EUROPEAN COMMISSION

Directorate General for Communications Networks, Content and Technology

eInfrastructure

ANNEX 1 (part A)

Research and Innovation action

NUMBER — 675191 — ESiWACE

Table of Contents

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 month[8]	Project exit month
1	DEUTSCHES KLIMARECHENZENTRUM GMBH	DKRZ	Germany	1	48
2	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	ECMWF	United Kingdom	1	48
3	CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE	CNRS-IPSL	France	1	48
4	MAX PLANCK GESELLSCHAFT ZUR FOERDERUNG DER WISSENSCHAFTEN E.V.	MPG	Germany	1	48
5	CENTRE EUROPEEN DE RECHERCHE ET DE FORMATION AVANCEE EN CALCUL SCIENTIFIQUE	CERFACS	France	1	48
6	BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION	BSC	Spain	1	48
7	SCIENCE AND TECHNOLOGY FACILITIES COUNCIL	STFC	United Kingdom	1	48
8	MET OFFICE	METO	United Kingdom	1	48
9	THE UNIVERSITY OF READING	UREAD	United Kingdom	1	48
10	SVERIGES METEOROLOGISKA OCH HYDROLOGISKA INSTITUT	SMHI	Sweden	1	48
11	NATIONAL UNIVERSITY OF IRELAND, GALWAY	ICHEC	Ireland	1	48
12	CENTRO EURO-MEDITERRANEO SUI CAMBIAMENTI CLIMATICI SCARL	CMCC	Italy	1	48
13	DEUTSCHER WETTERDIENST	DWD	Germany	1	48
14	SEAGATE SYSTEMS UK LIMITED	SEAGATE	United Kingdom	1	48
15	BULL SAS	Bull	France	1	48
16	ALLINEA SOFTWARE LIMITED	ALLINEA	United Kingdom	1	48

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	193.00	1	48
WP3	Usability	4 - MPG	120.00	1	48
WP4	Exploitability	7 - STFC	118.00	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	White paper on HPC roadmap for the ESM community	WP1	1 - DKRZ	Report	Public	36
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	WP2	13 - DWD	Report	Public	33
D2.2	Optimised community model code options tested on selected cases	WP2	6 - BSC	Other	Public	36
D2.3	OASIS version adapted to many- core architectures	WP2	5 - CERFACS	Other	Public	24
D2.4	XIOS version supporting GRIB2 format, including on-line diagnostics and adapted to many-core architectures	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
D3.1	ESiWACE Application Software Framework	WP3	4 - MPG	Report	Public	4
D3.2	Update of ESiWACE Application Software Framework (D3.1)	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	WP3	4 - MPG	Report	Public	14

Deliverable Number [14]	Deliverable Title	WP number [9]	Lead beneficiary	Type [15]	Dissemination level [16]	Due Date (in months) [17]
	Ensembles for Educational Purposes					
D3.5	How to select, configure and install ESM software stacks	WP3	6 - BSC	Report	Public	5
D3.6	Update Handbook for system administrators	WP3	6 - BSC	Report	Public	32
D3.7	Software Stack for ESM– A Specification:	WP3	6 - BSC	Report	Public	7
D3.8	A Report from the use of D3.5 for the 3rdE2SCMS	WP3	1 - DKRZ	Report	Public	16
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	12
D4.2	ESD middleware design	WP4	1 - DKRZ	Report	Public	12
D4.3	Final implementation of the ESD middleware	WP4	1 - DKRZ	Other	Public	36
D4.4	Final report on alternative tape usage	WP4	7 - STFC	Report	Public	36
D4.5	Final report on community support for semantic mapping	WP4	7 - STFC	Report	Public	36
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	Report	Public	6

Deliverable Number [14]	Deliverable Title	WP number [9]	Lead beneficiary	Type [15]	Dissemination level [16]	Due Date (in months) [17]
D5.5	Dissemination and Exploitation Plan (EP)	WP5	1 - DKRZ	Report	Public	12
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, 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 Roadmap for HPC for ESM
Task 1.4.1 will produce a roadmap for HPC for ESM [D1.2]. This will integrate input from other tasks and other work packages and will form a basis for the ESiWACE long-term sustainability. It will benefit from discussions at general Assemblies and the annual HPC workshops. It will be elaborated with input from the HPC task force and will benefit from exchanges with the supporting partners. The roadmap will address:
• The "place" of ESM applications in the HPC eco-system (Input from Task 1.1.3 and Task 1.2.2)
• A strategy for community tools (input from Task 1.1.2 / 1.1.3 and WP2 T2.1.2)
• How to cope with the changing HPC landscape and technology in the form of a programmers' guide for ESM (Input from WP2)
• A users' guide of best practices for ESM applications (basically prepared by WP3)

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.

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	White paper on HPC roadmap for the ESM community	1 - DKRZ	Report	Public	36
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 White paper on HPC roadmap for ESM modelling community: the deliverable will report on strategy of the weather and climate community on HPC,(DKRZ, R, PU, PM36). 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 : White paper on HPC roadmap for the ESM community [36]

White paper on HPC roadmap for the ESM community: The deliverable will report on strategy for the climate and weather community in using ESM on HPC

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
MS2	System Software Stack D3.7 handed over to T3.1.3 team	6 - BSC	7	Report published in the website
MS3	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	9 - UREAD	9	Summer School operational
MS4	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
MS6	Reports on user support, training, and integration of NEMO and EC-Earth community models	3 - CNRS-IPSL	30	Reports on user support, training, and integration of NEMO and EC-Earth community models, user-oriented development of XIOS and OASIS and Kick off Workshop on Metrics and Benchmark Strategies for dynamical ores, I/O Systems and coupling technologies Reports published on the intranet

Schedule of relevant Milestones

Milestone number [18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS7	Implementation of ESD middleware at STFC and CMCC	1 - DKRZ	30	ESD middleware finalised and rolled out
MS8	Prototypes of alternative storage backends	7 - STFC	30	Prototypes available
MS9	Prototype tape library for advanced tape subsystems	7 - STFC	30	Prototypes available

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

Objectives

• Centralize support of community models and tools and provide improved access to scientific and performance upgrades
• Coordinate efforts on performance inter-comparisons between community models and tools
• Provide support and training for community I/O tools and strategies for the future
• 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, METO, 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. 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. Performance analysis and inter-comparisons [Lead: CERFACS. Participants: ECMWF, BSC, Bull, CMCC, DKRZ, DWD, CNRS-IPSL, MetO, MPG, SMHI], Networking Activity
On the road to exascale, the performance of the I/O systems and coupling technologies is becoming a major challenge in weather and climate coupled models. This task will propose a set of metrics, with a focus on energy efficiency, for measuring the performances of parallel applications and will use them to benchmark I/O and coupling tools. A common kick-off workshop will offer training on HPC performance tools and will launch the preparation of the benchmarks. PRACE tier-0 and tier-1 systems are targeted to run the benchmark associated test cases, in the framework of the resource allocation dedicated to ESiWACE negotiated in WP1 T1.2.2. Results will be analysed in a final common workshop [D2.1 DWD] and a strategy to overcome main efficiency bottlenecks and to make appropriate choices for the future exascale platforms will be proposed.

Task 2.2.1 Standard efficiency metrics for parallel performance analysis [Lead: CMCC; Participants: BSC, SMHI]
In this task, the set of standard metrics defined in IS-ENES2 for ESM performance analysis will be extended to integrate energy efficiency for compute intensive tasks. The definition of a standard methodology will allow to perform the energy profiling of the applications and to identify the main sources of the energy cost (e.g. data moving on and off chip). Through the PAPI library, these measures will be enabled/disabled for specific routines and Extrae and Paraver, product developed by BSC, will then be used to visualise the corresponding data per routine.

Task 2.2.2 Benchmarking I/O servers and coupling technologies [Lead: DWD; Participants: CERFACS, Bull, CMCC, DKRZ, CNRS-IPSL, MetO, MPG, SMHI]
A first benchmark suite for coupling technologies has been defined in IS-ENES2 and is currently under implementation. In this task, the existing benchmark suite will be revised and the tests will be extended targeting new platforms with O(10K-100K) cores accessible during the ESiWACE longer timeframe. OASIS, OpenPALM, ESMF, XIOS and YAC will be considered. Benchmark suites for I/O libraries and servers will have to be built from scratch. The inter-comparison will include XIOS, ESMF and CDI-pio. EC-Earth will be used as a benchmark application.

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.

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. [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 benchmarked and compared to other couplers (see T2.2.2) 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, CNRS-IPSL, 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	37.00
3 - CNRS-IPSL	36.00
4 - MPG	3.00
5 - CERFACS	27.00
6 - BSC	22.00
8 - METO	3.00
10 - SMHI	20.00
11 - ICHEC	6.00
12 - CMCC	9.00
15 - Bull	20.00
Total	193.00

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	13 - DWD	Report	Public	33
D2.2	Optimised community model code options tested on selected cases	6 - BSC	Other	Public	36
D2.3	OASIS version adapted to many-core architectures	5 - CERFACS	Other	Public	24
D2.4	XIOS version supporting GRIB2 format, including on-line diagnostics and adapted to many-core architectures	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

Description of deliverables

D2.1: Final workshop report including benchmark results and software development strategy for future exascale platforms (DWD,R, PU, PM33). D2.2: Optimised community model code options tested on selected cases (BSC, OTHER, PU, M36). D2.3: OASIS version adapted to many-core architectures (CERFACS, OTHER, PU, PM24). D2.4: XIOS version supporting GRIB2 format, including on-line diagnostics and adapted to many-core architectures (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).

D2.1 : Final workshop report [33]

Final workshop report including benchmark results and software development strategy for future exascale platforms

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

Optimised community model code options tested on selected cases

D2.3 : OASIS version adapted to many- core architectures [24]

OASIS version adapted to many- core architectures

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

XIOS version supporting GRIB2 format, including on-line diagnostics and adapted to many-core architectures

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

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
MS2	System Software Stack D3.7 handed over to T3.1.3 team	6 - BSC	7	Report published in the website
MS3	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	9 - UREAD	9	Summer School operational
MS4	Establish governance rules for new ESiWACE developments	1 - DKRZ	12	Report published on the intranet
MS7	Implementation of ESD middleware at STFC and CMCC	1 - DKRZ	30	ESD middleware finalised and rolled out
MS8	Prototypes of alternative storage backends	7 - STFC	30	Prototypes available
MS9	Prototype tape library for advanced tape subsystems	7 - STFC	30	Prototypes available

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, 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 Oriol Mula-Valls, BSC (co-lead)

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 transport the environment to other PRACE and tier1 machines with more models than in 3rdE2SCMS.

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.

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 3rdE2SCMS [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 3rdE2SCMS. 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. An initial workshop on these topics will develop ideas for system architectures and design sketches for the appropriate software solutions, and will show paths to first implementations. Where systems already exist like Autosubmit29 (supported by IS-ENES2 and adhering to PRACE guidelines), Ophidia30, 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 it's 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	54.00
9 - UREAD	6.00
12 - CMCC	9.00
16 - ALLINEA	6.00
Total	120.00

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)	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	5

List of deliverables

Deliverable Number [14]	Deliverable Title	Lead beneficiary	Type [15]	Dissemination level [16]	Due Date (in months) [17]
D3.6	Update Handbook for system administrators	6 - BSC	Report	Public	32
D3.7	Software Stack for ESM– A Specification:	6 - BSC	Report	Public	7
D3.8	A Report from the use of D3.5 for the 3rdE2SCMS	1 - DKRZ	Report	Public	16
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), Version 2 (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, PM5) D3.6:Update Handbook for system administrators (D3.5) (BSC, R, PU, M32) D3.7: Software Stack for ESM– A Specification narrowing down D3.5 for the 3rdE2SCMS (BSC, R, PU, PM7) D3.8: Experiences with the ENES System Software Stack: A Report from the use of D3.5 for the 3rdE2SCMS (DKRZ, R, PU, PM16) 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) [30]

Update of ESiWACE Application Software Framework (D3.1) including recent experiences

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 [5]

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 [32]

Update Handbook for system administrators (D3.5)

D3.7 : Software Stack for ESM– A Specification: [7]

Software Stack for ESM– A Specification: Narrowing down D3.5 for the 3rdE2SCMS

D3.8 : A Report from the use of D3.5 for the 3rdE2SCMS [16]

Experiences with the ENES System Software Stack: A Report from the use of D3.5 for the 3rdE2SCMS

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
MS4	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 on user support, training, and integration of NEMO and EC-Earth community models, user-oriented development of XIOS and OASIS and Kick off Workshop on Metrics and Benchmark Strategies for dynamical ores, I/O Systems and coupling technologies Reports published on the intranet

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.
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.
4. Supporting semantic mapping between key weather and climate formats. While format interoperability exists between GRIB (primarily weather) and netCDF (primarily
climate) data, conversion currently involves information loss. ESiWACE will provide support for the community to define and populate a semantic mapping that will facilitate community interoperability.

Description of work and role of partners

WP4 - Exploitability [Months: 1-48]
STFC, DKRZ, ECMWF, 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, ECMWF] 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 and a performance optimized version for all data centres released in [D4.3]. 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].

Task 4.4 Semantic mapping between netCDF and GRIB [Lead: STFC, Participants: ECMWF] Joint research activity
One of the points of intersection between the operational weather and climate communities is data sharing, and differing data formats are a key barrier to exploitability - with both tools and semantic differences in content being significant issues. There are a number of tool building activities, but all are hampered by a lack of an agreed semantic cross walk between netCDF and GRIB, the two main formats involved. Such a semantic walk involves more than matching names, since key semantic concepts are built into coordinate definitions as well as parameter names. For example, the common concept 'maximum air temperature at 2m' requires knowledge of the time bounds to evaluate what maximum means, the definition of air temperature, and interaction with the vertical coordinate. While there is a one-to-one mapping between this concept in the two formats, there is not a two-way mapping in the way the information is implemented. The lack of a comprehensive map is both a barrier to data sharing between communities, and a barrier to evolution. Technology exists to maintain such a mapping, but the community needs to resource someone to populate, maintain and extend the mapping. This task will effectively contribute to a WMO action aiming to increase the support for such semantic mapping, and has been specifically identified by WMO colleagues as an important deliverable from ESiWACE. The

web presence to support this activity will be made available, with a report on activity and advances of the current status delivered as [D4.5].

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
2 - ECMWF	12.00
7 - STFC	33.00
12 - CMCC	10.00
14 - SEAGATE	18.00
Total	118.00

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	12
D4.2	ESD middleware design	1 - DKRZ	Report	Public	12
D4.3	Final implementation of the ESD middleware	1 - DKRZ	Other	Public	36
D4.4	Final report on alternative tape usage	7 - STFC	Report	Public	36
D4.5	Final report on community support for semantic mapping	7 - STFC	Report	Public	36

Description of deliverables

D4.1 Business model with alternative scenarios (DKRZ, R, PU, PM12). D4.2 ESD middleware design (DKRZ, R,PU, PM12). D4.3 Final implementation of the ESD middleware. (DKRZ, OTHER, PU, PM36). D4.4 Final report on alternative tape usage. (STFC, R, PU, PM36). D4.5 Final report on community support for semantic mapping (STFC,R, PU, PM36).

D4.1 : Business model with alternative scenarios [12]
Business model with alternative scenarios

D4.2 : ESD middleware design [12]
ESD middleware design

D4.3 : Final implementation of the ESD middleware [36]
Final implementation of the ESD middleware

D4.4 : Final report on alternative tape usage [36]
Final report on alternative tape usage

D4.5 : Final report on community support for semantic mapping [36]
Final report on community support for semantic mapping

Schedule of relevant Milestones

Milestone number [18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS2	System Software Stack D3.7 handed over to T3.1.3 team	6 - BSC	7	Report published in the website
MS3	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	9 - UREAD	9	Summer School operational
MS4	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 on user support, training, and integration of NEMO and EC-Earth community models, user-oriented development of XIOS and OASIS and Kick off Workshop on Metrics and Benchmark Strategies for dynamical ores, I/O Systems and coupling technologies Reports published on the intranet

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	Report	Public	6
D5.5	Dissemination and Exploitation Plan (EP)	1 - DKRZ	Report	Public	12

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, PM12) 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) [12]
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
MS2	System Software Stack D3.7 handed over to T3.1.3 team	6 - BSC	7	Report published in the website
MS3	Application Stack running at FMI/CSC, preparation of System Stack T3.2 Team	9 - UREAD	9	Summer School operational

Schedule of relevant Milestones

Milestone number [18]	Milestone title	Lead beneficiary	Due Date (in months)	Means of verification
MS4	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
MS6	Reports on user support, training, and integration of NEMO and EC-Earth community models	3 - CNRS-IPSL	30	Reports on user support, training, and integration of NEMO and EC-Earth community models, user-oriented development of XIOS and OASIS and Kick off Workshop on Metrics and Benchmark Strategies for dynamical ores, I/O Systems and coupling technologies Reports published on the intranet
MS7	Implementation of ESD middleware at STFC and CMCC	1 - DKRZ	30	ESD middleware finalised and rolled out
MS8	Prototypes of alternative storage backends	7 - STFC	30	Prototypes available
MS9	Prototype tape library for advanced tape subsystems	7 - STFC	30	Prototypes available

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	System Software Stack D3.7 handed over to T3.1.3 team	WP1, WP2, WP4, WP5	6 - BSC	7	Report published in the website
MS3	Application Stack running at FMI/ CSC, preparation of System Stack T3.2 Team	WP1, WP2, WP4, WP5	9 - UREAD	9	Summer School operational
MS4	Establish governance rules for new ESiWACE developments	WP1, WP2, WP3, WP4, WP5	1 - DKRZ	12	Report published on the intranet
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	WP1, WP3, WP4, WP5	3 - CNRS-IPSL	30	Reports on user support, training, and integration of NEMO and EC-Earth community models, user-oriented development of XIOS and OASIS and Kick off Workshop on Metrics and Benchmark Strategies for dynamical ores, I/O Systems and coupling technologies Reports published on the intranet
MS7	Implementation of ESD middleware at STFC and CMCC	WP1, WP2, WP5	1 - DKRZ	30	ESD middleware finalised and rolled out
MS8	Prototypes of alternative storage backends	WP1, WP2, WP5	7 - STFC	30	Prototypes available
MS9	Prototype tape library for advanced tape subsystems	WP1, WP2, WP5	7 - STFC	30	Prototypes available

1.3.5. WT5 Critical Implementation risks and mitigation actions

Risk number	Description of risk	WP Number	Proposed risk-mitigation measures
R1	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.
R2	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.
R3	Difficulty in achieving long-term sustainability.		The Business plan (see also Section 2.2.4) will be developed and continuously updated.
R4	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
R5	Scarcity of computing resources allocated by PRACE to ESiWACE:		ESiWACE will actively engage with PRACE to obtain allocations.

Risk number	Description of risk	WP Number	Proposed risk-mitigation measures
	impossibility to test benchmarks for I/O and couplers on tier-0 platforms on very high number of cores.		ESiWACE partner resources will be exploited for tests on O(10000) computational processor cores.
R6	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
R7	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.
R8	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.
R9	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.
R10	Insufficient fulfilment of the needs of the communities involved by the published recommendations.		The evolving versions of the specifications will be appropriately fine-tuned.
R11	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.

Risk number	Description of risk	WP Number	Proposed risk-mitigation measures
R12	Limited combinability of different scheduling approaches.		The number of approaches will be minimized, without imposing rules
R13	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.
R14	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.
R15	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	37	3	12	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	3	54	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	193	120	118	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	

1.4. Ethics Requirements

No ethics requirements indicated

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

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 **S**imulation of **W**eather **a**nd **C**limate in **E**urope

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

TABLE OF CONTENTS

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-of-the-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] http://www.etp4hpc.eu/
[2] http://www.prace-ri.eu/

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 high-end 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 coupled climate model EC-Earth[3], the ocean modelling platform NEMO[4] or the coupler module OASIS[5].
- 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.

[3] EC-Earth: http://ecearth.knmi.nl/
[4] NEMO: http://www.nemo-ocean.eu/
[5] OASIS: https://verc.enes.org/models/software-tools/oasis

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 EC-Earth models, OASIS coupler, XIOS[6] I/O server, Cylc[7] meta-scheduler), in response to a well-defined 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, trans-disciplinary considerations.

ESiWACE targets the convergent use of Earth System Modelling for weather and climate science (ESM)[8]. Global Earth System Models, the post-processing methods to handle the vast amounts of 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,

---

[6] XIOS:http://forge.ipsl.jussieu.fr/ioserver/wiki
[7] Cylc: http://cylc.github.io/cylc/
[8] **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.**

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 inter-comparisons, 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)[9]. 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[10]). It integrates the European climate modelling community, stimulates

[9] http://is.enes.org
[10] https://portal.enes.org/

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 IS-ENES2, 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 long-term 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 CNRS-IPSL, CMCC and Bull to improve the NEMO model, the cooperation of DKRZ, MPG and Bull to improve the ICON[11] 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)[12].

(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 E10-consortium[13]). ESiWACE partners established relationships with netCDF and GRIB technical teams and launched an early relationship with the HDF technical team[14]. 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 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.

---

[11] ICON: http://www.mpimet.mpg.de/en/science/models/icon.html
[12] HD(CP)2: http://hdcp2.zmaw.de/
[13] E10: http://www.exascale10.com
[14] 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/.

(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[15], HIRLAM[16], LACE[17], and COSMO[18] as well as existing and planned weather and climate European environmental (SC5-2014-1 "Earth systems" call under final evaluation) projects such as CRESCENDO, IMPULSE, PRIMAVERA, and the European Training Network Marie Skłodowska-Curie 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[19]) 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[20], WGNE[21], 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 IS-ENES2. 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

[15] ALADIN: http://www.cnrm.meteo.fr/aladin/
[16] HIRLAM: http://www.hirlam.org/
[17] LACE: http://www.rclace.eu/
[18] COSMO: http://www.cosmo-model.org/
[19] OBS4MIPS: https://www.earthsystemcog.org/projects/obs4mips/
[20] S2S: Subseasonal to Seasonal Prediction Project -http://www.wmo.int/pages/prog/arep/wwrp/new/S2S_project_main_page.html
[21] 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/

provision of community codes and the ability to identify and to prioritize user-driven requirements. 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 IS-ENES2 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 cross-cutting 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 ETP_4HPC[23].

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

[23] "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) two-sided 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.

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[24], the MetO, MPG and others aimed specifically at this user base.

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 non-

[24] The NZ National Institute of Water and Atmosphere

computer-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 time-to-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 non-traditional 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 achieve all four impacts 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-the-art 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 Group Committee (UGC) 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 web-platform 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 EU-funded 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**[25] 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 one-time 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.

[25] 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 Group Committee	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 time-scales 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 long-term 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
WP1	Governance, Engagement & long-term sustainability	CNRS-IPSL, Sylvie Joussaume	DKRZ, Joachim Biercamp
WP2	Scalability	ECMWF, Peter Bauer	CERFACS, Sophie Valcke
WP3	Usability	MPG Reinhard Budich	BSC, Oriol Mula-Valls
WP4	Exploitability	STFC Bryan Lawrence	DKRZ, Thomas Ludwig
WP5	Management & Dissemination	DKRZ, Joachim Biercamp	ECMWF Peter Bauer

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"** deals with the improvements in performance of model codes and tools (model coupler, I/O libraries) in order to augment their efficiency User training and support acting upon user requirements as well as provision of improved codes are an important follow-up ensuring 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-to-end 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 Performance analysis and inter-comparisons
- 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

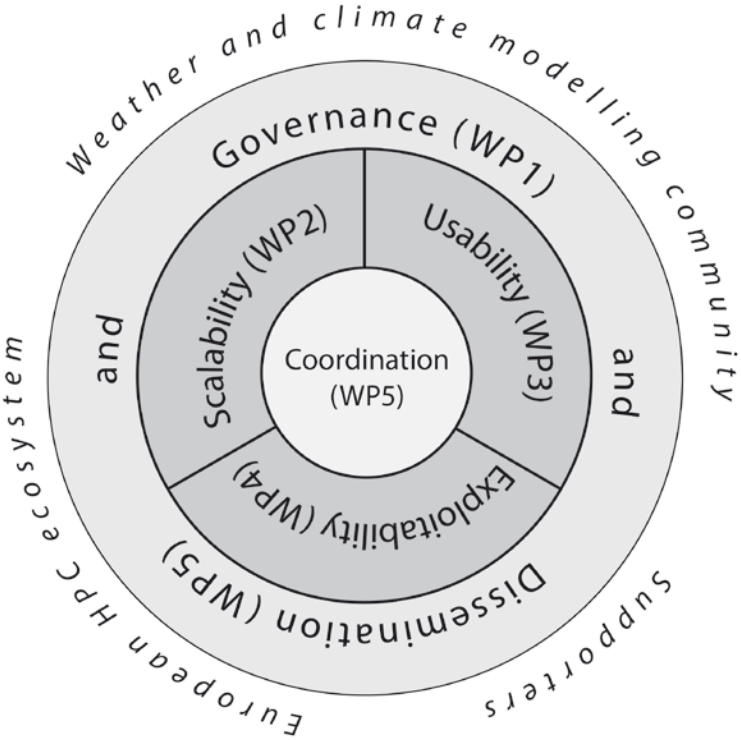

Table 3.1b Gantt chart: Timing of the different Work Packages. Deliverables, Milestones and Reports

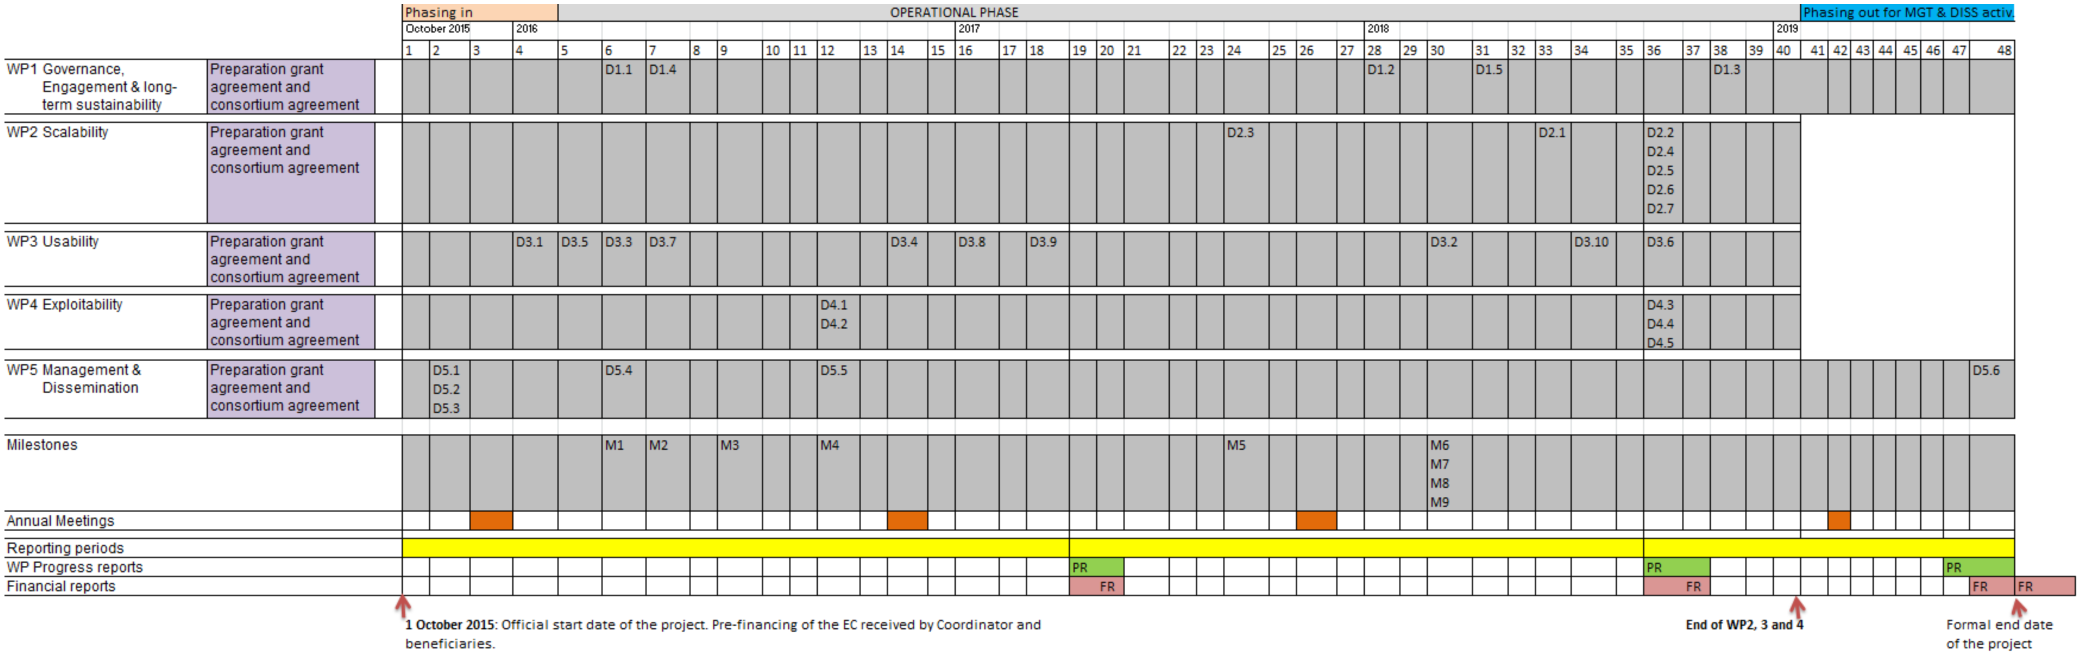

M = Milestone
D = Deliverable
WP = Work Package
PR = Progress Report
FR = Financial Report

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

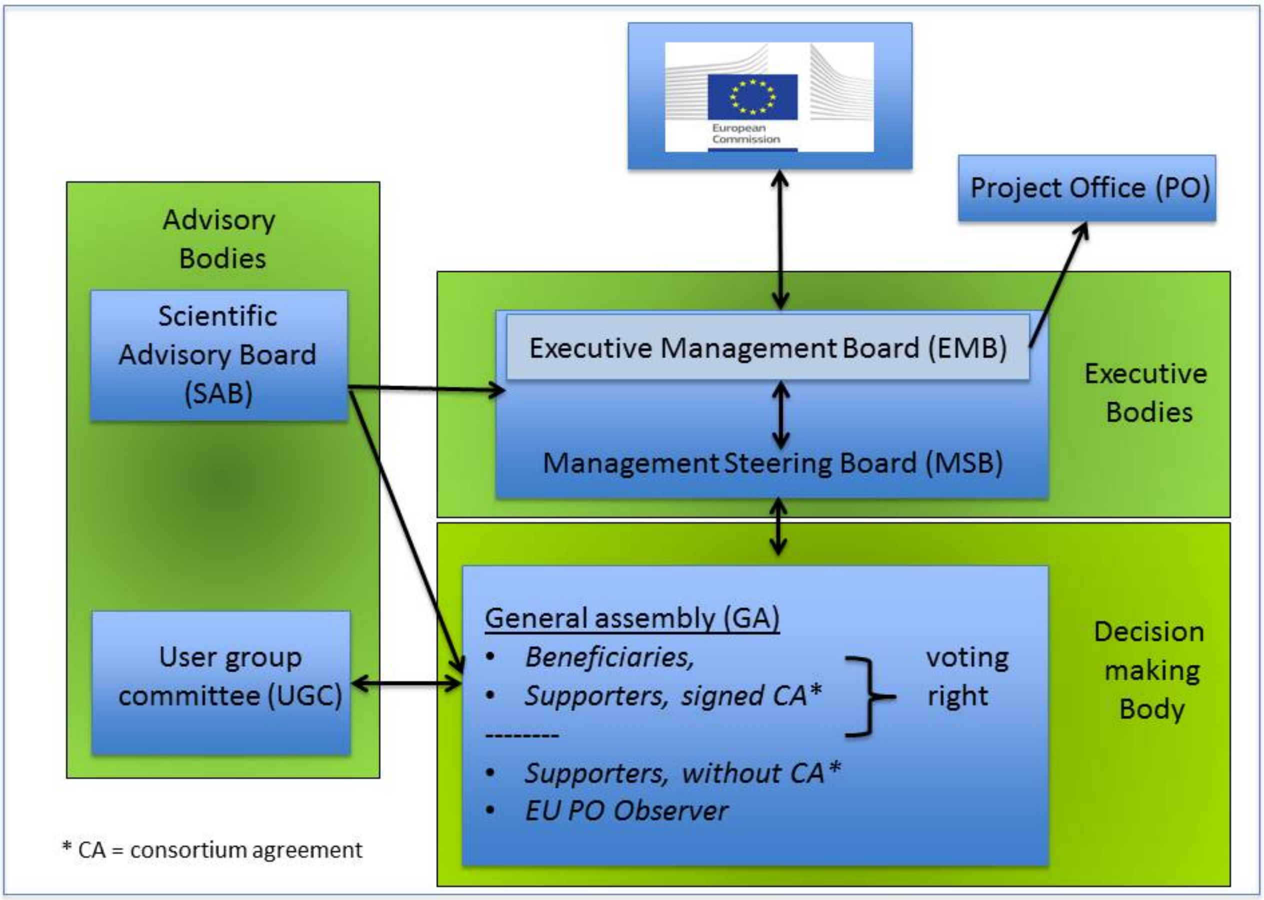

3.2.1 Description of the organisational structure and the decision-making

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 and the User Group Committee

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 and the user group committee. 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 bodies

We have foreseen two advisory bodies**: Scientific Advisory Board (SAB)** and **the User Group Committee (UGC).**

- 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.

- The **User Group Committee (UGC)** has the commitment to reflect the users' opinion, to propose upcoming topics and forward requests to the GA and / or the Executive Bodies. The UGC will be the key tool for the ESiWACE to get input from the wider scientific community and disseminate the results achieved by the ESiWACE. The **UGC** represents:
    - the members of the scientific weather and climate community interested in the outcomes of the ESiWACE,
    - The users of the shared (task 1.2) and newly developed software (task 1.3) identified in WP, which are not represented by beneficiaries in the consortium.
    - The representatives of the users from supporters.

  Supporting organisations within UGC will be invited to participate in COE workshops, general assemblies, and training activities.

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 Committee** for ESiWACE is central: we consider it crucial for bringing in new ideas, requests for new research and tasks forces to be installed. During the complete project runtime the UGC can address issues to MSB. 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, supporting partners and members of the user group committee 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 Medium-Range 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 compute 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 Limited-area 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.

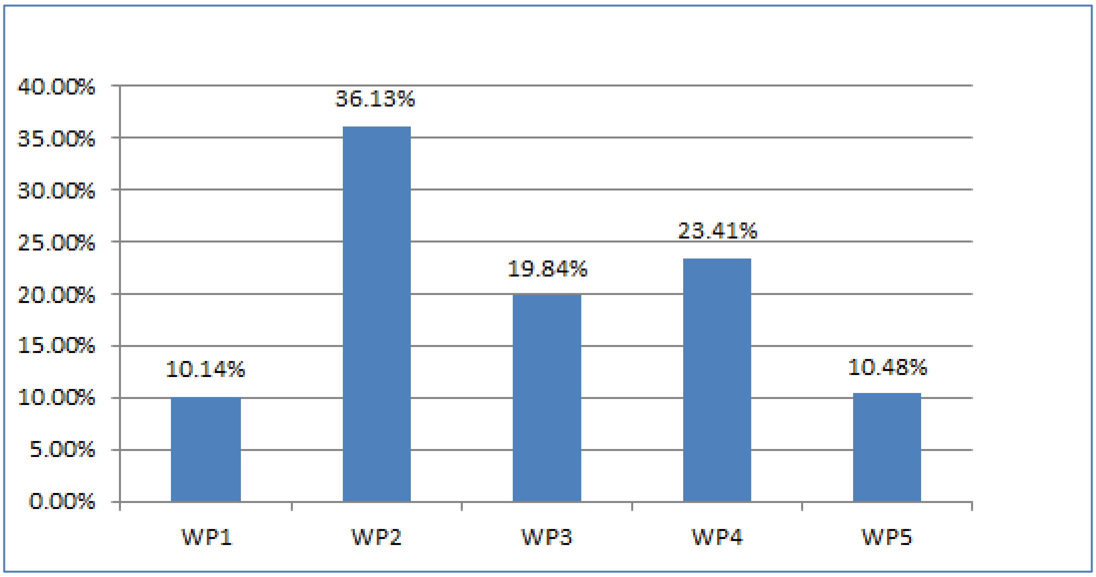

Chart 3.4 a: Distribution across the work packages of the direct costs. Figures are indicative.

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.

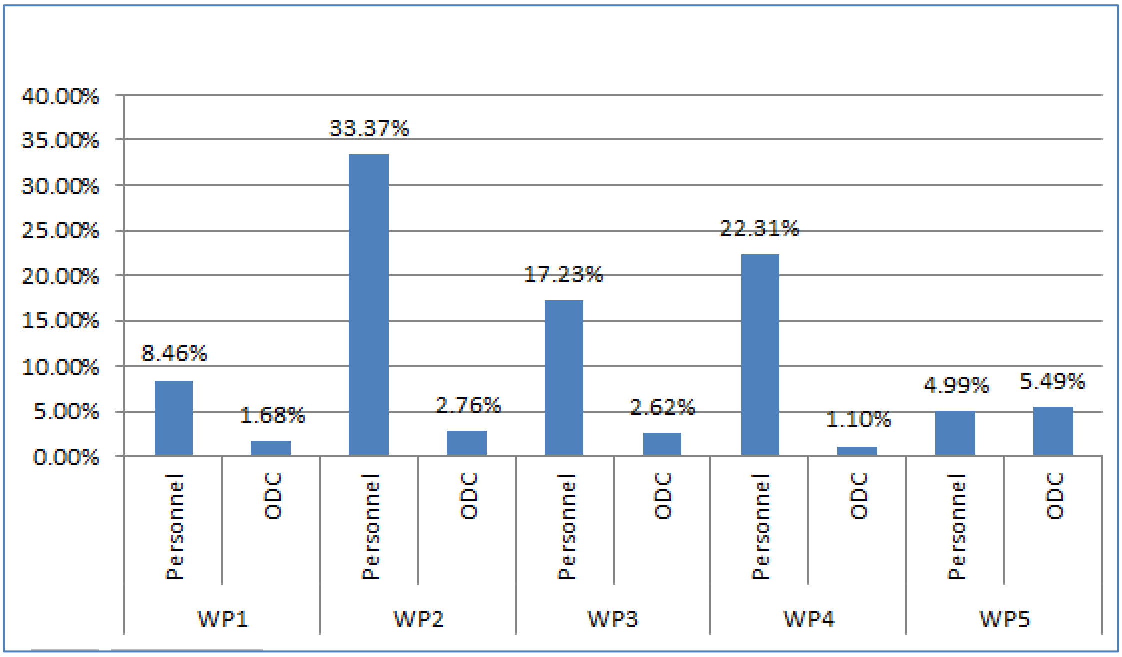

Chart 3.4 b: Distribution of the direct costs according to "personnel" and "other direct costs" (ODR) in the work packages. Figures are indicative.

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 [27]
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
I3	Integrated Infrastructures Initiative
ICON	Icosahedral non-hydrostatic (general circulation model)
IFS	Integrated Forecast System

[27] **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.**

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
UGC	User Group Committee
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-high-performance-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