatics, computational or data science".

Dr Jens Jensen is a group leader of the data services group in the scientific computing department. His primary responsibility is the group delivering data storage and database services to CEDA (and others), and he also leads the storage and data management group in GridPP, the UK part of the large hadron collider computing grid. With a background in maths, his research interests include large scale scientific data management for global collaborations, particularly the data security aspects and the need for globally trusted infrastructures and interoperable identity management. He is Area Director for security in the Open Grid Forum, and is currently leading the authentication and authorisation task force in EUDAT.

Esther Conway joined CEDA in 2011 as an Earth Observation Data Scientist providing level management support to ESA in the ESA LTDP initiative in addition to data curation and management activities at CEDA. Before transferring to CEDA she worked as an Analyst within STFC (2006 -2011) working on UK and European research projects aimed at strengthening the long-term reuse and exploitation of scientific data: SCIDIP- ES, ENSURE; SCAPE, PARSE INSIGHT; CASPAR; DCC – SCARP, BACI and FIDCUCIO. Her role on these projects involved the provision of research analysis, training and project management. The particular focus of her research was the development of process/information models and methods which support the creation and exploitation of scientific research assets. Responsibilities included writing of project proposal, papers and deliverables. She has also spent 4 years (2010-2014) providing consultancy and project managemnet expertise to the ESA LTDP programme.

Philip Kershaw is the Technical Manager in STFC-CEDA, responsible for the development of information and computing systems in support of CEDA. Phil joined CEDA in 2004 following work in a number of software-based projects within RAL. Prior to joining RALSpace Phil worked in the EO industry developing applications for image processing and image registration for EO data, leading to his present 15 years development experience as an experienced software engineer in EO and remote sensing, climate and space science. Through this work Phil has gained specialist knowledge in access control and security for federated systems, for example being a major contributor to the access control architecture for the Earth System Grid Federation, a globally distributed infrastructure deployed to support CMIP5. He is the Technical lead on academic CEMS, and leads the data modelling work in CHARME, an EU FP7 funded project aiming to link commentary metadata (e.g. annotations, supporting information about the data) and climate datasets.

Publications, and/or products, services or other achievements

1. J-C Andre, G Aloisio, J Biercamp, R Budich, S Joussaume, **B Lawrence**, S Valcke: *High Performance Computing for Climate Modeling*, J. Am. Met. Soc., Vol. 95, Issue 5 (May 2014), http://journals.ametsoc.org/doi/abs/10.1175/BAMS-D-13-00098.1
2. Mizielinski, M. S., Roberts, M. J., Vidale, P. L., Schiemann, R., Demory, M.-E., Strachan, J., Edwards, T., Stephens, A., **Lawrence, B. N**., Pritchard, M., Chiu, P., Iwi, A., Churchill, J., del Cano Novales, C., Kettleborough, J., Roseblade, W., Selwood, P., Foster, M., Glover, M., and Malcolm, A.: High-resolution global climate modelling: the UPSCALE project, a large-simulation campaign, Geosci. Model Dev., 7, 1629-1640, doi:10.5194/gmd-7-1629-2014, 2014
3. **Lawrence, B.N**., V.L. Bennett, J. Churchill, M. Juckes, P. Kershaw, S. Pascoe, S. Pepler, M. Pritchard, and A. Stephens. Storing and manipulating environmental big data with JASMIN. Proceedings of IEEE Big Data 2013, p68-75. doi:10.1109/BigData.2013.6691556
4. S Waddington, J Zhang, G Knight, **J Jensen**, R Downing, C Ketley: Cloud repositories for research data -- addressing the needs of researchers, Journal of Cloud Computing: Advances, Systems and Applications 2013, 2:13 (15 June 2013)
5. **J Jensen**: Clouds and Trust and All That. Norwegian Information Security Conference, Stavanger, Nov. 2013
6. **J Jensen** (STFC), P Mori (CNR), S Kindermann (DKRZ), M van de Sanden (SurfSARA): Implementing Community Security Policies for Trustworthy e/cyberinfrastructure, ISGC 2014 (March 2014)
7. **E Conway,** D Giaretta, S Lambert, B Matthews. Curating scientific research data for the long term : a preservation analysis method in context International Journal of Digital Curation 6 (2) (2011);
8. **E Conway**, B Matthews, S Lambert, D Giaretta, M Wilson, N Draper. Managing Risks in the Preservation of Research Data with Preservation Networks. International Journal of Digital Curation 7 (1) (2012).
9. **Esther Conway**, Sam Pepler, Wendy Garland, David Hooper, Fulvio Marelli, Luca Liberti, Emanuela Piervitali, Katrin Molch, Helen Glaves, and Lucio Badiali: Ensuring the Long Term Impact of Earth Science Data through Data Curaton and Preservation. ISQ / v.25 no.3 Fall 2013
(http://www.niso.org/apps/group_public/document.php?document_id=11598&wg_abbrev=isq)
10. D. Ross (ed): SCARF annual report (2014): https://epubs.stfc.ac.uk/work/12275693

Projects, and/or activities

- EU FP7 Project IS-ENES2 (https://verc.enes.org/ISENES2). IS-ENES2 is the second phase project of the distributed e-infrastructure of models, model data and metadata of the European Network for Earth System Modelling (**ENES**).IS-ENES2 combines expertise in climate modelling, computational science, data management and climate impacts.
- EU FP7 Project EUDAT (http://eudat.eu/) is building a data e-infrastructure which supports collaborations for, among others, climate modellers (ENES). STFC is leading the work package on scalability in the (now extended till March 2015) project as well as the task force on authentication and accounting
- SCIDIP-ES FP7 project (http://www.scidip-es.eu/scidip-es/), led by ESA, developed scalable digital preservations for Earth sciences data. It also investigated the use of (computing) clouds for trustworthy processing of data.
- The SCAPE project (www.scape-project.eu) finished at the end of September 2014. It developed a framework for workflows for the preservation of large scale science data, and the

workflows, automated processes, and quality assurance tools could be relevant to the present proposal.
- The CHARMe project, also FP7 (http://www.charme.org.uk/), aims to share climate knowledge through commentary annotations and linked data linked to the original datasets.

Significant infrastructure, and/or major items of technical equipment

CEDA hosts four major data centres: the British Atmospheric Data Centre, the NERC Earth Observation Data Centre, the IPPC Data Distribution Centre, and the Solar System Data Centre, as well as providing a hosting infrastructure to support analysis and exploitation of the data holdings.

The analysis infrastructure includes the JASMIN computing cluster, operated by SCD and owned by RAL Space. Based on VMWare vCloud, it is a significant HPC infrastructure for climate sciences, with 4K cores dedicated to analysis (not simulation) backed by a 12 PB Panasas system – the largest (known) in the world – it is capable of moving about 2 terabits per second. (See the SCARF report, its sister cluster, in references).

The analysis and data centre infrastructure is backed by the SCD datastore, based on Oracle/Sun/StorageTek tape robots. With a nominal nearline capacity of 170 PB, it serves as a dark archive, backup store, and working repository. With data volumes growing roughly exponentially over the years (by, very roughly, a factor 10 every five years), the growth in climate data is one of the key challenges.

Met Office (MetO)

About the institute

The Met Office (MetO) has been operating as a Trading Fund since 1996, originally as an Executive Agency of the UK Ministry of Defence (MoD). As part of a Machinery of Government change in July 2011 MetO became a Trading Fund within the Department for Business, Innovation and Skills (BIS). As the UK's national meteorological service, it provides a range of products and services to a large number of public and private sector organisations. It also represents the UK within the World Meteorological Organisation (WMO) and plays a prominent role in international meteorology.
MetO is one of the world's leading providers of environmental and weather-related services. It delivers proven weather related services for many different types of industry on a twenty-four hour basis. Many of these services are time critical. MetO is involved in many areas of research and development in the fields of atmospheric and oceanic sciences and observations. Its research and development activities aim to improve the accuracy of our weather forecast services and the efficiency with which they can be produced. This enables its customers to benefit from the progressive international advancement of weather forecasting techniques.
MetO provides the Met Office Hadley Centre Climate Programme which is supported by the Department of Energy and Climate Change (DECC), and the Department for Environment, Food and Rural Affairs (Defra). Their investment provides the core science on which Government can make decisions to help the UK become resilient to climate variability and change, benefit from opportunities for growth, and engage in international climate negotiations. For example, research findings from the programme help ensure cost-effective deployment of renewable energy, and a resilient future for the nation's infrastructure. To achieve this, the Hadley Centre needs a large

production facility to run complex multi-model integrations and ensembles of integrations as well as a resource for research and development. These models can run over periods of months and are time critical to meet deadlines for the customer and for the International Panel for Climate Change (IPCC) producing significant output that needs analysis over long periods of time.
MetO has a long experience in developing successful software infrastructures to support both Weather and Climate scientists and models including archive systems, user interfaces, build and configuration management systems.

Contribution to the specific project

MetO will contribute to the following tasks: Work Package 2, Tasks 2.2 and 2.3. Work Package 3, Task 3.3

Mr Mick Carter (male) has 31 years of experience working in Scientific Software Engineering for the Weather and Climate Communities at the Met Office (MetO). Mick is currently the Strategic Head of Scientific Software for the Met Office Hadley Centre for Climate Prediction. He is also the chair of the Technical Advisory Group for the Unified Model Partnership and a member of the HPC Taskforce for ENES.

David Matthews (male) has 17 years of experience working in Scientific Software Engineering at the Met Office (MetO). He leads the Modelling Infrastructure team containing 4 developers who contribute to development and support the CYLC meta-scheduler and are the primary developers of the Rose and FCM systems.

Professor Nigel Wood (male) has 28 years experience of numerical modelling of geophysical flows, including the design, analysis and implementation of various novel numerical methods. He is currently the Scientific Strategic Head of the Dynamics Research group, responsible for all aspects of the dynamical core of the Unified Model. He is the author of 76 peer reviewed journal papers and holds an honorary visiting professorship at the University of Exeter.

Dr. Michael Rezny (male) has over 20 years experience in High Performance Computing, mostly in the area of Earth Systems Modelling. His main areas of expertise are in performance analysis, benchmarking, and optimising applications running at large scale and he has worked in this role at three international HPC vendors: NEC, Cray, and SGI. He has also developed and presented numerous courses on these topics. More recently, he has been applying these skills at the Met Office on the GungHo / LFRic projects which are developing replacements for the current suite of Unified Models.

Mr Mike Hobson (male) is a computer scientist and software engineer with over 20 years of experience working on a wide variety of different projects including meteorological graphical display systems, defence projects and the Unified Model. He is currently sharing his time between working on the replacement for the Unified Model and working towards producing benchmarks for use with coupled systems.

Publications, and/or products, services or other achievements

1. Earth System Modelling – Volume 5. Tools for Configuring, Building and Running Models. Chapter 4, Configuration Management and Version Control in Earth System Modelling.
   Chapter Authors **Mick Carter and Dave Matthews**. Springer Briefs in Earth System Science Rupert Ford, Graham Riley, Reinhard Budich, Rene Redler.
2. Cascades, backscatter and conservation in numerical models of two-dimensional turbulence. John Thuburn, James Kent and **Nigel Wood**. 2014 QJ 140 pp 626-638

3. Runge-Kutta IMEX schemes for the Horizontally Explicit/Vertically Implicit (HEVI) solution of wave equations. Hilary Weller, Sarah-Jane Lock and **Nigel Wood**. 2013 JCP 252 pp 365-381
4. An inherently mass-conserving semi-implicit semi-Lagrangian discretisation of the deepatmosphere global nonhydrostatic equations. **Nigel Wood,** Andrew Staniforth, Andy White, Tom Allen, Michail Diamantakis, Markus Gross, Thomas Melvin, Chris Smith, Simon Vosper,
   Mohamed Zerroukat and John Thuburn. 2014 QJ 140 pp 1505-1520
5. Numerical analyses of Runge-Kutta implicit-explicit schemes for horizontally-explicit verticallyimplicit solutions of atmospheric models. Sarah-Jane Lock, **Nigel Wood**, and Hilary Weller. 2014 QJ 140 pp 1654-1669
6. Mixing properties of SLICE and other mass-conservative semi-Lagrangian schemes Kohei Aranami, Mohamed Zerroukat and **Nigel Wood**. 2014 QJ 140 pp 2084-2089
2. Using the UM dynamical cores to reproduce idealised 3D flows. N. J. Mayne, I. Baraffe, David M. Acreman, Chris Smith, **Nigel Wood**, David Skålid Amundsen, John Thuburn and David R. Jackson. 2013 GMDD 6 pp 3681-3741
3. The unified model, a fully-compressible, non-hydrostatic, deep atmosphere global circulation model, applied to hot Jupiters. ENDGame for a HD 209458b test case. N. J. Mayne, I. Baraffe, David M. Acreman, Chris Smith, Matthew K. Browning, David Skålid Amundsen, **Nigel Wood**, John Thuburn and David Jackson. 2014 Astronomy and Astrophysics 561
4. GungHo: A code design for weather and climate prediction on exascale machines
5. R. Ford, M.J. Glover, D.A. Ham, C.M. Maynard, S.M. Pickles, G.D. Riley and **N. Wood**. 2013 Submitted to Advances in Engineering Software
6. A mass restoration scheme for limited area models with semi-Lagrangian advection . Kohei Aranami, Terry Davies and **Nigel Wood**. 2014 To appear in the QJ

Projects, and/or activities

- EU FP7 METAFOR
- EU FP7 IS-ENES
- EU FP7 IS-ENES2
- EU FP6 PRISM

Significant infrastructure, and/or major items of technical equipment

The Met Office has recently secured a capital grant for HPC and related funding of £97M for HPC, tape based mass storage, a new computer building, networking and post processing infrastructure that will be installed in phases between 2015 and 2017.The HPC part is expected to grow beyond 350,000 cores and the archive in excess of 300 PBytes.

University of Reading (UREAD)

About the institute

The Department of Meteorology at the University of Reading (UREAD) is the largest in Europe with over 20 teaching staff, 80 research staff and around 50 PhD students. It has received the highest research rating of 5* in all UK Research Assessment Exercises, indicating an international reputation in all aspects of research. It is a member of Reading's Walker Institute for Climate System Research, established to promote integrative research across the University. This is reflected in the long-standing presence of staff from the UK Met Office, and the presence of the Natural Environment Research Council (NERC) funded National Centre for Atmospheric Science (NCAS) and the National Centre for Earth Observation (NCEO). The department also works closely with the European Centre for Medium-Range Weather Forecasts (ECMWF), which is located close to the University.

The Department hosts the Computational Modeling Services group of the UK's National Centre for Atmospheric Science (NCAS-CMS). The NCAS-CMS group at Reading provides modeling support the U.K. academic community for a wide range of climate and earth-system areas on several supercomputer platforms. The scope of support provided is wide ranging, covering areas as diverse as code management; model performance optimization; access to and management of high-performance compute and data services. It currently consists of 12 scientists, several of whom are on grants from NERC and other organizations, with a core of staff funded by NCAS. NCAS-CMS has strong links with the UK Met Office and works closely with them on many aspects of model infrastructure development and deployment. In addition to providing support NCAS-CMS is actively developing software for data processing, analysis, and visualization.

Contribution to the specific project

UREAD's contribution to ESiWACE is primarily in relation to WP3.

Complex modeling systems such as, for example, the UK Met Office Unified Model (UM), require a significant degree of supporting infrastructure to be in place in order to run successfully. Software for data pre-processing to satisfy model input requirements and software for output data postprocessing are just two essential components which must be installed, configured, and supported to run a numerical simulation. A given model itself requires a particular compute environment to function. Different models have evolved to require distinct compute environments, peripheral software, and workflow systems; NCAS-CMS will contribute to the design of a specification which will define a framework in which many modeling systems can be run seamlessly and we shall assist in implementing the specification.

The product will be installed by NCAS-CMS to provide the computational basis for the 3rd ENES summer school. NCAS-CMS has been responsible for installing and maintaining the UM on UK national supercomputers for many years along with additional software installation support and maintenance and consequently has the wealth of experience needed to efficiently complete this task.

Dr. Grenville Lister (male) is Head of NCAS-CMS. He leads a team of highly experienced computational scientists who among them have diverse and long standing expertise in complex numerical model installation, porting, trouble-shooting, and optimization. He was the computational support for the NERC-funded CASCADE project, and currently oversees the computational component of the NERC-funded SWAMMA project. Dr Lister leads the NCAS UM training course delivered biannually to the UK atmospheric research community. He was co-PI on a dCSE (www.hector.ac.uk) project to port and implement asynchronous IO in OpenIFS on the national HPC.

Professor Pier Luigi Vidale (male) is the Willis Chair of Climate System Science and Climate Hazards, Director of the Weather and Climate Hazards Laboratory at UREAD Department of Meteorology, and Senior Scientist at NCAS in Reading. He has over 18 years' international experience in the development of weather and climate models, focused on high-resolution and land surface modeling. Prof. Vidale has led, in partnership with Dr. M. Roberts (Met Office), four global high-resolution climate modeling programs in the UK, Europe and Japan: UK-HiGEM, UKJapan Climate Collaboration (UJCC), Partnership for Advanced Computing in Europe (PRACE)UPSCALE, and Joint Weather and Climate Research Programme (JWCRP) in Global HighResolution (HadGEM3-H). He currently leads a large consortium project, PAGODA, within

NERC's Changing Water Cycle program, and is a co-I in the EU's IS-ENES2 project. Before joining UREAD, he was co-I in two European regional climate programs (MERCURE and PRUDENCE) from which he has published highly-cited papers on changes in the variability of European summers. Prof. Vidale is currently a member of the Gung-Ho Exec, monitoring the development of the next-generation dynamical core for UK Earth System Models.

Publications, and/or products, services or other achievements, projects, and/or activities

- The NCAS-CMS web site (cms.ncas.ac.uk) summarizes activities undertaken by the group, at the level of both core national capability (NCAS funded) and externally funded projects, including model porting (to ARCHER, MONSooN (joint NERC/MO HPC), Polaris(Leeds), Mobilius(Southampton), HPC Wales), trouble-shooting (see NCAS-CMS helpdesk), and training.
- NCAS-CMS has developed UM-specific software tools widely used in the modeling community (xconv and xancil) and is actively developing CF-aware utilities (eg CFpython (http://cfpython.bitbucket.org/)) for use by more diverse research groups.
- NCAS-CMS delivers the PUMA (Providing UM Access) service to the academic community and MO partner. PUMA hosts several code repositories including the UM and OpenIFS, and is the system which manages remote job submission.
- NCAS-CMS has representation on several HPC management groups and committees, including the ARCHER Scientific Advisory Committee, the MONSooN Management Group, and the NERC HPC Steering Committee.

Swedish Meteorological and Hydrological Institute (SMHI)

About the institute

SMHI (http://www.smhi.se) is a government agency under the Swedish Ministry of Environment. SMHI offers products and services that provide organisations with important environmental information to support decision-making. The main fields include weather and climate forecasts/projections, industry-specific services, simulations and analyses. SMHI has a strong R&D focus. With climate research involving all of six research sections, including the Rossby Centre that is responsible for the development and application of regional and global climate models. In particular the Rossby Centre is active in the development of EC-Earth, being responsible for the development and release of the most recent generation, EC-Earth 3. The Rossby Centre also has extensive experience in the development and application of advanced regional climate models.

Contribution to the specific project

SMHI's contribution to this project is dedicated to WP1 and WP2 2.

SMHI will coordinate the efforts to provide community-wide access to the NEMO and EC-Earth models. This will include user support facilities as well as improvements for the scientific software development process. SMHI will contribute to the development of climate model performance metrics and provide performance benchmark results for the EC-Earth model. SMHI's will also assess the performance optimisations for the EC-Earth model. Finally SMHI will be analysing performance enhancements and maintain new developments in forthcoming EC-Earth releases.

Dr. Uwe Fladrich (male) is scientific software developer (education in applied mathematics and computer science) and one of the core developers of the EC-Earth model. His focus lies on efficient software development processes and numerical aspects of climate models. He is contributing to the EU project IS-ENES2, co-leading the work package JRA1 on multi-model, multimember high-resolution experiments.

Publications, and/or products, services or other achievements

1. Koenigk, T., L. Brodeau, R. Grand Graversen, J. Karlsson, G. Svensson, M. Tjernström, U. Willen, and K. Wyser, 2013. *Arctic Climate Change in the 21st Century in an Ensemble of AR5 Scenario Projections with EC-Earth.* Clim. Dyn. 40:2719-2743, doi:10.1007/s00382012-1505-y.
2. Smith D. M., Scaife A. A., Boer G. J., Caian M., Doblas-Reyes F. J., Guemas V., Hawkins E., Hazeleger W., Hermanson L., Ho C. K., Ishii M., Kharin V., Kimoto M., Kirtman B., Lean J., Matei D., Merryfield W. J., Müller W. A., Pohlmann H., Rosati A., Wouters B., Wyser K., 2013. *Real-time multi-model decadal climate predictions.* Climate Dynamics, 41(11-12), 2875-2888, doi:10.1007/s00382-012-1600-0
3. Hazeleger, W., V. Guemas, B. Wouters, S. Corti, I. Andreu-Burillo, F. J. Doblas-Reyes, K. Wyser and M. Caian, 2013. *Multiyear climate predictions using two initialisation Strategies.* Geoph. Res. Lett., DOI: 10.1002/grl.50355.

Projects, and/or activities

- EU FP7 project **IS-ENES2**: Infrastructure for the European Network for Earth System modelling - Phase 2
- EU FP7 project **CLIPC**: Climate Information Platform for Copernicus
- EU FP7 project **SPECS**: Seasonal-to-decadal climate Prediction for the improvement of European Climate Services
- EU FP7 project **EMBRACE**: Earth system Model Bias Reduction and assessing Abrupt Climate change

Significant infrastructure, and/or major items of technical equipment

- SMHI is providing the technical infrastructure for developer-developer and developer-user interaction for the EC-Earth model (the EC-Earth Development Portal)
- SMHI is one of the key laboratories developing and performing large model intercomparison exercises, such as CMIP5.

National University of Ireland Galway/Irish

Centre for High-End Computing (ICHEC)

About the institute

The Irish Centre for High-End Computing (ICHEC) is legally a centre within the National University of Ireland, Galway; PIC code 999978045.
ICHEC, founded in 2005, is Ireland's national high performance computer centre. Its mission is to provide High-Performance Computing (HPC) resources, support, education and training for researchers in third-level institutions and through technology transfer and enablement to support Irish industries large and small to contribute to the development of the Irish economy.

ICHEC works on code optimisation and development of climate and weather codes wirth academia and public organisations, in particular the EC-Earth climate model, where it is a consortium member, and the 'Harmonie' weather model with Met Éireann in the Hirlam consortium.

ICHEC has experience providing operational services for Met Éireann, the national weather service, since 2007. This involves redundant compute and computational scientist support as part of a scientific collaboration, where ICHEC scientists optimise and develop weather and climate codes on next-generation systems. This has recently expanded to include emergency dispersion modelling for the EPA (Environmental Protection Agency) and RPII (radiation), and Dept. of Agriculture (foot and mouth, disease dispersion); Met Éireann and ICHEC have also demonstrated flood forecasting for the Irish Office of Public Works.

ICHEC manages an Earth System Grid Federation (ESGF) portal for climate model data on behalf of the EC-Earth consortium, publishing data on behalf of 14 organisations; we have developed processing workflows and data management systems for this.

Contribution to the specific project

ICHEC's contribution to WP2 is dedicated to WP2 task 3 and 4.

ICHEC will work on the integration of GRIB2 format file output to the XIOS (XML I/O Server) library from IPSL. I/O is a major bottleneck for climate and weather codes, and ICHEC have worked with IPSL, integrating the XIOS library into the EC-Earth climate model, adding memory caching for scaling, and GRIB format writing.
ICHEC plans to add two components: use current and planned changes to XIOS by IPSL in memory layout to enable GRIB writing of large files. Currently GRIB output is limited by the need to do an in-memory transpose, requiring all of a dataset to be in memory simultaneously. Planned changes by IPSL in the server layer will make it possible to complete the transpose over multiple nodes, enabling high-resolution GRIB writing. IHEC as original author of the GRIB code will adapt the GRIB model for this.
Secondly, the current GRIB code is limited to lat-long and simple Gaussian outputs. There is no GRIB equivalent to the NetCDF unstructured grid outputs. We will work with partners (in particular ECMWF) to standardize the GRIB output for unstructured grids, and provide GRIB output in unstructured and icosahedral grids.
Dr Alastair McKinstry is a computational scientist with 15 years experience in HPC and Unix code optimization, originally in industry (Digital, Compaq, Oracle) and weather and climate code development.
He has worked on XIOS development within PRACE 2IP, leading development of the GRIB output, and memory caching to enable XIOS scaling on large low-memory-per-node HPC systems, and adding XIOS to the IFS atmosphere model within EC-Earth. At ICHEC Alastair leads environmental activities, optimizing user and community codes.

Publications, and/or products, services or other achievements

1. Colin O'Dowd, Saji Varghese, Damien Martin, Robert Flanagan, Darius Ceburnis, Jurgita Ovadnevaite, Giovanni Martucci, Jakub Bialek, Ciaran Monahan, Harald Berresheim, Aditya Vaishya, Zachary McGraw Grigas, **Alastair McKInstry** S Gerard Jennings, Baerbel Langmann, Tido Semmler, Ray McGrath (2011): The Eyjafjallajökull Ash Plume–Part 2: Simulating ash cloud dispersion with REMOTE Atmospheric environment
2. J Donners, C Basu**, A McKinstry**, M Asif, A Porter, Eric Maissonave, Sophie Valcke, Uwe Fladrich Performance Analysis of EC-EARTH 3.1, PRACE technical report, 2012
3. P.Nolan, A**. McKinstry**, Scaling Coupled Climate Models to Exascale: OpenACC-enabled EC-Earth3 Earth System Model, 2013, PRACE technical report
4. Shiyu Wang, Ray McGrath, **Alastair McKinstry**, Recent Irish Weather Extreme and Change of Extreme Precipitation Due to Climate Change, EPA Climate Change report, 2010

Projects, and/or activities

- EU FP7 **PRACE:** ICHEC was WP co-leader for climate code development in PRACE2IP,
- **CMIP5 EC-EARTH**: ICHEC managed the data publication for CMIP5 for the EC-Earth climate model in CMIP5, doing code development and optimization, running ensemble runs with Met Éireann, and post-processing ocean data from EC-Earth partners.

Significant infrastructure, and/or major items of technical equipment

- ICHEC's primary HPC facility is "Fionn", a 7680-core SGI ICE X / SGI UV 2000 system (147 Tflop peak) with additional accelerator and high-memory regions. This has 560 TB formatted Lustre storage, and multiple login nodes, including dedicated nodes for NWP and emergency service use
- ICHEC is also part of the eINIS collaboration, managing an Earth System Grid Federation (ESGF) node managing climate model data on 1 PB of storage based at DIAS in Dublin.

Centro Euro-Mediterraneo sui Cambiamenti Climatici (CMCC)

About the institute

The Euro-Mediterranean Center on Climate Change (CMCC; http://www.cmcc.it/) is a research center funded by the Italian Ministries of the Environment and Land Protection, of Education, of University and Research, and of Economy and Finance, that aims at furthering knowledge in the field of climatic variability, including causes and consequences, through the development of highresolution simulations and impact models. CMCC gathers the know-how from its funding Institutions (Istituto Nazionale di Geofisica e Vulcanologia, Università del Salento, Centro Italiano di Ricerche Aerospaziali, Università Ca' Foscari Venezia, Fondazione Eni Enrico Mattei, Università di Sassari, Università della Tuscia, Università degli Studi del Sannio), focusing on climate change issues and applications for environmental management. The mission of CMCC is also to encourage and foster collaboration among universities, national and international research institutions, local institutions and industrial sectors. CMCC represents, at the national and international scale, an institutional point of reference for decision makers, public institutions, as well as private and public companies seeking technical-scientific support. CMCC brings together highly qualified experts from different climate research areas in a single unique institution. The following eight research Divisions[24] work together in an interdisciplinary manner: ASC (Advanced Scientific Computing), CSP (Climate Simulations and Predictions), ECIP (Economic Analysis of Impact and Policy), IAFES (Impacts on Agriculture, Forests and Ecosystems Services), ODA (Ocean Modelling and Data Assimilation), OPA (Ocean Predictions and Applications), RAS (Risk Assesment and Adaptation Strategies), REMHI (Regional Models and Hydrogeological Impacts). Moreover, CMCC hosts in Lecce the Supercomputing Center (180 Tflops of computing power, 1.2 Petabyte on-line storage and 3 PetaBytes Archiving capacity).

The Advanced Scientific Computing (ASC) Division of CMCC carries out Research & Development activities on Computational Sciences applied to Climate Change. In particular, the Division works both on the scalability, the optimization and the parallelization of numerical models for climate change simulations and on the design and implementation of open source solutions addressing efficient access, analysis, and mining of large volumes of scientific data in the climate change domain. In this regard, CMCC provides a big data analytics framework (Ophidia), targeting parallel data analysis on large volumes of scientific/multidimensional data (e.g. multi-terabytes order datasets). The Ocean Modelling and Data Assimilation (ODA) Division focuses on the development of numerical models, methods of data assimilation, production of reanalysis and data sets for global marine forecasts and the study of the interactions between the physical and biogeochemical processes of oceans and the cryosphere in climate variability. The Climate Simulations and Predictions (CSP) Division deals with the development of models of the Earth system, the production of climate predictions and the realization of projections of climate change on scales which range from seasonal to centuries.

CMCC is member of the NEMO Consortium contributing to its System Team, whose main goal is, among the others, the optimizations of NEMO on the computers available in the Consortium. Moreover CMCC participates to the NEMO HPC working group. CMCC is also member of the European Network for Earth System Modelling (ENES) and partner of the Earth System Grid Federation (ESGF), providing access to 100TB CMIP5 data through its data node deployed at the CMCC Supercomputing Center. Finally, CMCC is partner of several EU FP7 and other national and international projects, working on the development of high resolution global and regional climate models, their parallel optimization on manycore clusters as well as on the development of efficient solutions for data management.

Contribution to the specific project

CMCC's contribution to ESiWACE is dedicated to Wp1, Wp2, WP3, WP4.

Prof. Giovanni Aloisio (male) is Full professor of Information Processing Systems at the Dept. of Innovation Engineering of the University of Salento, Lecce, Italy, where is leading the HPC laboratory. Former director of the Advanced Scientific Computing (ASC) Division at the EuroMediterranean Center on Climate Change (CMCC), he is now the Director of the CMCC Supercomputing Center and member of the CMCC Strategic Council. He is also member of the ENES (European Network for Earth System modelling) HPC Task Force. His expertise concerns high performance computing, grid & cloud computing and distributed data management. He has been involved into several EU grid projects such as GridLab, EGEE, IS-ENES1. He has been the responsible for ENES of the EU-FP7 EESI (European Exascale Software Initiative) project chairing the Working Group on Weather, Climate and solid Earth Sciences. He has also contributed to the

[24] *Here is reported the new CMCC organization starting from January 1st, 2015. The CMCC web site (http://www.cmcc.it) reports the old organization.*

IESP (International Exascale Software Project) exascale roadmap. He has been the chair of the European panel of experts on WCES that has contributed to the PRACE strategic document "The Scientific Case for HPC in Europe 2015-2020. Presently, he is coordinating CMCC activities into several EU FP7 projects such as EUBrazilCC, IS-ENES2, CLIP-C and the G8 ExArch. As CMCC, he also the coordinator of the OFIDIA (Operational FIre Danger preventIon plAtform) project in the context of the European Territorial Cooperation Program Greece-Italy 2007-2013. As University of Salento (PRACE Third Party), he is the responsible of the EU-FP7 EESI2 project, chairing the Working Group on Weather, Climate and solid Earth Sciences. He is the author of more than 100 papers in referred journals on high performance computing, grid computing and distributed data management.

Dr. Antonio Navarra (male) graduated in Physics in Bologna in 1980 and returned to Italy in 1986, after getting a Ph.D. at the Geophysical Fluid Dynamics Laboratory at Princeton University. He is Dirigente di Ricerca at the National Institute of Geophysics and Volcanology (INGV), where he carries out his activity in the field of the climate simulation with general circulation numerical models. He is now President of the CMCC Centro Euro-Mediterraneo sui Cambiamenti Climatici. The scientific interests of Dr. Navarra focus on the investigation of the dynamical mechanisms which control climate on the global scale, particularly regarding the natural climate variability of the atmosphere-ocean system on interannual, decadal and centennial scales. The general aim is to understand and document the main modes of climate variability on interannual and decadal scales (teleconnections) by means of statistical methods, numerical simulations and simplified models. These studies are the natural complement of the second activity field, which concerns the simulation and the evaluation of the climate changes using scenarios of the future climate. Dr. Navarra is also teaches in the PhD Program on “Science and Management of Climate Science” at Universita’ Ca’ Foscari, Venice. He is the author of several books and articles of general interest and contributes to national newspapers. He has written: El Nino, Truth and Myths of the Climate phenomenon of the Century (Avverbi, 1997); Weather Predictions (Il Saggiatore, 1996); and, with Andrea Pinchera, The Climate (Laterza, 2000).

Dr. Sandro Fiore (male) *Ph.D.*, is the Director of the Advanced Scientific Computing (ASC) Division of the Euro-Mediterranean Centre on Climate Change (CMCC). His research activities focus on distributed data management, data analytics/mining and high performance database management. He is Visiting Scientist at Lawrence Livermore National Laboratory (LLNL) working at PCMDI in the context of the Earth System Grid Federation. He is the P.I. of the Ophidia research project and he is involved in several EU FP7 (IS-ENES2, CLIP-C, EUBrazilCC) and international (ORIENTGATE, OFIDIA) projects. He is author of more than 60 papers in refereed books/journals/proceedings on parallel and distributed computing. He is editor of the book Grid and Cloud Database Management (Springer, 2011). He is ACM Member.

Dr. Silvia Mocavero (female) *Ph.D.*, is Scientist at the “Advanced Scientific Computing” (ASC) Division of the "Euro-Mediterranean Centre on Climate Change" (CMCC). She received a Ph.D. in 2006 in “Innovative Materials and Technologies” from the ISUFI at the University of Lecce, Italy. Her skills include high performance computing, distributed and grid computing. She is working on the analysis of scalability, optimisation and parallelisation of Earth System Models with a particular focus on NEMO as member of the System Team.

Dr. Simona Masina (female) *Ph.D.* (Princeton University) in Atmospheric and Oceanic Sciences, is Senior Researcher at Istituto Nazionale di Geofisica e Vulcanologia (INGV) and the Head of the “Ocean Modelling and Data Assimilation” (ODA) Division at CMCC. She has more than fifteen years of experience in the field of global ocean modelling, data assimilation and interactions between physical and biogeochemical processes in the climate system. During these years she has been involved in several national and international projects and has been principal investigator for many European projects. Since 2007 she teaches in the Ph.D. Programme in Science and Management of Climate Change at Ca’ Foscari University in Venice. She is member of the Governing Board of the “Italian Society for Climate Sciences” (SISC), and INGV representative at the "Commissione Oceanografica Italiana" (Italian IOC). She is also member of the CLIVAR Working Group on Ocean Model Development (WGOMD) and coordinator of the PRACE Project Ens4Ocean.

Dr. Silvio Gualdi (male) is Senior Scientist at CMCC, where he leads the "Climate Simulations and Predictions" Division (CSP). He holds a PhD in Geophysics and has more than 15 years of experience in climate modelling and simulations. During this period, he has contributed to the development of climate models (e.g., SINTEX, CMCC-Med, CMCC-CM) and seasonal-to-decadal prediction systems (e.g DEMETER, MERSEA, ENSEMBLES, COMBINE). He has been principal investigator and WP in several international projects.

Publications, and/or products, services or other achievements

1. **S. Fiore**, A. D'Anca, D. Elia, C. Palazzo, I. Foster, D. N. Williams, **G. Aloisio**, "Ophidia: a full software stack for scientific data analytics", IEEE Workshop on Big Data Principles, Architectures & Applications, HPCS2014, Bologna, Italy, July 21-25, 2014.
2. **S. Fiore**, C. Palazzo, A. D'Anca, I. T. Foster, D. N. Williams, **G. Aloisio**, "A big data analytics framework for scientific data management", IEEE BigData Conference 2013, Santa Clara, USA, 6-9 Oct. 2013, pp. 1-8, doi: 10.1109/BigData.2013.6691720.
3. **S. Fiore**, A. D'Anca, C. Palazzo, Ian T. Foster, Dean N. Williams, **Giovanni Aloisio**: Ophidia: Toward Big Data Analytics for eScience. ICCS 2013, June 5-7, 2013 Barcelona, Spain, ICCS, volume 18 of Procedia Computer Science, pp. 2376-2385. Elsevier, (2013).
4. L. Cinquini, D. Crichton, C. Mattmann, J. Harney, G. Shipman, F. Wang, R. Ananthakrishnan, N. Miller, S. Denvil, M. Morgan, Z. Pobre, G. M. Bell, C. Doutriaux, R. Drach, D. Williams, P. Kershaw, S. Pascoe, E. Gonzalez, **S. Fiore**, R. Schweitzer, The Earth System Grid Federation: An open infrastructure for access to distributed geospatial data, Future Generation Computer Systems, Volume 36, July 2014, pp. 400-417, ISSN 0167-739X, http://dx.doi.org/10.1016/j.future.2013.07.002.
5. Epicoco I, **Mocavero S**, **Aloisio G** (2012). The performance model for a parallel SOR algorithm using the red-black scheme. INTERNATIONAL JOURNAL OF HIGH PERFORMANCE SYSTEMS ARCHITECTURE, vol. 4, p. 101-109, ISSN: 1751-6528.
6. Epicoco I, **Mocavero S**, **Aloisio G** (2011). The NEMO oceanic model: Computational performance analysis and optimisation. In: Proceedings of the 2011 IEEE International Conference on High Performance Computing and Communications (HPCC 2011). Banff, Canada, Sep 2-4, 2011, p. 382-388, Los Alamitos, CA:IEEE Computer Society, ISBN: 978076954538-7, doi: 10.1109/HPCC.2011.56.
7. J. Dongarra, P. Beckman, T. Moore, P. Aerts, **G. Aloisio**, J. C. Andre, D. Barkai, J. Y. Berthou, T. Boku, B. Braunschweig, F. Cappello, B. M. Chapman, X. Chi, A. N. Choudhary, S. S. Dosanjh, T. H. Dunning, **S. Fiore**, A. Geist, B. Gropp, R. J. Harrison, M. Hereld, M. A. Heroux, A. Hoisie, K. Hotta, Z. Jin, Y. Ishikawa, F. Johnson, S. Kale, R. Kenway, D. E. Keyes, B. Kramer, J. Labarta, A. Lichnewsky, T. Lippert, B. Lucas, B. Maccabe, S. Matsuoka, P. Messina, P. Michielse, B. Mohr, M. S. Mueller, W. E. Nagel, H. Nakashima, M. E. Papka, D. A. Reed, M. Sato, E. Seidel, J. Shalf, D. Skinner, M. Snir, T. L. Sterling, R. Stevens, F. Streitz, B. Sugar, S. Sumimoto, W. Tang, J. Taylor, R. Thakur, A. E. Trefethen, M. Valero, A. van der Steen, J. S. Vetter, P. Williams, R. Wisniewski, K. A. Yelick: "The International Exascale Software Project roadmap". International Journal of High Performance Computing Applications (IJHPCA) 25(1): 3-60 (2011), ISSN 1094-3420, doi: 10.1177/1094342010391989.
8. Danabasoglu, G., Yeager, S. G., Bailey, D., Behrens, E., Bentsen, M., Bi, D., Biastoch, A., Bning, C., Bozec, A., Canuto, V. M., Cassou, C., Chassignet, E., Coward, A. C., Danilov, S., Diansky, N., Drange, H., Farneti, R., Fernandez, E., Fogli, P. G., Forget, G., Fujii, Y., Gries, S. M., Gusev, A., Heimbach, P., Howard, A., Jung, T., Kelley, M., Large, W. G., Leboissetier, A.,

Lu, J., Madec, G., Marsland, S. J., **Masina, S**., **Navarra, A.**, Nurser, A. G., Pirani, A., y Mlia, D. S., Samuels, B. L., Scheinert, M., Sidorenko, D., Treguier, A.-M., Tsujino, H., Uotila, P., Valcke, S., Voldoire, A., Wang, Q., (2014) North Atlantic simulations in Coordinated Ocean-ice Reference Experiments Phase II (CORE-II). Part I: Mean states. *Ocean Modelling,* 73, 76-107. doi:10.1016/j.ocemod.2013.10.005.

9. Cherchi, A., Annamalai, H., **Masina, S.**, **Navarra, A.** (2014) South Asian summer monsoon and the eastern Mediterranean climate: the monsoon-desert mechanism in CMIP5 simulations. Journal of Climate, 27(18), 6877-6903 doi:10.1175/JCLI-D- 13-00530.1.
10. Colleoni, F., **Masina, S.**, Cherchi, A., **Navarra, A.**, Ritz, C., Peyaud, V., and Otto-Bliesner, B. (2014) Modelling Northern Hemisphere ice sheets distribution during MIS5 and MIS7 glacial inceptions. Climate of the Past, 10 (1), 269-291.
11. Bellucci A., **S. Gualdi**, **S. Masina**, A. Storto, E. Scoccimarro, C. Cagnazzo, P. Fogli, E. Manzini and **A. Navarra**, (2013) Decadal Climate Predictions with a coupled OAGCM initialized with oceanic reanalyses. Climate Dynamics: Volume 40, Issue 5, Page 1483-1497.
12. Alessandri A., A. Borrelli, **S. Gualdi**, E. Scoccimarro, **S. Masina**, (2011) Tropical cyclone count forecasting using a dynamical Seasonal Prediction System: sensitivity to improved ocean initialization. Journal of Climate, Vol. 24, (12), 2963-2982, doi: 10.1175/2010JCLI3585.1 13. Alessandri A., A. Borrelli, **S. Masina**, A. Cherchi, S. **Gualdi, A. Navarra, P.** Di Pietro, A.F. Carril (2010) The INGV-CMCC Seasonal Prediction System: Improved Ocean Initial Conditions. Mon Wea Rev, Vol. 138, No. 7. 2930–2952.

Projects, and/or activities

- **EU FP7 Project EUBRAZILCC (2013-2016).** CMCC is working on the development of a PaaS level for big data analytics for climate change, exploiting scalable and dynamic VM-based solutions. CMCC is also leading the WP3 on the Federated Infrastructure of the project, as well as the Use Case 3 (WP5) on Biodiversity and Climate Change, working (i) on the use of observed and simulated data to better understand how climate change impacts on terrestrial biodiversity and (ii) on the implementation of a Science Gateway for scientific, multidimensional data analysis (climate and satellite data).
- **EU FP7 Project CLIP-C (2013-2016).** CMCC is working in this project to provide an interoperable interface (e.g. OGC Web Processing Service compliant) to analysis services for climate data and a caching systems for climate indicators.
- **INTERREG Project OFIDIA (2013-2015).** CMCC is providing a specialized data analysis platform for fire danger prevention, by working (i) on the implementation of specific fire danger indices and their integration into big data frameworks as well as (ii) on the development of a Science Gateway for fire danger data analysis. The main use cases concern the Apulia-Epirus cross-border regions.
- **EU FP7 Project IS-ENES (2009-2013) and its follow on IS-ENES2 (2013-2017).** The goal of IS-ENES is the development of a common climate and Earth system modelling distributed research infrastructure in Europe. IS-ENES2 further integrates the European climate modelling community, stimulates common developments of software for models and their environments, fosters the execution and exploitation of high-end simulations and supports the dissemination of model results to the climate research and impact communities. CMCC is working on the use of exascale to develop next generation climate models and on the data node monitoring infrastructure.
- **EU FP7 Project EESI1 (2010-2011) and its follow on EESI2 (2012-2015)**. The goal of these projects is the production of the European Roadmap on Exascale. CMCC is chairing the WP3/Task3.2 on Weather, Climatology and solid Earth Sciences.
- **G8 Project ExArch (2011-2014).** As partner of this project, CMCC has been involved in evaluating/extending/testing climate data analysis services/frameworks in use cases at large scale.
- **PRACE Preparatory Access "Optimisation of the NEMO oceanic model in the GLOB16 configuration" and its follow on 8th PRACE Project "ENS4OCEAN" (2014-2015).** CMCC

aims to produce a simulation of an eddy-resolving global ocean (NEMO at 1/16° horizontal resolution and 98 vertical levels), which will be the base for a real-time forecasting system able to provide on daily basis forecasts of global oceanographic parameters for the following 10 days. Starting from the hindcast for a few selected dates, an ensemble of perturbed forecasts will be also produced. This ensemble will form the basis for short-range predictability studies and for estimating ensemble-derived background-error covariances for further use in variational data assimilation experiments. The goal of the preparatory access has been the design and implementation of suitable strategies to overcome the scalability bottlenecks to improve the parallel efficiency.

- **EU FP7 Project MyOcean2 and MyOcean FollowOn (2012-2015)**. CMCC is working on the production of a set of validated global ocean re-analysis for the physical and biogeochemical state of the ocean over multidecal period. CMCC is leading the WP18 concerning the production and assessment of reprocessed satellite and in-situ data sets and re-analysis of the 3D ocean state for Global and European Regional Seas. CMCC participates also to WP4, WP10 and WP19 where contributes to the improvement of data assimilation schemes with the development of ensemble protocols to specify more realistic background error covariance models and analysis uncertainties.
- **EU FP7 Project ERA-Clim2 (2014-2016).** It is a collaborative research project funded by the European Union, with the goal of preparing input data and assimilation systems for a new global coupled reanalysis of the 20th century. CMCC will contribute to ERA-CLIM2 by developing and testing hybrid variational and ensemble ocean data assimilation systems and exploring the impact of coupled ocean-atmosphere model error covariances to correct nearsurface meteorological fields through ocean data assimilation.

Significant infrastructure, and/or major items of technical equipment

Infrastructure: The HPC infrastructure managed at the CMCC Supercomputing Centre is composed of a 960 cores IBM Power6 cluster (peak performance 18TFlops) and a 8000 cores Intel Xeon Sandy Bridge (peak performance 160TFlops). Part of this infrastructure will be used for running data management/analysis services and training activities.

Datasets: CMCC publishes about 100TB of climate simulations datasets in the CMIP5 federated data archive related to the following three models: CMCC-CM, CMCC-CESM, and CMCC-CMS.

Software: CMCC provides the Ophidia software, a cross-domain big data analytics framework for the analysis of scientific, multi-dimensional datasets. This framework exploits a declarative, serverside approach with parallel data analytics operators.

Deutscher Wetterdienst (DWD)

About the institute

The Deutscher Wetterdienst (DWD), which was founded in 1952, is as National Meteorological Service of the Federal Republic of Germany responsible for providing services for the protection of life and property in the form of weather and climate information. This is the core task of the DWD and includes the meteorological safeguarding of aviation and marine shipping and the warning of meteorological events that could endanger public safety and order. The DWD, however, also has other important tasks such as the provision of services to Federal and Regional governmental authorities, and the institutions administering justice, as well as the fulfilment of international commitments entered into by the Federal Republic of Germany. The DWD thus co-ordinates the meteorological interests of Germany on a national level in close agreement with the Federal

Government and represents the Government in intergovernmental and international organisations as, for example the World Meteorological Organization (WMO). Currently DWD has a total staff of about 2300 employees at more than 130 locations all over Germany. DWD's spectrum of activity is very wide and comprises of:

- Weather observation and forecasting around the clock,
- Climate Monitoring and modelling at local, regional and global scale,
- Development of precautionary measures to avoid weather-related disasters and to provide support for disaster control
- Advice and information on meteorology and climatology to customers,
- National and international co-operation in meteorological and climatological activities,
- Outlooks on possible future climatic conditions at local, regional and global scale,
- Research and development.

Contribution to the specific project

DWD's contribution to CoE Weather and Climate is dedicated to Task 1 in Work Package 2:
DWD will organize workshops on model I/O for Exascale simulations. It combines mid-term objectives, e.g. exploiting the available hardware bandwidth through I/O servers, with longer-term I/O developments beyond MPI. Among the design goals is the provision of a unified interface for model I/O which transparently maps between available file formats. As a side-effect of the project, the participating modelling groups will obtain the abilities to accurately analyze the model codes I/O performance, to predict their scalability and to design proper benchmarks mimicking larger models in reduced environments.

Dr. Florian Prill (male) is a senior scientist in the Numerical Modelling Department at DWD. He is responsible for the design, coding and optimization of the asynchronous, parallel I/O packages (for GRIB2 and NetCDF) which are part of the new non-hydrostatic global model ICON.

Publications, and/or products, services or other achievements

1. **Günther Zängl,** Daniel Reinert, Pilar Rípodas, Michael Baldauf. (2014) The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core. Quarterly Journal of the Royal Meteorological Society, n/a-n/a. Online publication date: 1-May-2014.
2. **Michael Baldauf,** Daniel Reinert, Günther Zängl. (2014) An analytical solution for linear gravity and sound waves on the sphere as a test for compressible, nonhydrostatic numerical models. Quarterly Journal of the Royal Meteorological Society, n/a-n/a. Online publication date: 1-Feb-2014.

Projects, and/or activities

- **G8 project ICOMEX** (Icosahedral-grid models for exascale Earth system simulations) from 2011 to 2014: Development of key components for global atmospheric models (parallel I/O, parallel internal postprocessing, compute optimization / GPU usage) in order to prepare them for applications on future massively parallel computer architectures. G. Zängl (DWD) was Project Leader of ICOMEX.

Significant infrastructure, and/or major items of technical equipment

The German weather service DWD has long-standing experience in the field of high performance computing, dating back as far as the 1960's. Recently, two independent Cray XC40 supercomputers have been installed at the German Meteorological Computing Centre in Offenbach, each system with more than 500 TFlop/s reaching rank 128 on the current Top500 list (Nov 2014).
Daily runs of a comprehensive numerical weather prediction suite produce more than 4 TByte of data/day. Meteorological data is moved between a 3 PByte file system and an archive with 50 PByte capacity and products are distributed to numerous institutions world-wide..

Seagate Systems UK Ltd (SEAGATE)

About the institute

Seagate is the worlds leading provider of Data Storage devices, equipment and services.
The organisation is a worldwide multi-national registered in Ireland (Seagate Technology plc) with more than 56,000 employees; the division of the organisation responsible for this project is Seagate Systems UK Ltd.
Seagate operates two primary divisions within its corporate operations, Seagate Technology develops and produces data storage devices including disk drives, solid state drives and solid state storage for use in applications from consumer to extreme performance HPC, a large facility is located in Northern Ireland. The newer division, Seagate Systems has been created following the acquisition of Xyratex Technology by Seagate in April 2014, combining this organisation with EVault and other internal Seagate systems activities to create a high capability storage systems supply organisation.
A key product line acquired from Xyratex and continuing with Seagate is the ClusterStor range of products, these are fully engineered data storage systems with all hardware, file systems software and system management provided. Systems are provided through our OEM or business partnerships including with BULL. These systems support some of the worlds most powerful supercomputers. The systems are installed or planned in a number of installations in Europe including ones in Met Office UK, EPCC, DKRZ and ECMWF with capacities of up to 45 Petabytes and >1.4 TB/s performance in future deployments.
The Seagate systems (ex Xyratex) group has around 500 engineers employed in creating Hardware and supporting Software for the Enterprise and High Performance Computing applications. Seagate owns the Lustre trademark and several of their engineers were involved in the original Lustre architecture and design.
Within Seagate Systems (UK) the Emerging Technology Group manages collaborative research activities within Europe and will work in concert with development engineering groups based in UK.

Contribution to the specific project

Seagate is happy to take part in ESiWACE as we see significant opportunity for business growth based around its success.

Seagate's skills have been harnessed to create a next generation object storage technology with capabilities well beyond any similar solutions on the market. Seagate will provide instances of this storage technology and develop native support software that enables the NETCDF and GRiB data formats to be efficiently stored and accessed, providing multiple 'views' of the stored data. Seagate will also contribute its deep skills and knowledge of current and future data storage technologies to assist the study of optimized data management by the community.
Seagate Systems is a key supplier of data storage systems (in partnership with Other Equipment Manufacturers) with installations of their equipment in a number of sites of partners within this proposal. We are keen to work much more closely, understanding the user needs and specific opportunities to create or tune systems to maximize the effectiveness of our systems in these user environments.

Andy Nimmo (male): Andy is a Principal Engineer in Systems Design and Systems Integration for Seagate. Andy holds a BEng (Hons) in Software Engineering and joined Seagate from Adaptive Computing in January 2014. He has 10 years experience working in both private and public sectors of the the HPC sector and has extensive experience from systems administrator level all the way up to system architecture, consultancy and security. After initially working as a QA Engineer focussing around networking and kernel comms on the ASCI Q project in 2003 he spent some time as a senior software engineer before moving to system management and workload scheduling back in HPC space. Since joining Seagate Andy has been chiefly involved with the next-generation High Availability project but more recently has been tasked with being in charge of systems integration for Seagate's next-gen systems product and is heavily involved with both scoping and the architecture of many aspects of this project.

Dr. Sai Narasimhamurthy (male): Sai is currently Staff Engineer, Seagate Research (formerly Lead Researcher, Emerging Tech, Xyratex) working on Research and Development for next generation storage systems (2010-). He has also actively led and contributed to many European led HPC and Cloud research initiatives on behalf of Xyratex (2010-). Previously(2005 - 2009) , Sai was CTO and Co-founder at 4Blox, inc, a venture capital backed storage infrastructure software company in California addressing IP SAN(Storage Area Network) performance issues as a software only solution. During the course of his doctoral dissertation at Arizona State University (2001-2005), Sai has worked on IP SAN protocol issues from the early days of iSCSI(2001). Sai also worked with Intel R&D and was a contributing participant in the first stages of the RDMA consortium (put together by IBM, Cisco and Intel) for IP Storage and 10GbE (2002). Earlier in his career, Sai worked as Systems Engineer with Nortel Networks through Wipro, India focussing on Broadband Networking solutions(2000-2001).

Malcolm Muggeridge (male): Malcolm is Sr Dir Engineering responsible for collaborative research at Seagate Systems UK. He joined Seagate through its acquisition of Xyratex in 2014 and was with Xyratex at its creation as a management buyout from IBM in 1994.
Malcolm has more than 38 years experience through his employment with IBM and Xyratex in the Technology, manufacturing, quality and reliability of Disk drives and Networked data storage systems and in recent years in HPC data storage, architecting and managing designs and new technologies across many products. More recently he has been focused on Strategic Innovation and Business development, Research & Technology. He is a steering board member of the ETP4HPC defining research objectives for future within Euorpe and is active in the Partnership board of the cPPP on HPC. He is a member of the UK eInfrastructure board with Special interest in HPC. Malcolm has a B.Eng degree in Electronics from Liverpool University.

Dr Nikita Danilov (male): Nikita Danilov is a Consultant Software Architect at Seagate. His work on storage started in 2001, when he joined Namesys to develop the reiserfs file system for Linux. Since 2004 he worked on Lustre in ClusterFS, later acquired by Sun. In 2009 he followed the original Lustre architect—Peter Braam—to the latter's new company Clusterstor to design and implement an exascale storage system, this technology was acquired by Xyratex and forms the basis of the NEXT system. He received a PhD in mathematical cybernetics from Moscow Institute of Physics and Technology.

Publications, and/or products, services or other achievements

As a commercial organisation Seagate does not generally submit material for publication in academic journals or to conferences however they do present publicly on selected technical aspects of the systems and solutions with presentations at major events such as Supercomputing, ISC and events such as Lustre developer conferences LAD and LUG. As stated earlier Seagate is an active supplier within this and other HPC use domains.

Projects, and/or activities

Seagate Systems is deeply involved with the strategy toward HPC in Europe with active membership of ETP4HPC and contribution to the Strategic Research Agenda.
Seagate Systems has been involved in a number of FP7 projects including leading IRMOS; creating Quality of Service capability for storage in 'real time' cloud systems and currently is a member of the DEEP-ER project particular focused on improved IO guidance mechanisms. The organisation also has involvement in a number of Research projects in the area of Optical Interconnects including PHOXTROT.
Seagate is also a supporter of educational activities with Early Stage Researcher development through the SCALUS project and recently awarded BIGSTORAGE Marie Curie, Initial training Network project.

Significant infrastructure, and/or major items of technical equipment

Seagate Systems has within its development operations some medium scale storage systems linked to small scale computational capabilities for the evaluation and test of new storage hardware and software. For this project this facility will be utilised to explore the characteristics of IO with new storage techniques.

Bull SAS (BULL)

About the institute

Bull is the trusted partner for enterprise data. The Group, which is firmly established in the Cloud and in Big Data, integrates and manages high-performance systems and end-to-end security solutions. Bull's offerings enable its customers to process all the data at their disposal, creating new types of demand. Bull converts data into value for organisations in a completely secure manner. Bull currently employs around 9,200 people across more than 50 countries, with over 700 staff totally focused on R&D. In 2013, Bull recorded revenues of €1.3 billion with a particularly

strong presence in the public, healthcare, finance, telecommunications, manufacturing and defence sectors.
Bull has organized its activities into two major segments, Data Management and Data Infrastructure, highlighting how technological issues are now so closely intertwined. Whether in a supercomputer, a mobile application or an embedded M2M solution, performance and security issues depend at least as much on hardware components as they do on software. Because the most significant innovations often arise from an ingenious combination of these two aspects, Bull has chosen to bring together all its R&D activities to support its entire business.
Closely attuned to changes in technology and business, Bull R&D is involved in all key areas of the Group's activities: high-performance computing, cyber-security, Cloud computing, information systems modernization, Big Data. In recent years, the Bull R&D labs have developed many major products that are recognized for their originality and quality. These include the bullx supercomputer, bullion servers for the private Clouds and Big Data, the Shadow intelligent jamming system designed to counter RCIEDs, the libertp tool for modernization of legacy applications and, most recently, hoox, the first European smartphone featuring native security. To explore new areas and develop tomorrow's solutions, today, Bull R&D is investing heavily in customers – with whom it has forged many successful technological partnerships – as well as in institutional collaborative programs (such as competitiveness clusters and European projects) and in partnerships with industry (Open Source, consortiums).
On the long-term, the objective of Bull is to create, with the CHANCE project, market differentiators. The main areas are the efficient support of applications and optimization, new job scheduler strategies, and better performance at node, interconnect and I/O level. With these differentiators, Bull aims at developing its market share in the supercomputer segment worldwide and especially in Europe. In its HPC development strategy, software is a key element, hence the CHANCE project is essential for Bull.

Contribution to the specific project

BULL will take part in WP1 and WP2.

The people involved in this project are located in the CEPP (Center for Excellence in parallel Programming). This center is based in Grenoble and its activity consists in porting, profiling and optimizing applications for their efficient use of a parallel computers. Their everyday duty requires a high expertise in the computer architectures, the trends in their evolutions, and the impact on the software. Today, the constraints on the hardware are no longer transparent for the software. To get benefits from the coming architecture, the software has to be deeply studied and modified. The notion of co-design, without being strictly mandatory becomes highly necessary.
Beside their high skills in parallel programming, most of the CEPP experts have a Ph.D. in a scientific domain: molecular dynamic, chemistry, oceanography, fluid dynamics, astrophysics, … Thanks to these two scientific pillars (HPC and science), the experts in the CEPP are able to understand the behavior of the application and also the goals and needs of the science.
The experts who will be involved in this project will be selected thanks to these two background. It is clear that people having oceanography background would bring more benefits to the project, but, some others experts may have a specific knowledge that would be necessary (co-processors, IO, network, …). Thus, this will not be one expert who will implement the work, but probably several different according to the needs.
However, an overall envelope can be agreed and the dedicated expert will be the best possible one.

Dr Xavier Vigouroux (male) after a Ph.D. from Ecole normale Superieure de Lyon in Distributed computing, worked for several major companies in different positions. He has now been working for bull for 9 years. He led the HPC benchmarking team for the first five years, then in charge of the "Education and Research" market for HPC at Bull, he is now managing de "Center for Excellence in Parallel Programming" of Bull.

Dr Cyril Mazauric (male) obtained his Ph.D. in Applied Mathematics from the University of Grenoble concerning “Data assimilation applied to flood modelling”. Since 2008, he is part of the Application and Performance Team that is dealing with applications performance commitments during bidding process. However, Cyril activities lead him to be part of research projects, and, more specifically around NEMO performance profile, its coupling with WRF … Finally, he has been contributor to the scalability improvement of Meteo France and DKRZ applications.

Franck Vigilant (male) received his Applied Mathematics M.S degree in Numerical simulation and Modelling Engineering in 2008 from the Joseph Fourier University of Grenoble and his M.S. degree in Process Engineering in 1999 from the National Polytechnique Institute of Grenoble. After a 7-year position as expert engineer at Philips semiconductors and NXP R&D, a IC manufacturer, he worked as expert developer at the National Institute for Computer Science and Control (INRIA). His work was related to data assimilation applied to ocean modelling. He is now HPC consultant at BULL in the Applications and Performances team for HPC business. His interests includes high performance computing for environmental sciences.

Enguerrand Petit (male) received a HPC Computer Master in 2014 from INRIA Bordeaux in Runtime team supervised by Denis Barthou and Olivier Aumage. His skills are centered on codes optimization on ARM and Intel processor in the European Mont Blanc project and the impact of the vectorization on energy consumption.Enguerrand also worked on optimizing NEMO as part of a project between Bull and Intel.

Publications, and/or products, services or other achievements

The CEPP has been involved in different **code optimizations**:

1. **Cardiff “fast start”**: when bull delivered the machine to Cardiff, some codes have to be ported to this new machines. This list had been selected by Cardiff staff and Bull had to make it available and efficient on the new machine.
2. **P3M - Université Reims Champagne Ardennes**: The problem here was to optimize a workload aiming at find the interaction between one ligand to thousands of protein. This method is called “inverse docking” and requires a huge computation resources. CEPP experts studied the implemented and optimized the throughput (by 9 times) and create a very scalable implementation.
3. CEPP experts are involved in a **DKRZ optimization service**. They have to optimize the DKRZ code of the newly acquired machine. This project is in his starting phase.
4. **NEMO optimization analysis** on Xeon Phi for intel.

Projects, and/or activities

Bull Application experts are involved in all the deals answered by Bull and requiring application level commitments: DKRZ, AEMET, Meteo France, …

Besides, they have worked on projects linked with NEMO (funded or as provided services): porting on Intel Xeon Phi, application tuning, IO analysis, scalability analysis, … The team was involved in the PULSATION project (funded by ANR in France). If we focus on the most three relevant projects

- Project 1: PULSATION Project: This project is funded by ANR French Agency. Its goal is to create a coupled (ocean atmosphere) multigrid simulation. This multigrid simulation makes it possible to focus on very precise zone with relevant effect (for instance el nino and el nina). Then, this precise zone is embedded in a coarse one for the holistic simulation.

- Project 2: DKRZ. DKRZ has selected bull has their supercomputer provider. Beside the material, they asked bull to work on the application to prepare the future evolution. The work load is around 2PY
- Project 3: NEMO Ported on Phi: Intel has asked bull to work on NEMO and port it on Phi. NEMO is very complex and the approach in this project is to focus on a very small part of the code (few functions, 10 loopnest) but makes it very efficicient on Phi

Significant infrastructure, and/or major items of technical equipment

Different supercomputers are part of the CEPP; they are selected according to the different needs: reproducibility, ability to be modified. In the CHANCE context, the experiments will be done on targeted hardware provided by the project. CEPP is used to integrate, modify and give access to the hardware, they can host it a reliable infrastructure (login nodes, storage, …).
For instance, we have pure CPU nodes, nodes with latest GPU and accelerators. Besides, in terms of interconnect, storage, software stack; we are able to build configurations with a lot of flexibility. This will bring to the project the ability to rely on an flexible infrastructure administrated by professional sysadmin.
The size of the machines evolves continuously, but, there are generally, around one hundred of CEPP nodes up and running.

Allinea Software Limited (ALLINEA)

About the institute

Allinea Software is a UK SME providing world-leading development tools and application performance analytics software and training for high performance computing (HPC).

Its headquarters and Research and Development groups are in the UK, and it has sales or technical operations in France, Germany, Canada and the USA. It has a customer base throughout the world, including over 75% of Europe’s leading HPC centres. Allinea's integrated profiling and debugging developer tools are relied on in fields ranging from climate modeling to astrophysics, and from computational finance to engine design - and together provide proven capability beyond Petascale computing.

Allinea provides training in the best practices of software development - including using debugging and profiling tools - enabling scientists to develop more maintainable and efficient software faster. This training is provided regularly including at PRACE, university and lab workshops.

Contribution to the specific project

Allinea will contribute to WP3.

Within the ESIWACE Allinea will bring its expertise in training of best practices for software engineering and development tools to HPC community and work with the climate community to help it develop excellent software to the benefit of the community and society.

Publications, and/or products, services or other achievements

Florent Lebeau (male) is an HPC Applications and Support Analyst at Allinea and develops and provides training in software tools. Being involved in HPC for many years, he has expertise in parallel programming, optimization, tools and training. Before joining Allinea, Florent graduated from the University of Dundee with an MSc in Applied Computing and has worked for CAPS entreprise, were he developed profiling tools for HMPP Workbench and provided training on parallel technology. Florent will contribute to WP 3.

4.2 Third parties involved in the project (including use of third party resources)

4.2.1 Subcontracting

Not applicable.

4.2.2 Work is performed by linked third parties (Art.14 of the Grant Agreement)

Centre National de la Recherche Scientifique (CNRS), beneficiary Nr. 2, foresees to have part of the work done by its linked third parties:

- Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA)
- Université Pierre et Marie Curie (UPMC)

Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA)
CEA is associated with CNRS within the Joint Research Unit LSCE which is part of IPSL.

CEA will be involved in WP2 for tasks concerning the I/O server XIOS, a software library dedicated to efficient IO management for climate models (Tasks 1.2, 2.2, 2.3, 3.2, 4.1).

The main staff of CEA to be involved in these tasks are two engineers: Yann Meurdesoif and Arnaud Caubel. They will lead the support and development associated with XIOS.

CEA as a third party to beneficiary number 2 CNRS will request the following costs for the task described (Annex 2 to the Grant Agreement):

- Direct personnel costs declared as actual costs: 32.000 Euro (corresponding to 4 personmonths)
- Other direct costs: 0 Euro
- Indirect costs: 8.000 Euro
- Total direct costs: 40.000 Euro

Université Pierre et Marie Curie (UPMC)
UPMC is associated with CNRS within the Joint Research Unit LOCEAN which is part of IPSL.

UPMC will contribute to the support and developments activities on the European ocean platform NEMO in WP2 (Tasks 1.1 and 3.1).

The main staff of UPMC to be involved in these tasks is: **Dr Sébastien Masson** (male) is a member of the NEMO Developing Team for 10 years and will participate to these activities.

UPMC as a third party to beneficiary number 2 CNRS will request the following costs for the task described (Annex 2 to the Grant Agreement):

- Direct personnel costs declared as actual costs: 8.000 Euro (corresponding to 1 personmonths)
- Other direct costs: 0 Euro
- Indirect costs: 2.000 Euro
- Total direct costs: 10.000 Euro

4.2.3 Third Parties providing in-kind contributions (Article 11 and 12)

Institució Catalana de Recerca i Estudis Avançats (ICREA)

ICREA, as a third party to beneficiary number 6 BSC, will provide in-kind contributions free of charge to BSC (Article 12 Grant Agreement).

ICREA is a foundation supported by the Catalan Government and guided by a Board of Trustees which aims to recruit top scientists for the Catalan R&D system: scientists capable of leading new research groups, strengthening existing groups, and setting up new lines of research.
Following the rules of ICREA, although the salary costs of Dr. Doblas-Reyes are paid by ICREA, he is assigned to physically work at the Earth Sciences Department of the BSC and considered a full member of the BSC. The terms and conditions of this cooperation between ICREA and BSC are reflected in a bilateral agreement between the two parties.
The beneficiary, BSC, is free to use these resources at will. They are therefore assimilated as "own resources" of the beneficiary, and will be charged to the project without being considered as a receipt. The cost will be declared by the beneficiary and it will be recorded in the accounts of the third party. These accounts will be available for auditing if required.

Key personnel
Prof. Francisco Doblas-Reyes (male) is an expert in the development of seasonal-to-decadal climate prediction systems and the head of the ES-BSC. He is involved in the development of the EC-Earth climate forecast system since its inception. He was an IPCC lead author in the Fifth Assessment Report, serves in WCRP and WWRP scientific panels, is a member of the ENES HPC Task Force, has participated in a number of FP4 to FP7 projects and is author of more than 100 peer-reviewed papers. He is shaping BSC's plans for the development of a weather and climate modelling service that brings the latest developments of HPC and Big Data research to the Earth science community, increasing at the same time the resilience of the European society to weather, air quality and near-term climate extremes.

5 Ethics and Security

5.1 Ethics

There are no ethical issues foreseen during the ESiWaCE project.

5.2 Security

ESiWaCE project will not involve activities or results raising security issues.
ESiWaCE project will not involve 'EU-classified information' as background or results.

Digitally sealed by the European Commission
Date: 2016.10.28 12:35:08 CEST

This document is digitally sealed. The digital sealing mechanism uniquely binds the document to the modules of the Participant Portal of the European Commission, to the transaction for which it was generated and ensures its integrity and authenticity.

Any attempt to modify the content will lead to a breach of the electronic seal, which can be verified at any time by clicking on the digital seal validation symbol.