Horizon 2020

Call: H2020-FETHPC-2016-2017
(FET Proactive – High Performance Computing)

Topic: FETHPC-02-2017

Type of action: RIA
(Research and Innovation action)

Proposal number: 800897

Proposal acronym: ESCAPE-2

Deadline Id: H2020-FETHPC-2017

Table of contents

How to fill in the forms

The administrative forms must be filled in for each proposal using the templates available in the submission system. Some data fields in the administrative forms are pre-filled based on the previous steps in the submission wizard.

1 - General information

Topic	FETHPC-02-2017
Call Identifier	H2020-FETHPC-2016-2017
Type of Action	RIA
Deadline Id	H2020-FETHPC-2017

Acronym	ESCAPE-2		
Proposal title*	Energy-efficient SCalable Algorithms for weather and climate Prediction at Exascale		
	Note that for technical reasons, the following characters are not accepted in the Proposal Title and will be removed: < > " &		
Duration in months	*36*		
Fixed keyword 1	*Mathematics for High-Performance Computing*	Add	
Fixed keyword 2	*Scalability*	Add	Remove
Fixed keyword 3	*Numerical analysis, simulation, optimisation, modelling tools, da*	Add	Remove
Fixed keyword 4	*High performance computing*	Add	Remove
Fixed keyword 5	*Scientific computing and data processing*	Add	Remove
Free keywords	*Weather and climate prediction, mathematics and algorithms, scalability, exascale*		

Proposal ID **800897** Acronym **ESCAPE-2**

Abstract

ESCAPE-2 will develop world-class, extreme-scale computing capabilities for European operational numerical weather and climate prediction, and provide the key components for weather and climate domain benchmarks to be deployed on extreme-scale demonstrators and beyond. This will be achieved by developing bespoke and novel mathematical and algorithmic concepts, combining them with proven methods, and thereby reassessing the mathematical foundations forming the basis of Earth system models. ESCAPE-2 also invests in significantly more productive programming models for the weather-climate community through which novel algorithm development will be accelerated and future-proofed. Eventually, the project aims at providing exascale-ready production benchmarks to be operated on extreme-scale demonstrators (EsD) and beyond. ESCAPE-2 combines cross-disciplinary uncertainty quantification tools (URANIE) for high-performance computing, originating from the energy sector, with ensemble based weather and climate models to quantify the effect of model and data related uncertainties on forecasting – a capability, which weather and climate prediction has pioneered since the 1960s.

The mathematics and algorithmic research in ESCAPE-2 will focus on implementing data structures and tools supporting parallel computation of dynamics and physics on multiple scales and multiple levels. Highly-scalable spatial discretization will be combined with proven large time-stepping techniques to optimize both time-to-solution and energy-to-solution. Connecting multi-grid tools, iterative solvers, and overlapping computations with flexible-order spatial discretization will strengthen algorithm resilience against soft or hard failure. In addition, machine learning techniques will be applied for accelerating complex sub-components. The sum of these efforts will aim at achieving at the same time: performance, resilience, accuracy and portability.

Remaining characters 57

Has this proposal (or a very similar one) been submitted in the past 2 years in response to a call for proposals under Horizon 2020 or any other EU programme(s)? ○ Yes ◉ No

Declarations

1) The coordinator declares to have the explicit consent of all applicants on their participation and on the content of this proposal.	☒
2) The information contained in this proposal is correct and complete.	☒
3) This proposal complies with ethical principles (including the highest standards of research integrity — as set out, for instance, in the European Code of Conduct for Research Integrity — and including, in particular, avoiding fabrication, falsification, plagiarism or other research misconduct).	☒

4) The coordinator confirms:

- to have carried out the self-check of the financial capacity of the organisation on http://ec.europa.eu/research/participants/portal/desktop/en/organisations/lfv.html or to be covered by a financial viability check in an EU project for the last closed financial year. Where the result was "weak" or "insufficient", the coordinator confirms being aware of the measures that may be imposed in accordance with the H2020 Grants Manual (Chapter on Financial capacity check); or	○
- is exempt from the financial capacity check being a public body including international organisations, higher or secondary education establishment or a legal entity, whose viability is guaranteed by a Member State or associated country, as defined in the H2020 Grants Manual (Chapter on Financial capacity check); or	◉
- as sole participant in the proposal is exempt from the financial capacity check.	○

5) The coordinator hereby declares that each applicant has confirmed:

- they are fully eligible in accordance with the criteria set out in the specific call for proposals; and	☒
- they have the financial and operational capacity to carry out the proposed action.	☒
The coordinator is only responsible for the correctness of the information relating to his/her own organisation. Each applicant remains responsible for the correctness of the information related to him/her and declared above. Where the proposal to be retained for EU funding, the coordinator and each beneficiary applicant will be required to present a formal declaration in this respect.	

According to Article 131 of the Financial Regulation of 25 October 2012 on the financial rules applicable to the general budget of the Union (Official Journal L 298 of 26.10.2012, p. 1) and Article 145 of its Rules of Application (Official Journal L 362, 31.12.2012, p.1) applicants found guilty of misrepresentation may be subject to administrative and financial penalties under certain conditions.

Personal data protection

The assessment of your grant application will involve the collection and processing of personal data (such as your name, address and CV), which will be performed pursuant to Regulation (EC) No 45/2001 on the protection of individuals with regard to the processing of personal data by the Community institutions and bodies and on the free movement of such data. Unless indicated otherwise, your replies to the questions in this form and any personal data requested are required to assess your grant application in accordance with the specifications of the call for proposals and will be processed solely for that purpose. Details concerning the purposes and means of the processing of your personal data as well as information on how to exercise your rights are available in the privacy statement. Applicants may lodge a complaint about the processing of their personal data with the European Data Protection Supervisor at any time.

Your personal data may be registered in the Early Detection and Exclusion system of the European Commission (EDES), the new system established by the Commission to reinforce the protection of the Union's financial interests and to ensure sound financial management, in accordance with the provisions of articles 105a and 108 of the revised EU Financial Regulation (FR) (Regulation (EU, EURATOM) 2015/1929 of the European Parliament and of the Council of 28 October 2015 amending Regulation (EU, EURATOM) No 966/2012) and articles 143 - 144 of the corresponding Rules of Application (RAP) (COMMISSION DELEGATED REGULATION (EU) 2015/2462 of 30 October 2015 amending Delegated Regulation (EU) No 1268/2012) for more information see the Privacy statement for the EDES Database).

Proposal ID **800897** *Acronym* **ESCAPE-2**

List of participants

#	Participant Legal Name	Country
1	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	United Kingdom
2	DEUTSCHES KLIMARECHENZENTRUM GMBH	Germany
3	MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV	Germany
4	EIDGENOESSISCHES DEPARTEMENT DES INNERN	Switzerland
5	BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION	Spain
6	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES	France
7	LOUGHBOROUGH UNIVERSITY	United Kingdom
8	INSTITUT ROYAL METEOROLOGIQUE DE BELGIQUE	Belgium
9	POLITECNICO DI MILANO	Italy
10	DANMARKS METEOROLOGISKE INSTITUT	Denmark
11	FONDAZIONE CENTRO EURO-MEDITERRANEO SUI CAMBIAMENTI CLIMATICI	Italy
12	BULL SAS	France

2 - Administrative data of participating organisations

PIC	***Legal name***
999916741	*EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS*

Short name: ECMWF

Address of the organisation

Street SHINFIELD PARK

Town READING

Postcode RG2 9AX

Country United Kingdom

Webpage www.ecmwf.int

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person ....yes

Non-profit ....yes

International organisation ....yes

International organisation of European interest ......yes

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status.................... unknown

SME self-assessment .... unknown

SME validation sme.................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Research Department

☐ not applicable

☒ Same as organisation address

Street: SHINFIELD PARK

Town: READING

Postcode: RG2 9AX

Country: United Kingdom

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **ECMWF**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Peter**

Last name: **Bauer**

E-Mail: **peter.bauer@ecmwf.int**

Position in org.: Deputy Director

Department: Research Department

☐ Same as organisation

☒ Same as organisation address

Street: SHINFIELD PARK

Town: READING

Post code: RG2 9AX

Country: United Kingdom

Website: www.ecmwf.int

Phone 1: +441189499080

Phone 2: +xxx xxxxxxxxx

Fax: +44 118 986 9450

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Daniel	THIEMERT	daniel.thiemert@ecmwf.int	+441189499024
Nils	Wedi	nils.wedi@ecmwf.int	
Ben	Brown	ben.brown@ecmwf.int	
Laura	Drion	laura.drion@ecmwf.int	

PIC	***Legal name***
998692310	*DEUTSCHES KLIMARECHENZENTRUM GMBH*

Short name: DKRZ

Address of the organisation

Street BUNDESSTRASSE 45A

Town HAMBURG

Postcode 20146

Country Germany

Webpage http:/www.dkrz.de

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status.................................03/11/2008 - no

SME self-assessment .... unknown

SME validation sme.......................................03/11/2008 - no

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Application Support-Anwendungen

☐ not applicable

☒ Same as organisation address

Street: BUNDESSTRASSE 45A

Town: HAMBURG

Postcode: 20146

Country: Germany

Dependencies with other proposal participants

Character of dependence	***Participant***	
Is controlled by	MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN E	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **DKRZ**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Joachim** Last name: **Biercamp**

E-Mail: **biercamp@dkrz.de**

Position in org.: Head of Application Support Unit

Department: Application Support ☐ Same as organisation

☒ Same as organisation address

Street: BUNDESSTRASSE 45A

Town: HAMBURG Post code: 20146

Country: Germany

Website: www.dkrz.de

Phone 1: 0049 40460094314 Phone 2: *+xxx xxxxxxxxx* Fax: *+xxx xxxxxxxxx*

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Katja	Brendt	brendt@dkrz.de	
Chiara	Bearzotti	chiara.bearzotti@gmail.com	
Kerstin	Fieg	fieg@dkrz.de	

PIC	***Legal name***
999990267	*MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV*

Short name: MPI-M

Address of the organisation

Street	HOFGARTENSTRASSE 8
Town	MUENCHEN
Postcode	80539
Country	Germany
Webpage	www.mpg.de

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....no

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status..................................05/04/2016 - no

SME self-assessment ....05/04/2016 - no

SME validation sme..31/10/2008 - no

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name	Scientific Computing Laboratory
	☐ not applicable
	☐ Same as organisation address
Street	Bundesstr. 53
Town	Hamburg
Postcode	D-20146
Country	Germany

Dependencies with other proposal participants

Character of dependence	*Participant*	
Controls	DEUTSCHES KLIMARECHENZENTRUM GMBH	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **MPI-M**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Mr.

Sex: (●) Male ( ) Female

First name **Reinhard** Last name **Budich**

E-Mail **reinhard.budich@mpimet.mpg.de**

Position in org.: Reponsible for Strategic IT Partnerships

Department: Scientific Computing Laboratory ☐ Same as organisation

☐ Same as organisation address

Street: Bundesstr. 53

Town: Hamburg Post code: D-20146

Country: Germany

Website: https://www.mpimet.mpg.de/en/home/

Phone 1: +494041173-369 Phone 2: +xxx xxxxxxxxx Fax: +494041173-298

PIC	***Legal name***
990186089	*EIDGENOESSISCHES DEPARTEMENT DES INNERN*

Short name: MSWISS

Address of the organisation

Street Inselgasse 1

Town BERN

Postcode 3003

Country Switzerland

Webpage www.edi.admin.ch

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....no

Enterprise Data

SME self-declared status.................................. unknown

SME self-assessment .... unknown

SME validation sme.. unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name	Federal Institute of Meteorology and Climatology MeteoSwiss	☐ not applicable
	☐ Same as organisation address	
Street	Operation Center 1	
Town	Zurich-Flughafen	
Postcode	8058	
Country	Switzerland	

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **MSWISS**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Oliver** Last name: **Fuhrer**

E-Mail: **oliver.fuhrer@meteoswiss.ch**

Position in org.: Team lead

Department: Numerical Prediction ☐ Same as organisation

☐ Same as organisation address

Street: Operation Center 1

Town: Zurich-Flughafen Post code: 8058

Country: Switzerland

Website: http://www.meteoswiss.ch/

Phone 1: +41 58 460 93 59 Phone 2: +xxx xxxxxxxxx Fax: +xxx xxxxxxxxx

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Carlos	Osuna	carlos.osuna@meteoswiss.ch	

PIC	***Legal name***
999655520	*BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION*

Short name: BSC

Address of the organisation

Street Calle Jordi Girona 31

Town BARCELONA

Postcode 08034

Country Spain

Webpage www.bsc.es

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status.................................01/03/2005 - no

SME self-assessment .... unknown

SME validation sme....................................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Earth Sciences Department

☐ not applicable

☐ Same as organisation address

Street: c/ Jordi Girona 29

Town: Barcelona

Postcode: 08034

Country: Spain

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **BSC**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Mr.

Sex: (●) Male ( ) Female

First name **Kim** Last name **Serradell Maronda**

E-Mail **kim.serradell@bsc.es**

Position in org. Computational Earth Sciences Group Leader

Department Earth Sciences Department ☐ Same as organisation

☐ Same as organisation address

Street c/ Jordi Girona 29

Town Barcelona Post code Spain

Country Spain

Website www.bsc.es

Phone 1 +34 934134051 Phone 2 *+xxx xxxxxxxxx* Fax +34 934137721

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Mar	Rodríguez	mar.rodriguez@bsc.es	

PIC	***Legal name***
999992401	*COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES*

Short name: CEA

Address of the organisation

Street RUE LEBLANC 25

Town PARIS 15

Postcode 75015

Country France

Webpage www.cea.fr

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status.................................01/10/2008 - no

SME self-assessment .... unknown

SME validation sme.......................................01/10/2008 - no

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: CEA Saclay, DEN/DANS/DM2S/STMF/LGLS ☐ not applicable

☐ Same as organisation address

Street: N/A

Town: Gif-sur-Yvette

Postcode: 91191

Country: France

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **CEA**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Mr.

Sex: (•) Male ( ) Female

First name **Adrien** Last name **Bruneton**

E-Mail **adrien.bruneton@cea.fr**

Position in org.: Research Engineer

Department: DEN/DANS/DM2S/STMF/LGLS ☐ Same as organisation

☐ Same as organisation address

Street: N/A

Town: Gif-sur-Yvette Post code: 91191

Country: France

Website: www.cea.fr

Phone 1: +33 1 69 08 65 93 Phone 2: +xxx xxxxxxxxx Fax: +33 1 69 08 68 86

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **LU**

PIC	***Legal name***
999990752	*LOUGHBOROUGH UNIVERSITY*

Short name: LU

Address of the organisation

Street ASHBY ROAD

Town LOUGHBOROUGH

Postcode LE11 3TU

Country United Kingdom

Webpage www.lboro.ac.uk

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....unknown

International organisation of European interest ......unknown

Secondary or Higher education establishment ....yes

Research organisation ....yes

Enterprise Data

SME self-declared status.................................19/04/1966 - no

SME self-assessment .... unknown

SME validation sme.......................................19/04/1966 - no

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Wolfson School of Mechanical and Manufacturing Engineering ☐ not applicable

☒ Same as organisation address

Street: ASHBY ROAD

Town: LOUGHBOROUGH

Postcode: LE11 3TU

Country: United Kingdom

Dependencies with other proposal participants

Character of dependence	Participant	

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: ◯ Male ◉ Female

First name: **Joanna**

Last name: **Szmelter**

E-Mail: **j.szmelter@lboro.ac.uk**

Position in org.: Reader

Department: Wolfson School of Mechanical and Manufacturing Engineering

☐ Same as organisation

☒ Same as organisation address

Street: ASHBY ROAD

Town: LOUGHBOROUGH

Post code: LE11 3TU

Country: United Kingdom

Website: http://www.lboro.ac.uk/departments/meme/staff/joanna-szmelter/

Phone 1: +44 (0)1509 227580

Phone 2: *+xxx xxxxxxxxx*

Fax: *+xxx xxxxxxxxx*

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **RMI**

PIC	***Legal name***
996699154	*INSTITUT ROYAL METEOROLOGIQUE DE BELGIQUE*

Short name: RMI

Address of the organisation

Street AVENUE CIRCULAIRE 3

Town BRUXELLES

Postcode 1180

Country Belgium

Webpage www.meteo.be

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....unknown

International organisation of European interest ......unknown

Secondary or Higher education establishment ....unknown

Research organisation ....yes

Enterprise Data

SME self-declared status.................................. unknown

SME self-assessment .... unknown

SME validation sme....................................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Meteorological and Climatological Research

☐ not applicable

☒ Same as organisation address

Street: AVENUE CIRCULAIRE 3

Town: BRUXELLES

Postcode: 1180

Country: Belgium

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **RMI**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Piet** Last name: **Termonia**

E-Mail: **piet.termonia@meteo.be**

Position in org.: Head of Department

Department: Meteorological and Climatological Research ☐ Same as organisation

☒ Same as organisation address

Street: AVENUE CIRCULAIRE 3

Town: BRUXELLES Post code: 1180

Country: Belgium

Website: www.meteo.be

Phone 1: +32 2 373 0638 Phone 2: *+XXX XXXXXXXXX* Fax: +32 2 373 6762

PIC	***Legal name***
999879881	*POLITECNICO DI MILANO*

Short name: POLIMI

Address of the organisation

Street PIAZZA LEONARDO DA VINCI 32

Town MILANO

Postcode 20133

Country Italy

Webpage www.polimi.it

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....yes

Research organisation ....yes

Enterprise Data

SME self-declared status..................................10/05/2016 - no

SME self-assessment .... unknown

SME validation sme..25/02/2009 - no

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Dipartimento di Matematica

☐ not applicable

☐ Same as organisation address

Street: Via E. Bonardi 9

Town: Milano

Postcode: 20133

Country: Italy

Dependencies with other proposal participants

Character of dependence	***Participant***	

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Prof.
Sex: (•) Male ( ) Female

First name: **Luca**
Last name: **Bonaventura**

E-Mail: **luca.bonaventura@polimi.it**

Position in org.: Associate Professor

Department: Mathematics
[ ] Same as organisation

[ ] Same as organisation address

Street: Via E. Bonardi 9

Town: Milano
Post code: 20133

Country: Italy

Website: http://www1.mate.polimi.it/~bonavent/

Phone 1: +39-02-23994600
Phone 2: +xxx xxxxxxxxx
Fax: +39-02-23994568

PIC	***Legal name***
999509438	*DANMARKS METEOROLOGISKE INSTITUT*

Short name: DMI

Address of the organisation

Street Lyngbyvej 100

Town KOBENHAVN

Postcode 2100

Country Denmark

Webpage www.dmi.dk

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....yes

Legal person ....yes

Non-profit ....yes

International organisation ....unknown

International organisation of European interest ....unknown

Secondary or Higher education establishment ....unknown

Research organisation ....yes

Enterprise Data

SME self-declared status.................... unknown

SME self-assessment .... unknown

SME validation sme.................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: The Department of Research and Development

☐ not applicable

☒ Same as organisation address

Street: Lyngbyvej 100

Town: KOBENHAVN

Postcode: 2100

Country: Denmark

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **DMI**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Kristian Pagh** Last name: **Nielsen**

E-Mail: **kpn@dmi.dk**

Position in org.: Physicist/Meterologist

Department: The Department of Research and Development ☐ Same as organisation

☒ Same as organisation address

Street: Lyngbyvej 100

Town: KOBENHAVN Post code: 2100

Country: Denmark

Website: http://research.dmi.dk/staff/all-staff/kpn/

Phone 1: +45 3915 7403 Phone 2: *+xxx xxxxxxxxx* Fax: +45 3927 1080

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Bent Hansen	Sass	bhs@dmi.dk	

PIC	***Legal name***
999419422	*FONDAZIONE CENTRO EURO-MEDITERRANEO SUI CAMBIAMENTI CLIMATICI*

Short name: CMCC

Address of the organisation

Street VIA A IMPERATORE 16

Town LECCE

Postcode 73100

Country Italy

Webpage www.cmcc.it

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....no

Legal person .... yes

Non-profit ....yes

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....yes

Enterprise Data

SME self-declared status.................................11/05/2005 - no

SME self-assessment .... unknown

SME validation sme....................................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: ASC - Advanced Scientific Computing Division ☐ not applicable

☒ Same as organisation address

Street: VIA A IMPERATORE 16

Town: LECCE

Postcode: 73100

Country: Italy

Department 2

Department name: ODA - Ocean modeling and Data Assimilation Division ☐ not applicable

☐ Same as organisation address

Street: via M. Franceschini, 31

Town: Bologna

Postcode: 40128

Country: Italy

Dependencies with other proposal participants

Character of dependence	*Participant*	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **CMCC**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: ◯ Male ◉ Female

First name: **Silvia** Last name: **Mocavero**

E-Mail: **silvia.mocavero@cmcc.it**

Position in org.: Researched at ASC - Advanced Scientific Computing Division

Department: ASC - Advanced Scientific Computing Division ☐ Same as organisation

☒ Same as organisation address

Street: VIA A IMPERATORE 16

Town: LECCE Post code: 73100

Country: Italy

Website: www.cmcc.it

Phone 1: +39 0832 297304 Phone 2: +xxx xxxxxxxxx Fax: +xxx xxxxxxxxx

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Giulia	Galluccio	giulia.galluccio@cmcc.it	+390243986856

PIC	***Legal name***
996058081	*BULL SAS*

Short name: BULL

Address of the organisation

Street RUE JEAN JAURES 68

Town LES CLAYES SOUS BOIS

Postcode 78340

Country France

Webpage www.bull.com

Legal Status of your organisation

Research and Innovation legal statuses

Public body ....no

Legal person .... yes

Non-profit ....no

International organisation ....no

International organisation of European interest ......no

Secondary or Higher education establishment ....no

Research organisation ....no

Enterprise Data

SME self-declared status.................................. unknown

SME self-assessment .... unknown

SME validation sme....................................... unknown

Based on the above details of the Beneficiary Registry the organisation is not an SME (small- and medium-sized enterprise) for the call.

Department(s) carrying out the proposed work

Department 1

Department name: Bull BDS – Center for Excellence in Parallel Programming (CEPP)

☐ not applicable

☐ Same as organisation address

Street: 1, rue de Provence

Town: Echirolles

Postcode: 38432

Country: France

Dependencies with other proposal participants

Character of dependence	***Participant***	

Proposal ID **800897** *Acronym* **ESCAPE-2** *Short name* **BULL**

Person in charge of the proposal

The name and e-mail of contact persons are read-only in the administrative form, only additional details can be edited here. To give access rights and basic contact details of contact persons, please go back to Step 4 of the submission wizard and save the changes.

Title: Dr.

Sex: (•) Male ( ) Female

First name: **Xavier**

Last name: **Vigouroux**

E-Mail: **xavier.vigouroux@atos.net**

Position in org.: Head of CEPP

Department: Bull BDS – Center for Excellence in Parallel Programming (CEPP)

☐ Same as organisation

☐ Same as organisation address

Street: 1, rue de Provence

Town: Echirolles

Post code: 38432

Country: France

Website: https://atos.net/en/products/high-performance-computing-hpc/hpc-s

Phone 1: +33 6 86 28 50 81

Phone 2: +xxx xxxxxxxxx

Fax: +xxx xxxxxxxxx

Other contact persons

First Name	***Last Name***	***E-mail***	***Phone***
Erwan	RAFFIN	erwan.raffin@atos.net	
Medur	SRIDHARAN	medur.sridharan@bull.net	

Proposal ID **800897** *Acronym* **ESCAPE-2**

3 - Budget for the proposal

No	Participant	Country	(A) Direct personnel costs/€	(B) Other direct costs/€	(C) Direct costs of sub-contracting/€	(D) Direct costs of providing financial support to third parties/€	(E) Costs of inkind contributions not used on the beneficiary's premises/€	(F) Indirect Costs / € (=0.25(A+B-E))	(G) Special unit costs covering direct & indirect costs / €	(H) Total estimated eligible costs / € (=A+B+C+D+F +G)	(I) Reimburse-ment rate (%)	(J) Max.EU Contribution / € (=H*I)	(K) Requested EU Contribution/ €
1	Ecmwf	UK	552039	62000	0	0	0	153509,75	0	767548,75	100	767548,75	767548,75
2	Dkrz	DE	189000	15000	0	0	0	51000,00	0	255000,00	100	255000,00	255000,00
3	Mpi-m	DE	202500	10000	0	0	0	53125,00	0	265625,00	100	265625,00	265625,00
4	Mswiss	CH	334800	17000	0	0	0	87950,00	0	439750,00	100	439750,00	439750,00
5	Bsc	ES	162000	24000	0	0	0	46500,00	0	232500,00	100	232500,00	232500,00
6	Cea	FR	264024	21174	0	0	0	71299,50	0	356497,50	100	356497,50	356497,50
7	Lu	UK	267200	20000	0	0	0	71800,00	0	359000,00	100	359000,00	359000,00
8	Rmi	BE	299988	10000	0	0	0	77497,00	0	387485,00	100	387485,00	387485,00
9	Polimi	IT	180000	25000	0	0	0	51250,00	0	256250,00	100	256250,00	256250,00
10	Dmi	DK	215802	4026	32215	0	0	54957,00	0	307000,00	100	307000,00	307000,00

Proposal ID **800897** *Acronym* **ESCAPE-2**

11	Cmcc	IT	87500	16500	0	0	0	26000,00	0	130000,00	100	130000,00	130000,00
12	Bull	FR	173040	21355	0	0	0	48598,75	0	242993,75	100	242993,75	242993,75
	Total		2927893	246055	32215	0	0	793487,00	0	3999650,00		3999650,00	3999650,00

4 - Ethics issues table

1. HUMAN EMBRYOS/FOETUSES		Page
Does your research involve Human Embryonic Stem Cells (hESCs)?	○ Yes ◉ No	
Does your research involve the use of human embryos?	○ Yes ◉ No	
Does your research involve the use of human foetal tissues / cells?	○ Yes ◉ No	
2. HUMANS		**Page**
Does your research involve human participants?	○ Yes ◉ No	
Does your research involve physical interventions on the study participants?	○ Yes ◉ No	
3. HUMAN CELLS / TISSUES		**Page**
Does your research involve human cells or tissues (other than from Human Embryos/ Foetuses, i.e. section 1)?	○ Yes ◉ No	
4. PERSONAL DATA		**Page**
Does your research involve personal data collection and/or processing?	○ Yes ◉ No	
Does your research involve further processing of previously collected personal data (secondary use)?	○ Yes ◉ No	
5. ANIMALS		**Page**
Does your research involve animals?	○ Yes ◉ No	
6. THIRD COUNTRIES		**Page**
In case non-EU countries are involved, do the research related activities undertaken in these countries raise potential ethics issues?	○ Yes ◉ No	
Do you plan to use local resources (e.g. animal and/or human tissue samples, genetic material, live animals, human remains, materials of historical value, endangered fauna or flora samples, etc.)?	○ Yes ◉ No	
Do you plan to import any material - including personal data - from non-EU countries into the EU?	○ Yes ◉ No	
Do you plan to export any material - including personal data - from the EU to non-EU countries?	○ Yes ◉ No	
In case your research involves low and/or lower middle income countries, are any benefits-sharing actions planned?	○ Yes ◉ No	
Could the situation in the country put the individuals taking part in the research at risk?	○ Yes ◉ No	

7. ENVIRONMENT & HEALTH and SAFETY		Page
Does your research involve the use of elements that may cause harm to the environment, to animals or plants?	◯ Yes ◉ No	
Does your research deal with endangered fauna and/or flora and/or protected areas?	◯ Yes ◉ No	
Does your research involve the use of elements that may cause harm to humans, including research staff?	◯ Yes ◉ No	

8. DUAL USE		Page
Does your research involve dual-use items in the sense of Regulation 428/2009, or other items for which an authorisation is required?	◯ Yes ◉ No	

9. EXCLUSIVE FOCUS ON CIVIL APPLICATIONS		Page
Could your research raise concerns regarding the exclusive focus on civil applications?	◯ Yes ◉ No	

10. MISUSE		Page
Does your research have the potential for misuse of research results?	◯ Yes ◉ No	

11. OTHER ETHICS ISSUES		Page
Are there any other ethics issues that should be taken into consideration? Please specify	◯ Yes ◉ No	

I confirm that I have taken into account all ethics issues described above and that, if any ethics issues apply, I will complete the ethics self-assessment and attach the required documents. ☒

How to Complete your Ethics Self-Assessment

5 - Call specific questions

Extended Open Research Data Pilot in Horizon 2020

If selected, applicants will by default participate in the Pilot on Open Research Data in Horizon 2020[1] , which aims to improve and maximise access to and re-use of research data generated by actions.

However, participation in the Pilot is flexible in the sense that it does not mean that all research data needs to be open. After the action has started, participants will formulate a Data Management Plan (DMP), which should address the relevant aspects of making data FAIR – findable, accessible, interoperable and re-usable, including what data the project will generate, whether and how it will be made accessible for verification and re-use, and how it will be curated and preserved. Through this DMP projects can define certain datasets to remain closed according to the principle "as open as possible, as closed as necessary". A Data Management Plan does not have to be submitted at the proposal stage.

Furthermore, applicants also have the possibility to opt out of this Pilot completely at any stage (before or after the grant signature). In this case, applicants must indicate a reason for this choice (see options below).

Please note that participation in this Pilot does not constitute part of the evaluation process. Proposals will not be penalised for opting out.

We wish to opt out of the Pilot on Open Research Data in Horizon 2020.	(●) Yes	( ) No

If opting out please indicate the reason(s) for not being able to participate in the Pilot:

- the project does not generate any data	☒
- to allow the protection of results (e.g. patenting)	☐
- incompatibility with the need for confidentiality linked to security	☐
- incompatibility with privacy/data protection	☐
- achievement of the project's main aim would be jeopardised	☐
- other legitimate reasons	☐

Further guidance on open access and research data management is available on the participant portal: http://ec.europa.eu/research/participants/docs/h2020-funding-guide/cross-cutting-issues/open-access-dissemination_en.htm and in general annex L of the Work Programme.

[1] *According to article 43.2 of Regulation (EU) No 1290/2013 of the European Parliament and of the Council, of 11 December 2013, laying down the rules for participation and dissemination in "Horizon 2020 - the Framework Programme for Research and Innovation (2014-2020)" and repealing Regulation (EC) No 1906/2006.*

Energy-efficient SCalable Algorithms for weather and climate Prediction at Exascale

ESCAPE2

Research and Innovation Action

Responding to: FETHPC-02-2017 Transition to Exascale Computing [Subtopic e) Mathematics and algorithms for extreme scale HPC systems and applications working with extreme data]

List of participants

Participant No	**Participant organisation name**	**Short name**	**Country**
1 (Coordinator)	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	ECMWF	International
2	DEUTSCHES KLIMARECHENZENTRUM GMBH	DKRZ	DE
3	MAX-PLANCK-GESELLSCHAFT ZUR FÖRDERUNG DER WISSENSCHAFTEN EV	MPIM	DE
4	EIDGENÖSSISCHES DEPARTEMENT DES INNERN	MSWISS	CH
5	BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION	BSC	ES
6	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES	CEA	FR
7	LOUGHBOROUGH UNIVERSITY	LU	UK
8	INSTITUT ROYAL METEOROLOGIQUE DE BELGIQUE	RMI	BE
9	POLITECNICO DI MILANO	POLIMI	IT
10	DANMARKS METEOROLOGISKE INSTITUT	DMI	DK
11	FONDAZIONE CENTRO EURO-MEDITERRANEO SUI CAMBIAMENTI CLIMATICI	CMCC	IT
12	BULL SAS	BULL	FR

Coordinator: Dr Peter Bauer, ECMWF, peter.bauer@ecmwf.int

Contents

1 Excellence

At the 2015 Paris Climate Conference, leaders from 194 countries of the world unanimously acknowledged the serious threat posed by anthropogenic emissions of greenhouse gases[1]. Society must now become resilient to changes in climate over coming decades. Most importantly, this will require quantitative estimates of the changing character of climate extremes. Such extremes include not only exceptional weather events such as violent wind storms and flash floods, but also persistent anomalies in planetary–scale circulation patterns, which lead to pervasive flooding in some regions and seasons, and long-lived drought and extremes of heat in others. However, providing such information at present goes beyond the capability of even the most advanced weather prediction and climate simulation models and will require a step change in the quality of global weather and climate prediction. In particular, the resolution of these models must increase to a level which allows both ocean eddies and individual cloud systems to become represented explicitly (Bony et al. 2015, Schneider et al. 2017). This is the most promising approach for obviating long-standing model biases.

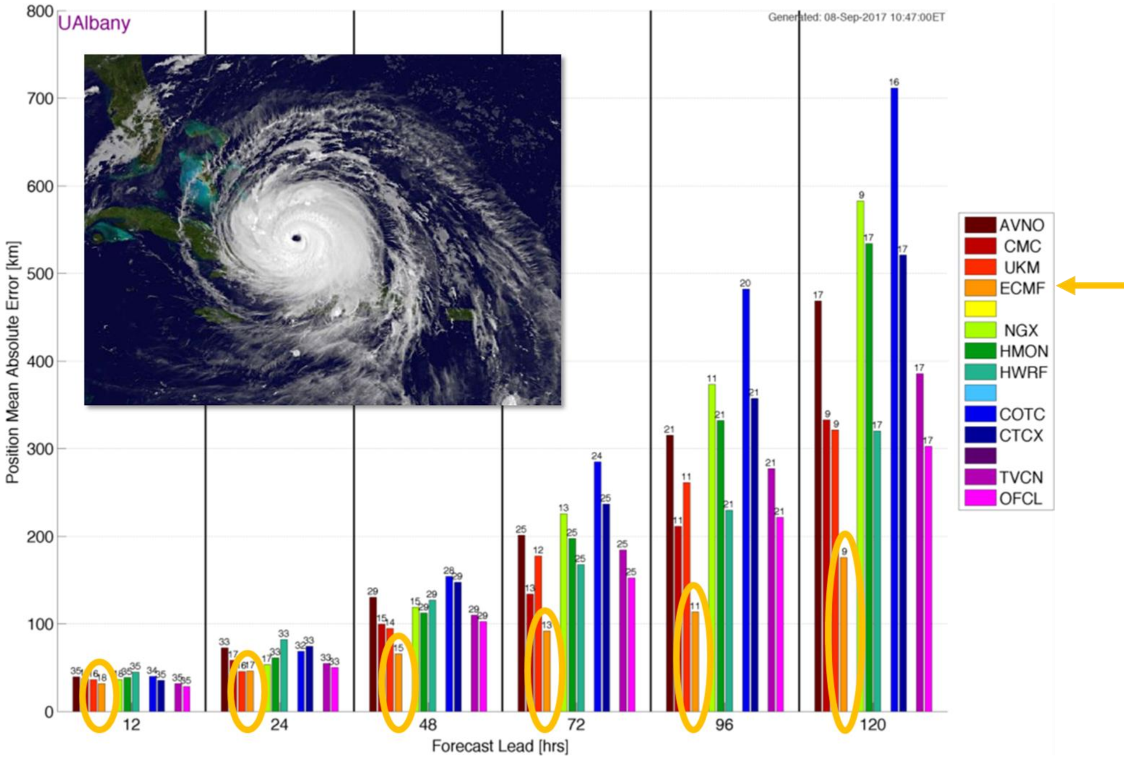

Figure 1: Mean absolute track forecast error for Hurricane Irma of different forecast models over the period of 30 August to 8 September 2017. ECMWF's forecast (in orange, labelled ECMF) indicates the lead the European model has over other centres in the US, Canada and the UK (Figure courtesy U Albany).

The recent episode of extreme tropical cyclone events affecting the Caribbean and the US clearly demonstrates the need for accurate and reliable prediction capability. Hurricane Harvey alone is estimated to have caused 180 billion US Dollar damage[2], and hurricane Irma is considered the largest and most powerful hurricane occurring in the recorded history of the Atlantic basin.

[1] https://www.ipcc.ch/report/ar5/syr/

[2] http://fortune.com/2017/09/03/hurricane-harvey-damages-cost/

Figure 1 shows the mean absolute track error forecasts for Irma from various centres around the world. Accurate track errors are crucial for guiding evacuation plans and risk management. The ECMWF forecasts have been by far the most accurate (orange bars and circles) not only for Irma. However, errors still amount to several hundreds of kilometres for forecasts beyond a few days, and only substantial improvements to key model drivers like spatial resolution and complexity will eventually help to achieve sufficient skill.

Only by advancing Earth system modelling and harnessing exascale technologies to substantially accelerate weather and climate prediction models this scientific challenge can be met. **This clearly poses an exascale computing challenge to our community** (Bauer et al. 2015).

The recognition of this fact explains why weather and climate prediction has produced outstanding momentum for the development of new mathematical methods and numerical algorithms towards exascale computing and extreme data handling in recent years. Leadership in the ESCAPE[3] project and ESiWACE[4] centre of excellence as well as partnership in projects like CRESTA[5], NextGenIO[6] and the recently started EuroEXA[7] co-design effort are proof that weather and climate prediction drive key areas of development for the high-performance computing community already now, and why it represents a must-have application for next-generation hardware configurations introduced by extreme-scale demonstrators eventually leading to European exascale platforms in the future.

ESCAPE-2 paves the way for the European weather and climate prediction community and beyond towards the exascale. This will be achieved by "connecting the dots" (see Section 1.3): **Developing bespoke and novel mathematical and algorithmic concepts, combining them with proven methods**, and thereby reassessing the mathematical foundations forming the basis of Earth system models. Moreover, ESCAPE-2 invests in significantly more productive programming models for the weather-climate community through which novel algorithm development will be accelerated and future-proofed. Eventually, the project aims at **providing exascale-ready production benchmarks to be operated on extreme-scale demonstrators and beyond**. Moreover, ESCAPE-2 combines **cross-disciplinary uncertainty quantification tools** (URANIE) for high-performance computing, originating from the energy sector, with ensemble based weather and climate models to quantify the effect of model and data related uncertainties on forecasting – a capability, which weather and climate prediction has pioneered since the 1960s.

[3] Energy-efficient Scalable Algorithms for Weather Prediction at Exascale (www.hpc-escape.eu)
[4] Centre of Excellence in Simulation of Weather and Climate (www.esiwace.eu)
[5] Collaborative Research into Exascale Systemware, Tools and Applications (www.cresta-project.eu.)
[6] Next Generation I/O for the Exascale (www.nextgenio.eu)
[7] Co-designed Innovation and Systems for Resilient Exascale Computing in Europe: From Applications to Silicion (www.euroexa.com)

1.1 Objectives

ESCAPE-2 will develop world-class, extreme-scale computing capabilities for European operational numerical weather and climate prediction, and provide the key components for representative benchmarks to be deployed on extreme-scale demonstrators and beyond.

To achieve this goal, ESCAPE-2 has set five top-level objectives:

Table 1: ESCAPE-2 top-level objectives.

Objective	Achieved by	Addressed in work package
1. Combine frontier research **on mathematics and algorithm development** and extreme-scale, high-performance computing applications with novel hardware technology: → to design scientifically flexible and sustainable weather and climate prediction systems	- implementing **data structures and tools supporting parallel computation** of dynamics and physics on **multiple scales and multiple levels**; - combining **highly-scalable spatial discretization** with **proven large time-stepping techniques** to optimize **time-to-solution**; - applying **machine learning** for accelerating complex sub-components - combining multi-grid tools, iterative solvers, and overlapping computations with flexible-order spatial discretization to **strengthen algorithm resilience** against soft or hard failure.	WP1
2. Develop and apply a **domain-specific language (DSL) concept for the weather and climate** community: → to maximize flexibility, programmability and performance portability to heterogeneous hardware solutions across different weather and climate models	- defining a weather and climate **DSL concept** for a comprehensive set of models; - developing and demonstrating **an open source toolchain for code adaptation and performance portability** to different hardware architectures; - sustaining community-wide **code usability and maintainability** beyond the lifetime of the project.	WP1, WP2
3. Establish **weather and climate model benchmarks** based on world class European prediction models: → to enable deployment on energy efficient and heterogeneous HPC architectures, in particular Extreme-scale Demonstrators (EsD)	- setting up a **benchmark hierarchy** of representative models for atmosphere and ocean together with machine-wide simulated, domain specific workflows; - incorporating seamlessly novel and **disruptive numerical algorithms and mathematics**; - ensuring portability through a weather and climate **domain-specific language**; - representing (world-) **leading European weather and climate prediction models** for both atmosphere and ocean.	WP1, WP2, WP3

4. Develop a cross-disciplinary Verification, Validation, Uncertainty Quantification (**VVUQ) framework** → to establish exascale-ready verification and uncertainty quantification tools for weather and climate prediction and beyond.	- implementing highly **non-linear and multi-dimensional** weather & climate dwarfs in a cross-disciplinary VVUQ framework; - estimating mathematical, numerical and data parameter related **uncertainties** on simulation performance; - demonstrating **VVUQ capability** across a hierarchy of high-dimensional modelling systems.	WP4, WP3
5. Produce an **open-source software framework**: → to accelerate mathematical algorithm development, foster continued leadership of European weather and climate prediction models and sustain commercialisation of weather-dependent innovative products and services in Europe	- providing an **open-source DSL toolchain software** and support beyond the project lifetime **to sustain & accelerate novel algorithm development and ensuring performance portability** to emerging HPC hardware; - reforming fundamentally the way weather and climate modelling is performed and verified; - supporting formally incremental upgrades from **novel mathematical and algorithmic concepts** into large-scale weather and climate model legacy codes; - disseminating and providing training on novel code development concepts **through direct public engagement, stakeholder communication and through the ESiWACE Centre of Excellence**; - providing sustainability for **Copernicus services** that deliver vital information for European society.	WP1, WP2, WP3, WP4, WP5

Why do we focus on these objectives?

Weather and climate simulations have been intimately connected with progress in supercomputing since the first numerical forecast was made about 65 years ago. **Today's** operational prediction centres run global model forecasts down to 9 km resolution into the medium range (3-30 day forecasts), multiple forecasts (ensembles) at 18 km resolution up to one month ahead, and climate simulations at 100 km resolution for decades and centuries.

The initial conditions are produced with complex mathematical algorithms permitting the assimilation of 60 million observations per day. While the number of observations is expected to grow more slowly (one order of magnitude), model resolution and complexity are expected to increase much further (three orders of magnitude) over the next decade. Model complexity includes chemical processes in the atmosphere, and fully coupled simulations with ocean, sea-ice and land surfaces. The uncertainty of small-scale processes with increased resolution naturally requires the use of suitably perturbed simultaneous forecasts, and likely analyses, to be executed in parallel. **Weather and climate prediction models are expected to converge towards so-called Earth-system models** representing, in great detail, all relevant physical and (bio-)chemical processes. Just like today's weather forecasting models, Earth-system models will be run in ensemble mode to produce predictions of future states and their uncertainties.

The key challenge is to perform simulations at a rate of at least 1 forecast year / 1 wall-clock day with resolutions of 1 km (National Research Council 2012). ESCAPE-2 will pave the way towards this goal through the combination of novel mathematics and algorithms and new programming models exploiting the parallel processing capability offered by new processor technology.

The current forecast production rate capabilities are illustrated in Figure 2 suggesting that - even with perfect scalability - acceleration by at least a factor of O(100) is required while acquiring and operating an O(100) bigger HPC system with today's technology is clearly not affordable.

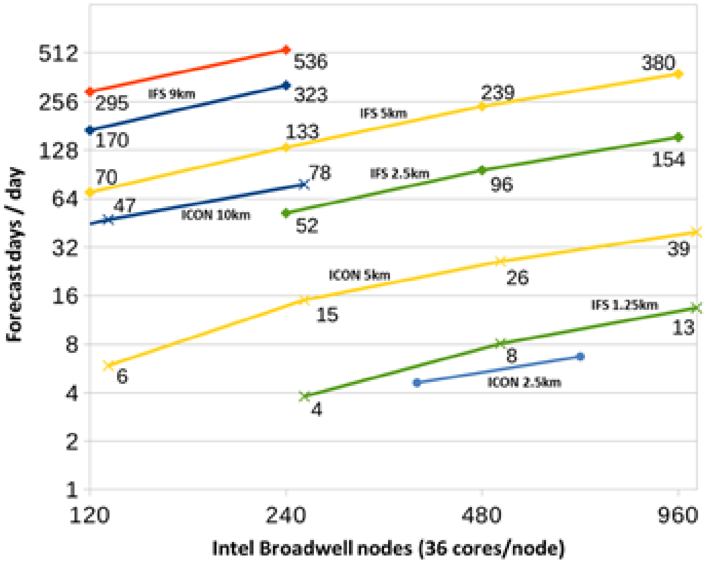

Figure 2: Time-to-solution tests with ECMWF's IFS[8] and MPIM/DWD's[9] ICON[10] for various spatial resolutions using Intel Broadwell node allocations (36 core) on ECMWF's Cray XC-40. The simulations were run within the ESiWACE project. At resolutions of 1 and 2.5 km both IFS and ICON were run with only half the operational vertical levels. At 9 km horizontal grid-spacing, IFS was run with both single and double precision (red vs. blue curves), all other runs were performed at single precision. The operational target is at least one year/day including I/O; however, all runs here did not account for I/O.

Already today, numerical weather and climate prediction and, more generally, computing centres face substantial challenges due to the rising cost of energy associated with running complex high-resolution forecast models on more and more processors and the likelihood that Moore's law will soon reach its limit, with microprocessor feature density (and performance) no longer doubling every two years. But the biggest challenge to state-of-the-art computational weather and climate prediction services arises from their own software productivity shortfall. The application software at the heart of all services throughout Europe is not equipped to efficiently adapt to the rapidly evolving heterogeneous hardware provided by the supercomputing industry. **If this challenge is not addressed, it will have dramatic negative consequences for weather and climate prediction and the associated European service infrastructure.**

ESCAPE-2 joins leading European weather and climate prediction models, namely IFS (atmosphere), NEMO[11] (ocean) and ICON (atmosphere and ocean) to comprehensively represent the state-of-the-art of Earth system science.

The first key ingredient for obtaining both time-to-solution and energy-to-solution is to **combine novel mathematical methods and algorithms for reducing data movement and near-perfect scalability** with proven large time-step forward-in time methods. The second key ingredient is provided by **hierarchical multi-level and multi-grid tools** to optimally address the challenging multi-scale character of atmospheres and oceans. The third ingredient is based on **machine learning techniques** aiming at self-learning statistical information exploration outside the critical forecast path, and thus providing a trade-off between accuracy and time-to-solution acceleration within the critical path. The fourth key ingredient is the development of an **open-source, weather and climate prediction domain-specific language** (DSL) **toolchain** that ensures performance portable access to

---

[8] Integrated Forecasting System

[9] Deutscher Wetterdienst

[10] ICOsahedral Non-hydrostatic model

[11] Nucleus for European Modelling of the Ocean

the parallel processing power in multiple heterogeneous HPC architectures, while maintaining programmability and sustainability of large weather and climate model codes by expert scientists. The fifth corner stone is provided by a **generic, cross-disciplinary VVUQ framework**.

The final ESCAPE-2 outcome – enabling the transition to exascale for weather and climate prediction – will be the **foundation of a representative domain-specific High Performance Climate and Weather (HPCW) benchmark**. HPCW will include much more relevant and realistic, near-operational workloads than available from general-purpose benchmarks such as HPL (High Performance LINPACK[12]) and High Performance Conjugate Gradient Benchmark (HPCG[13]).

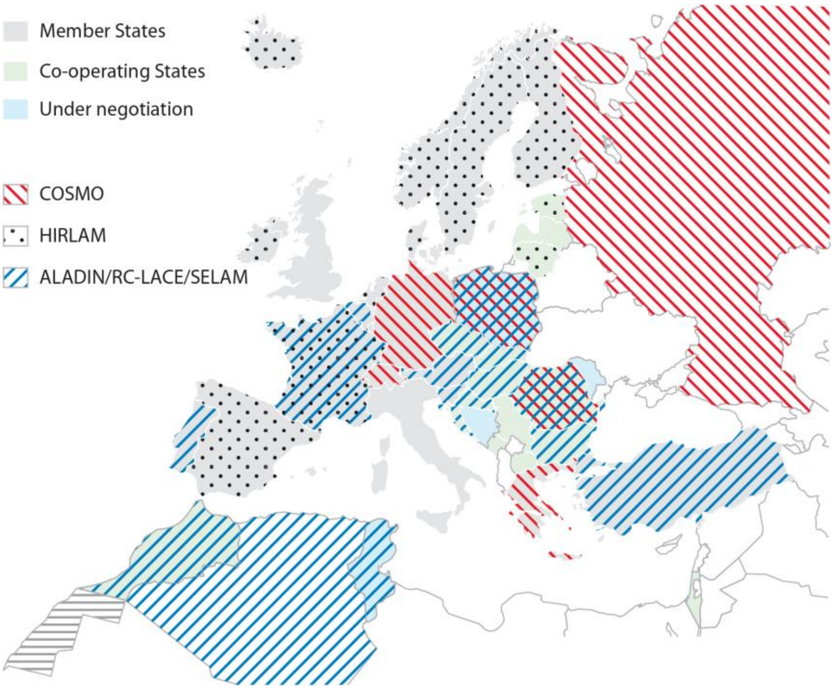

Figure 3:Weather and climate prediction consortia represented in ESCAPE-2: ECMWF, the world leading weather prediction centre, its member and co-operating states, the associated regional consortia ALADIN and HIRLAM, and the MPIM/DWD ICON consortium (to be fully implemented in 2022) targeting global and regional weather and climate prediction.

ESCAPE-2 builds on the success of the ESCAPE project that has **pioneered the concept of weather and climate dwarfs** and has laid the foundation by extending a domain-specific language to unstructured grids used by global weather and climate models. Weather and climate dwarfs allow to break down accumulated and required knowledge from legacy codes into functional units for which tailored numerical solutions can be designed and adapted to emerging processor technology for optimal performance at scale.

ESCAPE-2 feeds into ESiWACE that disseminates the future evolution of weather and climate models to an operational community. This concept is rapidly emerging as an accepted development template for the entire weather and climate prediction community in the future.

With ESCAPE-2, ECMWF is leading a consortium composed of weather and climate prediction centres operating at global as well as European regional scales (Figure 3), with university institutes performing research on numerical methods and novel code optimization techniques, and HPC centres with vast experience in scalable code development and diverse processor technologies. A leading European HPC provider will act as the benchmark recipient producing performance assessment for different architectures. Due to the focus on operational prediction systems as well as climate services there is a direct benefit for all countries shown in Figure 3 which all operate prediction systems based on the ESCAPE-2 models. This includes the European Commission's

---

[12] https://www.top500.org/project/linpack/
[13] http://www.hpcg-benchmark.org/

Copernicus Climate Change and Atmospheric Monitoring Services (C3S and CAMS), both being hosted by ECMWF on behalf of the Commission.

1.2 Relation to the work programme

ESCAPE-2 responds to the call FETHPC-2–2017 – Transition to Exascale Computing with a specific focus on sub-topic **e) 'Mathematics and algorithms for extreme-scale HPC systems and applications working with extreme data'**. This topic was introduced with the European Technology Platform for HPC Strategic Research Agenda Update (ETP4HPC SRA[14]) to acknowledge the fundamental role that mathematics and algorithms play in making efficient use of future HPC technologies for real-world applications.

By recognising the dilemma of producing performance-portable mathematics and algorithms for an entire community and the need to optimally exploit a variety of emerging hardware technologies, ESCAPE-2 will also partly address sub-topic **a) 'High productivity programming environments for exascale'**. ESCAPE-2 will provide weather and climate DSL concepts, which will be crucial for being able to assess new and disruptive algorithms in a general way across different models and HPC platforms.

ESCAPE-2 addresses the call by targeted research and innovation actions to ensure continued European world-leadership in weather and climate prediction. ESCAPE-2 focuses on a unique application area with enormous socio-economic relevance for which the successful step towards exascale computing will depend exactly on advances in the areas covered by sub-topic (e). **ESCAPE-2 will therefore leverage the transition from peta- to exascale computing** for an entire **community** and thus **combine world-leading European applications with the emerging leadership of European technology.** The ETP4HPC SRA defines the need for HPC to address the grand societal and economic challenges in Europe. By providing enabling technology for world-leading European weather and climate prediction capability, ESCAPE-2 directly contributes to (i) safety of the population in case of natural or industrial disasters, (ii) secure, clean and efficient energy, (iii) supply of raw materials, resource efficiency, and climate action.

ESCAPE-2 **addresses the aims and scope of FETHPC-2-2017 through its top-level objectives** (Section 1.1), particularly advancing the '**exascale and extreme scale**', '**ease of use**' and '**efficiency**' as defined in the overall technical research priorities of the SRA.

A detailed overview of how ESCAPE-2 targets the specific topics of the call - with reference to the project's work packages and tasks - is given in Table 2 below.

Table 2: ESCAPE-2 relation to work programme topics of FETHPC-2-2017.

Work Programme Topic	Work Programme Text	ESCAPE-2 Relation
Specific challenge	"Take advantage of the full capabilities of exascale computing, in particular through high-productivity programming environments,…,supercomputing for extreme data and emerging HPC use modes, mathematics and algorithms for extreme scale HPC systems for existing or visionary applications, including data-intensive and extreme data applications in scientific areas such as physics, chemistry, biology, life sciences,	ESCAPE-2 will advance Europe's world-leading weather and climate prediction capability to the exascale era through fundamental mathematical and algorithmic research. The step towards exascale readiness will be made by substantially enhancing the time-to-solution effectiveness of highly scalable, flexible-order spatial discretizations, introducing fault tolerant algorithms

[14] http://www.etp4hpc.eu/strategy/strategic-research-agenda/

	materials, climate, geosciences, etc."	supported by hierarchical multigrid tools and controlled sensitivity to numerical precision, introducing surrogate neural network models, and evaluating parameter uncertainty in a generic trans-disciplinary VVUQ framework. ESCAPE-2 will comprehensively represent extreme data applications in the scientific area of weather & climate prediction by incorporating algorithmic motives from leading European atmosphere and ocean models into large-scale model benchmarks ready for tests on pre-exascale systems, in particular Extreme-scale Demonstrators. **WP1**: **Task 1.5, WP3***:* **Task 3.5**
e) Mathematics and algorithms for extreme scale HPC systems and applications working with extreme data	"Mathematical methods, numerical analysis, algorithms and software engineering for extreme parallelism." "Novel and disruptive algorithmic strategies should be explored to minimise data movement as well as the number of communication and synchronisation instances in extreme computing."	ESCAPE-2 will combine well-established large-timestep forward-in-time computing with highly scalable, flexible order spatial discretization with minimal data movement. Moreover, ESCAPE-2 will introduce hierarchical fault tolerant algorithmic concepts for local failure recovery. **WP1**: **Tasks 1.1, 1.2, 1.3**
	"Multi-scale, multi-physics and extreme data."	Neural networks as surrogate models for selected physical processes will be developed with appropriate very large training data sets of ECMWF's meteorological data archive and specific highly-accurate models. **WP1**: **Task 1.4, WP3: Task 3.3**
	"Parallel-in-time methods may be investigated to boost parallelism of simulation codes across a wide range of application domains."	Parallel-in-time methods will not be directly explored in ESCAPE-2 as they are not readily applicable or sufficiently mature for highly nonlinear applications with extremely high numbers of degrees of freedom as found in complex weather prediction and climate simulations. However, the multi-grid methods developed in ESCAPE-2 will provide a useful infrastructure for future investigations on time-parallelism, and a dual time-stepping approach will be explored to boost performance. **WP1**: **Task 1.3**
	"Taking into account data-related uncertainties is essential for the acceptance of numerical simulation in decision making; a unified European VVUQ (Verification Validation and Uncertainty Quantification) package for Exascale computing should be provided by improving methodologies and solving problems limiting usability for very large computations on many-core configurations."	A cross-disciplinary VVUQ package employed on diverse HPC platforms and applied to weather and climate prediction will be pioneered in ESCAPE-2 that combines an existing industrial set-up (URANIE from the energy sector) with ensemble based methods used in operational weather prediction to quantify the impact of input data and

	"Quantification of uncertainties and noise." "Access to the VVUQ techniques for the HPC community should be facilitated by providing software that is ready for deployment on supercomputers."	parameter uncertainties on forecast predictions. The impact of uncertainty and noise will be investigated at different levels of complexity (weather & climate dwarfs and models) within the VVUQ framework. **WP4**: **Tasks 4.1-4.6**
a) High productivity programming environments for exascale	"Development of more productive programming models and environments." "To provide simplified development and to ensure the maintainability of domain-specific languages (DSLs), DSL frameworks are required which target a general-purpose stable programming model and runtime."	Through the development of a DSL for programming weather and climate models, ESCAPE-2 will achieve performance portability, accelerate algorithm development, and increase domain-scientist productivity through a higher level of abstraction, increased conciseness of the programmed mathematical algorithm and the application of newly developed front-end parser and a reliable DSL toolchain operating on a comprehensive intermediate representation of the science code. Continued interoperability with legacy code is an important design feature. **WP2**: **Tasks 2.1-2.4**
	"Easier combination of different programming models, and using increased intelligence throughout the programming environment. Managing data transfers, data locality and memory management, including support for heterogeneous and reconfigurable systems Dealing with inter-application dynamic load balancing and malleability, adapting to changes in the number of processors." "APIs, runtime systems and the underlying libraries should support auto-tuning for performance and energy optimisation."	The DSL toolchain will be able to use domain knowledge to automatically manage and optimise data transfers such as inter-node communication, data locality through automatic loop fusion and caching and transparently support different hardware targets such as accelerators. Performance models can be integrated into the toolchain that can reorder and fuse user computations in order to maximise the data locality present in the algorithms. The work will build on defined domain decomposition strategies and DSL on-node optimisations. The DSL toolchain includes optimization steps that facilitate auto-tuning. **WP2**: **Tasks 2.1-2.4**

1.3 Concept and methodology

1.3.1 Overall concept

Europe's proven and fastest algorithms for numerical weather and climate prediction, including both ocean and atmosphere components, represent substantial financial and intellectual investments, often in the form of up to millions of lines of code. Substantial progress has been made in improving the scalability of these codes which has enabled Europe's first ever prototype global kilometre-scale simulations. This capability is expected to eliminate climate uncertainty associated with parametrised processes such as clouds and convection (Schneider et al. 2017).

But fundamental changes are required to enable optimisation and adaptation for emerging HPC architectures as well as to obtaining further increases in parallelism towards exascale. A particular predicament that substantially slows down novel mathematical algorithm development is the need to adapt to conflicting HPC optimisation requirements. The ESCAPE project has addressed this issue based on so-called weather and climate dwarfs that allow to focus adaptation and optimisation on very specific compute and data movement patterns. Dwarfs represent the wide range of different domain specific algorithms that are key components of weather and climate models, and they contribute significantly to the overall workload. Focussing on dwarfs has clearly facilitated algorithm development.

However, porting and maintaining models on emerging HPC architectures continues to be a significant concern and has substantially delayed the development of novel algorithms. As a result, time- and energy-to-solution have not accelerated at a rate producing affordable exascale simulations of weather and climate any time soon.

ESCAPE-2 recognises that disruptive changes are not only required for making algorithm choices but also for defining software development pathways (Figure 4).

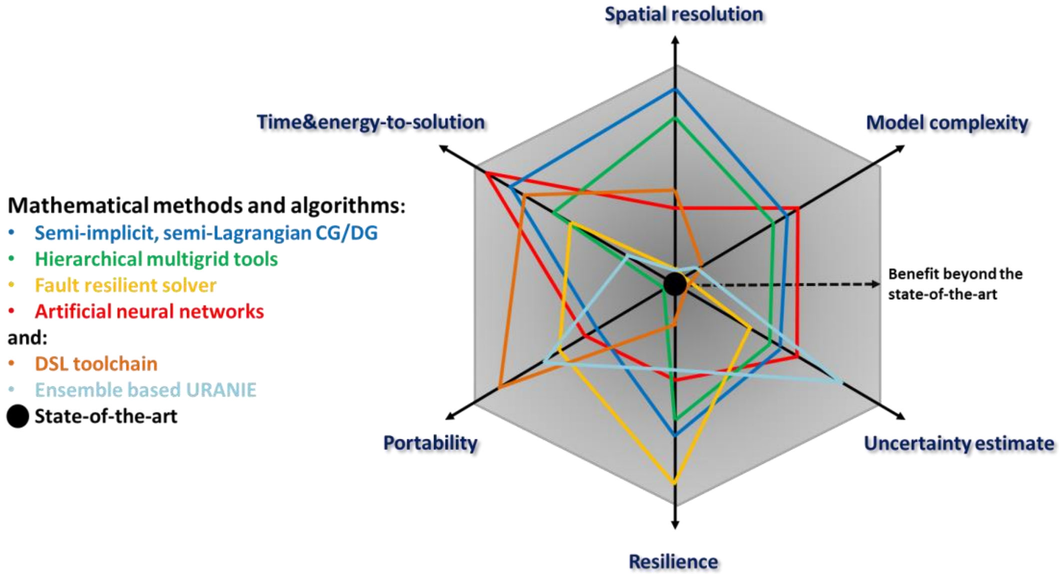

Figure 4: Aiming at the combination of performance, accuracy, resilience and portability for weather and climate prediction through novel mathematical and algorithmic methodologies, programming and uncertainty estimation.

ESCAPE-2 seeks a more fundamental and domain-specific solution **to accelerate mathematical algorithm development and to accelerate the associated research-to-operation (R2O) transition process** by including the compatible back-integration of optimised algorithms into larger scale legacy codes and relevant HPC benchmarks. DSLs are a promising avenue to close the emerging software productivity gap, and have received increased interest in the last decade. Thus the overall concept in ESCAPE-2 for enabling the transition of weather and climate prediction to the exascale is **fundamentally based on a multi-fold approach:**

1. Develop large time-step advection algorithms for highly-scalable, communication-minimising, high-order CG/DG discretisations with very large parallel-domain decompositions and potentially higher flop-rates, resulting in closer-to-peak HPC use. This work is necessary to achieve a **breakthrough for currently used algorithms in weather and climate prediction that are highly scalable, but inefficient in either or both time- and energy-to-solution** and where acceleration through hardware optimisation is likely to be insufficient. The proposed work is also a significant stepping stone towards the potential future use of time-parallel algorithms that would promise further scalability gains, but which are not sufficiently mature for practical applications in weather and climate prediction today.
2. Provide **robustness and resilience to HPC system delays or failures, to changes in precision, and to changes in the algorithmic concepts affecting accuracy**. This is developed in a dual approach: On the one hand, resilience is likely improved by suitably modifying existing iterative

Krylov-solvers through introducing pipelining concepts (Sanan et al. 2016) and by further developing multi-grid algorithms and pairing these with overlapping and one-sided communication concepts. On the other hand, a **generalised VVUQ approach is pursued** in this proposal that will share concepts between the weather and climate community and other disciplines (atomic energy and alternative energies) **via the URANIE platform.** Weather and climate prediction is already at the forefront of highly scalable and time-critical computer applications running routinely on up to O(10^5 - 10^6) computer cores, and is routinely evaluating the uncertainty associated with their predictions through ensemble forecasts. ESCAPE-2 will perform VVUQ-type analyses on high performance computing facilities, assessing scalability and generalisation of URANIE beyond its initial design. The project will provide a comprehensive VVUQ package for addressing parameter uncertainty in algorithms, and for potentially measuring the impact of (hard or soft failure) resilient algorithms on the estimated overall forecast prediction uncertainty.

3. **Explore machine learning techniques** as they have the potential to make weather forecasts substantially fasters. These techniques are very unlikely to replace the physical modelling from basic principles due to their many degrees of freedom and strong non-linearities, but they are **appropriate tools for accelerating models by replacing time consuming and uncertain parts of the model physics**, for instance radiation, cloud microphysics and convection parameterisations (Johnson et al. 2015). Radiative transfer poses a problem of specific computational complexity as the radiative intensity (the radiance) varies in three spatial dimensions, two directional dimensions and the spectral dimension. If the polarisation state is accounted for, the problem becomes further complicated. Solving this problem fully is not feasible for current O(10 km) resolution global models, and much less so at O(1 km) resolution. Here, training artificial neural networks (ANN) on accurate radiative transfer computations will be explored in ESCAPE-2 to achieve more accurate and faster computations (Stamnes et al. 2017; Fan et al. 2017). This concept will be **also included in the VVUQ framework** to estimate the impact of using ANN on overall forecast uncertainty.

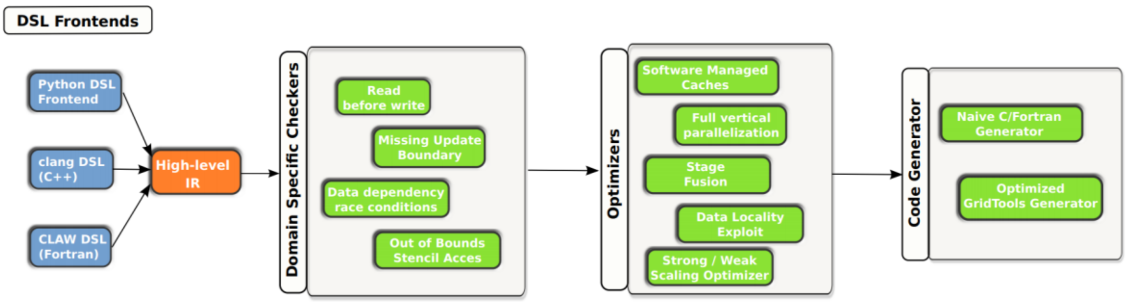

Figure 5: Design of a modular DSL toolchain for Weather & Climate dwarfs.

4. Formally decouple the development of key algorithmic motives from the generation of optimized, machine-dependent code at the highest possible level by **defining and developing a DSL toolchain** (Figure 5) and **applying** it to a **comprehensive list of algorithmic motives (dwarfs) in weather and climate prediction**. This bridges the gap between HPC optimisation/porting experts and novel mathematical algorithm developers and is achieved through the community development and buy-in of an **easy-to-learn front-end to a DSL in which domain scientists express key mathematical and algorithmic concepts,** as represented by weather and climate models dwarfs or full models. The development includes the definition and implementation of **a high-level intermediate representation (HIR),** which provides a specification of a comprehensive list of algorithmic motives in a way that is **agnostic to the programming language** or to the weather and climate prediction model it is applied in. The latter enables a separation of concerns in the development of **subsequent code analysis** and

code generation steps for HPC architecture-specific optimisations. The resulting code is compiled with standard and vendor supplied compilers.

Experience with large-scale code adaptation to new HPC platforms and HPC procurements, exploring novel architectures, reveals that the availability of meaningful and maintainable benchmarks representative of operational numerical weather and climate prediction system workloads is essential. The evaluation and implementation of new programming and parallelization paradigms, given the accelerating progress in HPC technologies, challenges present benchmarks and prohibits the effective uptake of new mathematics and algorithms (**WP1**) and programming models (**WP2**) in performance assessments. In order to ensure the relevant representation and extension of the successful dwarf concept, new dwarfs will be developed as well as a hierarchy of benchmarking components representing the key elements in the work flow of weather and climate prediction.

Based on this ESCAPE-2 will establish a **High Performance Climate and Weather (HPCW) benchmark** for (pre)-exascale application of climate and weather codes that can be demonstrated on emerging HPC platforms. The participating reference models, from which algorithmic motives will be derived, include the global atmosphere models ICON co-developed by DWD and MPIM, the IFS of ECMWF, the ocean model 'ICON-Ocean' and the community ocean model NEMO). In addition, an entirely new workload simulator (Kronos) developed within the NextGenIO project will be employed to assess the performance of (simulated) model benchmarks at scale.

1.3.2 Links with national and international research activities

The ESCAPE-2 consortium members are involved in several national, international and European projects that provide relevant input to ESCAPE-2. ESCAPE-2 builds on the innovative outcomes of ESCAPE but ingests methodologies, code and knowledge from the projects listed below, in addition to the substantial operational weather forecasting and climate projection experience of the project partners.

Table 3: Recent projects with partner involvement relevant for ESCAPE-2.

Project	Outputs provided to ESCAPE-2	ESCAPE-2 partners involved
ESCAPE	Weather & climate dwarfs for energy-efficient, scalable algorithms towards exascale; accelerator implementations; DSL concepts.	ECMWF, MSWISS, RMI, DMI, LU, BULL
PantaRhei	Mathematical and numerical technology for the next generation of weather and climate models.	ECMWF
NextGenIO	I/O and work flow performance simulator for benchmarking.	ECMWF
EuroEXA	Optimizations for ARM based HPC platforms, including the use of DSL toolchains.	ECMWF, BSC
ESiWACE	Variable precision algorithms, km-scale weather and climate model benchmarks.	DKRZ, MPIM, ECMWF, BSC, BULL, CMCC
URANIE	Generic VVUQ framework.	CEA
PASC	GridTools DSL; DSL concepts and developments.	MSWISS
GNCS	Semi-implicit, semi-Lagrangian DG algorithms.	POLIMI
ISENES-2 (ended)	Benchmarking of coupled climate models.	DKRZ, MPIM, MSWISS, BSC, CMCC

1.3.3 Overall methodology

The overall approach and methodology adopted in ESCAPE-2 is based on **three key principles** relevant for a successful and highly technical project that combines elements of information technology, frontier research on mathematical algorithms, and uncertainty quantification in time-critical applications:

1. Provide pathways for further increases in scalability and algorithm resilience (against hard and soft failures) by eliminating existing limitations on time- and energy-to-solution for forward-in-time algorithms based on highly-scalable high-order CG/DG[15] discretizations and by exploiting bespoke hierarchical multi-level/multi-grid algorithms, as well as deep learning techniques.
2. Bridge the gap between HPC experts and novel mathematical algorithm developers, and enable future co-design through a common, software-engineering driven, domain-specific language concept that facilitates, combines and manages portability and optimisation of current and future key algorithms for emerging HPC architectures across a wide range of domain specific applications.
3. Build an open-source, trans-disciplinary and high-performance European VVUQ package forming the basis for exhaustive resilience and uncertainty testing, for new developments with general community benefit, and for ensuring future business exploitation by European weather and climate services.

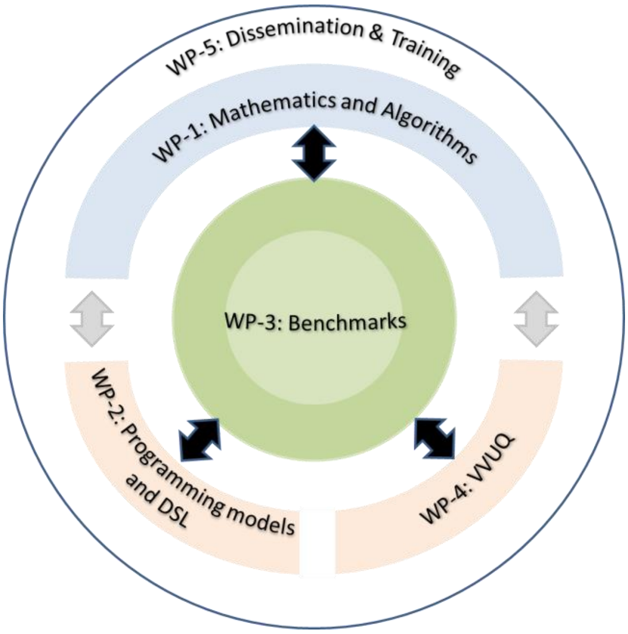

Figure 6: Interaction between ESCAPE-2 work packages WP1 – WP5.

The project activities are organised around **four technical and two cross-cutting work packages (WPs)**, the latter being dedicated to the dissemination and training, and to the overall project and risk management. Training will focus on filling the above mentioned gap between science and hardware specific code levels by employing the newly developed DSL toolchain to key algorithmic motives and applying it to achieve performance portability of novel mathematical concepts as developed in ESCAPE-2. The work package interactions are illustrated in Figure 6, with **WP1** providing a diverse range of relevant algorithmic motives (weather and climate dwarfs) for defining and working with a comprehensive domain-specific language (DSL) toolchain (**WP2**). The subsequent integration in benchmarks for weather and climate prediction, and the progressive development of a domain-specific **HPCW** benchmark as well as back-integration of the output of the DSL toolchain into models is done in **WP3**. VVUQ concepts to both weather and climate dwarfs as well as full prediction system workloads are applied in **WP4**. Dissemination and training (**WP5**) will encompass elements from all technical work packages.

[15] Continuous Galerkin / Discontinuous Galerkin

Aim and main features of the different work packages are defined as:

WP1: Algorithms and Mathematics

Aim: **WP1** will develop mathematical methods and implement advanced and disruptive algorithms suitable for extreme-scale parallelism that achieve major improvements in the accuracy, efficiency, fault-tolerance, and scalability of dynamical cores and of physical parametrizations for next-generation weather and climate prediction models. Moreover, **WP1** will extract and provide a range of relevant algorithmic motifs (weather and climate dwarfs) as a prerequisite for other work packages. These will include key algorithms of advection, time-stepping methodologies, and of physical parametrizations, representative for leading European weather and climate models.

Approach and methodology: A significant contribution to this work package will build on weather and climate dwarfs developed in ESCAPE. In order to widen the spectrum of algorithmic concepts used in Earth system modelling, contributions will be extracted from the ocean models NEMO and ICON-ocean, and the radiation physical parametrizations in the context of Artificial Neural Networks (ANN). Additionally, newly developed dwarfs will cover a semi-Lagrangian higher-order, DG-approach and a fault-tolerant implementation of an iterative elliptic solver.

Suitability of the research approach: **WP1** will cover a wide range of relevant algorithms used in weather and climate sciences. It facilitates the comprehensive definition of a user-friendly and widely applicable DSL toolchain as well as the definition of highly-relevant benchmarks for vendors and HPC hardware developers.

Measures for Success of the Work Package/ KPIs:

- Weather and climate dwarfs delivered, covering at least eight different algorithmic motifs (dwarfs).
- Development of a large time-step, highly-scalable, higher-order method.
- Proof-of-concept of the ANN approach.

WP2: Programming Models & Domain-specific Languages

Aim: **WP2** will define, develop, and apply a DSL toolchain applicable to a comprehensive list of weather and climate dwarfs. The code adaptation and code generation via the DSL toolchain will be demonstrated for a number of representative and fundamentally different mathematical algorithms and horizontal discretizations. Moreover, **WP2** will develop and promote APIs and generic interfaces across the DSL toolchain in order to improve reusability and inter-operability, and leverage code adaptation to emerging HPC architectures.

Approach and methodology: **WP2** will build heavily on developments and expertise from the ESCAPE project and leverage existing open-source technology. Dwarfs are used as vehicles to design, prototype and demonstrate the newly developed DSL toolchain, in particular in the design of the DSL front-end. Following the development of a high-level DSL specification involving community-wide experts in the process, a DSL front-end to parse and translate into high-level intermediate representation (HIR) is designed. Subsequent parts of the modular DSL toolchain will be developed based on existing, open-source technologies such as Atlas, GridTools, and clang. Finally, the DSL toolchain will be applied to a wide range of weather and climate dwarfs.

Suitability of the research approach: Community-wide involvement in the iterative development of a high-level DSL frontend language (including the associated design decisions) ensures a comprehensive coverage of requirements and is key to achieving broad support and adoption. Building on existing, open-source technologies and integration with existing DSL efforts in Switzerland will ensure exploitable results within the scope of the ESCAPE-2 project.

Measures for Success of the Work Package/ KPIs:

- DSL implementation of at least five dwarfs with increased readability and equal or better performance as compared to their reference implementation on any of the three targeted hardware architectures (x86 multi-core, NVIDIA Tesla, Intel Xeon Phi).

WP3: Weather & Climate Benchmarks

Aim: WP3 will develop a hierarchy of benchmarking components representing the key elements in the workflow of weather and climate prediction systems and re-integrate and test code adaptations generated from the DSL toolchain. This work will establish a representative High Performance Climate and Weather benchmark (HPCW). **HPCW** will serve as a benchmark for (pre)-exascale applications of climate and weather codes and will facilitate communication with HPC hardware developers and vendors. The value of HPCW will be demonstrated using the range of available hardware architectures (Table 4).

Approach and methodology: WP3 will define the HPCW based on representative Earth system models from which key dwarfs will have been extracted (in **WP1**). Moreover, **WP3** will ensure reliable and automatic verification through developing routines that check the correctness of the benchmark execution when different software implementations or different hardware options (cf. Table 4) are explored. Several known approaches will be implemented and evaluated following the methodological developments in **WP1** and **WP2**, and also producing an evaluation option in the VVUQ framework of **WP4**. The HPCW benchmark establishes a comprehensive set of test cases and models featuring a number of representative algorithmic motives, as well as system-sized workloads. The workload simulator Kronos (developed by the NextGenIO project) will be employed to create realistic operational scenarios for executing multiple workloads within a single benchmarking environment. This level of benchmarking is entirely new, and it will allow exploring the effect of complex resource contention not observable if single workloads are executed in isolation.

Suitability of the research approach: WP3 closely involves HPC centres and the leading European infrastructure vendor to ensure the suitability of the benchmark design as both user and vendor requirements with respect to HPC benchmark use and relevance will be addressed. In addition, leading European models provide a comprehensive suite of current and future requirements.

Measures for Success of the Work Package/ KPIs:

- Number of selected algorithmic motives (dwarfs) for which successful back-integration of DSL-toolchain generated code of in models can be demonstrated.
- Number of delivered components of the HPCW benchmark.
- Number of performance analyses (per pair benchmark code/hardware system).
- Number of downloads of HPCW components from portal.

WP4: Verification, Validation, Uncertainty Quantification (VVUQ)

Aim: WP4 will develop a generic European VVUQ package for weather and climate simulations that is deployable on supercomputers and that prepares workloads of pre-exascale computations on many-core configurations. The VVUQ package will be demonstrated for both dwarf and full forecasting system workloads, and scenarios will be explored with optimized case performance based on the available VVUQ methodologies. **WP4** aims at confronting ensemble-based and other methodologies to improve VVUQ practices and to produce a generic VVUQ framework for climate simulation at the European level.

Approach and methodology: WP4 builds on the URANIE project that has been started at CEA to understand and quantify uncertainties in numerical simulations. Although URANIE was originally developed to focus on one of the four main CEA areas of research and development, namely the nuclear energy and modelling of multi-physics phenomena in a nuclear plant, its scope of application has been broadened as soon as it became an OpenGL platform and has widened its user community beyond the historical partners. The URANIE platform will be enhanced to capitalise and disseminate the approaches learned from the weather and climate community to other science disciplines and use cases. The weather and climate prediction community has a long experience with the quantification of uncertainties through ensemble methods. Ensemble Prediction methods are targeted in **WP4** by running several instances of the same simulation with modified initial conditions, and to analyse the produced results. The work will start from a small, but representative configuration and target a full, more realistic system. In order to provide a VVUQ software ready for deployment, first developments and tests will be performed on the CEA and BSC HPC clusters. The final version, dealing with a production-level ensemble weather prediction model will be deployed on the ECMWF cluster and potentially other platforms (Table 4).

Suitability of the research approach: A two-way continuous exchange between both energy and weather and climate prediction communities is anticipated to benefit from their respective expertise, fostering substantial cross-disciplinary exchange of ideas and methodologies. The result is anticipated to benefit other research communities.

Measures for Success of the Work Package/ KPIs:

- Providing a VVUQ package for the weather and climate prediction, and energy community with at least two successful deployments (and executions of the full-sized system) at computing clusters, demonstrating the usability and generality of selected VVUQ concepts across different HPC systems.

WP5 and WP6: Dissemination, Training, Project Management

Aim: Prompting a paradigm shift in the understanding and use of DSLs and their impact on weather and climate model design and application, **WP5** will focus on training, support and dissemination activities. In addition, **WP5** will actively promote the dissemination and use of weather and climate dwarfs and the application of the HPCW benchmark in co-design, model development, as well as in training and education. **WP6** coordinates the project and ensures that its innovation actions, objectives and impacts will be delivered.

Approach and methodology: **WP5** will provide the public web portal, confluence interactive development pages for remote working and exchange between partners, and the provision of a common software development and exchange platform suitable for rapid deployment and developments in a distributed environment. A particular focus of **WP5** is to ensure adequate training and dissemination of the novel concepts. Training of early-career scientists through use of novel concepts will foster community acceptance and showcase the achievable acceleration of developer productivity when applying the DSL toolchain developed in ESCAPE-2.

The **WP6** management structures will coordinate and ensure to:

- set-up and maintain a structure, procedures and tools that will allow a coherent and efficient technical and administrative management of the project;
- keep the project on time and within the assigned budget;
- identify and manage risks and solve problems;
- identify opportunities for improved results and collaboration;
- coordinate the interactions between work packages and partners;
- provide and manage working procedures ensuring transparency within the team and for the EC;
- manage quality assurance.

Suitability of the research approach: The common software development platform based on Atlassian tools hosted at ECMWF has proven successful in the project ESCAPE and will be continued in this project. The additional benefit derives from the continued support and promotion of ESCAPE outcomes together with the exchange and novel mathematical and algorithmic development of ESCAPE-2 in a common development environment. The early exposure of the novel development concepts based on a DSL toolchain that generates code is crucial for domain scientists, both experts and early career. This will provide important feedback that can influence design choices thus avoiding costly redesign at a later project stage.

Measures for Success of the Work Package/ KPIs:

- Successful provision of all relevant communication tools between project partners, stakeholders and the public.
- Successful interaction with domain-scientists and early career researchers in the form of at least two dissemination and discussion workshops and 1 summer school.

Table 4: HPC systems and architectures available for ESCAPE-2 benchmark assessment.

Partner	HPC and other architectures	Use in ESCAPE-2
ECMWF	• Cray XC-40, 2x 3550 Intel Broadwell nodes, 36 cores/node, 2.1 GHz + 32 KNL 7210 nodes, 64 cores/node, 1.3 GHz • 34 Intel Haswell nodes, 24 cores/node, 2.6 GHz + 2x NVIDIA K80/nodes	Dwarf adaptation, Benchmarking
MSWISS	Using CSCS infrastructure: • Cray XC-40/50 (Piz Daint) 5320 hybrid nodes (Intel® Xeon® E5-2690 v3 @ 2.60GHz (12 cores, 64GB RAM) and NVIDIA® Tesla® P100 16GB) and 1431 multi-core nodes (Intel® Xeon® E5-2695 v4 @ 2.10GHz (18 cores, 64/128 GB RAM) • Cray XC-40 (Grand Tavé) 164 hybrid nodes (64 cores Intel(R) Xeon Phi(TM) CPU 7230 @ 1.30GHz) • CS Storm (Piz Kesch) 24 fat hybrid nodes (2x Intel Haswell E5-2690v3 2.6 GHz 12-core plus 8x NVIDIA Tesla K80)	DSL development, dwarf adaption
DKRZ/ MPIM	• bullx DLC 720, 1550 Intel Haswell nodes 24 cores/node, 2.5 GHz + 1750 Intel Broadwell nodes 36 cores/node 2.1 GHz • 21 nodes (Broadwell/Haswell) with NVidia Kepler and Maxwell GPUs	DSL tests, benchmarking, dwarf adaptation

BSC	• Mare Nostrum 4: Intel Xeon Platinum 8160 CPU with 24 cores each, 2.10 GHz, for a total of 48 cores per node, 100 Gbit/s Intel Omni-Path
	• Future Emerging Technologies clusters (part of Mare Nostrum 4). Available during 2018-2019:
	  ○ cluster IBM POWER9 processors and NVIDIA Tesla GPUs
	  ○ cluster Intel Knights Hill processors
	• cluster formed of 64 bit ARMv8 processors (technologies from the Japanese Post-K supercomputer)
		Performance analysis, benchmarking, VVUQ analysis

CEA	• Curie thin nodes (TGCC, GENCI) - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR. 77,000 cores, with 308TB RAM • Cobalt (CCRT, CEA) - Bullx DLC 720, Xeon E5-2680v4 14C 2.4GHz, Infiniband EDR. 38,528 cores, with 182TB RAM • Occigen (CINES, GENCI) - Bullx DLC, Xeon E5-2690v3 12C 2.6GHz, Infiniband FDR. 50,544 cores, with 202TB RAM • In 2018, access to the successor of Curie: Intel Xeon Skylake-EP processors and Intel Xeon Phi with a total of 124,700 cores. 400TB RAM. • CEA local cluster: Xeon E5 2680 V2 (2.8GHz) 60 nodes, 1200 cores, with some fat nodes @512Go RAM	Performance analysis, benchmarking, VVUQ analysis
LU	• Hydra: 2460-core 64-bit Intel Xeon cluster supplied by Bull, 103TB storage, with 185 total compute nodes • Athena: 1 PB file store, supplied by ClusterVision and Huawei, with Power8 large memory nodes; 512 compute nodes in Huawei X6000 quad-node chassis, each 28 cores, two Intel Xeon E5-2680v4, 128 GB of memory. 1 Petabyte disk storage; 10 TB SSD GPFS filestore; EDR Infiniband high-performance internal network; 5 OpenPOWER compute nodes, each with 20 cores and 1 TB of memory; two Nvidia GP100 GPGPU cards, via NVlink.	Dwarf development and testing

POLIMI	Using CINECA infrastructure: • Marconi A1 (Broadwell): 1.512 nodes, 2 x 18-cores Intel Xeon E5-2697 v4 (Broadwell) at 2.30 GHz 36 cores/node, 54.432 cores in total, 128 GB/node, 3.5 GB/core RAM, 2 PFlop/s peak performance • Marconi A2 (Knights Landing): 3.600 nodes, 1 x 68-cores Intel Xeon Phi 7250 CPU (Knights Landing) at 1.40 GHz, 244.800 cores in total 16 GB/node of MCDRAM and 96 GB/node of DDR4, 11 PFlop/s peak performance • Marconi A3 (Skylake):1.512 + 792 nodes, 2 x 24-cores Intel Xeon 8160 CPU (Skylake) at 2.10 GHz, 72.576 + 38.016 cores in total, 192 GB/node of DDR4, 7.00 PFlop/s peak performance	Dwarf adaptation, Benchmarking
CMCC	• Athena (Sandy Bridge): 482 nodes, 2 x 8-cores Intel Xeon E5-2670 (Sandy Bridge) at 2.6 GHz 16 cores/node, 7.712 cores in total, 64 GB/node, 4 GB/core RAM, 160 Flops/s peak performance	DSL tests, benchmarking
BULL	Bull France facilities: • Genji: medium scale system based on standard x86 CPUs, Nvidia GPUs, Intel Xeon Phis, ARM; one large in-memory server MESCA-2 • Manny: large scale system based on last generations of CPUs and possibly accelerators	Performance analysis, benchmarking, optimization

1.3.4 ESCAPE-2 gender dimension

As in science, technology, engineering, maths and medicine (STEMM) subjects in general, ESCAPE-2 has a large bias towards male scientists. ESCAPE-2 aims to balance this by active recruitment endorsing the principles of the European Charter for Researchers and Code of Conduct for the Recruitment of Researchers. All partners have working environments that are welcoming for female candidates. ESCAPE-2 will support researchers with children or other dependants. Engagement with target audiences, and when designing information and services, will be gender neutral. Loughborough University's commitment to advancing women's careers in STEMM subjects has been recognized by the Athena SWAN awards scheme, with partner LU holding double bronze medal awards and ESCAPE-2 supports this scheme to encourage talented female scientists.

1.4 Ambition

1.4.1 State of the art and progress beyond

State-of-the-art and main limitations

Weather and climate prediction models are converging in their development paths, as they include increasingly more Earth system modelling complexity by including atmospheric composition, dynamic ocean, sea-ice and more complex land-surface interactions. Climate modelling mostly aims to address major uncertainties in climate sensitivity (Schneider 2017), which leads to the requirement of cloud-resolving climate simulations for addressing uncertainties with respect to shallow clouds and their impact on the Earth's radiation balance. **The big challenge is to efficiently (in terms of energy- and time-to-solution) run ensembles of Earth-system models at kilometre-scale resolutions on the European HPC infrastructure within the next 5 years**.

Time criticality in weather applications matches well with reasonable run times of climate applications, e.g. about one simulated year per (wallclock) day. In recent years the need to increase the efficiency and parallelism of Earth science applications has been addressed by porting parts or all of the codes to specific accelerators (e.g. GPUs, Gysi et al. 2015), by assessing and exploiting precision requirements (Vana et al. 2017), and by substantially changing numerical and mathematical algorithms and associated discretizations on the sphere, towards nearest-neighbour rather than global communication patterns (Zängl et al. 2015; Müller et al. 2015).

However, as has been recently expressed in a US position paper (Carman et al. 2017), despite these efforts, there is a performance wall to overcome, due to limited parallelism in parts of crucially required Earth system modelling tasks, and due to the development of HPC architectures less suited for Earth system modelling. Consequently, one of the fundamental outstanding issues in Earth system modelling is to successfully combine stringent time-to-solution and energy-efficiency requirements. This is because highly-scalable, low-communication (per time-step) methods require excessive numbers of cores to break even in time-to-solution with established globally-communicating methods that can use 10-100 times larger time-steps. This has implications on energy efficiency and is unattractive for computing centre managers, because the most scalable solution is not necessarily the most energy efficient to achieve the desired result.

But substantial progress can be made with model developments that can adapt to a variety of conflicting choices in terms of data structure arrangements, vendor-specific compiler optimization, and the need to include bespoke management of memory hierarchies. Novel memory (hardware) technologies offer also substantial opportunities, e.g. NVRAM technologies explored by ESCAPE-2 partner ECMWF in the H2020 project NextGenIO. Moreover, compiler availability and support of mixed Fortran/C/C++ codes, as well as MPI, OpenMP, OpenACC and other parallelization strategies require substantial extensions which are currently difficult to manage with million-of-lines operational codes. There have been positive developments towards a domain-specific language development to address part of the porting and associated multi-platform scalability challenge with the development of GridTools[16]. The latter provides a set of libraries for expressing distributed memory implementations of structured and unstructured grid applications, e.g. stencils, in a hardware agnostic high-level DSL. However, the application programming interface (API) of the GridTools DSL requires advanced C++ knowhow and this may limit productivity of domain scientists in a community historically expressing algorithmic motives in Fortran. This has been improved in the H2020 project ESCAPE through linking GridTools with the Atlas library (Deconinck et al. 2017).

However, there is an urgent need for more accessibility beyond expert level, through source-to-source translation and a high-level DSL front-end in order to translate code that is close to the equation into code that is HPC optimal. This is likely to be only achievable in a domain-specific way, and the ESCAPE-2 partners MSWISS, ECMWF and DKRZ will contribute substantial expertise to the progress in this area. There are similar developments in the US developing the Kokkos C++ library[17], and at the UK MetOffice, who are developing a specific Fortran-to-Fortran source-to-source translation called PsyClone[18] for their future forecasting system (LFRIC/Gung-Ho). More recently a dialogue has emerged how specific compilers such as CLAW[19] could be used for the benefit of the community overall. The latter is explored in the H2020 project EuroExa.

Moving beyond the state-of-the-art

For achieving progress in weather and climate prediction it is crucial to overcome the performance wall for ensembles of Earth-System models at kilometre-scale resolutions. This is not a matter of simply increasing computing power measured in theoretical flop rates. Addressing Europe's grand challenges, especially under continued budgetary constraints for many European countries, and limited energy availability for computing, it is imperative to identify, apply and implement flexible software engineering design principles for current and future weather and climate models. Our vision is illustrated in Figure 7, where domain science and multi-disciplinary abstractions are separated by a formal interface that facilitates rapid and simultaneous developments of the key weather and climate model algorithms together with hardware adaptation avoiding conflicting choices made in one or the other.

[16] http://www.pasc-ch.org/projects/2013-2016/grid-tools
[17] https://github.com/kokkos
[18] https://github.com/stfc/PSyclone
[19] https://github.com/C2SM-RCM/claw-compiler

The **four key reasons** why ESCAPE-2 will substantially leverage the transition to exascale computing for the key application area of weather and climate prediction, and the associated socio-economic impacts in Europe and worldwide are by:

1. **Combining world-leading mathematical and algorithmic expertise in efficient forward-in-time computing to enhance algorithmic robustness and resilience to failure at scale, and minimize time- and energy-to-solution for highly scalable, high-order algorithms**, thus opening novel pathways for subsequent further advances in scalability and potentially time-parallelism.
2. **Defining an accessible and community-wide domain-specific language concept, and implementing this in key stake-holder application weather and climate dwarfs to accelerate mathematical algorithm development.**
3. **Feeding the mathematical and programming developments into weather and climate prediction community benchmarks facilitating co-design and providing application specific benchmarks for performance.**
4. **Providing an open-source, trans-disciplinary and (exa-)scalable VVUQ package for European research.**
5. **Synthesizing the complementary skills of all project partners. The challenges can only be met through synergies with supercomputing centres, leading HPC vendors, research institutes and universities, and key weather and climate application stake holders.**

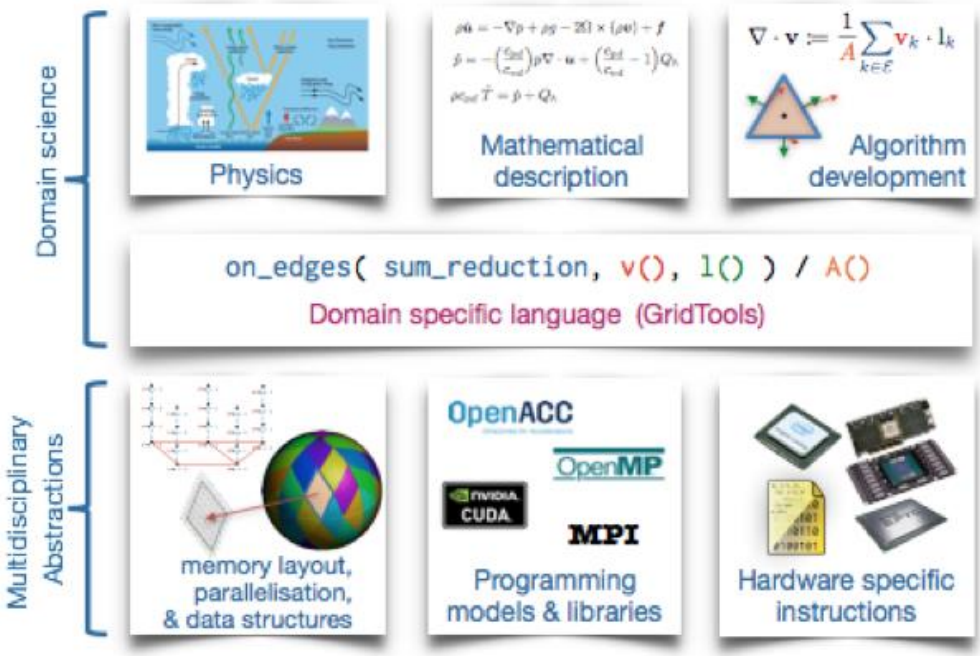

Figure 7: Future vision of mathematical algorithm development and programming levels realizing the separation of concerns in weather and climate prediction.

1.4.2 Innovation potential

The shortcomings of currently well-established benchmarks such as HPL (High Performance LINPACK) for the prediction of a real system's performance have been widely recognized and major efforts have been made to develop benchmarks that are more representative of the performance of a broad set of real applications (e.g. High Performance Conjugate Gradient Benchmark, HPCG). However, the specific aspects of weather and climate applications are still not well represented, for example the ratio between computing and I/O or the wide diversity of computing bottlenecks related to data communication across the network, memory size and bandwidth. The HPCW benchmark to be developed in ESCAPE-2 is therefore particularly attractive for HPC architecture procurements, and acceptance and performance tests run by weather and climate centres.

It is anticipated that HPCW will receive widespread acceptance beyond ESCAPE-2 and beyond the weather and climate domain as a vehicle to interact with HPC hardware developers in Europe and worldwide. Computer systems that are configured and optimized for achieving good performance for the HPCW benchmark will inevitably enhance the performance of operational climate and weather models, thus enhancing productivity and efficiency in one of Europe's most important HPC application areas.

The project outcome provides substantial HPC market opportunities for SMEs to develop bespoke technologies addressing specifically weather and climate dwarfs. The DSL toolchain, and the HIR in particular, enable SMEs to accelerate their hardware adaptation substantially quicker and more independently from full weather and climate applications when transitioning ‘from lab to market’. ESCAPE-2 aims to raise the Technology Readiness Level (TRL) for the DSL toolchain from a level where basic technological components exist (TRL4) to a prototype system demonstration in a weather and climate prediction context (TRL6-7).

ESCAPE-2, as a Research and Innovation Action, clearly focuses on the earlier stages, namely “activities aiming to establish new knowledge and/or to explore the feasibility of a new or improved technology, product, process, service or solution[20]”.

Consequently, an important deliverable of ESCAPE-2 is **dissemination and training**. With the anticipated **paradigm shift in the understanding and use of novel technologies** and their impact on Earth-system model design and applications, it is very important "to channel knowledge, creativity, and technology into innovative, internationally competitive products and services that respond to societal needs[21]". This remains particularly true for the future of operational weather and climate centres and related Copernicus services in Europe.

2 Impact

2.1 Expected impacts

ESCAPE-2 will impact European excellence for efficiently employing exascale high-performance computing in support of one of the largest societal impact areas, namely high-resolution weather and climate forecasting. To this end, achieving kilometre-scale global simulation capability represents the top scientific challenge[22], with ensembles of km-scale simulations delivering unprecedented uncertainty quantification that requires and effectively uses exascale computing resources and associated big data handling.

Weather- and climate-related disasters have caused $2.4 trillion in economic losses and nearly 2 million deaths globally since 1971 as a result of hazards such as droughts, extreme temperatures, floods, tropical cyclones and related health epidemics, according to a new report by WMO[23]. Notably, the United Nation’s Global Assessment Report on Disaster Risk Reduction 2013[24] concluded

---

[20] http://ec.europa.eu/research/participants/data/ref/h2020/other/wp/2016-2017/annexes/h2020-wp1617-annex-ga_en.pdf
[21] European Commission’s Director General Report on Research and Innovation, http://ec.europa.eu/research/innovation-union/pdf/state-of-the-union/2012/innovation_union_progress_at_country_level_2013.pdf
[22] http://www.nature.com/news/climate-forecasting-build-high-resolution-global-climate-models-1.16351
[23] http://newsroom.unfccc.int/nature-s-role/wmo-report-the-escalating-impacts-of-climate-related-natural-disasters/
[24] http://www.unisdr.org/we/inform/publications/33013

that direct and indirect losses from natural hazards of all kinds have even been underestimated by at least 50%. More precise forecasts and their uncertainty bounds in both time and space are critical for human activities and concerns such as travel, health, work, and safety. While there is no control about the weather itself, its impact on society through forecasting and preparation has been drastically reduced, with advances in predictive skill enabling more timely decisions. The WMO report quotes benefit-cost-ratios between 1:4 and 1:36 of investments in improved weather and hydrological forecasting for disaster prevention and early warning. Leading national meteorological services in Europe have estimated the benefit-cost-ratios for their investments in HPC being 1:8 (Météo-France) and 1:20 (UK Met Office).

A new risk is the changing characteristics (frequency, location, severity) of weather and climate related hazards since natural climate variability is now exacerbated by long-term, human-induced climate change. The socio-economic impact of disasters given climate change is likely to be escalating because of their increasing frequency and severity and the growing vulnerability of human societies. Investment in forecasting systems that provide reliable and timely warning is therefore critical. The above economic assessments conclude that these investments pay for themselves many times over.

Future improvements in predictive skill for both weather and climate will originate from enhanced spatial resolution, the better representation of more complex physical and chemical processes, from the coupling between atmosphere, land surface and oceans, cryosphere and biosphere, and from a better characterization of forecast uncertainty through ensembles. All of these require fundamentally different approaches to computing and data management given constraints on power and storage consumption that will represent impediments to progress from present-day petascale to exascale computing. The development of new paradigms is outside the realm of individual operational services and requires internationally coordinated research and funding, a need that is addressed by ESCAPE-2's research focus and constellation of domain representative partners.

ESCAPE-2 will directly benefit the consortia representing the countries shown in Figure 3. ESCAPE-2 will further produce direct impact on both Copernicus Atmospheric Monitoring Service (CAMS) and Copernicus Climate Change Service (CCCS), as they rely - among others - on the world-leading European modelling systems IFS, ICON and NEMO.

2.1.1 Impacts listed in the work programme

Table 5 lists the impact of ESCAPE-2 with respect to the specific topics outlined in the work programme.

Table 5: ESCAPE-2 Expected Impacts.

Contribution to the realisation of the ETP4HPC Strategic Research Agenda, strengthened European research and industrial leadership in HPC technologies:
As the only forecasting centre ECMWF has assumed full membership status of the ETP4HPC since 2014 with the aim to generate a concerted European approach to produce sustainable excellence in HPC for weather and climate prediction. ESCAPE was the first project targeted to address specific aspects of the ETP4HPC research agenda and its realisation in multi-scale numerical modelling of the atmosphere. ESCAPE-2 elevates this effort to a much wider community, the full Earth system and across a wider set of representative models, and with substantial technical enhancements towards fulfilling the ETP4HPC SRA's multi-dimensional HPC vision. ***Technical Research Priorities (and milestones)*:** - **HPC system architecture and components:** ESCAPE-2 prepares for diverse options to use specialized compute units but does not contribute directly to hardware development. Based on existing and pre-release technology to be made available by BULL and HPC centres associated with the ESCAPE-2 consortium partners, the performance portability and programmability of

ESCAPE-2 benchmarks will be tested, and performance (with focus on energy efficiency and time-to-solution) will be assessed based on defined metrics. Hence, ESCAPE-2 will directly impact on the application side of co-design. Moreover, compilers are an essential part of HPC installations. In providing weather & climate dwarfs and the HPCW to compiler developers, systematic and routine testing of vendor provided compilers against archetypical domain algorithms will aid and accelerate compiler development cycles and add robustness together with enhanced customer satisfaction.

Relevant SRA-2 milestones: M-ARCH-1, M-ARCH-7.

- **Programming environment:** ESCAPE-2 contributes directly to the programming environment by adopting the effectiveness of domain-specific languages for enhancing productivity, accelerating development cycles and achieve performance portability in terms of computing and energy efficiency of key algorithmic components. The design and implementation of a weather and climate domain-specific language (DSL) concept based on the tools introduced in ESCAPE will bridge the chasm between highly complex heritage codes and software layers with substantial hardware specific design elements. This development is considered crucial for enabling the mathematical and algorithmic developments to be (a) useable across different models and (b) applicable to and portable between existing and future hardware technologies. The collaboration with BULL as a partner introduces an interface to compiler design in support of hardware abstraction towards a flexible management of data locality and concurrency. ESCAPE-2's weather and climate DSL design will dramatically transform the implementation and adaptation efficiency of weather and climate prediction applications throughout the FET programme's co-design phase.

  Relevant SRA-2 milestones: M-PROG-API-1, M-PROG-API-2, M-PROG-API-5.

- **Energy and resilience:** Enhancing energy efficiency in weather and climate prediction is essential when approaching global kilometre-scale simulations and under stringent operational time constraints. ESCAPE combined a paradigm change for the relevant algorithms with a concept for employing specialized hardware in a heterogeneous environment for dedicated tasks dealing with the resolved flow (model dynamics) and unresolved processes (physical parameterizations). New approaches to enhancing time-to-solution at the same time as energy-to-solution (e.g. energy per forecast) represents a key objective of ESCAPE-2. ESCAPE-2 proposes the development of novel numerical techniques that combine highly effective large-time-step advection with highly scalable, flexible order spatial discretization, thus minimizing communication and enhancing data locality without compromising time-to-solution. The definition of metrics and the employment of generic VVUQ tools will achieve community-wide applicability by providing a detailed quantification of performance portability achieved through domain-specific language implementations.

  ESCAPE-2 directly addresses resilience with hierarchical concepts for fault tolerant solvers that support application resilience during large-scale parallel simulations under strict weather and climate work schedule constraints. The solvers will be tested by implementing a fault detection scheme and iterative data recovery schemes preserving the numerical performance of the solver.

  Relevant SRA-2 milestones: M-ENR-MS-2, M-ENR-FT-6, M-ENR-AR-7, M-ENR-AR-8.

- **Mathematics and algorithms for extreme-scale HPC systems:** ESCAPE-2 aims to deliver a breakthrough in time-to-solution effectiveness of highly scalable, flexible-order spatial discretizations, introducing fault tolerant algorithms supported by hierarchical multigrid tools and a controlled sensitivity to numerical precision, as well as introducing surrogate neural network models by essentially moving training periods outside the critical path and by transforming low-flop operations typical in physical parametrizations to efficient matrix-multiply operations. Connecting and combining these techniques, ESCAPE-2 will directly address the software gap between complex hardware and complex applications through its focus on advancing energy efficient algorithmic building blocks optimized for data flow, data locality and communication patterns across processors. Weather and climate dwarfs pioneered in ESCAPE are emerging as an accepted development template for the entire weather and climate prediction community. Performance portability for emerging hardware is a second corner stone that will ensure sustainable development productivity of software cycles with complex weather & climate codes. The developments will impact the European science community by advancing productivity and showcasing performance portability with world-leading and highly complex forecasting models. These models are at the core of operational service providers and the ESCAPE-2 developments will affect (a) science implementation roadmaps and (b) future HPC procurements throughout the community as the community's workloads approach the exascale era.

Relevant SRA-2 milestones: M-ALG-1, M-ALG-2, M-ALG-8, M-ALG-9.

Extreme-scale Demonstrators:

- ESCAPE-2 will play an important role in defining a key European weather and climate prediction application benchmark (HPCW) for Extreme-scale Demonstrators. As shown in Figure 8, ESCAPE-2 will further develop the dwarf concept pioneered in ESCAPE and the Kronos workload simulator to generate ready-to-use applications for co-design projects (e.g. EuroEXA, NextGenIO) and Extreme-scale Demonstrators. The inclusion of the ICON and NEMO models and the establishment of a weather & climate specific DSL concept --- to allow the implementation of novel mathematical concepts and algorithms across models and hardware platforms --- prepares the weather and climate applications for deployment on the Extreme-scale Demonstrators. The combined outcomes of ESCAPE, NextGenIO, EuroEXA and ESCAPE-2 will be readily available in phase B of the demonstrators.

Ecosystem at large - stakeholders and initiatives:

- **European Extreme Data and Computing Initiative:** Weather and climate prediction represents a dedicated application area within the current EXDCI project (its work package 3), co-led by an ESCAPE-2 partner institute (CMCC). ESCAPE-2 will impact the definition of the science case for the weather and climate community and act as a focal point for the transition of community models to the exascale with the centre of excellence (ESiWACE) as a dissemination hub. Note that ESCAPE-2 will be the only core development project supporting this transition within FET.
- **Centres of Excellence in Computing Applications:** ESCAPE-2 develops user-driven application components that provide scalable benchmarks for weather and climate prediction. This will be used to support the definition of the use cases that represent the grand science challenges addressed by the weather and climate prediction centre of excellence ESiWACE (and its potential successor). ESCAPE-2 partners comprise the ESiWACE co-leading institutes (DKRZ and ECMWF), and key partners (MPIM, CMCC, BSC, BULL). ESCAPE-2 will be instrumental in defining the scope of community models supported by future centres of excellence acting on behalf of the weather and climate prediction community.

ETP4HPC SRA, Completing the value chain:

- ECMWF combines advanced research and operational applications which benefits both the application and service layers spanned by ECMWF (including Copernicus services), its member states and ESCAPE-2 project partners as they represent a significant portion of the European weather and climate forecasting community. The push-through of the envisaged ESCAPE-2 developments follows the same impact route. While the ETP4HPC SRA focuses its recommendations on the industrial impact, a similar value-chain template applies to environmental application and service provision.

Successful transition to practical exascale computing for the addressed specific element of the HPC stack:

ESCAPE-2 will play a central role in advancing European weather and climate prediction capabilities to the exascale through the unique combination of transformational research on mathematics and numerical methods supported by a powerful domain-specific language concept. This combination is considered crucial for achieving, at the same time, performance and performance portability across hardware technologies providing exascale computing capability.

The concept of weather and climate dwarfs pioneered in ESCAPE will be extended across models towards a community representative benchmark integration. Dwarfs are easy to handle for rapid development cycles and identify key bottlenecks preventing the full application from running efficiently on exascale systems. Models and benchmarks in ESCAPE-2 will be scalable to the full size of HPC systems available to the ESCAPE-2 consortium. The benchmarks will serve as application demonstrators in the ESiWACE centre of excellence (and its potential successor) for quantifying the computability of the exascale science challenge. The Kronos workload simulator developed in NextGenIO is included in the ESCAPE-2 benchmark. ESCAPE-2 dwarfs, models, and ultimately the HPCW benchmark will represent weather and climate applications in the EuroEXA co-design project and the anticipated Extreme-scale Demonstrators (EsD) - again representing realistic workloads on full size machines with close to real-world applications (Figure 8). In the project itself, BULL will serve as the test-bed provider for establishing the porting and adaptation process. Enhancing energy efficiency is a prime target for ESCAPE-2 since operational prediction imposes strict runtime requirements on product delivery with direct implications on flop and watt rates. Concerning energy efficiency and resilience, the algorithmic methods developed in ESCAPE-2 aim at enhancing numerical accuracy and stability across a diverse range of heterogeneous hardware options. This is of key importance to future-proof key European applications and supporting the continued development of alternative energy-efficient hardware options.

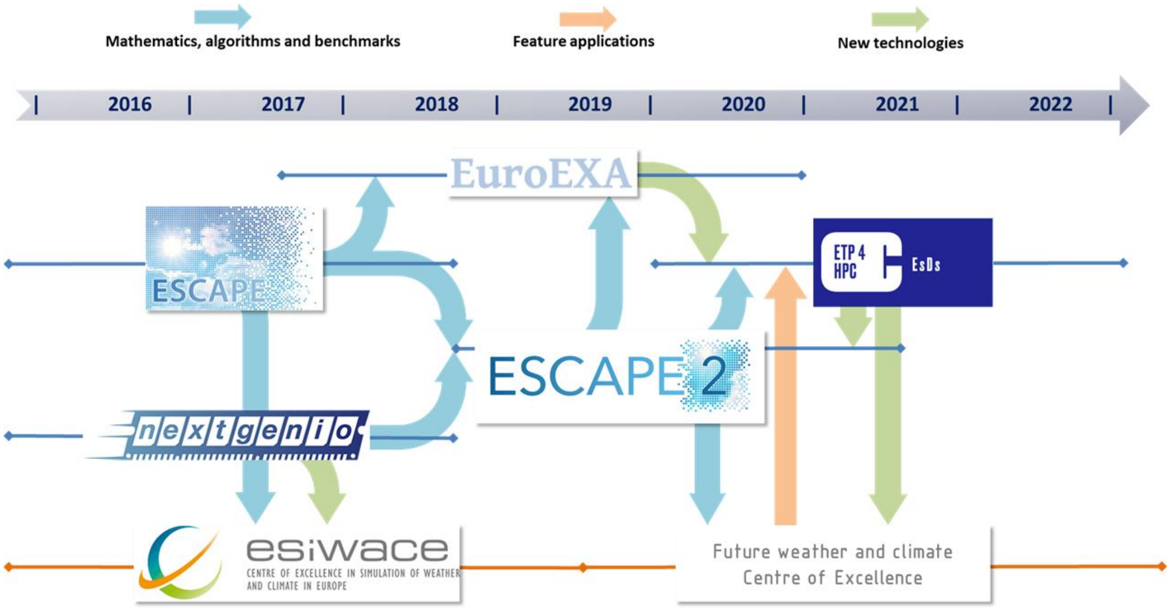

Figure 8: Roadmap from technological weather and climate development projects (ESCAPE, NextGenIO) towards full-scale application demonstrators in ambitious co-design projects (EuroEXA) and, eventually, Extreme-scale Demonstrators (EsD). The centres of excellence (ESiWACE and potential successor) disseminate the results from research and development projects to the wider community. Arrows indicate input provided by one project for another.

Covering important segments of the broader and/or emerging HPC markets, especially extreme-scale HPC systems

ESCAPE-2 will support emerging extreme-scale markets by investigating specific models of heterogeneous architectures for one of the most challenging scientific applications. The specific complexity of weather and climate forecast production workflows, engrained in the heritage of existing (legacy) applications, requires a new approach for often conflicting programming models required to efficiently use more conventional and novel low-energy processor types. Specific to Earth-system prediction are the fundamentally large memory requirements, the combination of global and local stencil-type operations, and the very high flop rates needed for completing high-resolution complex sets of equations in prescribed strict wall-clock time limits with highest possible reliability. ESCAPE-2 will produce a template for mathematics/algorithms in concert with domain specific programming models to be run on a diverse range of HPC architectures applicable to a wide community in weather and climate forecasting.

The involvement of industrial partners such as BULL – currently providing the bulk of the European HPC systems – and other vendors through a network of accessible HPC infrastructures (Table 4), indicates the significant economic potential for HPC in weather and climate prediction. The complementary setup of research and innovation activities in this community (see Figure 8), with a tight link to innovative hardware developments through co-design projects (e.g. EuroEXA, NextGenIO), supports the important role of weather and climate prediction for defining the development of excellence in European HPC and data handling technology, as well as the role of ESCAPE-2 in defining a representative and community-wide application benchmark.

Impact on standards bodies and other relevant international research programmes and frameworks:

The most relevant international research programmes supporting weather and climate prediction are WWRP and WCRP, which are administered by WMO. WWRP aims to advance society's ability to cope with high impact weather through research focused on improving the accuracy, lead time and utilization of weather prediction while the WCRP mission is to facilitate analysis and prediction of Earth-system variability and change for use in an increasing range of practical applications of direct relevance, benefit and value to society. Both programmes foster future high-resolution model development to which scalability is recognized as a major impediment[25]. The recently published WWRP strategy includes, for the first time, a specific action area on new technologies which includes HPC. This has been due to ECMWF's involvement in several WWRP and WCRP coordination activities and highlights the importance of this consortium's involvement in promoting the European HPC strategy in international (here WMO) programmes.

The collaboration with vendors and the important research on weather and climate domain-specific languages in ESCAPE-2 offers the opportunity for programming model co-design, adding leverage at the boundary between hardware and application. The project will establish whether it is possible to retain performance portability with an open-source, technically supported DSL toolchain including traceable source-to-source translation, thus promoting a unified approach in the community. Programming models are seen as a key cornerstone to unlocking performance and performance portability at the same time. ESCAPE-2 is the only project where the responsibility for this development is strongly supported by the entire weather and climate prediction community.

ESCAPE-2 will pioneer a new type of domain-specific benchmark standard (**HPCW**) that offers more realistic hardware performance assessment compared to existing general-purpose benchmarks such as HPL and HPCG. Notably, key weather and climate prediction codes operate at 5% sustained performance or less, and HPCW prepares the way to establish benchmark standards that vendors can work with to extract efficiency rates that are higher than today.

European excellence in mathematics and algorithms for extreme parallelism and extreme data applications to boost research and innovation in scientific areas such as physics, chemistry, biology, life sciences, materials, climate, geosciences, etc.:

[25] http://www.wmo.int/pages/prog/arep/wwrp/new/documents/HighRes_NWP.pdf

ESCAPE-2 conducts European frontier research in algorithm development, by pooling the best European expertise and international partners in large-time-step advection algorithms and highly-scalable, communication-minimizing discretizations with very large domain decompositions. This work provides a significant stepping stone towards the use of time-parallel algorithms, which today are not sufficiently mature for any practical applications in weather and climate prediction. ESCAPE-2 is designed to overcome the current barrier that exists for algorithms in weather and climate that are highly scalable, but inefficient in both time- and energy-to-solution. This work has the potential to substantially influence the future European research agenda for geosciences and beyond (e.g. such technology would readily apply to other potentially time-critical applications such as seismology or biomedical applications where DG-approaches are considered).

With ESCAPE-2, today's proven and fastest algorithms for numerical weather prediction and climate, including both ocean and atmosphere components, will be made portable and optimisable for emerging HPC via an easy-to-learn front-end to a domain-specific language toolchain. This approach will formalise an interface between HPC optimisation/porting experts and novel mathematical algorithm developers beyond the community through the separation of concerns.

ESCAPE-2 will introduce the concept of surrogate neural network models into time-critical applications. By essentially moving extensive training periods – utilising the world's largest meteorological archive as well as specialised expert models (on radiation) – outside the time-critical path in the workflow, and subsequently transforming low-flop operations typical in physical parametrizations with efficient matrix-multiply operations that apply the training period results, substantial accelerations of specific key components in Earth system modelling are expected.

Lastly, with weather and climate prediction being at the forefront of highly important, highly scalable and immensely time-critical computer applications running routinely on up to $O(10^5 - 10^6)$ computer cores, ESCAPE-2 will explicitly address real-world, time-critical application resilience. Through suitably modifying existing mathematical algorithms as well as developing novel mathematical concepts, ESCAPE-2 will substantially enhance European excellence in critical application resilience.

The four themes combine to substantially leverage developments on the route to exascale by overcoming key transition barriers, and in so doing ESCAPE-2 will also facilitate continued European excellence in mathematical and algorithmic developments.

2.1.2 Further impacts

Table 6 lists the impact of ESCAPE-2 with respect to topics that go beyond those outlined in the work programme.

Table 6: Further impacts.

Strengthening Competitiveness and Growth:
Weather and climate prediction capability sits at the beginning of the socio-economic value and decision-making chain, with downstream beneficiaries in water, food and energy resource management, disaster prevention and risk mitigation. Given the anticipated climate change effects and the increasing probability of weather extremes with unprecedented global effects on Europe, much enhanced prediction capabilities are vital for European society. The emerging markets for resource management and climate change adaptation technology at national and European level will inevitably benefit from these capabilities. Therefore, sustaining and advancing the existing world-leadership in numerical weather prediction to a level that it is fit for the upcoming decades where climate change impacts will fundamentally influence European economies is paramount for ensuring sufficient growth and competitiveness.
Improving innovation capacity and the integration of new knowledge:
The multi-lateral engagement of cutting-edge developments with applied science is built into

ESCAPE-2 through the strong partnership between world leading global and regional weather and climate prediction centres, with universities leading mathematical research in Europe, with the leading French national institute for research, development and innovation fostering a wide network of industrial partners, and with the key European HPC provider. Subsequently, all downstream collaborations from ESCAPE-2 partners and adjacent services will benefit from the project's developments. These are in particular (i) the national hydro-meteorological services, (ii) their public and private sector customers, (iii) the Copernicus Atmospheric Monitoring, Climate Change Services and Emergency Management Services, (iv) the associated industrial collaboration networks. ESCAPE-2 will enable these sectors to maintain productive and competitive despite the future computing challenges, will facilitate the integration of new knowledge efficiently, while enhancing the capabilities of their decision making processes during and beyond the lifetime of the project.

The outcomes of ESCAPE-2 will produce models with higher computational fidelity than have ever been used before. ESCAPE-2 outcomes will feed directly into the prediction systems of partner operational centres, namely the ECMWF and Météo-France global as well as the HIRLAM and ALADIN regional operational prediction systems, DWD and ICON consortium weather prediction centres, MPIM and EC-Earth climate prediction activities, and the Copernicus Atmospheric Monitoring Service’s model core, the Copernicus Climate Change Service contribution to global reanalyses and seasonal forecasting, and the Copernicus Emergency Management Services' flood prediction capability. ECMWF’s key role in these activities makes ESCAPE-2 an investment into the future of the Copernicus services and an effective vehicle of cross-benefit between programmes across Europe and beyond.

The ESCAPE-2 strategy also provides a template for the wider weather and climate prediction community and will offer benefit to related Horizon 2020 projects dealing with climate topics mostly under the Societal Challenges 2, 5 and 7. Other projects and modelling groups will be able to build upon the advances made by ESCAPE-2, which will strengthen the innovation capacity of Europe. The substantial number of significant letters of support (attached) indicate the large potential ESCAPE-2 offers to the European community and the interest of the non-European, international community in supporting ESCAPE-2 teams.

ESCAPE-2 will further contribute to the integration of new knowledge by training the next generation of scientists accordingly using the well-established ECMWF training courses and by adding modules to curriculae at two universities and one computing centre. Selected young scientists will be offered the opportunity to participate in working visits at ECMWF, project meetings, workshops and general assemblies, present their results, prepare discussions and chair sessions, and organize a dedicated summer school. Involving early-career scientists in ESCAPE-2 will guarantee a safe transfer of knowledge and prepare the ground for ambitious projects in the future.

Definition of the future path for the world’s leading weather and climate consortium:

ECMWF is an international organisation supported by 34 member and co-operating states. Its foundation principle was the implementation of a centre of excellence for weather forecasting and supercomputing including the largest meteorological data archive in the world, and the notion of economy of scale provided by strengthening and broadening of capability created by scientists from across Europe and beyond through ECMWF. Since the production of the first operational forecast in 1979, ECMWF has been the world leader in global forecasting. The ECMWF IFS has been extended to the world’s highest resolution global data assimilation and forecast system, and supports regional modelling by the regional consortia (like ALADIN and HIRLAM) of European and North-African NMSs as an option under the existing IFS framework. This 40-year effort has been joined by regional consortia representing groups of countries. The resulting limited-area models share the same code infrastructure as IFS (Bénard et al. 2010, Seity et al. 2011). The ICON (ocean and atmosphere model) development was coordinated between MPIM and DWD in support of global and limited area applications. The limited area version will replace the existing COSMO model shared by eight countries in central Europe, Russia and Israel. The NEMO ocean model is a European community ocean model framework used by ECMWF and Met Office for operational weather forecasting and EC-Earth, UK Met Office, CMCC, and IPSL for climate prediction. Laying the foundation of the next-generation forecasting systems with ESCAPE-2 will immediately benefit global and regional forecasting capabilities shared by the vast majority of European countries. This benefit will be directly propagated into services operating downstream, namely climate adaptation, energy resource and water management, crop prediction, disaster reduction and health.

Internationally, ESCAPE-2's pioneering developments will support Europe's competitiveness compared to significant efforts currently pursued in the US by NOAA and NCAR, in Japan by JMA and Riken, in China by CMA, South Korea by KMA and by Environment Canada aiming at their next-generation extreme-scale Earth-system Modelling. The core developments of ESCAPE-2 will impact climate prediction capabilities and will therefore be of immediate relevance to IS-ENES[26] that coordinates the European infrastructure and HPC strategy for Earth-system modelling and European contributions to the IPCC. The fundamental technical developments in ESCAPE-2 will be propagated into the ESiWACE centre of excellence for further dissemination and maintenance.

Foundation for future global model usage in European Commission Copernicus services:

ECMWF is the operator of the recently established Copernicus Atmospheric Monitoring Service (CAMS) and the Copernicus Climate Change (C3S) by delegation agreement with the EC. The IFS model is used in the CAMS global monitoring and forecasting component, and the ECMWF contribution to the C3S global reanalysis and seasonal forecast production. IFS is also the atmosphere component in EC-Earth[27] while NEMO contributes the ocean component. Further, operational ECMWF model output supports the product chain established by the European Flood Awareness System (EFAS). ICON is the new generation of MPIM-DWD global climate models with both atmospheric and ocean components, both of which are employed in ESCAPE-2. ICON and EC-Earth will participate in future Coupled Model Intercomparison Projects (CMIP) that form the foundation of the assessments of the Intergovernmental Panel on Climate Change (IPCC). The CMIP climate projections also provide key input for C3S. The development of key components in the future IFS and ICON through ESCAPE-2 will thus benefit and strongly impact the user communities spanned by both CAMS and C3S as well as the European Flood Awareness System (EFAS; being part of the Copernicus Emergency Management Service).

2.1.3 Barriers, obstacles and other conditions with relevance to the impact

Weather and climate prediction provides invaluable societal and economic benefits and the development of the next-generation forecasting systems crucially depends on future hardware options and software solutions to allow sustainable energy-efficient Earth-system modelling. The magnitude of this task requires a broad engagement of operational services, industry, academia and cutting-edge research supported by the EC.

ESCAPE-2 represents the first concerted effort to address these challenges in the weather and climate community at European level, but a longer-term commitment is to be envisaged and sought from all stakeholders to overcome the biggest challenge for computational numerical weather prediction since its initiation and to provide a sustained return for the European society.

2.2 Measures to maximise impact

To maximise the impact of the project results, we have chosen a combination of measures for: **Dissemination**, **Communication, Training and Exploitation**. These activities will involve **all consortium partners** and their respective staff, including researchers.

[26] https://verc.enes.org/ISENES2
[27] https://www.ec-earth.org

- A **Dissemination and Communication strategy** will be fully developed in the early stage of the project (**WP5**). Focus on dissemination and communication activities in all work packages ensures that the full impact of the project is achieved. An online platform - an integral part of the plan - will be complemented with a comprehensive programme of offline activities, training, and dissemination activities 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 context. **Communication** foresees **a number of tools** for the implementation of our strategy. **Suitably framed messages delivered through suitable dissemination tools** will help us publicize our work in such a way that the consortium will profit and help us to 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. We have planned to tailor our messages for a full range of potential users and uses: for each audience, we will work on a distinct strategy using targeted messages, means and language, working at the right level according to the identified needs (local, regional, national).
- **Exploitation (WP6)** will involve the development of suitable exploitation routes within **WP6** combined with a pro-active innovation management approach. The exploitation plan will be developed early on in the project (**M12**) with the aim to inform the further development in terms of focus, also informed through the community interaction. The exploitation plan will be revisited towards the end of the project to identify and refine the activities to be implemented beyond the lifetime of the project.

Table 7: Overview of strategic plans for dissemination, exploitation, communication and software/data management.

	Plan for Dissemination and Communication	Project Portal and Software Collaboration Platform	Exploitation Plan
Objective	Enable and facilitate implementation of ESCAPE-2 outputs for European weather and climate prediction community	Provide metrics, data and software to facilitate numerical experimentation inside/outside ESCAPE-2	Enable long-term sustainability and uptake of ESCAPE-2 results
Target audiences	Dissemination: - ESCAPE-2 project partners - Weather and climate prediction centres - Scientific Community - EC (as a multiplier) Communication: - General public - Scientific community - HPC industry - ETP4HPC - WMO programmes - EC (as a multiplier)	ESCAPE-2 project partners EC (as a multiplier)	- ESCAPE-2 project partners - Weather and climate prediction centres - Scientific Community - EC (as a multiplier) - Scientific community - HPC industry - ETP4HPC - WMO programmes
Instruments	ESCAPE-2 website ESCAPE-2 reports ESCAPE-2 Wiki ESCAPE-2 workshops Training events, summer school Scientific publications	ESCAPE-2 website ESCAPE-2 software collaboration platform ESCAPE-2 suites, data	ESCAPE-2 website ESCAPE-2 software collaboration platform

Access	Public	2-tier (project, public)	2-tier (project, public)
Responsibility	WP5	WP5	WP6

Since ESCAPE-2 develops novel numerical algorithms and new DSL concepts for accelerating algorithm development, the communication, dissemination and output management activities will be shared by POLIMI and ECMWF, representing the numerical developer and application communities, respectively. Additionally, in the consortium we can count on institutions with **in-house capacity** for science communication.

2.2.1 Dissemination and exploitation of results

Scientific and technical results of ESCAPE-2 will be made available to the public and disseminated regularly through project reports and technical documentation, publications in the scientific literature and conference proceedings, and most importantly publications on the project website. Engagement with the university and education sector will be crucial both for promoting the work of ESCAPE-2 and for ESCAPE-2 to benefit from scientific and technical developments taking place outside the consortium. Dissemination, output management and exploitation will use a set of key instruments. These instruments support a dialogue rather than a unilateral distribution of information. Within the project, the interplay between development and application with strong involvement from the industrial partners will be crucial.

- ***ESCAPE-2 community portal***: A dedicated portal will be implemented comprising public and project compartments. The public compartment will represent a showcase for information on key results and demonstration projects produced by ESCAPE-2. The Plan for Dissemination and Communication (**D5.1**) will define milestones for public dissemination. The public site will also include a blog-function allowing external users to communicate with the team. The project compartment will provide limited access for project partners as well as the EC, and direct access to all material generated in the course of the project as well as quick looks at the project status. The project site will also include an online exchange function to track the communication between partners on global project topics. Further, a project Wiki will be established as the central documentation tool. Further, a branch maintained by students will be established to link to classes, relevant course material, and training events organized in the framework of ESCAPE-2, to be integrated in the established ECMWF training courses[28].
- ***Dissemination workshops***: Through its European wide scalability programme ECMWF has established a strong engagement with its member states, the weather and climate prediction community and HPC vendors in the general area of scalability. This momentum will be maintained by the organisation of a regular ECMWF workshop on scalability, in which the ESCAPE-2 activities will feature prominently through separate sessions (cf. dissemination workshops I and II). The workshop's working groups represent a forum for interaction between the ESCAPE-2 team and the community. The bi-annual ECMWF HPC Workshop[29] also highlighted the scalability theme in 2016, and ESCAPE-2 will be featured in future events. In the beginning of the project a particular emphasis will be on expert workshops discussing and defining a community representative DSL intermediate representation and related concepts. This work will continue in a dedicated working group (see also **WP2**).
- ***Training courses***: The training effort in ESCAPE-2 will be organized in two stages. In the first stage, the ECMWF training course that provides centralized European facilities to the community, will be extended by modules on scalability in general and numerical and programming solutions for addressing the Weather & Climate Dwarfs in particular. The

[28] https://www.ecmwf.int/en/learning/training

[29] https://www.ecmwf.int/en/learning/workshops-and-seminars/17th-workshop-high-performance-computing-meteorology

university and HPC centre partners will add modules based on this material in their curriculae to support the next generation of scientists bridging between applied and computing science. A summer school dedicated to ESCAPE-2 will be hosted by partner POLIMI introducing young scientists to the subject and introducing novel development concepts (such as the DSL toolchain). The timing of this summer school is therefore targeting **M18**. Event organizations will be carried out with early career scientists and students under supervision of ESCAPE-2 to enhance student engagement and transparency at universities.

- ***Committees***: The representation of ECMWF and project partners in international committees will be used as a channel for disseminating ESCAPE-2 results and output in the weather and climate prediction communities (mostly through WMO programmes WWRP and WCRP) as well as forums at the science - technology interface, namely ETP4HPC. ECMWF and its partners are strongly represented in these communities. This ensures an efficient push-through of the ESCAPE-2 outcome to the relevant target groups.

Dissemination and exploitation within ESCAPE-2*:* This will be managed in WP5 and WP6. The central platform will be the ESCAPE-2 project website and its interactive facilities for the exchange of information, software and data. Repositories for project documentation (plans, progress tracking, reports, financial information), scientific outcome (papers, reports, documentation, conference proceedings), data (sample input and output), and general Wiki-type interactive tools will be established. Exchange forums for discussion between partners will be linked to the respective repositories.

WP5 will establish a technical and scientific collaboration platform comprising a source code repository and a bug tracking system to facilitate distributed software development. The planned workflow model for software development collaboration is illustrated in Figure 9. The central web-platform, collaboration tools such as the wiki, and the collaboration platform on software development will be maintained by ECMWF. The ESCAPE-2 meeting schedule is described in WP5 and Section 3.2.

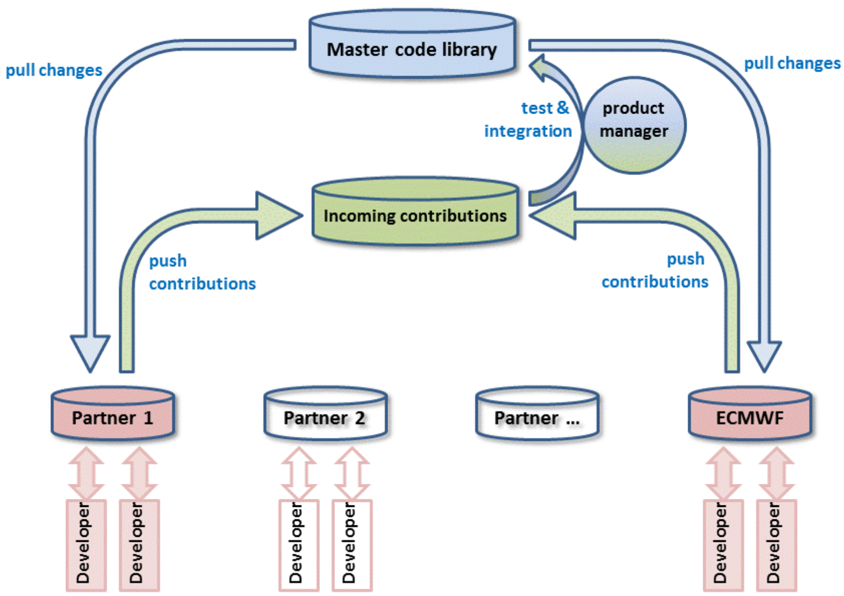

Figure 9: Project collaboration platform on software development.

Dissemination and exploitation beyond ESCAPE-2: The products of ESCAPE-2 will comprise reports, graphical displays and datasets. Reports will be viewable and downloadable from the public pages of the central ESCAPE-2 website. Graphical products will be publicly viewable either directly from the central website or from partner websites to which users will be directed from the central site. Many of the pages served by partners will be given look-and-feel similar to that of those served by the central site. The central site will be hosted by ECMWF to ensure that the site benefits from the resilience built into the ECMWF web system.

In the course of the project, external parties such as the scientific community, operational weather and climate prediction centres and hardware/software vendors will be provided with substantial amounts of information that they can utilize to track progress and make optimal use of the information provided by ESCAPE-2. A shortened version of the project plan suitable for a more general readership will be prepared based on the contractual Description of Work that, in turn, will be based on this proposal document. It will be made publicly available on the project website. The regular project reports will include accessible, self-contained summaries of progress, and these too will be made available as separate publicly available documents on the website. A 'news' item on the project website will draw attention to recent results and their potential impact on stakeholders.

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.

- **Open access to peer-reviewed scientific publications**: Open access will be granted to all scientific publications resulting from ESCAPE-2 with a combination of golden and green open access. We will make use of institutional and topic repositories for making our publications available. ESCAPE-2 scientists, as EC grant recipients, will make their best efforts to 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.
- **Open access to software and tools**: Software developed in ESCAPE-2 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 ESCAPE-2 is fundamentally open. Data, documentation and training tools will be made freely available, and software produced by ESCAPE-2 will be managed through an Apache-2 open source license where possible. A limited set of code parts is based on IFS heritage code and this will be managed by an OpenIFS-type software license[30] already provided by ESCAPE, which is limited in time, only permits non-commercial research use, and limits redistribution. The OpenIFS software licence also regulates that any further developments on OpenIFS regulated software may be commercially exploited by ECMWF or its member states in future. The OpenIFS software license has been drafted in agreement with the ECMWF Member States. These software elements are already subject to multiple ownership rights in several instances and thus sufficient experience exists within the consortium to resolve any issues that might arise.
- **Open access to data**: Since ESCAPE-2 is a software and application development project the rules and procedures governing data management do not apply.

The procedures above will be monitored by the Project Office. This will ensure coordination and standardization through a code exchange inventory and licensing register. The code management will monitor the availability of software to all project partners and ensure that formats and

[30] https://software.ecmwf.int/wiki/display/OIFS/Licensing

documentation are in line with the established guidelines. The code management protocol will be established at the kick-off General Assembly and will include representatives from each work package. It will collect and maintain the required code documentation. All partners will share the knowledge and methodologies developed within the project for achieving the expected results during the project. IPR protection will be addressed in the Consortium Agreement with a strong emphasis on open access data policies.

Dissemination measures *in the closing phase* of the project: The final report of the project (**D6.3**) will include a plan for the use and dissemination of foreground, 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 EC 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. A final project booklet collecting all project publications will be produced at the end of the project. The booklet will be made available for download on the website.

Dissemination measures *after the closure* of the project: After the official end of the project, the foreground of the project will be available as a web-based archive for all interested actors. The domain name of the project website will be assigned to ECMWF. The website archives all documentation related to the project, including publications, and will be accessible for 5 years after the end of the project.

2.2.2 Communication Activities

Defining the target audience is important to produce impact outside ESCAPE-2. The following audiences have been identified:

- The first audience are the (scientific) project users, which are potentially direct beneficiaries of the project and are represented for instance by other weather and climate prediction centres inside and outside Europe.
- The second category comprises the potential industrial users, which are organisations that can benefit from the use of the output that ESCAPE-2 will develop. Examples of this are hardware vendors that have commercial interest in weather and climate prediction customer requirements.
- The third audience contains more general stakeholders such as ECMWF member states, academia, international committees, and the general public. The online platform and the social media managed by **WP5** will provide the tools to build and maintain an engaged community around the project.
- The general public, a wider audience with a non-scientific background.

The communication with these target audiences will be managed as shown in Table 8.

Table 8: Overview of ESCAPE-2 dissemination target groups and tools (this list is not exhaustive).

Target audience	Communication	Responsibility
Scientific project users	- Peer-reviewed scientific papers (J. Comp. Phys., Bull. - Amer. Meteor. Soc.) - Project workshops, conferences (ECMWF/NCAR HPC) - ESCAPE-2 portal - Link with related H2020 activities (e.g. Centres of Excellence)	All partners
Potential industrial users (hardware, software vendors)	- Targeted publication material - Presence at inter-disciplinary conferences (WMO WGNE meetings, ClimathNet[31] Network)	ECMWF, DKRZ, BULL

[31] http://www.climathnet.org/

	- Participation in advisory boards (e.g. WMO, WWRP, JSC, WGNE) supporting best practices and standards - ETP4HPC Board (ESCAPE-2 partner members) - ESCAPE-2 website - Link with related H2020 activities (FET-HPC, Centres of Excellence)	
ECMWF Member States, International committees	- ECMWF committee documents - Strategy papers, white papers - ESCAPE-2 website	ECMWF
General public	- General information material - Demonstration examples - ESCAPE-2 website	All partners

3 Implementation

3.1 *Work plan — Work packages, deliverables*

3.1.1 Overall work plan and structure

The work plan builds on the interaction between the different work packages by further extending and applying the concept of weather and climate dwarfs that encapsulate and modularise key algorithmic motifs for effectively simulating weather and climate (**WP1**). Underpinning these developments, ESCAPE-2 will invest in porting approaches and higher-level data-structure and language abstractions (via weather and climate specific DSL) to leverage performance portability across a diverse range of HPC architectures (**WP2**). Existing and novel mathematical and algorithmic developments are incorporated in domain-specific high-performance benchmarks (**HPCW**) representing a comprehensive range of world-leading European Earth-system models (**WP3**). This work is supported by a generic, cross-disciplinary and (exa-)scalable VVUQ framework (**WP4**). Results and ESCAPE-2 outcomes are disseminated to industry, programming environment and compiler developers, other stakeholders, and ensuring future recruitment opportunities by adequate training and education (**WP5**). Finally appropriate project planning and management is ensured (**WP6**).

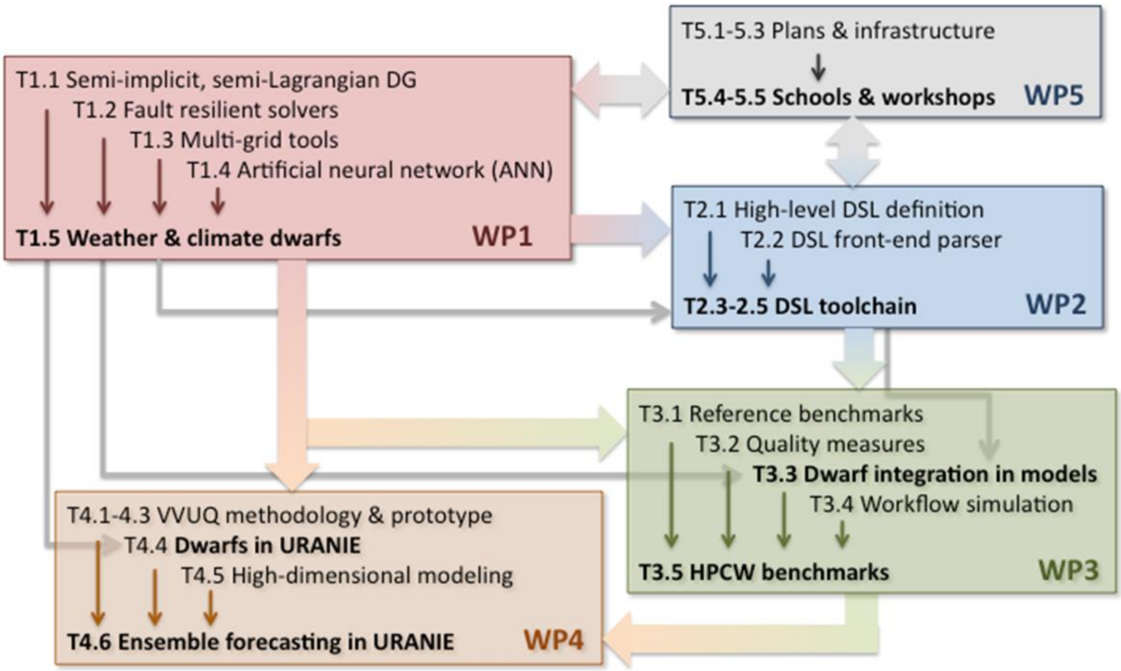

Figure 10: PERT chart showing dependencies between work packages and tasks. Arrows indicate "provides input to".

3.1.2 Timing of work packages

The dependencies and timings of the different work packages are presented in Figure 10 (PERT chart) and Figure 11 (Gantt chart), the project work packages are detailed in Table 10, and all deliverables are listed in Table 11. The major project milestones are listed in Table 12.

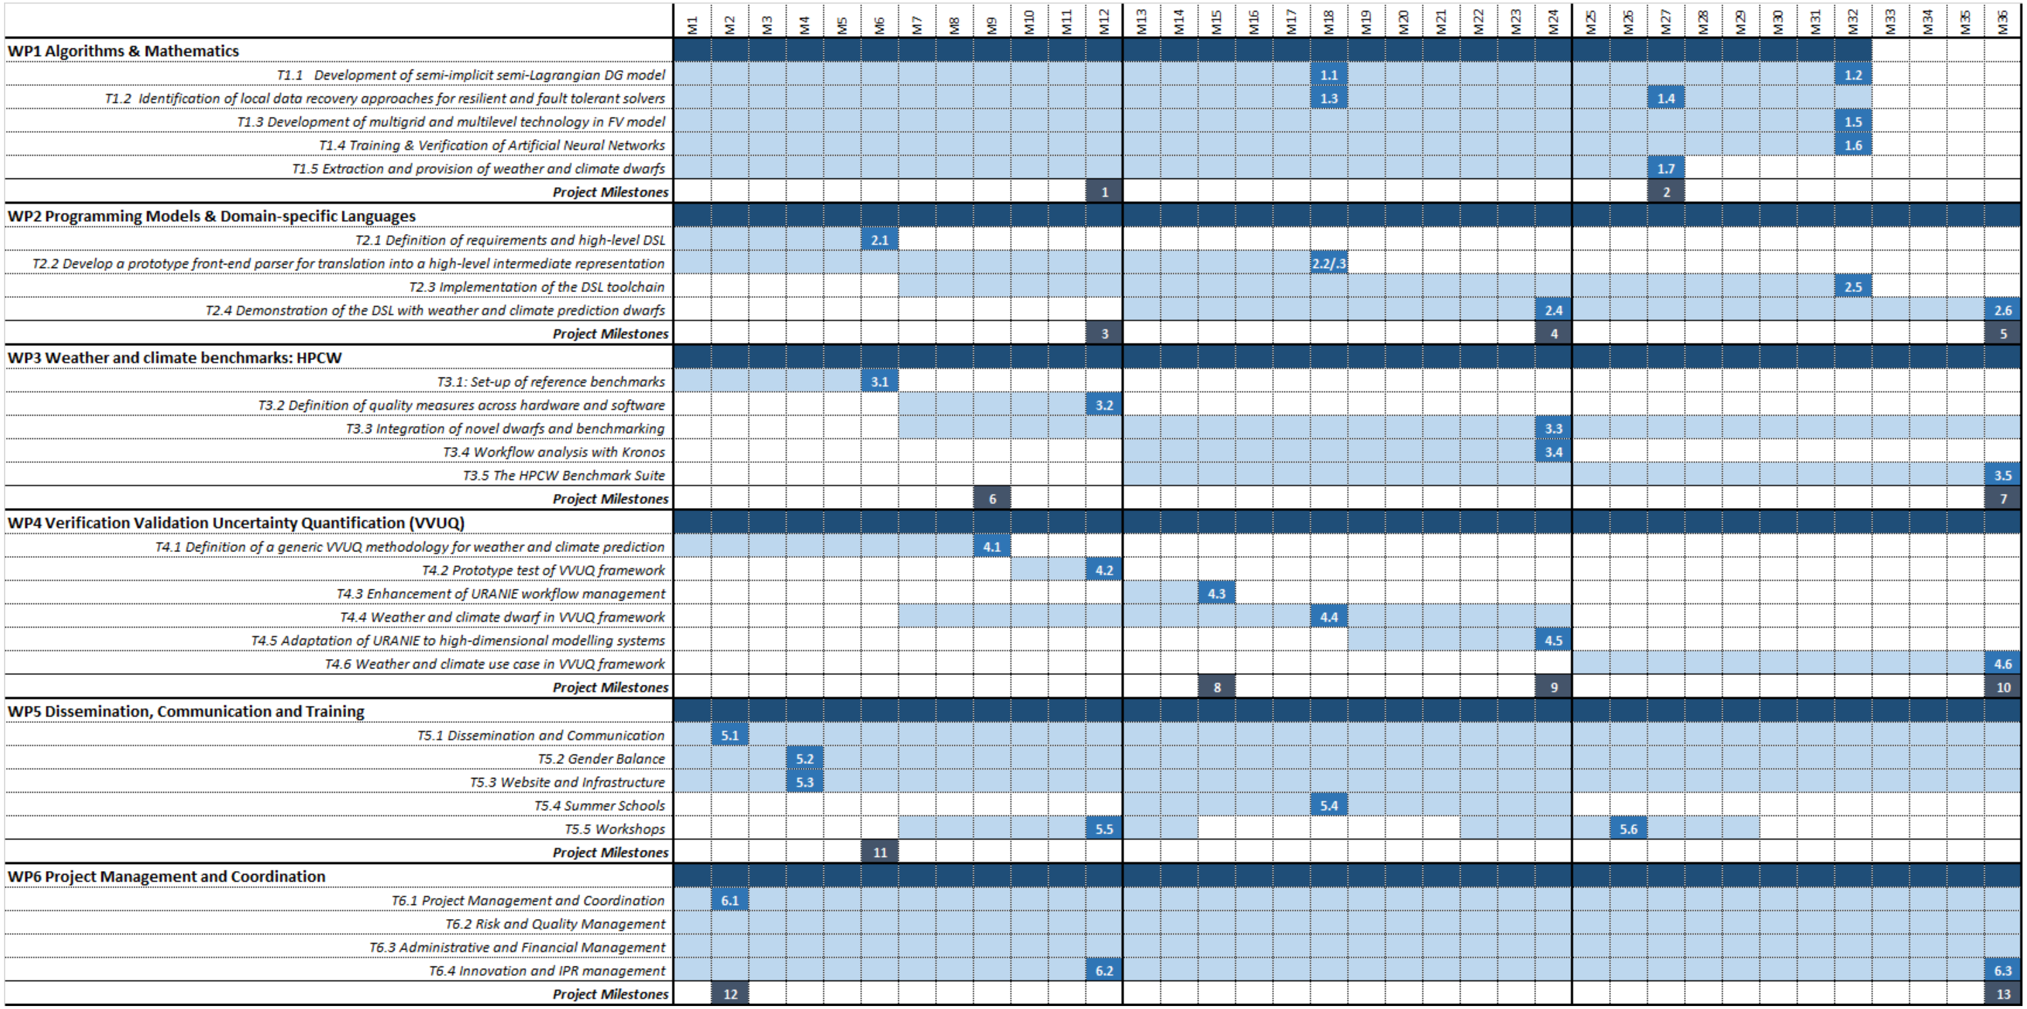

Figure 11: ESCAPE-2 Gantt chart showing time line of work packages (WP) and tasks (blue), milestones and points where deliverables will be provided (white).

3.1.3 Table 10: Work package description.

Work package number	1	**Lead beneficiary**							POLIMI, LU				
Work package title	Algorithms & Mathematics												
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	**TOTAL**
Person months/ participant:	10	1	1	0	0	0	38	0	32	36.3	1	0	**119.3**
Start month	1						**End month**			36			

Objectives

WP1 is directly related to **ESCAPE-2 key objective no. 1**, and produces output for **WP2**, **WP3** and **WP4**. **WP1** aims to:

- Develop mathematical methods and implement advanced and disruptive algorithms suitable for extreme-scale parallelism that achieve major improvements in the accuracy, efficiency, fault-tolerance, and scalability of dynamical cores and of physical parametrizations for next-generation weather and climate prediction models.
- Extract and provide a range of relevant algorithmic motifs (weather and climate dwarfs) for defining and working with a comprehensive domain-specific language (DSL) toolchain (**WP2**), for subsequent integration into benchmarks of weather and climate prediction models (**WP3**), and for applying VVUQ concepts (**WP4**).

Description of Work

WP1 is led by POLIMI (Luca Bonaventura) and co-led by LU (Joanna Szmelter)

Task 1.1 Development of semi-implicit semi-Lagrangian DG model (Lead: POLIMI, Partners: ECMWF)

A semi-implicit, semi-Lagrangian (SISL) Discontinuous Galerkin (DG) discretization for the fully compressible Euler equations on hexahedral meshes will be implemented for the first time in a natively parallel environment and equipped with local data recovery procedures. The SISL approach has been shown to allow for substantially longer time steps also in the DG framework, but its application has so far been limited to tests of academic interest only. The implementation to be developed in the framework of this project will instead allow to run simulation in realistic configurations. It will allow the flexible choice of the local element degree, so as to enhance resolution locally for different applications. Different hexahedral meshes obtained from the extrusion of unstructured quadrilateral meshes will be tested. Experiments with the coupling to simplified physical parameterizations and with parallel load balancing of adaptive versions will be carried out. To test the communication patterns induced by the SL method a SL DG dwarf will first be produced (**D1.1, POLIMI**). In the later stages of the development it is planned to incorporate the fault tolerant solvers developed in **Task 1.2**, and to combine the DG spatial discretization with SL advection in a semi-implicit time-stepping framework (**D1.2, POLIMI**). An Eulerian semi-implicit discretization is foreseen as alternative in case the communication pattern highlighted by the SL DG dwarf were unsatisfactory.

Task 1.2 Identification of local data recovery approaches for resilient and fault tolerant solvers (Lead: POLIMI, Partners: LU, ECMWF)

This task will identify resilient implementations of the DG and finite-volume (FV) solvers, based on current research on local data recovery approaches and on fault tolerant solvers. Resilient computation approaches will be examined to identify local data recovery suitable for weather and climate applications (**D1.3, POLIMI**). This will be followed by a fault tolerant implementation of

Generalized Minimum-Residual (GMRES) methods (**D1.4, POLIMI**) suitable for application in weather and climate prediction. Multi-grid solvers and matrices resulting from higher-order SI DG discretizations implemented in **Task 1.1** and **Task 1.3** will also be used in this context. The impact of floating point precision reduction on the main families of linear solvers considered will also be analysed jointly with ECMWF.

Task 1.3 Development of multigrid and multilevel technology in FV model (Lead: LU, Partners: ECMWF)

Recent developments and evaluation of the FV Module (FVM) of the IFS have shown that the accuracy of simulation it can offer is at least comparable to that of the spectral transform based predictions by IFS. However, to provide adequate speed in the operational prediction environment, the Finite Volume Module needs to be further optimised and take advantage of new algorithmic and computing solutions. The developments proposed in this task are designed to increase the computational speed of the module with new algorithmic concepts but with full preservation of the accuracy and structure of the module as well as its compatibility with the IFS.

Specifically, this task will generalise the multi-grid apparatus initiated in ESCAPE for the preconditioning of a class of Krylov-solvers in the FVM of ECMWF's IFS. In particular, the developments will exploit coarse mesh predictions for perturbation variables and the regularization for stiff physics. Specific accelerations of the time-to-solution performance are going to be addressed by a dual time-stepping approach and a multiple-grid MPDATA advection implementation, both forming pre-requisites for exploring parallel-in-time schemes. The developments will be based on ECMWF's Atlas library and tested for canonical weather and climate prediction test cases. The developments will be applicable to the FVM of IFS (**D1.5, LU**) and extended to the hexahedral meshes (**Task 1.1**).

Task 1.4 Training & Verification of Artificial Neural Networks (Lead: DMI, Partners: ECMWF)

This task will develop the training and verification of Artificial Neural Networks (ANN) for the parametrization of radiation, which is the most costly single physical parameterization in weather and climate prediction models (also used in **WP4** for **Task 4.4**). The following steps will be undertaken (**D1.6, DMI**):

- Define the ANN method for use in weather and climate prediction for radiation processes.
- Train the ANN based on relevant datasets from ECMWF archive (ECMWF hosts the largest meteorological data archive in the world) and a full (and precise) radiation transfer model.
- Provide a preliminary validation of the ANN approach on a specific weather prediction benchmark before proceeding to **Task 3.3** in **WP3.**
- Adapt and benchmark the ANN for different HPC architectures.

The principal tool that will be used to create training data sets will be the AccuRT radiative transfer model described in Chapter 10 of Stamnes et al., 2017. This radiative transfer code is well tested and has been used in many applications to create training data sets required to develop ANNs as illustrated by several examples provided in Chapter 10 of Stamnes et al. (2017) and in a recent publication. Furthermore, empirical high quality radiation datasets will be used for training the ANN specifically for situations with complex cloud fields.

Task 1.5 Extraction and provision of weather and climate dwarfs (Co-Lead: POLIMI&LU, Partners: ECMWF, DMI, MPIM, DKRZ, CMCC)

This task will extract and provide a range of relevant algorithmic motifs (weather and climate prediction dwarfs) for defining and working with the comprehensive DSL toolchain (**WP2**), for subsequent integration into benchmarks of weather and climate prediction models (**WP3**), and for applying VVUQ concepts (**WP4**). These weather and climate prediction dwarfs will be extracted from highly representative weather and climate prediction models (IFS, ALADIN/HIRLAM - HARMONIE, ICON, NEMO), providing a comprehensive list of key algorithms and discretizations (**D1.7, POLIMI&LU**):

- Advection: reduced octahedral grid, semi-Lagrangian finite-difference (extracted from IFS-spectral

atmosphere).
- Advection: reduced octahedral grid, Eulerian finite-volume (extracted from IFS-FVM atmosphere).
- Advection: tri-polar grid, Eulerian finite-difference (extracted from NEMO ocean model).
- Advection: icosahedral mesh, Eulerian finite-volume/element combination (extracted from ICON-ocean model).
- Time-stepping: semi-implicit, iterative 3D elliptic solver (extracted from IFS-FVM atmosphere).
- Time-stepping: horizontally explicit, vertically implicit solver (extracted from ICON-ocean model).
- Physical parametrization: radiation (extracted from HARMONIE atmosphere and from ICON atmosphere).
- Physical parametrization: cloud microphysics scheme (extracted from IFS atmosphere).

In order to facilitate the preparatory work related to DSL the advection reduced octahedral grid modules will be pre-released to WP2 within the first 6 months of the project.

Novel dwarfs as developed in **Task 1.1, Task 1.2** and **Task 1.4** will complement the list of representative schemes, in particular (**D1.8, POLIMI&LU**):
- Advection: hexahedral grid, semi-Lagrangian DG,
- Time-stepping: Fault-tolerant GMRES solver; semi-implicit DG solver,
- Physical parametrization: ANN radiation.

Deliverables

D1.1 A SL DG dwarf (POLIMI, OTHER, CO, 18)

D1.2 A SI-SL DG prototype dynamical core (POLIMI, OTHER, CO, 32)

D1.3 A report on identification of local data recovery approaches suitable for weather and climate prediction applications (POLIMI, R, PU, 18)

D1.4 Fault tolerant implementations of GMRES (POLIMI, OTHER, CO, 27)

D1.5 Multigrid and multilevel technology based IFS-FVM model (LU, OTHER, CO, 32)

D1.6 Training & preliminary validation of the ANN for the physical parametrization of radiation (DMI, OTHER, CO, 32)

D1.7 Extract and provide a list of representative algorithmic motives in the form of weather and climate dwarfs (POLIMI&LU, OTHER, CO, 9)

D1.8 Complement the list of dwarfs with novel concepts (POLIMI&LU, OTHER, PU, 27)

Partner Roles	
PM	will co-lead this work package, is responsible for delivering D1.1, D1.2, D1.3, D1.4 and coordinate delivering D1.7, D1.8
LU	will co-lead this work package, is responsible for delivering D1.5, contribute to D1.3, D1.4. and coordinate delivering D1.7, D1.8
ECMWF	will contribute to D1.1, D1.2, D1.5, D1.7, D1.8
DMI	is responsible for delivering D1.6, contribute to D1.7
MPIM	will contribute to D1.7
DKRZ	will contribute to D1.7
CMCC	will contribute to D1.7

WP Dependencies	
WP2	Delivery of dwarfs for developing a DSL toolchain (including a comprehensive definition of an intermediate representation; application of the DSL toolchain to the weather and climate dwarfs developed in **WP1**.
WP3	Provision of expertise to support the (re-)integration of dwarfs developed in **WP1** and modified

	in **WP2** for the development and testing of benchmarks.
WP4	Delivery of selected dwarfs for the application of VVUQ.
WP5	Provision of dwarfs for the summer school practical training.

Work package number	2		**Lead beneficiary**						MSWISS, MPIM			
Work package title	Programming Models & Domain-specific Languages											
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	**TOTAL**
Person months/ participant:	10	11	18	33	2	0	0	3	0	0	7	0	**84**
Start month	1					**End month**			36			

Objectives

WP2 is directly related to **ESCAPE-2 key objective no. 2**, it uses input from **WP1** and produces output for **WP3** and **WP4**. **WP2** aims to:

- Define, develop, and apply a domain-specific language (DSL) toolchain applicable to a comprehensive list of algorithmic motives (dwarfs) in weather and climate prediction.
- Demonstrate code adaptation and code generation via the DSL toolchain for a number of representative and fundamentally different mathematical algorithms and horizontal discretizations.
- Develop and promote APIs and standard interfaces across the DSL toolchain in order to improve reusability and inter-operability, and leverage code adaptation to emerging HPC architectures.

Description of Work

WP2 is led by MSWISS (Oliver Fuhrer) and co-led by MPIM (Reinhard Budich)

Task 2.1 Definition of requirements and high-level DSL (Lead: MSWISS, Partners: MPIM, DKRZ, ECMWF, CMCC, RMI)

A key element to facilitate broad acceptance and adoption of using a DSL toolchain that translates code written by the weather and climate prediction scientist to a machine readable code adapted and optimized for a specific HPC architecture is a comprehensive and high-level DSL front-end that formally separates the task of analysis and synthesis. A key aspect of the proposed DSL toolchain is the High-level Intermediate Representation (HIR), which provides a specification of the algorithmic motif in a way that is agnostic to the programming language as well as the weather and climate prediction model used by the individual scientist. In this task the DSL front-end will parse a high-level, domain-scientist readable DSL and translate it into the HIR. The former is close to the algorithm and equations, the latter is abstract and machine readable.

In this task we define an expressive higher-level language that avoids "boiler plate" code and provides a concise syntax increasing the readability of a comprehensive list of numerical algorithms. This task will collect and describe the required computational patterns that a DSL for weather and climate codes should cover, based on the different dwarfs provided by **WP1**. The outcome will be a definition of the high-level DSL that allows to syntactically express these dwarfs in a concise manner.

Specifically, this task will define the requirements and develop a first definition of the DSL language (**D2.1, MSWISS**) that supports a comprehensive spectrum of algorithmic motives present in state-of-the-art weather and climate prediction models. This will be achieved by broad engagement and involvement of representative model development experts in a joint **working group**, covering a wide spectrum of dwarfs, i.e. physical parameterizations, advection, solvers and different grid discretizations. The working group will organize a workshop that will be open to key developers other than ESCAPE-2 partners.

Task 2.2 Develop a prototype front-end parser for translation into a high-level intermediate representation (Lead: MSWISS, Partners: ECMWF, DKRZ)

This task will develop a prototype parser (building heavily on existing open-source tools such as clang, CLAW, python) for the DSL that model developers can use to prototype and evaluate the language definition (**D2.2, MSWISS**) for their specific applications of interest. It is expected that domain scientists and model developers use the new high-level DSL with the front-end parser to generate a HIR. The HIR forms the basis for the rest of the DSL toolchain developed in **Task 2.3**, and it will be defined in this task (**D2.3, MSWISS).** The HIR allows us to decouple the front-end DSL language from the rest of the compiler toolchain by allowing other front-ends to generate HIR and leverage the DSL toolchain developed in **Task 2.3** for optimization and code generation. As an important step towards the standardization of tools, the DSL front-end may be replaced as long as, ultimately, there is a translation into the HIR. This will allow us to compile codes efficiently on multiple architectures by using the toolchain proposed in **Task 2.3**. The early availability of the front-end parser will allow an iterative refinement and improvement of the DSL language together with the domain scientists and model developers and directly involve them in required trade-offs and design choices. Consequently, an iterative process between **Task 2.2** and **Task 2.1** is anticipated. Furthermore, appropriate training in the use of the parser will be organized in *WP5* to support the acceptance and usability of the toolchain.

Task 2.3 Implementation of the DSL toolchain (Lead: MSWISS, Partners: ECMWF, DKRZ)

This task will develop the modular DSL toolchain (**D2.5, MSWISS**) (consisting of front-end, transformations, code generator; see Section 1.3). It will build on existing, open-source technologies (Atlas, GridTools, clang, etc.) which have partly been developed within the ESCAPE project and deliver a modular implementation with standardized interfaces that allows for easy inter-operability with other approaches that have been developed. The HIR is passed on to the rest of the toolchain. A set of safety checkers will signal errors or warnings to the user about bad or ill-formed computational implementations (e.g. race conditions, access to uninitialized fields). They will provide a high level of safety since they can protect codes from typical errors common in parallel implementations using standard programming models like OpenMP or OpenACC. Furthermore, a set of performance model driven optimizers will reorganize the code, group or fuse operations in order to increase data locality and generate an efficient implementation targeting a specific computing architecture. Halo-exchanges and/or lateral boundary conditions will be auto-generated from the horizontal dependencies in the user code. Finally the code generators (back-ends) will produce code in standard general purpose programming languages (C++/CUDA) that can be compiled with industry-strength compilers and will target three different hardware architectures (x86 multi-core, Intel Xeon Phi, NVIDIA GPU).

Task 2.4 Demonstration of the DSL with weather and climate prediction dwarfs (Lead: MPIM, Partners: MSWISS, ECMWF, CMCC, DKRZ, BSC, RMI)

This task will demonstrate and verify the usability of the newly developed DSL toolchain for code generation. The task will translate a specified set of representative dwarfs delivered by **WP1** to the new high-level DSL and apply the DSL toolchain to generate code for multiple hardware targets (**D2.4, MPIM**). These dwarfs are tested and verified against their original versions and subsequently supplied to **WP3** for re-integration (**Task 3.3**). An in-depth evaluation of the DSL in terms of usability and performance will be the outcome, including a comparison between the original code and the DSL version and their behaviour on different systems (**D2.6, MPIM**).

Deliverables

D2.1 High-level domain-specific language (DSL) specification (MSWISS, R, PU, 6)

D2.2 DSL front-end to parse DSL into high-level intermediate representation (HIR) (MSWISS, OTHER,

	CO, 12)
D2.3	High-level intermediate (HIR) representation specification (MSWISS, R, PU, 12)
D2.4	Demonstration of domain specific language toolchain for selected weather and climate dwarfs (MPIM, R, PU, 24)
D2.5	Implementation of comprehensive domain specific language toolchain (MSWISS, OTHER, PU, 32)
D2.6	Report on the usability and performance of applying DSL toolchains in weather and climate prediction (MPIM, R, PU, 36)

Partner Roles	
CMCC	Contributes to D2.1, D2.4, D2.6
DRKZ	Contributes to D2.1, D2.2, D2.3, D2.4, D2.5, D2.6
ECMWF	Contributes to D2.1, D2.2, D2.3, D2.4, D2.6
MPIM	Co-leads WP2; is responsible for D2.4, D2.6; contributes to D2.1, D2.4
MSWISS	Co-leads WP2; is responsible for D2.1, D2.2, D2.3, D2.5, D2.6; contributes to D2.6
BSC	Contributes to D2.1, D2.6
RMI	Contributes to D2.1, D2.6

WP Dependencies	
WP1	Delivers dwarfs for specification of DSL and DSL adaptation.
WP3	**WP2** delivers dwarfs rewritten using the DSL to **WP3** (**Task 3.3**) **WP2** receives feedback and expertise on performance testing of dwarfs and the building/enhancement of optimizers.

Work package number	3	**Lead beneficiary**							DKRZ, BULL			
Work package title	Weather and climate benchmarks: HPCW											
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	**TOTAL**
Person months/ participant:	5	15	8	2	6	0	0	13	1	0	4.5	22	**76.5**
Start month	1					**End month**	36						

Objectives

WP3 is directly related to **ESCAPE-2 key objective no. 3**, it ingests input from **WP1** and **WP2** and produces output for **WP4**. **WP3** aims to:

- Develop a hierarchy of benchmarking components representing key elements in the workflow of weather and climate prediction systems.
- Establish a representative benchmark suite (the High Performance Climate and Weather benchmark, **HPCW**) based on these components for (pre)-exascale applications of climate and weather codes.
- Demonstrate the value of HPCW using a wide range of available hardware architectures.

Description of Work

WP3 is led by DKRZ (Joachim Biercamp) and co-led by BULL (Xavier Vigouroux and Erwan Raffin)

Task 3.1 Set-up of reference benchmarks (Lead: DKRZ, Partners: BULL, CMCC, MPIM, ECMWF, RMI, DMI)

The goal of this task is to establish reference HPCW benchmark set-ups (version-0, v0) across two tiers of the benchmark hierarchy. The reference benchmarks are selected so as to include all dwarfs, which will undergo methodological developments in **WP1** and **WP2**. **Task 3.1** will define control versions against which to compare the performance of the upgrades across the range of the available hardware infrastructure (see Table 4). The final set of HPCW benchmarks will be defined and assessed towards the end of the project (**Task 3.5**, see Table 9).

HPCW v0-benchmarks will include two tiers:

- Models: ICON-OCE finite volume/element combination (FV/FE) and NEMO finite difference (FD) ocean dynamical cores, IFS FVM and DG atmosphere dynamical core.
- Forecasting systems: Kronos workload simulator.

In addition, model benchmarks for which no developments will be performed (e.g. IFS atmosphere (spectral transforms, ST)) will be included in the evaluation as stable suites. The ALADIN ensemble prediction system will be set up for the assessment of quality measures in **Task 3.2**, and for the weather and climate VVUQ use case in **Task 4.6**.

Task 3.1 will perform:

1. set-up of v0-benchmarks on selected hardware;
2. definition of first-order correctness measures for initial tests;
3. definition of representative workload for the respective benchmarks;
4. establishment of baseline performance metrics for v0-benchmarks (e.g. time/energy to solution).

Output will be the installation of the above set of v0-benchmarks on the available hardware (**D3.1, BULL**).

Task 3.2 Definition of quality measures across hardware and software (Lead: RMI, Partners: BULL, DKRZ)

For addressing the challenge of developing reliable and portable verification routines that ensure the correctness of benchmark execution two basic scenarios will be addressed:

Technical modifications: when the same code is run on different HPC systems (selected from Table 4), which implies the usage of different software environments (e.g. compilers, libraries) or the code has been modified without changing the scientific formulation of the model (e.g. by performance optimization or by implementing new techniques like DSL). In these cases the results will not be bit-identical and thus we need to evaluate if the model is still correct in the sense that it produces the same climate. This will be done using ensemble-based consistency tests as described by Baker et al. (2015) and Milroy et al. (2017). The method is already used by NCAR for similar purposes in their new Ensemble Consistency Test (ECT) tool[32].

Scientific modifications: when two different formulations of the same model or model component need to be compared, namely codes that have been run before and after undergoing mathematical and algorithmic changes, different radiation schemes, or different ocean and atmosphere models. In this case we need to evaluate and compare the scientific quality of the results. This includes case-based comparison in well-defined sensitivity study environments and statistical assessments based on larger experiment samples. The scientific evaluation will be coordinated with **WP4**. In **Task 3.2** the uncertainties in a meteorological and numerical sense will be quantified and a subset of relevant variables and/or model parameters will be identified and extracted that drive benchmark variability. These will be used to train the meta-model of URANIE used in **WP4** where the tests to quantify uncertainties will be carried out in **Task 4.4** and **Task 4.6**. In particular, the scientific evaluation will be

[32] https://www2.cisl.ucar.edu/news/new-tools-verifying-cesm-simulation-outputs

performed between different models, i.e. ICON vs NEMO ocean and IFS vs ICON atmosphere models. The use of single-precision instead of double-precision arithmetics will be used to test its impact on reproducibility.

Output will be the assessment of the above evaluation tools based on the configurations defined in **Task 3.1** (**D3.2, RMI**). The correctness evaluation will feed into the performance assessment in **Task 3.3** and the VVUQ assessment in **Task 4.6**.

Task 3.3 Integration of novel dwarfs and benchmarking (Lead: BULL, Partners: MSWISS, LU, POLIMI, RMI, MPIM, CMCC, ECMWF, DKRZ)

This task will establish the benchmarks defined in **Task 3.1** (see table below in **Task 3.5**) with input from dwarf development on mathematics and DSL provided by **WP1** and **WP2**. These benchmarks will constitute versions 0.n (v0.n), where 'n' denotes sequential update versions.

Specifically, this task will:

1. evaluate and select dwarfs and DSL implementations to be integrated in benchmarks (see Table 9 below).
2. integrate dwarfs in relevant benchmarks and ALADIN ensemble on available hardware;
3. verify correctness of results across hardware options employing verification and evaluation tools specified in **Task 3.2.**

Output will be an intermediate installation and initial correctness assessment of the above set of v0.n-benchmarks on the available hardware (**D3.3, BULL**). The above sequence of steps will be progressively applied as the developments from **WP1** and **WP2** reach maturity.

This task will further prepare demonstration cases to be used by **WP4** for the quantification and modelling of uncertainty. These cases will be based on the ALADIN ensemble prediction system with both v0 and v0.n as the individual version upgrades represent a source of uncertainty.

Task 3.4 Workflow analysis with Kronos (Lead: ECMWF, Partner: MPIM, ECMWF)

The Kronos workload simulator has been developed by ECMWF for the NextGenIO project. The simulator can generate synthetic workloads that are characteristic for operational weather and climate prediction systems taking progressively into account the specific requirements of the advanced approaches and solutions developed within ESCAPE-2. Kronos can (a) integrate selected dwarfs and models into its synthetic application execution framework, (b) synthesise dwarfs by mimicking their workload in terms of memory usage and communication patterns.

Task 3.4 develops Kronos to:

- Synthesise the model benchmarks and act as a high-level control suite that allows execution and analysis of workloads of increasing scale and complexity. In this phase Kronos generates synthetic applications that reproduce the behaviour of model benchmarks from their profiling data. Kronos will provide an additional degree of flexibility that allows exploring complex resource contentions that only appear when a mix of workloads of increasing size make use of shared resources (compute, network and storage). The definition of Kronos benchmark configurations is derived from **Task 3.1**.
- Define scaled-up schedules of model benchmarks. To scale the workload up, a model of workload scaling will be developed. Key workload parameters are taken into account (number of concurrently executing workloads, sequencing, specific scaling factors, etc.) and an estimation of the scaled-up workload is automatically generated. This step will make use, and further develop, the Machine Learning algorithms already embedded in Kronos.
- Use dwarfs as realistic application components within the workload. Up to this point Kronos has only made use of light-weight highly-controlled kernels, whose performance characteristics are both strictly known and well specified. In order to use dwarfs instead of synthetic applications, the following features will be developed:
    - Add a higher degree of flexibility of the Kronos execution framework.
    - Add a novel monitoring and data collection structure to interpret different (and not strictly

controlled) types of behaviour of the applications. This work will be coordinated with **Task 3.2**.

The output of this task will be synthesized versions of Tier-1 models within the Kronos framework including scaled-up workload schedules and selected integrations of real dwarfs in the synthetic execution framework of Kronos (**D3.4, ECMWF**). Kronos will be an open-source software provided by the NextGenIO project.

Task 3.5 The HPCW Benchmark Suite (Lead: DKRZ, Partner: BULL, BSC, ECMWF, CMCC)

This task will assemble and disseminate the components of the final version of the benchmark suites, denoted as version-1 (v1). The list of benchmarks is shown in Table 9. This task includes packaging of codes, provision of run-time scripts and environments, and of all data needed as input and for evaluation. It will leverage experience from procurements done by the partners and also from the IS-ENES2 and ESiWACE projects. Rules and guidelines for set up and evaluation of benchmarks based on the tools defined in **Task 3.2** will be provided.

Table 9: HPCW benchmarks

HPCW benchmark tier	Specification	Options for novel developments to be included
Models	ICON ocean FV NEMO ocean FD IFS atmosphere FV IFS atmosphere DG IFS atmosphere ST ICON atmosphere FV	Mathematics (finite-difference, time stepping), DSL Mathematics (time stepping), DSL Mathematics (discretization, time stepping, fault tolerance), DSL Mathematics (discretization, time stepping, fault tolerance), DSL N/A (only as reference) Mathematics (neural networks), DSL
Systems	Kronos workload simulator	Simulating the above

Best practice standards for HPCW benchmark dissemination and application will include:

- complete and portable stand-alone source codes,
- simple license request procedure for use and modification for benchmarking purposes,
- portal for software and documentation downloads,
- physically meaningful experiments with typical workload characteristics,
- resolution-cases and parallelism choices that cover the spectrum of computational capabilities of systems contained in the current and near-future HPC landscape (e.g., TOP 500 list),
- correctness criteria,
- performance and energy metrics permitting full cost evaluation.

The suites will be tested on a variety of available system architectures. The performance for computing, communication, memory access, I/O patterns, and power efficiency will be assessed. Profiling tools will include BSC's analysis tools Extrae and Paraver, and Bull's energy and performance measurement software. The analysis will provide a range of performance metrics, it will identify the main bottlenecks for time/energy-to-solution, and investigate the individual performance the **WP1** and **WP2** developments have contributed when comparing HPCW v0 and v1.

Our long term goal is to establish HPCW as a reference that is used beyond the scope and lifetime of the ESCAPE-2 project. We expect the HPCW to grow as model development evolves. The framework for maintaining and promoting HPCW in the long term will be set up in collaboration with the EU funded Centre of Excellence ESiWACE.

Output of this task will be a portal from which benchmark software, associated input data, and documentation can be accessed (**D3.5, DKRZ**).

Deliverables

D3.1 Set of v0-reference benchmarks (BULL, OTHER, CO, 6)

D3.2 Assessment of the evaluation and verification tools based on the v0-benchmark configurations

	(RMI, R, PU, 12)
D3.3	Installation and initial correctness assessment of the above set of v0.n-benchmarks on the available hardware (BULL, OTHER, CO, 24)
D3.4	Synthesized tier-1 model benchmarks by Kronos and dwarfs replacing selected model components in Kronos (ECMWF, OTHER, CO, 24)
D3.5	Full HPCW suite v1.0 (DKRZ, OTHER, PU, 36)

Partner Roles	
RMI	Verification and validation tools. Demonstration cases.
DKRZ	Lead of WP. Responsible for definition, assembly and dissemination of HPCW benchmark suite. Contribute to the definition of ICON based benchmarks.
BULL	Co-lead of WP. Lead of task 3.3. Responsible of D3.1 and D3.3. Port and execute HPCW and tools to different hardware platforms available to CEPP and partner sites. Provide performance and energy measures jointly with BSC. HPCW test with innovative tools.
BSC	Provide performance measurement and profiling analysis for the HPCW benchmark suite on different machines.
ECMWF	a) Kronos. c) Link to WP1 (contribution to Task 3.3). d) FVM reference and upgrade path.
CMCC	Contribute to set-up the HPCW v0-benchmarks (NEMO). Integrate and validate novel DSL NEMO dwarf in tier-1 model benchmark suite. Contribute to provide performance measurement and profiling analysis for the HPCW benchmark suite (NEMO).
MPIM	ICON (ocean, atmosphere) references and upgrade paths.
MSWISS	Link to WP2 (through contribution to Task 3.3).

WP Dependencies	
WP1	Delivery of dwarfs with novel mathematical and algorithmic concepts for integration in benchmarks.
WP2	Delivery of dwarf DSL implementation for integration in benchmarks.
WP4	Delivery of selected benchmarks for the application of VVUQ; provision of uncertainty measures for VVUQ from ALADIN ensemble benchmark.
WP5	Provision of benchmarks for dissemination to wider public; further promotion of benchmarks in ESiWACE centre of excellence.

Work package number	4	Lead beneficiary						CEA, BSC				
Work package title	Verification Validation Uncertainty Quantification (VVUQ)											
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	TOTAL
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	
Person months/ participant:	5	0	0	0	27	34	0	20	0	0	0	0	86
Start month	1					End month		36					

Objectives

WP4 is directly related to **ESCAPE-2 key objective no. 4**, it ingests input from **WP1** and **WP3**. Based on the URANIE platform, **WP4** aims to:

- Develop a generic European VVUQ package for weather and climate simulations that is deployable on supercomputers preparing workloads of pre-exascale computations on many-core configurations.
- Demonstrate the VVUQ package for both dwarf and full forecasting system workloads and optimize the performance of the VVUQ package for the use cases based on the available VVUQ methodologies.
- Enhance the URANIE platform to capitalise and disseminate the approaches learned from the weather and climate community to other science disciplines and use cases.

Description of Work

WP4 is led by CEA (Adrien Bruneton) and co-led by BSC (Kim Serradell)

Task 4.1 Definition of a generic VVUQ methodology for weather and climate prediction (Lead: CEA, Partners: ECMWF, RMI)

This task will define how to adapt the existing URANIE VVUQ framework to weather and climate prediction models and how to exploit the vast experience of ensemble based forecasting for refining the current URANIE set-up for the benefit of other applications. The specific actions in this task comprise:

- the definition of the URANIE software adaptation necessary to interface with and run weather and climate dwarfs, benchmarks and full sized suites (such as the ALADIN ensemble);
- the choice of methods for reducing the degrees of freedom from weather and climate models to those manageable in the VVUQ framework (e.g. EOF, Singular Vectors, spectral truncation, but also the computation of integrated quantities such as total kinetic energy and mass or spectral norms, etc.);
- the implementation of support for the various input/output formats commonly used in weather and climate prediction codes (e.g. NetCDF, GRIB) and of the relevant I/O processing in URANIE;
- the definition of quality measures and scientific evaluation metrics required to perform verification and validation;
- the characterisation of weather and climate model sensitivities, e.g. into initial uncertainties, parameter uncertainties, and non-linear feedbacks;

The outcome of this task will be the set-up of a common and generic framework for the subsequent VVUQ assessments that will be applied throughout **WP4** (**D4.1, CEA**). This task will organize a workshop to support the design of the generic VVUQ framework.

Task 4.2 Prototype test of VVUQ framework (Lead: CEA, Partners: ECMWF, BSC)

This task will introduce a reduced-complexity weather and climate prediction pseudo-model to produce a first-order test with the VVUQ framework defined in **Task 4.1**. The pseudo-model follows the upgraded Lorenz-95 methodology (Thornes et al. 2017) that represents non-linear processes with three distinct scales of variability (mimicking convection, synoptic and planetary wave scale motion) and which is limited to O(500) degrees of freedom. The advantage of using a pseudo-model is that it can be integrated in URANIE as it is, and that the parameter space has a simple structure (low dimensionality, isotropy).

The task will:

- run the Lorenz-95 model in URANIE;
- perform a sensitivity analysis to assess how input data and key parameters drive results;
- perform an uncertainty analysis given pre-defined error bounds for input data and key parameters;
- test the resulting output on a statistically significant set of simulations providing a stable uncertainty quantification (specification of an input law) and sensitivity analysis (Morris-analysis, Sobol-indices, etc.).

As a by-product, this task will lead the way to:

- help to build a common language across communities;
- help to setup the environment for automated sensitivity testing in URANIE;

The outcome of this task is a report summarizing the assessment of the URANIE VVUQ performance for the Lorenz-95 pseudo-model, and a comparison against other existing VVUQ approaches for this model (**D4.2, CEA**).

Task 4.3 Enhancement of URANIE workflow management (Lead: BSC, Partners: CEA, ECMWF)

At present, URANIE can run multiple instances of the same code with different input parameters in a concurrent fashion. This is handled by a simple workflow management implementation where the pool of allocated resources is inspected once and the code is then executed on the available computing facility. To allow adequate scaling the workflow management will be reviewed to efficiently deal with increasing I/O requirements and for running many codes in parallel on large node allocations. This task will therefore:

- Implement a new workflow management system, the Autosubmit tool, which is the workflow management tool developed at BSC for operating large-scale and diverse applications on HPC systems.
- Enhance the I/O management of URANIE to process large data workloads across multiple tasks efficiently without affecting the scalability of the application.
- Perform tests and produce performance assessments using the BSC tools Extrae and Paraver, and to identify potential bottlenecks in the processing.
- Testing with node allocation patterns beyond today's limits of URANIE.

Output of this task will be an upgraded workflow management environment for URANIE that allows running large-scale applications in parallel and on large HPC node allocations (**D4.3, BSC**). A performance analysis will be produced.

Task 4.4 Weather and climate dwarf in VVUQ framework (Lead: CEA, Partners: RMI, ECMWF, BSC, RMI)

Based on the experience in **Task 4.2**, a weather and climate prediction dwarf (Physical parametrization: radiation scheme extracted from HARMONIE atmosphere) from **WP1** will be used to demonstrate the next level of the generic VVUQ framework offered by URANIE, as defined in **Task 4.1**. The radiation dwarf is more challenging for URANIE than the Lorenz pseudo-model due to its increased number of input parameters, its numerical complexity and its computing cost. The dwarf will be run in the URANIE framework, and the URANIE VVUQ methodology will be applied in the same way as in **Task 4.2** (cf. list of bullets).

This task will further explore the impact of reducing numerical precision on the accuracy and performance of the dwarf output. URANIE will be used to drive an arbitrary precision floating point arithmetic emulator that allows a flexible emulation of precision produced by software or introduced by hardware. The search for the optimal level of arithmetic precision for model components or individual model parameters can be automated when the impact of a local precision reduction is compared against typical uncertainty ranges. Those ranges will be obtained from the URANIE analysis, and the quality measures defined in **Task 3.2 of WP3** will be used as a metric to assess the impact of precision reduction. The same procedure will be applied for the Neural Network version developed in **Task 1.4**.

Output of this task will be a report summarizing the assessment of the URANIE VVUQ performance for the weather and climate radiation dwarf and its Neural Network version, including a detailed assessment of the impact of reduced numerical precision (**D4.4, CEA**).

Task 4.5 Adaptation of URANIE to high-dimensional modelling systems (Lead: CEA, Partners: ECMWF, BSC)

In the past, URANIE has proven to be efficient on simulations exhibiting up to 50 input parameters. Input parameter categories are (a) parameters in physical models that drive the output of the model significantly hence produce strong sensitivity, (b) physical input state that predictions use as initial conditions. The latter can assume very large dimensions as they initialize the three-dimensional state of the atmosphere. In weather and climate, the latter corresponds to the initial state of the atmosphere and thus can easily have on the order of 10^8 to 10^{10} degrees of freedom. Both categories are associated with errors in weather and climate prediction. URANIE is currently able to deal with category (a) but not (b). For category (a), URANIE needs to be extended to deal with larger numbers of physical model parameters, the capability for category (b) needs to be introduced from scratch. The former capability is already going to be touched upon for larger parameter sets in previous tasks.

Hence Task 4.5 aims at:

- enhancing the number of parameters for (a), by materialising the feedback from **Task 4.2** and **Task 4.4**
- and introducing (b) for very large input datasets.

Consequently, this task will:

- Assess current URANIE capabilities for a high-dimensional input parameter space.
- Implement and test degree-of-freedom reduction techniques based on the approaches defined in **Task 4.1.** Assess and benchmark the performance of such techniques and define limits of dimensionality beyond which the URANIE framework will not be further workable.
- Review and optimise the URANIE code and workflow management for this level of high-dimensional setting, employing the BSC performance tools established for URANIE in **Task 4.3**.

Output of this task will be the implementation and assessment of degree-of-freedom reduction techniques for URANIE and code enhancements to cope with high-dimensional input parameters (**D4.5, CEA**).

Task 4.6 Weather and climate use case in VVUQ framework (Lead: RMI, Partners: CEA, BSC, ECMWF)

This task will capitalize on developments from previous tasks to investigate the use of the VVUQ framework within the ALADIN ensemble system. Given the complexity of the model setup (a typical model setup would have O(450x540x65) grid-points plus prognostic parameters), the parameter space for the VVUQ assessment for input data category (a) will be limited to those of the radiation dwarf. For input parameter category (b) data reduction techniques developed in **Task 4.5** will be employed. The quality of simulations subject to parameter changes will be compared against the ensemble spread of the model (e.g. by evaluating how many prognostic parameters deviate by more than one or two standard deviations from the ensemble mean). Here, the ensemble spread will provide an estimate for forecast error.

This task will:

- Run the ALADIN ensemble system in the URANIE VVUQ framework;
- Employ the input parameter category (a) perturbation approach developed in **Task 4.4**;
- Apply the data reduction techniques to input data category (b) developed in **Task 4.5**;
- Compare parameter sensitivities against estimates of forecast errors that are given by the ensemble spread;

- Perform a full VVUQ assessment for an entire weather and climate prediction work load; - Assess the computing performance of the full chain. Output of this task will be the run of full-system sized ensemble simulations in URANIE, the assessment of the corresponding VVUQ indicators, and finally the assessment of the corresponding computing performance (**D4.6, RMI**).

Deliverables
D4.1 Definition of common and generic framework for subsequent VVUQ assessments (CEA, R, PU, 9)
D4.2 Assessment of the URANIE VVUQ performance for the Lorenz-95 pseudo-model (CEA, R, CO, 12)
D4.3 Upgraded workflow management environment for URANIE that allows running large-scale applications in parallel and on large HPC node allocations (BSC, OTHER, CO, 15)
D4.4 Assessment of the URANIE VVUQ performance for the weather and climate radiation dwarf (CEA, R, CO, 18)
D4.5 Implementation and assessment of degree-of-freedom reduction techniques for URANIE, and relevant code enhancements (CEA, R, CO, 24)
D4.6 Full-system sized ensemble forecasts within URANIE framework, assessment of the corresponding VVUQ indicators, and assessment of the computing performance (RMI, OTHER, PU, 36)

Partner Roles	
CEA	Co-leads WP2. Contributes to all deliverables D4.1, D4.2, D4.3, D4.4, D4.5 and D4.6, and is responsible for D4.1, D4.2, D4.4 and D4.5
BSC	Co-leads WP2. Contributes to deliverables D4.3, D4.4, D4.5 and D4.6, and is responsible for D4.3
ECMWF	Contributes to deliverables D4.1, D4.2, D4.4, D4.5 and D4.6
RMI	Contributes to deliverables D4.1, D4.4, D4.5 and D4.6, and is responsible for D4.6

WP Dependencies	
WP1	Delivery of weather and climate radiation dwarf, and potential upgrades arising from mathematical and algorithmic developments.
WP3	Delivery of ALADIN ensemble HPCW benchmark and input on evaluation and verification metrics provided by general benchmark testing.
WP5	Provision of demonstrator case and technical enhancements for URANIE that can be exported to other applications.

Work package number	5	**Lead beneficiary**						ECMWF				
Work package title	Dissemination, Communication and Training											
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	**TOTAL**
Person months/ participant:	10	0	0	0	1	1	2	0	2	0	0	1	**17**
Start month	1					**End month**			36				

Objectives
WP5 directly addresses **top-level objectives no. 2, 4 and 5** defined **in Section 1.1**. **WP5** aims to:
• **Prompt a paradigm shift in the understanding and use of novel technologies** and their

impact on weather and climate model design and application.

- **Promote the dissemination and use of weather and climate dwarfs** as well as the **HPCW benchmark** to vendors and data centres for co-design, model development, as well as in training and education.

Description of Work

WP5 is led by ECMWF (Daniel Thiemert) and co-led by partner POLIMI (Luca Bonaventura).

Key tasks of ESCAPE-2 comprise **training, support, dissemination and communication activities**. The aim of **WP5** is to promote the understanding and use of novel technologies and their impact on weather and climate model design and application.

To achieve a large audience, **maintenance, dissemination and communication beyond the project network and lifetime**, the project will promote the dissemination of the weather and climate dwarfs and HPCW benchmarks through the ESCAPE-2 website portal administered by ECMWF. The project partners ECMWF and POLIMI will develop regular workshops and a summer school. The EC's e-infrastructures Centre of Excellence for Performance Optimisation and Productivity (POP[33]) and ESiWACE would support continued benchmarking and code optimization efforts, coordinate common tool developments, and strengthen the technical basis of future weather and climate simulation capabilities in the longer term.

An important objective is thus to **facilitate ESCAPE-2's dissemination and communication** established through five tasks:

Task 5.1 Dissemination and Communication (Lead: ECMWF, Partners: BULL)

In the first stage, **Task 5.1** will develop the plan for Dissemination and Communications including Media Outreach **(D5.1, ECMWF).** The plan will identify avenues and targets for communication, dissemination and exploitation both within and outside the ESCAPE-2 Consortium including communication strategy and brand guidelines.

The second stage of this task will then see the implementation of this plan, including publications, webinars, talks, fairs, newsletters, press releases, etc. Further details on the dissemination and communication activities are already presented in Section 2.2.

Task 5.2 Gender Balance (Lead: LU)

Task 5.2 will develop and implement the plan to ensure women's participation in the developments and applications as active members of research/engineering teams (**WP1** - **WP4**). It will ensure promotion of gender equality and diversity in the activities involved in the outreach and training activities that will be built in **WP5** by ensuring equal gender participation in training videos and webinars, by promoting success stories about women through the project website, and, by promoting partner activities that implement equal opportunity policies in their organisations (**D5.2, LU**).

Task 5.3 Website and Infrastructure (Lead: ECMWF)

To establish efficient interaction between project partners and for communicating with external stakeholders**,** this task will establish the **web portal, the software collaboration platform** including bug tracking facilities, **and the software repository**. In addition, **collaborative tools** such as a wiki will be established and maintained (**D5.3, ECMWF**).

Task 5.4 Summer Schools (Lead: POLIMI)

A summer school will be organised to facilitate early dissemination of outcomes to early-career scientists in the weather and climate prediction communities (**D5.4, POLIMI**). The school will take place after the representative weather and climate prediction dwarfs are released (**D1.7**) and tested by other partners. The main goals of the school will be:

[33] https://pop-coe.eu/

1. introduce novel mathematical techniques implemented in **Task 1.1**, **1.2**, **1.3**, **1.4**,
2. provide an overview of domain-specific languages developed in **WP2**,
3. introduce the generic unified European VVUQ package for weather and climate simulations developed in **WP4,**
4. provide hands-on training to using of the newly developed dwarfs and benchmarks.

Task 5.5 Workshops (Lead: ECMWF)

To support the dissemination activities and interact with the scientific community, this task will organise two dissemination workshops, one at **M12** (**D5.5, ECMWF**), and one at the end of the project (**D5.6, ECMWF**). Whilst the latter aims to provide the results of the project to the community, the former aims to integrate the community into the activities of the ESCAPE-2 project through take-up of feedback from outside the consortium during the workshop.

Deliverables

D5.1 Plan for Dissemination and Communication (ECMWF, R, PU, M2)

D5.2 Plan for Gender Balance (PGB) (LU, R, PU, M4)

D5.3 Web Portal including Software Collaboration Platform (ECMWF, DEC, PU, M4)

D5.4 Summer school (POLIMI, DEC, PU, M18)

D5.5 Dissemination workshop I (ECMWF, DEC, PU, M12)

D5.6 Dissemination workshop II (ECMWF, DEC, PU, M26)

Partner Roles	
ECMWF	will lead this work package and is responsible for delivering Tasks 5.1, 5.3 and 5.5
LU	will lead Task 5.2
POLIMI	will lead Task 5.4
BULL	will take part in Task 5.1 with a focus on industrial aspects

WP Dependencies	
WP1-4	All technical WPs provide input to the dissemination, communication and training activities.

Work package number	6		**Lead beneficiary**					ECWMF				
Work package title	Project Management and Coordination											
Participant number	1	2	3	4	5	6	7	8	9	10	11	12	
Short name of participant	ECMWF	DKRZ	MPIM	MSWISS	BSC	CEA	LU	RMI	POLIMI	DMI	CMCC	BULL	**TOTAL**
Person months/ participant:	15	1	1	1	0	1	0	0	1	0	0	1	**21**
Start month	1					**End month**			36				

Objectives

The general objective of **WP6** is to coordinate the project and ensure that its innovation actions, objectives and impacts will be delivered. Specific project management objectives are:

- 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 interactions between work packages and partners;
- to provide and manage working procedures ensuring transparency within the team and for the EC;
- to manage quality assurance.

Description of Work

WP6 is led by ECMWF (Peter Bauer - Coordinator and Daniel Thiemert - Project Manager).

Task 6.1 Project Management and Coordination (Lead: ECMWF)

Task 6.1 will provide the project management and coordination for the ESCAPE-2 project. The work will comprise project initiation, including setting up the management structure such as General Assembly, Executive Board, project office etc. The necessary communication channels including mailing lists, and document management system will be implemented to support the project work.

The task will also be responsible for the organisation of project meetings including kick-off meeting and annual project board meetings.

In addition, progress monitoring will occur through regular Executive Board meetings (online, monthly) as well as quarterly progress reports from partners and work packages which will also aid the risk and quality management (see **Task 6.2**).

Finally, **Task 6.1** will also serve the scientific coordination of the project, ensuring that all partners and work packages are aligned and working towards the goal of the project.

Task 6.2 Risk and Quality Management (Lead: ECMWF)

Based on the initial risk table already identified in Table 13, this task will maintain a risk register (part of **D6.1, ECMWF**) that is regularly updated (through quarterly progress reports described in **Task 6.1**) or on an ad-hoc basis, as required. The task will thus actively monitor existing and emerging risks and identify mitigation measures.

Task 6.2 will also perform the quality management. A quality manual (also part of **D6.1, ECMWF**) will detail the Key Performance Indicators and will provide target measures for these. The manual will also provide the quality review process for each deliverable of the project, through a rigorous internal review process (at least two reviewers for each deliverable). The quality manual will also contain templates for the project, including deliverables, presentations, posters, etc.

Task 6.3 Administrative and Financial Management (Lead: ECMWF, M1-M36)

Task 6.3 provides the administrative and financial management for the ESCAPE-2 project. This includes management of potential amendments to Grant Agreement and Consortium Agreement, periodic reporting, management of the project finances (distribution of funds, financial claims, etc.) as well as organisation of periodic reviews. The task will also ensure liaison and communication with the European Commission.

Task 6.4 Innovation and IPR management (Lead: ECMWF)

Management of knowledge and innovation is an integral part of our project. We will focus on the role and synergies between the areas of competence of the partners, their experience, capabilities, and on how partners will protect, share, manage IPR capital actual exploitation.

The **consortium agreement** will be set up for regulating the ownership and access to key knowledge (IPR, data etc.) and scientific foreground, among other things, after the communication of the approval of the project by the EC and **before the signature of the Grant Agreement with the EC**.

A strategy for the knowledge management, protection and for the exploitation of results, the **Exploitation Plan** (**D6.2, ECMWF**), will be defined for the consortium in the early stage of the project based on the principles explained 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.

At the *end of the project* a **revised version of the Exploitation plan** (**D6.3, ECMWF**) will be drafted for providing best practices in capturing and assessing the Intellectual Property and providing measures for exploitation for the phase after the closure of the project. ESCAPE-2 will provide **open access to peer-reviewed scientific publications** through a combination of golden open access and green open access (cf. Section 2.2 for details on software IPR and data).

Deliverables
D6.1 Quality Manual, Risks Register: Manual describing procedures for managing quality during project. Risk register tracking type, level, responsibility, mitigation measures (ECMWF, R, CO, 2) D6.2 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 (ECMWF, R, CO, 12) D6.3 Exploitation Plan after closure: Strategy for defining measures for exploitation "after the project" phase, providing evidence of best practices in capturing and assessing IPR (ECMWF, R, PU, 36).

Partner Roles	
ECMWF	will lead the work package and deliver all its tasks
WP (Co) leaders	Will contribute to the management of the project and their WPs

WP Dependencies	
All	WP1 ensures that all WPs perform according to contract.

Table 10: List of work packages

Work package No	Work Package Title	Lead Participant No	Lead Participant Short Name	Person-Months	Start Month	End month
WP1	Algorithms & Mathematics	9	POLIMI	119.3	1	36
WP2	Programming Models & Domain-specific Languages	4	MSWISS	84	1	36
WP3	Weather and climate benchmarks: HPCW	2	DKRZ	76.5	1	36
WP4	Verification Validation Uncertainty Quantification (VVUQ)	6	CEA	86	1	36
WP5	Dissemination, Communication and Training	1	ECMWF	17	1	36
WP6	Project Management and Coordination	1	ECMWF	21	1	36
			Total person-months	403.8		

Table 11: List of deliverables.

Deliverable (number)	Deliverable name	Work package number	Short name of lead participant	Type	Dissemination level	Delivery date (in months)
D1.1	A semi-Lagrangian DG dwarf	WP1	POLIMI	OTHER	CO	M18
D1.2	A semi-implicit, semi-Lagrangian DG prototype	WP1	POLIMI	OTHER	CO	M32

	dynamical core					
D1.3	A report on identification of local data recovery approaches suitable for weather and climate prediction applications	WP1	POLIMI	R	PU	M18
D1.4	Fault tolerant implementations of GMRES	WP1	POLIMI	OTHER	CO	M27
D1.5	Multigrid and multilevel technology based IFS-FVM model	WP1	LU	OTHER	CO	M32
D1.6	Training & preliminary validation of the ANN for the physical parametrization of radiation	WP1	DMI	OTHER	CO	M32
D1.7	Extract and provide a list of representative algorithmic motives in the form of weather and climate dwarfs	WP1	POLIMI&LU	OTHER	CO	M9
D1.8	Complement the list of dwarfs with novel concepts	WP1	POLIMI&LU	OTHER	PU	M27
D2.1	High-level domain-specific language (DSL) specification	WP2	MSWISS	R	PU	M6
D2.2	DSL front-end to parse DSL into high-level intermediate representation (HIR)	WP2	MSWISS	OTHER	CO	M12
D2.3	High-level intermediate (HIR) representation specification	WP2	MSWISS	R	PU	M12
D2.4	Demonstration of domain specific language toolchain for selected weather and climate dwarfs	WP2	MPIM	R	PU	M24
D2.5	Implementation of comprehensive domain specific language toolchain	WP2	MSWISS	OTHER	PU	M32
D2.6	Report on the usability and performance of applying DSL toolchains in weather and climate prediction	WP2	MPIM	R	PU	M36
D3.1	Set of v0-reference benchmarks	WP3	BULL	OTHER	CO	M6
D3.2	Assessment of the evaluation and verification tools based on the v0-benchmark configurations	WP3	RMI	R	PU	M12
D3.3	Installation and initial correctness assessment of the above set of v0.n-benchmarks on the available hardware	WP3	BULL	OTHER	CO	M24
D3.4	Synthesized Tier-1 models by Kronos and dwarfs replacing selected model components in Kronos	WP3	ECMWF	OTHER	CO	M24
D3.5	Full HPCW suite v1.0	WP3	DKRZ	OTHER	PU	M36
D4.1	Definition of common and generic framework for subsequent VVUQ assessments	WP4	CEA	R	PU	M9
D4.2	Assessment of the URANIE VVUQ performance for the Lorenz-95 pseudo-model	WP4	CEA	R	CO	M12
D4.3	Upgraded workflow management environment for URANIE that allows running large-scale applications in parallel and on large HPC node allocations	WP4	BSC	OTHER	CO	M15
D4.4	Assessment of the URANIE VVUQ performance for the weather and climate radiation dwarf	WP4	CEA	R	CO	M18
D4.5	Implementation and assessment of degree-of-freedom reduction techniques for URANIE, and relevant code enhancements	WP4	CEA	R	CO	M24
D4.6	Full-system sized ensemble forecasts within URANIE framework, assessment of the corresponding VVUQ	WP4	RMI	OTHER	PU	M36

	indicators, and assessment of the computing performance					
D5.1	Plan for Dissemination and Communication		ECMWF	R	PU	M2
D5.2	Plan for Gender Balance		LU	R	PU	M4
D5.3	Web Portal including Software Collaboration Platform		ECMWF	DEC	PU	M4
D5.4	Summer school		POLIMI	DEC	PU	M18
D5.5	Dissemination workshop I		ECMWF	DEC	PU	M12
D5.6	Dissemination workshop II		ECMWF	DEC	PU	M26
D6.1	Quality Manual, Risks Register: Manual describing procedures for managing quality during project. Risk register tracking type, level, responsibility, mitigation measures		ECMWF	R	CO	M2
D6.2	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		ECMWF	R	CO	M12
D6.3	Exploitation Plan after closure: Strategy for defining measures for exploitation “after the project” phase, providing evidence of best practices in capturing and assessing IPR		ECMWF	R	PU	M36

3.2 Management structure, milestones and procedures

3.2.1 Project organisation

This section describes the management of the project at the Consortium level. The objective is to keep the project on-track by defining and applying control and reporting procedures. At the same time, the consistency of and the cooperation among the activities in different work packages will be ensured.

A **Consortium Agreement** (CA) will define and complete any points not covered by the EC Grant Agreement. The CA will be signed before the project commences. The CA includes in particular details about:

- The organisation of the consortium, as described below.
- The financial distribution on the basis of each participant's effort and activity type.
- Procedures for changes in the consortium composition.
- IPR and exploitation: definition of the background brought by all participants and related access.
- Rights and rules for joint ownership, and access rights to project results for participants and 3rd parties.
- Dissemination rules for managing confidentiality and approving public presentations and publications.

The structure that has been adopted for ESCAPE-2 follows the DESCA Model Consortium Agreement for Horizon 2020 projects, and specific roles and decision-making responsibilities have been assigned accordingly here. Further details about all components are given below and in Figure 12.

The management body of ESCAPE-2 is composed of the General Assembly, the Executive Board, and the Project Office, aiming to:

- stay transparent and flexible at each stage,
- keep the structure as light as possible,
- provide working procedures offering full transparency for the participants and the EC,
- support the efficient cooperation between WPs,
- ensure a good link between technology, science and application,

- maintain control of the overall project while exploring synergies and encouraging creativity.

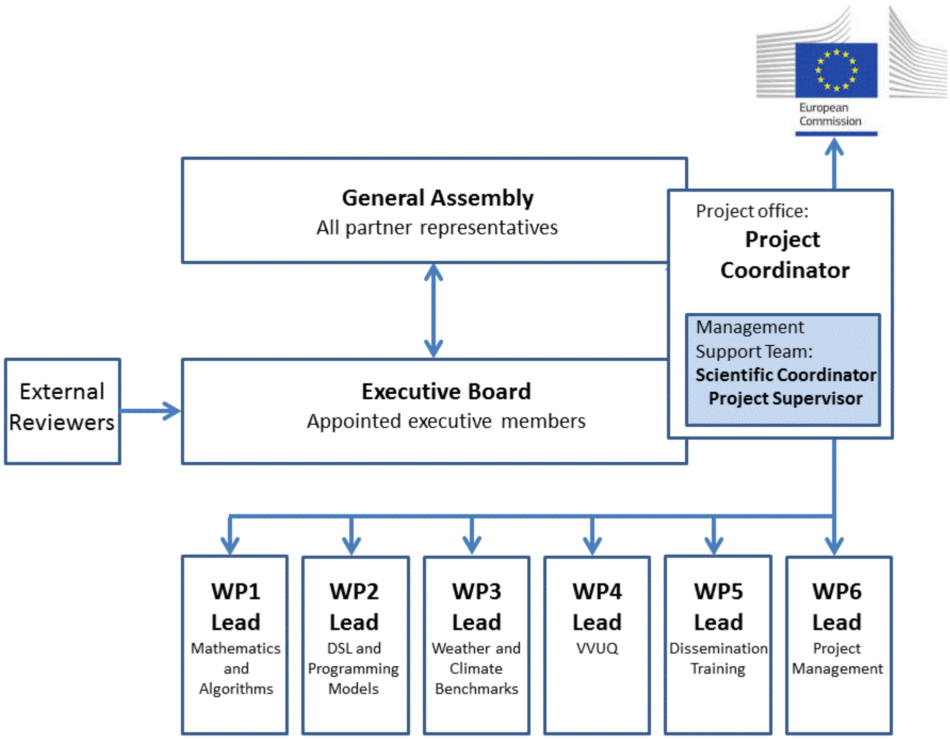

Figure 12: ESCAPE-2 project management structure.

The **General Assembly** (GA) is the ultimate decision making body for ESCAPE. This will be chaired by the Project Coordinator and will be composed of representatives from all partner organisations. The GA is the decision making body. The responsibilities of the GA include:

- making the strategic decisions about the progress of the project,
- making the strategic decisions in case of conflict between beneficiaries,
- making the strategic decisions concerning IPR,
- modifying the Consortium Agreement if necessary,
- modifying the composition of the Consortium, according to the provisions of the Consortium Agreement and the rules imposed by the EC, if necessary,
- and approving major redistributions of budget if necessary.

The GA will convene, when necessary, during the bi-annual meetings of the Executive Board or on request by one third of its members. The members of the GA are the official representatives of the beneficiaries as appointed by their organisation. They are responsible for the management of the technical and/or scientific team working in their organisation and the timely production of the reports and results for which they have responsibility. The partner representatives report to the management of their organisation, to the GA for overall contractual and technical issues and to the WP leaders for issues specific to the progress of the WP's involving their organisation.

The **Executive Board** (EB) will be the supervisory body ensuring successful execution of the project, and will be accountable to the GA. The board will be chaired by the Project Coordinator. The EB is the coordination body and its members will be appointed by the GA. The EB will convene twice per year. EB meetings will be chaired by the Project Coordinator. Responsibilities of the EB include the:

- monitoring of the effective and efficient implementation of the project,
- preparation of meetings, proposal of decisions and preparation of the agenda for GA meetings,
- assurance that the project is compliant with the CA and, if required, proposal of modifications of the CA,
- provision of advice to the GA on rearrangement of tasks and budget, if applicable.

The EB will be supported by two **external reviewers** (Prof. Rupert Klein, FU Berlin, and Prof. Frank Giraldo, Naval Postgradual School; see letters of commitment in Section 6 and Table 14). This review function will be enabled on demand by the EB.

The **Project Office** (PO) will conduct the day-to-day management of ESCAPE-2 on behalf of the GA. The PO executes the decision made by the EB to which it reports. Members are the Project Coordinator, the Scientific Coordinator and the Project Supervisor. The PO is supported by the administrative staff of the Coordinator's organisation. The Project Office is responsible for the:

- management of planning and practical decision making on the progress of the project,
- submission of deliverables to the EC and coordination of follow-up actions,
- supervision of the dissemination of project output within the consortium and the external communication,
- organisation of committee meetings,
- coordination of activities between the WPs,
- management of dissemination activities and IPR issues,
- management of the budget,
- management of the risks,
- management of the quality control.

The members of the PO are contributed by the lead partner ECMWF. ECMWF has an extensive track-record in coordinating and participating in complex multidisciplinary projects.

The **Project Coordinator** (PC) is responsible for the overall coordination of the project and will act as the point of contact for the EC. The PC has ultimate responsibility for ensuring that the project delivers what is expected. The PC is also in charge of the management of innovations originating from the project.

The **Scientific Coordinator** (SC) is responsible for overseeing the scientific content of the project and ensuring that all work packages are contributing to the outcomes and scientific excellence of ESCAPE-2.

The **Project Supervisor** (PS) will have overall responsibility for the efficient coordination of the project in terms of resources and time schedule, risks, quality, and innovation management. The PS acts on behalf of the PC and makes all the decisions relating to the contractual provisions. The PS's role is to ensure the smooth operation of the project in accordance with the provisions of the CA. The PS reports to EC (via the PC) and GA for financial and contractual issues, and to the EB for technical and management issues. The PS is also responsible for the maintenance of the risks register. The PS will organise the project reviews in collaboration with the partners and will be responsible for compiling the final report. From the first review onwards, the PS will be responsible for issuing the "cross-reference" document, which will trace how recommendations made by the EC's Project Officer and the reviewers are taken into account.

The **Work Package Leaders** (WPL; and their co-leaders) have a responsibility to ensure delivery of their WP objectives and deliverables, working closely with PC and SC to support the outcomes of ESCAPE-2 as a whole. They are the coordinators of the different work packages. Each WP has its own leader, objectives and resources. The WPL will define the "Quality Plan" (QP) for the work of the teams involved. This QP will define the acceptance criteria for each deliverable. These activities are part of the WP's in terms of resources and responsibility and are conducted by the partners contractually involved in each WP. The WPLs control the progress of the WP.

Milestones relevant to the ESCAPE-2 project have been identified in Table 12 below.

Table 12: List of milestones.

Milestone number	Milestone name	Related work package(s)	Due date (in month)	Means of verification
MS1	Dwarfs extracted and provided to other work packages	WP1	M12	D1.7 submitted
MS2	Novel concept dwarfs developed and provided to other work packages	WP1	M27	D1.1, D1.4, D1.6 submitted
MS3	High-level DSL agreed upon and prototype parser available	WP2	M12	D1.1, D1.2, D1.3 submitted
MS4	DSL toolchain applied to a set of dwarfs	WP2	M24	D2.4 submitted
MS5	Proof of usability and advantages of the DSL toolchain approach	WP2	M36	D2.5, D2.6 submitted
MS6	Benchmarks version-0	WP3	M9	D3.2 submitted
MS7	Benchmarks version-1 with general access through public portal	WP3	M36	D3.3, D3.4, D3.5 submitted
MS8	URANIE VVUQ adaptation to weather and climate prediction and adaptation to large HPC systems	WP4	M15	D4.1, D4.3 submitted
MS9	Assessment and implementation of degree-of-freedom reduction techniques for URANIE VVUQ platform	WP4	M24	D4.5 submitted
MS10	Full-sized weather and climate prediction suite demonstration with URANIE VVUQ platform	WP4	M36	D4.6 submitted
MS11	Dissemination and Communication activities planned and implementation started	WP5	M6	D5.1, D5.2 D5.3 submitted
MS12	Project Initiated and set-up	WP6	M2	D6.1 submitted
MS13	Project completed	WP6	M36	All deliverables submitted, final report provided

3.2.2 Project management procedures

Decision-making process and conflict resolution

The project is steered and guided following a top-down approach, while issues are raised and problems resolved following a bottom-up approach. The details of the procedures will be described in the CA. Its general principles are:

- The GA is the ultimate decision making body.
- Unless otherwise specified decisions are made by applying a simple majority of the members. In case of equal votes, the vote of the Chairman is decisive.
- In case of serious failing of one participant to fulfil their commitments to the extent that the project is jeopardized, the GA will have the right to expel this participant. This decision, as well as the one to include a new partner, will be made requiring approval by 2/3 of the GA members.
- The project could be affected by conflicts of various types: strategic, technical, shortage of resources, and others. WP Leaders, PC and PS should anticipate the emergence of such conflicts, and discuss the best way to resolve them with the partners. The project employs a multi-level escalation strategy. Conflicts at work level should first be reported to WPLs. If not resolved at this level, problems and possible solutions will be formally discussed at the EB-level, moderated by the PS, where, if necessary, a vote could take place. If no unanimous decision can be found,

the EB will report the issue to the GA where the final decision will be taken in a meeting moderated by the PS.

Risk Management

The EB is in charge of managing the risks while the PS maintains the risks register, in order to prevent any deviation from the plans. The risks register contains the following information for each risk identified:

- risk description (fact or event which could jeopardize the correct functioning of the project),
- work package/task involved,
- probability (almost certain, likely, moderate, unlikely, rare),
- impact (insignificant, minor, moderate, major, catastrophic),
- risk response (type and description),
- responsibility,
- due date,
- status (open, closed).

At each EB meeting, the register will be analysed and the values of each parameter updated. With the help of this information, the EB will be able to make the most appropriate decisions. The PS will ensure that the risks register is permanently up to date and that the recommended measures are in place in due time.

The ESCAPE-2 project risks and appropriate mitigation measures are listed in Table 13.

Table 13: ESCAPE-2 risks.

Description of risk (indicate level of likelihood: Low/Medium/High)	WP(s) involved	Proposed risk-mitigation measures
Novel disruptive algorithms do not produce expected accuracy and efficiency or their implementation is delayed. (D1.1, D1.5) (medium).	WP1	For D1.1 and D1.5 several algorithmic strategies are possible. An iterative implementation of alternative algorithmic options and experimentation with a trade-off between numerical accuracy and computing efficiency would be performed. If performance of D.1.1 is unsatisfactory, the decision can be taken to provide a final deliverable for the DG model based on an Eulerian IMEX approach instead while similarly work can proceed with the standard IFS-FVM if D1.5 does not produce the expected efficiency.
The scientific novelty of the training and validation of the ANN based model (D1.6) poses a natural technical risk due to complexity of task the task especially taking into account the large number of potentially relevant datasets (medium).	WP1	In case deliverable D.1.6 is delayed, the decision can be taken not to proceed with adaptation and benchmarking for different HPC architectures.
Provision of weather and climate dwarfs not shared effectively within consortium (D1.6) (low).	WP1	D1.6 provides a range of dwarfs. This lowers the risk since the work in WP2-WP5 can proceed even with a selection of the proposed dwarfs. Meetings and communication with the related partners involved in D1.6 and tasks in WP2-WP5 lowers the risk and the tasks

		involved allow flexibility to change the selection of required dwarfs.
The technical complexity and novelty of (D1.1, D1.2, D1.4 and D1.5) poses a risk that they will not be ready for D1.8 (medium).	WP1	All of the proposed novel concepts could provide benefits for the weather and climate dwarfs. The successful advancements can complement the list of dwarfs (D1.8) but any delay would not have a critical implication at this stage of research.
Fail to define a common DSL applicable to a wide range of weather and climate dwarfs. If the requirements from the different model developers cannot be aligned sufficiently or the different algorithmic motifs cannot be mapped onto a common denominator, it might be unfeasible to define a single, acceptable high-level DSL frontend for all possible dwarfs (medium).	WP2	1. Define more than one high-level DSL frontend that can be translated into the high-level intermediate representation (HIR) and thus still leverage a large part of the DSL toolchain. 2. Restrict the applicability of the DSL to a specific subset of domain-scientists or algorithmic motifs.
Insufficient performance of DSL toolchain. The main thrust of the work in ESCAPE-2 is on defining a common high-level DSL frontend and we rely on existing open-source technology (GridTools) for delivering performance on the different hardware architecture targets. New dwarfs not considered so far may not perform well (medium).	WP2	1. Adapt back-ends for the new algorithmic motifs in order to deliver performance. 2. Accept performance penalty and conduct performance comparison for dwarfs that have a good match for the existing backend implementations.
Hardware not available as expected due to providers change in their roadmap will reduce the number of HPC architectures available for testing during project duration. (D3.5) (medium)	WP3	As soon as information on roadmap modifications is publicly available, the list of target HPC systems should be revised and some system should be excluded from Task 3.3 and D3.3.
Software environments incompatibility (due to compilers, libraries, operating systems, etc.) can lead to the incapacity to re-integrate dwarfs into larger code on some HPC systems. (D3.3, D3.5) (high)	WP3	The potential incompatibilities should be monitored since the beginning of the project. Anyway, the list of targeted HPC systems and associated software environment should be revised according to incompatibilities and some system should be excluded from D3.3. Incompatibilities will be documented in D3.5.
Disagreement on a common VVUQ framework definition. If the NCWP and URANIE community fail to jointly define a common VVUQ framework, it will prevent the relevant usage of URANIE on NWCP components, and the fruitful interpretation and discussion of the URANIE results (low).	WP4	1. Organize further workshop sessions where each participant exposes its own vision and reach a common agreement 2. Ask an external partner to contribute to this definition, helping project partners to find a common agreement.
Implementation of NWCP file format support in the URANIE platform. URANIE and the various NWCP software components need to be able to communicate through common file formats otherwise it will impossible to apply the URANIE	WP4	Amend the dwarfs or models that have to be run in URANIE so that they are able to deal with plain ASCII files in input/output, which URANIE handles perfectly well.

methodology on those components (low).		
New workflow management not implemented in URANIE. The new workflow will allow an efficient run of the NWCP components on a significant number of nodes, and not having it might compromise those runs (medium).	WP4	Review the current URANIE workflow management to optimize it and gain in performance.
Failing to run the ALADIN ensemble benchmark in URANIE. The ALADIN ensemble HPCW benchmark is significantly more involved than the Lorentz model analysed in task 4.2. URANIE might not be able to deal with such a complex system, even with the reduction techniques developed in Task 4.5 (medium).	WP4	1. Enhance the dimension reduction techniques to better fit the ALADIN complexity. 2. Reduce the parameter space dimensionality by agreeing on not varying certain parameters (expert decision).
Personnel involved or recruited not able to fulfil tasks (low).	All WP	Monitoring by the PO, and implementing adjustments within each organisation through GA.
Partner leaving the consortium (low).	All WP	Some of the competencies are partially overlapping, introduction of new partner.
Underperforming partners (medium).	WP6	Close contact between WP Leaders and Project Leader, short feedback loops and personal contact.

Innovation Management: Effective innovation management within this project will require an overview of the project in its entirety and for this reason the PC will be responsible for the process of innovation management. The management of innovation activities is integral to ESCAPE, as already indicated in Section 2.1. By nature of the structure of the consortium and the work packages within the project, the technical, scientific and application aspects of innovation are addressed and combined. There is expertise regarding all these aspects represented in the consortium. Through the PC and within the management structure already identified above these elements will be brought together and will ensure that innovative products are developed and exploited as a result of the project.

The PC, with support from the SC will produce an overview of the WPs and the outcomes of ESCAPE-2 to provide effective management and therefore exploitation of these both during and after the project.

At each stage of the project, the tasks associated with innovation management will be slightly different. The interfaces between elements of the value-chain (Section 1.3) form critical stages of innovation management because they define the hand-over between hardware and software as well as between different software levels**.**

- The innovation management will form an integral part of the Dissemination and Exploitation Plan (**D5.1**) which will be defined at an early stage in the project.
- At a more mature stage of the project, this plan will be adapted to take into account best practice methods of maximising 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 for the intellectual property exploitation after the project.

WP1, **WP2**, **WP3** and **WP4** include feedback options to other work packages. To ensure that the project responds to feedback in the most efficient way, the timings of the deliverables and milestones of the work packages have been planned to allow sufficient time to incorporate feedback and development. In this way the project will be responsive to any opportunities that may arise. The EB will also ensure that any opportunities are addressed and incorporated if applicable.

Quality Management: The quality control management involves the product description and quality expectations of key deliverables, and a review and acceptance procedure. In case of software, the quality control process also comprises output from code testing and validation of results. The procedures for the quality control process of deliverables will be described in the Quality Manual (**D5.1**) taking into account Technology Readiness Levels (TRL; Section 1.3). When a deliverable is ready for review, it will be forwarded to the PO, which will verify its general compliance before forwarding it to the EB. The EB can request revision of the deliverable before approval.

Intellectual Property Management: The CA that will be produced and signed after acceptance of the project will define and complete any points not covered by the EC Grant Agreement. It will formalize project management procedures but will also consider Intellectual Property Rights (IPR).

The innovation management will manage the strategy and the concrete actions of the consortium for the protection, exploitation and dissemination (Section 2.1) of the results of the project. This will include software licensing issues.

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. The innovation management will be in charge of tracking and proposing commercialisation and exploitation strategies (including software licensing issues) to the whole consortium. When needed, he/she will give advice to the parties concerned about the ownership, access rights, legal and commercial implications, patents, publications, copyrights, etc.

Administrative management: To ensure the transparency of the project management, all reports will be made available on the ESCAPE-2 website. Only documents or part of documents concerning individuals will be kept confidential. Two main reporting types are identified:

- *Reporting to the EC*: contractually, regular management and financial reports have to be submitted to the EC. The PS will draft the reports in due time and submit them to the PC for approval.
- *Internal Reporting*: in order to control the project, each partner will submit a monthly performance sheet to the PS. The PO will publish a monthly dashboard displaying the progress of the project and a management report will be presented during each EB meeting and annexed to the minutes.

Financial management: The budgets and advance payments will be distributed according to the provisions of the CA and the decisions of the GA. Budget forecasts will be established as an input for EB meetings. The information on participant expenses will be gathered following the reporting procedure and information from the management report.

3.3 Consortium as a whole

3.3.1 Partners

The ESCAPE-2 consortium (Figure 13) comprises one international European organisation, three national meteorological centres, two university, two HPC centres and one hardware vendors. This composition ensures that the relevant expertise is well covered by the consortium:

- the scientific and technical expertise along the value chain of novel accelerator hardware, programming environment, weather and climate application software, and performance/energy efficiency diagnostics;
- the representation of the world-leading European global and key regional weather centres translating the benefits of ESCAPE-2 into substantial societal benefit.

This combination of technical and scientific excellence as well as experience will support early and sustained progress towards the primary objective of redefining a leading application towards energy efficiency required for the exascale. This objective cannot be fulfilled with a less experienced team. Each partner brings a particular expertise needed for the successful development and deployment of the project's outcome. Several partners have already worked together and there is substantial experience working in an international and multi-disciplinary context thus ensuring strong synergies within the team and an efficient implementation of actions.

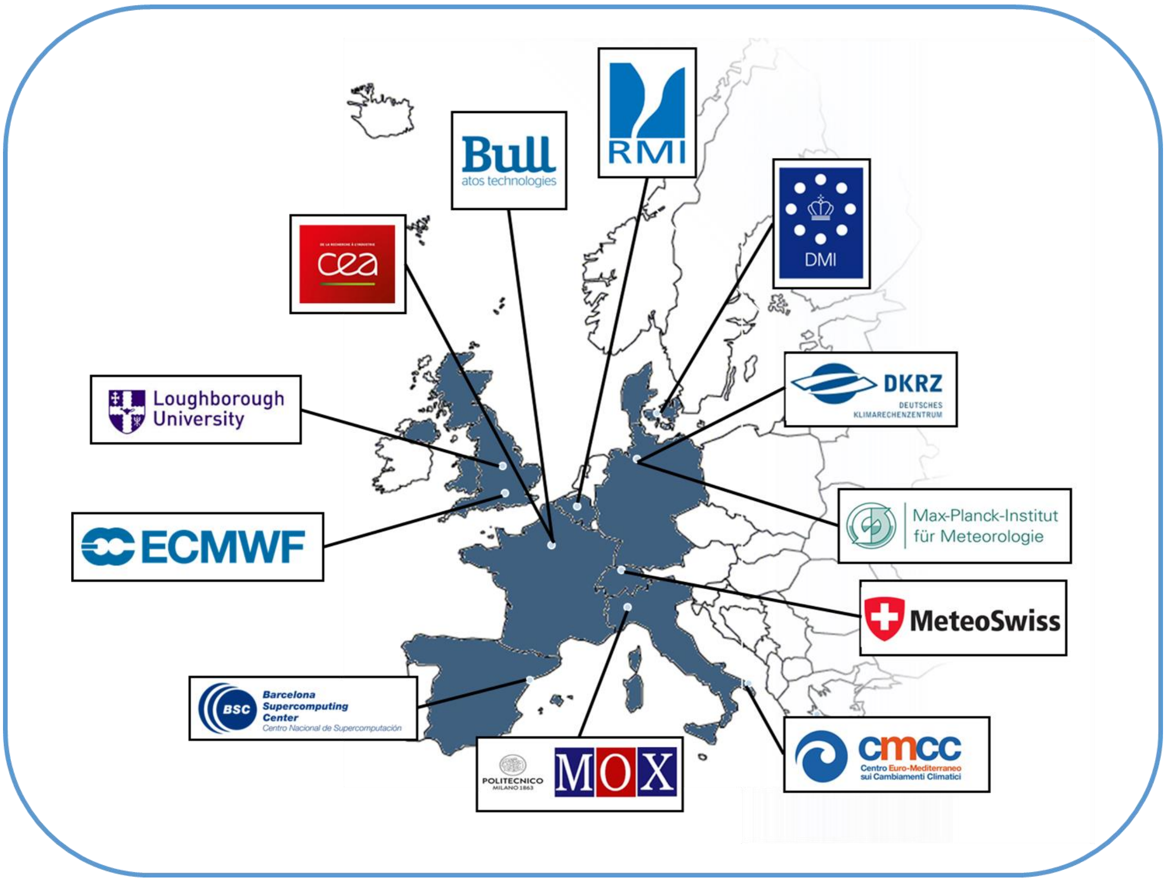

Figure 13: ESCAPE-2 partnership.

In particular:

- The consortium is coordinated by the world-leading global weather prediction centre, which was founded by its member states to focus European excellence in a single organisation. Moreover, ECMWF not least through its participation in CRESTA[34], has a proven record in co-design with HPC vendors. ESCAPE-2 provides a crucial ingredient for extending this excellence to the next Earth-system model generation (including the European community climate model EC-Earth) enabled by exascale computing. ECMWF will also be the operator of the Copernicus Atmospheric Monitoring and Climate Change services, which will directly benefit from ESCAPE-2.
- A number of regional European numerical weather prediction consortia, are represented in the consortium through participation of RMI, DMI, and MSWISS. Therefore the transfer of computational performance from global to regional level is built in.
- The tight collaboration between MPIM and DKRZ for developing and running the ICON weather and climate prediction models, and their business relationship with BULL, ensure low-risk R2O transfers. DKRZ and ECMWF co-coordinate the ESiWACE centre of excellence so that the dissemination of ESCAPE-2 results to the wider community is built in.
- MSWISS has obtained substantial experience in radical code refactoring in the course of the adaptation of COSMO to hybrid architectures in collaboration with industry and HPC centres. LU and POLIMI contribute academic excellence in the areas of code adaptation, energy efficiency diagnostics and numerical methods.

---

[34] https://www.cresta-project.eu/

- BULL have substantial experience in the provision of heterogeneous architectures, adaptation of hardware and software solutions for the weather and climate modelling communities, and in maintaining close collaboration with research in support of future hardware generations.
- Two of the consortium members (BULL, ECMWF) are full members of ETP4HPC ensuring the realization of the Scientific Research Agenda in this application area.
- Project management, dissemination of results and communication are centralized at ECMWF that has a long-term record in these areas and will ensure a highly professional level is reached.

Figure 13 illustrates the geographical distribution of partners and Figure 14 shows the distribution of expertise between partners as required by the ESCAPE-2 project. This indicates that the consortium provides an excellent complementarity of expertise for achieving the objectives of this project.

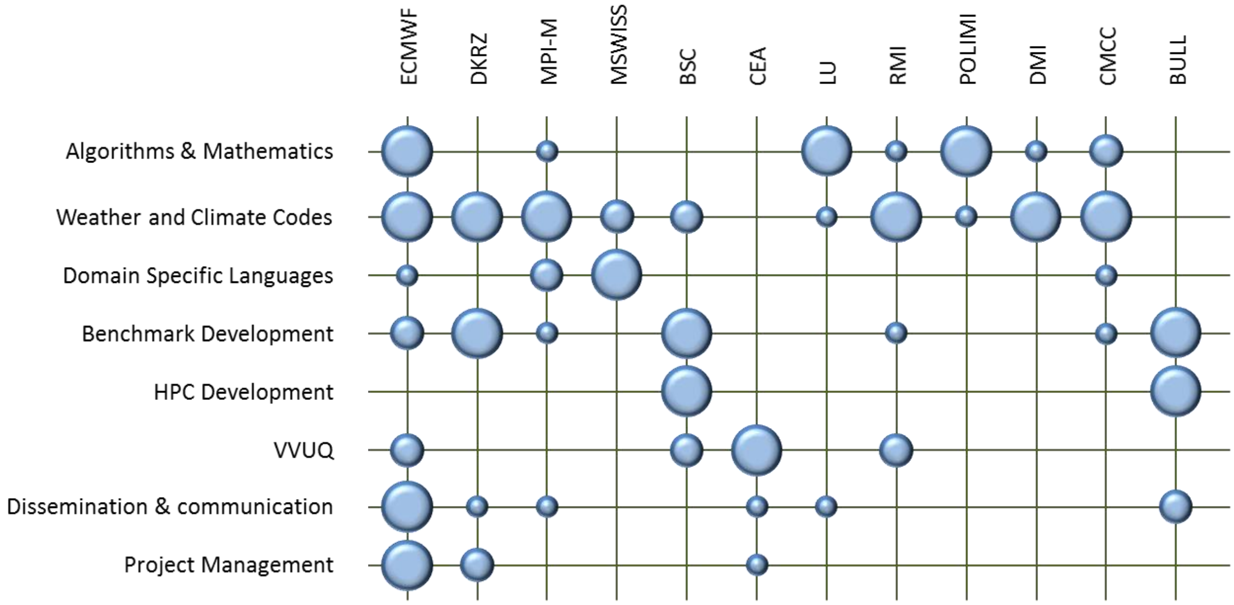

Figure 14: ESCAPE-2 Partner competency map.

3.3.2 External support

The project will be further strengthened through the support of a strong community of climate and weather related institutions, HPC centres and WMO, called “supporters”. The supporters are legal national and international entities, which are not beneficiaries but wish to support the project through complementary activities. These will be addressed in WP5 through dissemination and communication activities, and the supporters will participate in workshops and reviews.

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 Executive Board. A list of our supporters is given in Table 14 and letters of support and commitment dedicated to ESCAPE-2 can be found in Section 6.

Table 14: List of external ESCAPE-2 supporters.

Institute	Representative
Europe:	
Swiss National Supercomputing Centre	Prof. *Thomas Schulthess*, Director
Eidgenössische Technische Hochschule (ETH), Switzerland	Prof Dr. *Christoph Schär*, Institute for Atmospheric and Climate Science
Deutscher Wetterdienst (DWD), Germany	Prof. Dr. *Sarah Jones*, Director Research
SMHI, Sweden/ EC-EARTH	Dr *Ralf Doescher*, Chair of EC-EARTH Steering Committee
Niels-Bohr Institute, U Copenhagen, Denmark	Prof. Dr. *Eigil Kaas*

Freie Universitaet Berlin, Germany	Prof. Dr.-Ing. *Rupert Klein*, Institute for Mathematics
Alfred Wegener Institute, Germany	Prof. Dr. *Thomas Jung*, Head of Climate Dynamics
University Oxford, UK	Prof. Dr. *Tim Palmer*, Royal Society Research Professor in Climate Physics
IPSL, France	Dr. *Sylvie Joussaume*, Chair of ENES Board
University of Helsinki, Finland	Prof. *Heikki Jaervinen*, Professor of Meteorology
Netherlands eScience Center/ Wageningen University	Prof. Dr. *Wilco Hazeleger*, Chair Climate Dynamics
United States:	
Naval Postgraduate School, Department of Applied Mathematics	Prof. *Frank Giraldo*, Professor of Applied Mathematics
Argonne National Laboratory	Dr. *Paul Messina*, Director of DOE's Exascale Computing Project
University of Tennessee	Prof. *Jack Dongarra*, University Distinguished Professor of Computer Science
Princeton University, USA	Prof. *V. Balaji*, Head of Modelling Systems Group, Cooperative Institute for Climate Science
National Center for Atmospheric Research, USA	Dr. *Anke Kamrath*, Director Computational and Information Systems Laboratory
University Corporation of Atmospheric Research (UCAR), USA	Dr. *John Michalakes*, UCAR/ CPAESS visiting scientist, NRL Marine Meteorology Division
International:	
World Meteorological Organization, Switzerland	Dr. *Paolo Ruti*, Chief World Weather Research Program
Asia:	
Global Scientific Information and Computing Center / Dept. Mathematical and Computing Sciences, Tokyo Institute of Technology, Japan	Dr. *Satoshi Matsuoka*, Professor / Division Director of HPC Infrastructure

3.4 Resources to be committed

The overall budget is outlined in the table below. In total, 403.8 PM will be deployed during the 3 years project time, averaging 11.22 persons per project month. The budget of 3,999,650.00€ (Table 15) requested EC contribution is within the budget expectations indicated by the European Commission in the call text, and due to the complexity of the project as well as the uniqueness of its offerings requiring the involvement of a variety of partners, this is considered appropriate by the Consortium.

Table 15: ESCAPE-2 budget

			(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)	(I)	(J)	(K)
	Participant	Country	Direct Personnel costs, €	Other Direct costs, €	Direct costs of subcontracting, €	Direct costs of providing financial support to third parties, €	Costs of inkind contributions not used on the beneficiary's premises, €	Indirect costs, € (= 0.25*(A+B-E))	Special unit costs covering direct & indirect costs	Total estimated eligible costs (= A+B+C+D+F+G)	Reimbursement rate	Max grant , € (=H*I)	Requested grant, €
P1	ECMWF	UK/Int	552,039.00	62,000.00	-	-	-	153,510.00	-	7,677.00	100%	767,548.75	767,548.75
P2	DKRZ	DE	189,000.00	15,000.00	-	-	-	51,000.00	-	255,000.00	100%	255,000.00	255,000.00
P3	MPI-M	DE	202,500.00	10,000.00	-	-	-	53,125.00	-	265,625.00	100%	265,625.00	265,625.00
P4	MSWISS	CH	334,800.00	17,000.00	-	-	-	87,950.00	-	439,750.00	100%	439,750.00	439,750.00
P5	BSC	ES	162,000.00	24,000.00	-	-	-	46,500.00	-	232,500.00	100%	232,500.00	232,500.00
P6	CEA	FR	264,024.00	21,174.00	-	-	-	71,299.50	-	356,497.50	100%	356,497.50	356,497.50
P7	LU	UK	267,200.00	20,000.00	-	-	-	71,800.00	-	359,000.00	100%	359,000.00	359,000.00
P8	RMI	BE	299,988.00	10,000.00	-	-	-	77,497.00	-	387,485.00	100%	387,485.00	387,485.00
P9	POLIMI	IT	180,000.00	25,000.00	-	-	-	51,250.00	-	256,250.00	100%	256,250.00	256,250.00
P10	DMI	DK	215,802.00	4,026.00	32,215.00	-	-	54,957.00	-	307,000.00	100%	307,000.00	307,000.00
P11	CMCC	IT	87,500.00	16,500.00	-	-	-	26,000.00	-	130,000.00	100%	130,000.00	130,000.00
P12	BULL	FR	173,040.00	21,355.00	-	-	-	48,598.75	-	242,993.75	100%	242,993.75	242,993.75
Total			**2,927,893.00**	**246,055.00**	**32,215.00**	**-**	**-**	**793,487.25**	**-**	**3,239,778.25**		**3,999,650.00**	**3,999,650.00**

Table 16 and Table 17 outline the efforts allocated to each partner and WP as well as the breakdown of "Other Direct Cost" for Partners whose Other Direct Cost exceed 15% of the Personnel Cost.

Table 16: Summary of staff effort.

	WP1	WP2	WP3	WP4	WP5	WP6	Total Person-Months per Participant
1 - ECMWF	10.0	10.0	5.0	5.0	**10.0**	**15.0**	**55**
2 - DKRZ	1.0	11.0	**15.0**	0.0	0.0	1.0	**28**
3 - MPIM	1.0	18.0	8.0	0.0	0.0	1.0	**28**
4 - MSWISS	0.0	**33.0**	2.0	0.0	0.0	1.0	**36**
5 - BSC	0.0	2.0	6.0	27.0	1.0	0.0	**36**
6 - CEA	0.0	0.0	0.0	**34.0**	1.0	1.0	**36**
7 - LU	38.0	0.0	0.0	0.0	2.0	0.0	**40**
8 - RMI	0.0	3.0	13.0	20.0	0.0	0.0	**36**
9 - POLIMI	**32.0**	0.0	1.0	0.0	2.0	1.0	**36**
10 - DMI	36.3	0.0	0.0	0.0	0.0	0.0	**36.3**
11 - CMCC	1.0	7.0	4.5	0.0	0.0	0.0	**12.5**
12 - BULL	0.0	0.0	22.0	0.0	1.0	1.0	**24**
Total Person Months	**119.3**	**84**	**76.5**	**86**	**17**	**21**	**403.8**

Table 17: 'Other direct cost' items (travel, equipment, infrastructure, goods and services, large research infrastructure).

10 CMCC	Cost (€)	Justification
Travel	16,500.00	Participation in meetings between project partners, participation to GA, participation to dissemination events
Equipment	0.00	
Other goods and services	0.00	
Total	16,500.00	

In terms of effort, Figure 15 below outlines the distribution between partners in terms of percentages. Research Partners make up 46% of the effort which shows the research nature of the ESCAPE-2 proposal. Nevertheless, industry represent 10% of the effort, while Operational Centres with research activities are responsible for 44%. Finally, the final chart shows that the bulk of efforts is spent on research activities (91%), with 9% focused on project management and dissemination/ exploitation.

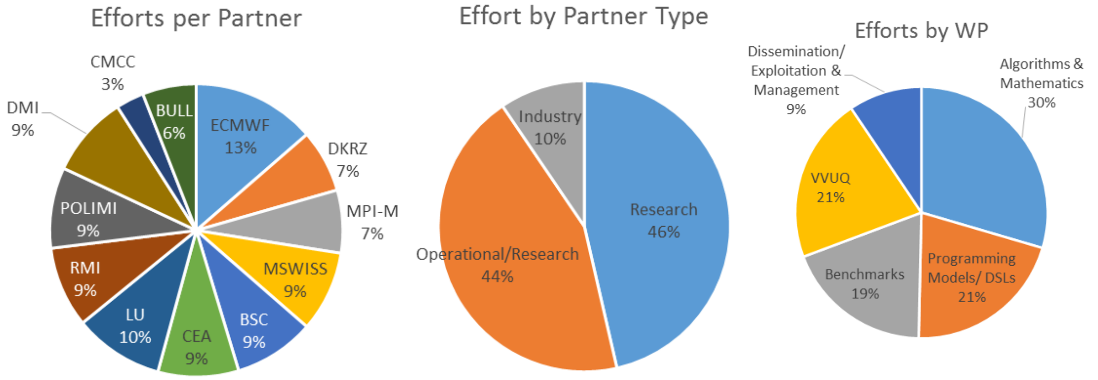

Figure 15: Effort distribution by partner, domain and work topic.

References

Baker, A. H., D. M. Hammerling, M. N. Levy, H. Xu, J. M. Dennis, B. E. Eaton, J. Edwards, C. Hannay, S. A. Mickelson, R. B. Neale, D. Nychka, J. Shollenberger, J. Tribbia, M. Vertenstein, and D. Williamson, “A new ensemble-based consistency test for the Community Earth System Model (pyCECT v1.0)”. Geosci. Model Dev., 8, 2829–2840, doi:10.5194/gmd-8-2829-2015, 2015.

Bénard, P., Jozef Vivoda, Ján Mašek, Petra Smolíková, Karim Yessad, Ch Smith, Radmila Brožková, and J-F. Geleyn, "Dynamical kernel of the Aladin–NH spectral limited-area model: Revised formulation and sensitivity experiments". Quarterly Journal of the Royal Meteorological Society 136, no. 646 (2010): 155-169, 2016.

Bauer, P., A. Thorpe, and G. Brunet, “The quiet revolution of numerical weather prediction”. Nature 525, 47–55 (03 September 2015) doi:10.1038/nature14956, 2015.

Bony, S., B. Stevens, D. M. Frierson, C. Jakob, M. Kageyama, R. Pincus, T. G. Sheperd, S. C. Sherwood, A. P. Siebesma, A. H. Sobel, M. Watanabe, and M. J. Webb, “Clouds, circulation and climate sensitivity”. Nature Geoscience, doi:10.1038/ngeo2398, 8(4), 261-268, 2015.

Carman, J., T. Clune, F. Giraldo, M. Govett, B. Gross, A. Kamrath, T. Lee, D. McCarren, J. Michalakes, S. Sandgathe, T. Whitcomb “Position paper on high performance computing needs in Earth system prediction”. National Earth System Prediction Capability. https://doi.org/10.7289/V5862DH3, 2017.

Deconinck, W., P. Bauer, M. Diamantakis, M. Hamrud, C. Kühnlein, P. Maciel, G. Mengaldo, T. Quintino, B. Raoult, P.K. Smolarkiewicz, Nils P.Wedi, “Atlas: A library for numerical weather prediction and climate modelling”. Computer Physics Communications, 2017, avail. online, 2017.

Fan, Y., W. Li, C. K. Gatebe, C. Jamet, G. Zibordi, T. Schroeder, and K. Stamnes, “Atmospheric correction and aerosol retrieval over coastal waters using multilayer neural networks”. Remote Sens. Environ., doi:10.1016/j.rse.2017.07.016, 199, 218-240, 2017.

Gysi, T., C. Osuna, O. Fuhrer, M. Bianco and T. C. Schulthess, "STELLA: a domain-specific tool for structured grid methods in weather and climate models". SC15: International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, pp. 1-12. doi: 10.1145/2807591.2807627, 2015.

Johnson, J., Z. Cui, L. A. Lee, J. P. Gosling, A. M. Blyth, and K. S. Carslaw, “Evaluating uncertainty in convective cloud microphysics using statistical emulation”. J. Adv. Model. Earth Syst., doi:10.1002/2014MS000383, 7, 162–187, 2015.

Milroy, D.J., Baker, A.H., Hammerling, D.M., Jessup, E.R.: “Nine time steps: ultra-fast statistical consistency testing of the Community Earth System Model (pyCECT v3.0)”. Geoscientific Model Development Discussions. In review. (DOI: 10.5194/gmd-2017-49), 2017.

Müller, A., M.A. Kopera, S. Marras, L.C. Wilcox, T. Isaac, F.X. Giraldo, “Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA”. arXiv:1511.01561 [cs.DC], 2016.

National Research Council. “A National Strategy for Advancing Climate Modeling”. Washington, DC, The National Academies Press. https://doi.org/10.17226/13430, 2012.

Sanan, P., Schnepp, S. M., & May, D. A., “Pipelined, Flexible Krylov Subspace Methods”. SIAM Journal on Scientific Computing, 38(5), C441-C470, 2016.

Schmidt, G. A., D. Bader, L. J. Donner, G. S. Elsaesser, J. C. Golaz, C. Hannay, A. Molod, R. B. Neale, and S. Saha, “Practice and philosophy of climate model tuning across six US modeling centers”. Geoscientific Model Development, doi:10.5194/gmd-10-3207-2017, 10(9), 3207-3223, 2017.

Seity, Y., P. Brousseau, S.Malardel, G. Hello, P. Benard, F. Bouttier, C. Lac, and V. Masson, “The AROME-France Convective-Scale Operational Model”. Monthly Weather Review, 139, 976-991. doi: http://dx.doi.org/10.1175/2010MWR3425.1, 2011.

Schneider, T., J. Teixeira, C. S. Bretherton, F. Brient, K. G. Pressel, C. Schär and A. P. Siebesma “Climate goals and computing the future of clouds”. Nature Climate Change, 7, 3–5, doi:10.1038/nclimate3190, 2017.

Stamnes, K., G. E. Thomas, and J. Stamnes, “Radiative Transfer in the Atmosphere and Ocean”, 2nd Edition, Cambridge University Press, Cambridge, UK, 2017.

Thornes, T., P. Düben, and T.N. Palmer, T., “On the use of scale-dependent precision in Earth System modelling”. Q.J.R. Meteorol. Soc, 143: 897–908. doi:10.1002/qj.2974, 2017.

Vana, F., P Duben, S Lang, T Palmer, M Leutbecher, D Salmond, G Carver, “Single Precision in Weather Forecasting Models: An Evaluation with the IFS”. MONTHLY WEATHER REVIEW 145, p495-502, 2017.

Zängl, G., D. Reinert, P. Ripodas, M. Baldauf, “The ICON (ICOsahedral Non-hydrostatic) dynamical core”, Quarterly Journal of the Royal Meteorological Society, Volume 141, Issue 687 January 2015 Part B Pages 563–579, 2015.

Table of Contents

4 Members of the consortium

4.1 Participants

01 - European Centre for Medium-range Weather Forecasts (ECMWF)

About the Organisation

The European Centre for Medium-Range Weather Forecasts (ECMWF) is an international organization supported by 34 States: 22 Members (Austria, Belgium, Croatia, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Luxembourg, the Netherlands, Norway, Portugal, Serbia, Slovenia, Spain, Sweden, Switzerland, Turkey and the United Kingdom) and 12 Co-operating Members (Bulgaria, Czech Republic, Estonia, the former Yugoslav Republic of Macedonia, Hungary, Israel, Latvia, Lithuania, Montenegro, Morocco, Romania and Slovakia). ECMWF's principal objectives are the preparation, on a regular basis, of medium-range and long-range weather forecasts for distribution to the meteorological services of the Member States, the development of scientific and technical research directed to the improvement of these forecasts, and the collection and storage of appropriate meteorological data ECMWF's strategy includes the principal goals to provide reliable forecasts of severe weather across the medium-range and high-quality, near-surface forecast products focusing on areas such as precipitation, wind and temperature. ECMWF's computer facility includes supercomputers, archiving systems and networks.

In 2013, ECMWF initiated a Programme on Scalability that aims at developing the next-generation forecasting system addressing the challenges of future exascale high-performance computing and data management architectures. The Programme is required to optimize system performance allowing ECMWF to fulfil its strategy within the expected resource constraints.

Role in this Project

ECMWF will coordinate the ESCAPE-2 project, both in terms of administration and scientifically, and lead WPs 5 and 6. ECMWF will contribute to the development of dwarfs (WP1), DSL specification (WP2), provide KRONOS (WP3), and support the VVUQ aspects (WP4).

Expertise with Relevance to this Project and Role

ECMWF has extensive competence in operating complex global forecasting suites on high-performance computers and in transitioning top-level science from research to operations exploiting innovative approaches in computing science to fulfil the tight runtime and delivery constraints required by Member States. ECMWF coordinates the ESCAPE project, upon which this proposal builds. In addition, ECMWF has coordinated a number of other EC-funded projects.

Key Persons

Dr Peter Bauer (Male) - joined ECWMF in January 2000. Before becoming Deputy Director of Research he headed the Model Division that comprises the physical and numerical aspects of numerical weather prediction and the Satellite Section at ECMWF. Before joining ECMWF, he was leading a DLR research team on satellite meteorology in Cologne, Germany. His background covers physical modelling, data assimilation and satellite remote sensing. During his career, he was awarded research fellowships by NOAA and NASA, and a science award by DLR. He is the author and co-author of 100 peer-reviewed scientific journal papers, and his publications have an h-index of 41. He is a member of several scientific advisory committees at the international level (WMO, ESA, EUMETSAT) and has extensive experience with managing international research projects. At ECMWF, his current duties also include the management of the transition of new model cycles from research to operations and he is the manager of the recently launched Scalability Programme. He coordinates the ESCAPE project and co-coordinates the ESiWACE centre of excellence.

Dr Nils Wedi (Male) - joined ECMWF in 1995. He received his PhD degree from the Ludwig-Maximilians-Universität München. His career at ECMWF encapsulates a diverse range of work both technical and scientific. He has been a Principal Scientist since 2010 and is the Head of Earth System Modelling at ECMWF that addresses all aspects of scientific and computational performance relating to ECMWF's forecast model and the ensemble forecasting system. He develops strategies to secure the scalability of the model on future high-performance computing systems. He is the scientific coordinator of the European H2020 project ESCAPE to address the challenges of rising energy cost for computing towards affordable, exascale high performance simulations of weather and climate, and he is a member of the WMO working group on numerical experimentation (WGNE).

Dr Wedi will involve and lead members of this section at ECMWF, supporting and driving forward the development of scalable and energy-efficient mathematical algorithms for Earth system modelling.

Dr Daniel Thiemert (Male) - joined ECMWF in 2015 as Project Manager in the Research Department. Before joining ECMWF, he was working for the University of Reading as Senior Researcher and Project Manager, being responsible for the management and coordinator of large scale European and national projects in the area of intelligent systems. He obtained his masters in computer science from the Anhalt University of Applied Sciences, Germany, and his PhD in Computer Science from the University of Reading, UK. At ECMWF he is responsible for the project management of the EU-funded projects ESCAPE and CHE, as well as for managing externally funded project acquisition.

Relevant Publications (up to 5)

1. Energy-efficient SCalable Algorithms for weather Prediction at Exascale, Bauer, Peter; Wedi, Nils; Baldauf, Michael; Benard, Pierre; Fuhrer, Oliver; Kulczewski, Michal; McKinstry, Alastair; Messmer, Peter; New, Nick; HansenSass, Bent; Szmelter, Joanna; Termonia, Piet; Vigouroux, Xavier, Source: Impact, Volume 2017, Number 1, January 2017, pp. 69-71(3).
2. Atlas: A library for numerical weather prediction and climate modelling, W. Deconinck, P. Bauer, M. Diamantakis, M. Hamrud, C. Kühnlein, P. Maciel, G. Mengaldo, T. Quintino, B. Raoult, P.K. Smolarkiewicz, Nils P.Wedi, Computer Physics Communications, 2017, avail. online.
3. Representing the Earth surfaces in the Integrated Forecasting System: Recent advances and future challenges, G Balsamo, A Agustì-Panareda, C Albergel, A Beljaars, S Boussetta, E Dutra, T Komori, S Lang, J Muñoz-Sabater, F Pappenberger, P de Rosnay, I Sandu, N Wedi, A Weisheimer, F Wetterhall, E Zsoter, ECMWF Research Department Technical Memorandum 729, ECMWF Reading, UK, 2014.
4. Bauer, Peter, Alan Thorpe, and Gilbert Brunet. "The quiet revolution of numerical weather prediction." Nature 525.7567 (2015): 47-55.

Recent Projects (up to 5)

- Collaborative Research into Exascale Systemware, Tools & Applications (CRESTA). FP7-287703, 2011-2014.
- PantaRhei. FP7/2012/ERC Advanced Grant agreement no. 320375, 2013-2018.
- Energy-efficient SCalable Algorithms for weather Prediction at Exascale (ESCAPE). H2020-671627, 2015-2018
- Excellence in SImulation of Weather and Climate in Europe (ESIWACE), H2020=675191, 2015-2019
- Co-designed Innovation and System for Resilient Exascale Computing in Europe: From Applications to Silicon (EuroExa), H2020-754337, 2017-2021

Major Hardware/ Infrastructure available

ECMWF maintains a multi-petaflops supercomputer facility which is designed for operational resiliency featuring two Cray XC40 systems and independent Cray Sonexion storage systems. Each subsystem consists of 20 Cray XC40 cabinets equipped with Intel Broadwell processors and around 3,600 dual-socket compute nodes per system, a number of Cray Development and Login nodes and around 10 petabytes of Lustre storage with the ability to cross mount the Lustre file systems between the halls.

02 - Deutsches Klimarechenzentrum GmbH (DKRZ)

About the Organisation

DKRZ, the German Climate Computing Centre, is a national service provider which constitutes an outstanding research infrastructure for model-based simulations of global and regional climate and the investigation of the processes in the climate system. DKRZ's principal objectives are provision of adequate computer performance, data management, and service and support to use these tools efficiently. DKRZ operates one of the largest supercomputers in Germany and provides its more than 1000 scientific users with the technical infrastructure needed for the processing and analysis of huge amounts of data from climate simulations. This also includes training and support for related application software and data processing issues. DKRZ participates in many national and international projects aiming to improve the infrastructure for climate modeling. Through its research group on scientific computing DKRZ is linked to the Department of Informatics of the University of Hamburg. DKRZ is a non-profit and non-commercial limited company with four shareholders. MPG (Partner 3) holds 55% of the shares of DKRZ (see http://www.dkrz.de/about-en/Organisation/gesellschafter for more references). The dependency relationship has been declared in the Part 2 – Administrative data of participating organisation of this application form Contribution to the specific project

Role in this Project

WP1:

- participate in the extraction of ICON related climate dwarfs

WP2:

- participate in the definition of requirements and design of a high level DSL
- support of DSL development
- participate in the application of DSL to ICON related dwarfs
- participate in the evaluation of DSL usability and testing of adapted dwarfs

WP3:

- DKRZ is workpackage leader of WP3
- contribute to the definition of ICON based benchmarks
- participate in the re-integration of DSL versions of ICON based dwarfs into benchmarks
- DKRZ is responsible for definition, assembly and dissemination of HPCW benchmark suite.

Expertise with Relevance to this Project and Role

DKRZ as a leading computer centre has a long history and experience in benchmarking and evaluating tier-1 super computers based on the performance of real climate applications, which will feed into WP3 (benchmarking)

DKRZ supports scientists in using a broad spectrum of climate codes and thus has leading expertise in analysing performance and structure of climate models. In addition to knowledge of HPC programming in general and ICON development in particular, DKRZ staff has experiences with rapid prototyping of recursive descent parsers using high level scripting languages (perl, pyhton) and beyond prototyping with C++. This knowledge fits to the role of participating in defining/designing (T2.1) and implementing (T2.2) a DSL prototype front-end parser and to the role of supporting the modular DSL toolchain development (T2.3). DKRZ's experience on extracting the ICON communication kernel from the full ICON model into a separate dwarf and extraction of the radiation kernel from ECHAM6 can be applied to the tasks of extraction (T1.5) and re-integration (T3.3) of ICON dwarfs

Key Persons

Dr. Joachim Biercamp (**Male**) holds a PhD in Physical Oceanography. He is leading the Application department of DKRZ. As part of his responsibilities, he coordinated the procurement and benchmarking of the last three generations of DKRZ super computers all ranked within the TOP 30 of the ToP500 list: a NEC SX6, an IBM Power6 system and the current bullx system, which has a peak performance of 3.6 PetaFLOPS and a file system of world leading 55 PetaBYTES. Joachim is involved in several national and international projects dealing with infrastructure for climate modeling. In particular he is the coordinator of H2020 Center of Excellence ESiWACE and is a member of the steering committee of the German project HD(CP)^2 aiming at development and operation of a cloud resolving version of the ICON model which is used for both, climate research and numerical weather prediction.

Jörg Behrens (**Male**) received his Diploma in Physics from the University of Göttingen, Germany and now works in the Application department of DKRZ mainly on performance optimization of climate models and has 20 years of experience with developing scalable parallel Fortran applications in material and climate science. He is one of the authors of the YAXT library, which simplifies efficient MPI communication. He also has defined and extracted the ICON communication kernel from the full ICON model into a separate dwarf (FP7 project IS-ENES-II) and has supervised the definition and extraction of the radiation kernel from ECHAM6 (BMBF project PALMOD).

Irina Fast (**Female**) graduated in Meteorology at the Freie Universität Berlin, Germany. She gained expertise in different fields of applied Earth System modelling, for example implementation of reliable workflows for performance of high-resolution long-term climate simulations and management of extensive output data as well as processing, scientific analysis and visualisation of the model data. As a member of the help desk staff she supports users by porting ESMs to HPC systems operated at DKRZ taking into account computational performance aspects. She was involved into the selection and adaptation process of the real application benchmarks for the last HPC system procurement at DKRZ and the activities on benchmarking coupled climate models within the IS-ENES2 project.

Mathis Rosenhauer (**Male**) received his PhD from the Faculty of Mechanical Engineering of RWTH Aachen University. He has 20 years of experience with highly parallel simulation codes for fluid mechanics and climate science. At DKRZ he supports the efficient use of computing and storage resources. Since three generations of DKRZ super computers, he is involved in design and execution of the benchmarking process.

Relevant Publications (up to 5)

1. Crossing the Chasm: How to develop weather and climate models for next generation computers?
   Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S., Osuna, C., Porter, A., Serradell, K., Valcke, S., Wedi, N., and Wilson, S.: , Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-186, in review, 2017.
2. CPMIP: measurements of real computational performance of Earth system models in CMIP6
   Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Caubel, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., and Wright, G.: , Geosci. Model Dev., 10, 19-34, https://doi.org/10.5194/gmd-10-19-2017, 2017
3. Report on the suite of base benchmarks and analysis of performance on available platforms
   G. Aloisio, C. Basu, J. Behrens, J Biercamp, A. Caubel, I. Fast, S. Fiore, M.-A. Foujols, P.G. Fogli, M. Hanke, S. Masina, S. Mocavero, P. Neumann, H. Struthers, S. Valcke
   Deliverable 10.4 of FP7 project IS-ENES2
   https://is.enes.org/documents/deliverables/is-enes2_d10-4_report-on-the-suite-of-base-benchmarks-and-analysis-of-performance-on-available-platforms/view

Recent Projects (up to 5)

- ESiWACE: Excellence in Simulation of Weather and Climate in Europe;
  Role of DKRZ: coordinator;
  H2020 project, Center of Excellence, GA 675191;
  www.esiwace.eu
- HD(CP)2: High definition clouds and precipitation for advancing climate prediction;
  Role of DKRZ: Work Package leader for "Modelling", Member of Steering Committee
  Funded by German Federal Ministry of Education and Research)
  www.hdcp2.eu
- IS-ENES2: Infrastructure for the European Network of Earth System Modelling
  Role of DKRZ: Work package leader "Performance benchmarks for coupled climate models"
  FP7 project, GA 312979
  is.enes.org

Major Hardware/ Infrastructure available

- bullx DLC 720, 1550 Intel Haswell nodes 24 cores/node, 2.5 GHz + 1750 Intel Broadwell nodes 36 cores/node 2.1 GHz
- 21 additional nodes (Broadwell/Haswell) with NVidia Kepler and Maxwell GPUs

03 - Max-Planck-Institut für Meteorologie (MPI-M)

About the Organisation

The Max Planck Society (MPG) is Germany's most successful research organization established in 1948. The 83 Max Planck Institutes and facilities (as of 2017) conduct basic research in the service of the general public in the natural sciences, life sciences, social sciences, and the humanities. Continuous renewal in the institutional structure preserves the scope the Max Planck Society needs to react quickly to pioneering scientific developments. The MPG contributes to the project with the institute for meteorology.

Max Planck Institute for Meteorology (MPI-M, http://www.mpimet.mpg.de) is dedicated to fundamental climate research. The overall mission of MPI-M is to understand how chemical, physical, and biological processes, as well as human behavior contribute to the dynamics of the Earth system, and specifically how they relate to global climate changes. The institute comprises three departments: The Atmosphere in the Earth System; The Land in the Earth System and The Ocean in the Earth System and hosts independent research groups focused on Fire in the Earth System, Forest Management in the Earth System, Sea Ice in the Earth System, Stratosphere and Climate as well as Turbulent Mixing Processes in the Earth System.

Role in this Project

Partner, Co-Lead of WP-2

WP1

- Extraction and providing ICON-OCE and PSRad dwarfs

WP2

- participation in the requirements capture and design of a high level DSL
- participation in the application of DSL to ICON related dwarfs
- participation in the evaluation of DSL usability and testing of adapted dwarfs
- demonstration of the application of the domain specific language toolchain to the above dwarfs

WP3

- contribution to the definition of ICON based benchmarks
- implementation of reference global configurations of the above dwarfs for the HPCW v0-benchmarks, integration, validation and benchmark of the novel DSL OCE and Rad dwarfs in the tier-1 benchmark suite
- contribution to the definition and assembly of the HPCW benchmark suite.

Expertise with Relevance to this Project and Role

MPI-M develops the ICON based MPI-ESM2 and uses it to understand and project climate and carbon cycle dynamics, in particular under anthropogenic forcings. MPI-ESM will be used in the sixth phase of the coupled model intercomparison project (CMIP6) in concerted action with many other world leading climate centres, including groups developing six other European ESMs (e.g., EC-Earth, IPSL-CM, UKESM) involved in projects like ESiWACE or IS-ENES. A strength of the MPI-M team in ESCAPE2 is the access to the knowledge from the joint efforts of MPI-Ms fundamental science departments land, ocean, and atmosphere on linking processes and feedbacks in the Earth System, supported by the MPI-M Scientific Computing Laboratory (SCLab). We will utilize in-house expertise on software infrastructure development (a group led by B. Stevens, including Luis Kornblueh and Leonidas Linardakis) and on ocean model development (a group led by P. Korn with in the ocean department of MPI-M, led by J. Marotzke).

Key Persons

Prof. Dr. Bjorn Stevens (Male) is a director at the Max-Planck-Institute for Meteorology where he leads the Atmosphere in the Earth System Department and is a professor (§17) at the University of Hamburg. Prior to moving to Hamburg Dr. Stevens was a full professor of Dynamic Meteorology at the University of California of Los Angeles. He received a PhD in Atmospheric Science in 1996 from the Colorado State University in Ft Collins CO, and holds a Bachelor and Masters of Science in Electrical Engineering from Iowa State University. Prof. Stevens' research blends modeling, theory and field work to help articulate the role of aerosols, clouds and atmospheric convection in the climate system. He has made pioneering contributions to both understanding and modelling of mixing and microphysical processes and their impact on the structure and organization of clouds. Likewise, his contribution to an understanding of how clouds respond to warming, and how radiative forcing responds to aerosol perturbations, has proven fundamental to the present comprehension of the susceptibility of Earth's climate to perturbations. Prof. Stevens served as a lead-author of Chapter 7, "Cloud and Aerosols" for the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. He is the lead principal investigator for the HD(CP)2 project, High Definition Clouds and Precipitation for Climate Prediction, a national project supported by the Germany Ministry of Education and Research. He serves on a number of international advisory boards, has served as editor of leading journals in his field and has been honored by a number of awards, and fellowships.

Reinhard Budich (Male) is an Oceanographer active in Climate Modelling and related infrastructures for over 25 years. He was the Technical Coordinator of the FP7 project IS-ENES1, and work package leader in ISENES2. He was running the COSMOS network as a project manager and Director of the FP5 PRISM project. Since 2001 to 2015 he has also been running the ENES office in Hamburg, Germany. He is also involved in the EUDAT2020 project. He is ICT-Representative of MPI for Meteorology and responsible for the HPC strategy of this institute as the leader of the group “Strategic IT partnerships” within MPI-M’s Scientific Computing Laboratory.

Peter Korn (Male) received his PhD in Mathematics from University Erlangen-Nürnberg, he leads the group "Applied Mathematics and Computational Physics" at MPI-M. The group performs research on the partial differential equations of the ocean with emphasis on mathematical, numerical and computational physics aspects. The work of the group is embedded in MPI-M's integrated project ICON. The group leader is coordinating the ocean model development within the department "The Ocean in the Earth System" (OES) and across several groups, thereby providing the ocean departments contribution to the development of the ICON-based MPI-ESM2 at MPI-M as well as to the total ICON modelling system.

Relevant Publications (up to 5)

1. P. Korn, Formulation of an unstructured grid model for global ocean dynamics, J. Comp. Phys. 339 (2017), 525-552
2. P. Korn, S. Danilov, Elementary dispersion analysis of some mimetic discretizations on triangular C-grids, J. Comp. Phys. 330 (2017), 156-172
3. P. Korn, Analysis of a Coupled Atmosphere-Ocean Data Assimilation Problem using Derivative-Based Cost Functionals, submitted
4. Bonaventura, L., Redler, R. & Budich, R. (2012). Earth system modelling 2: Algorithms, code infrastructure and optimization.Springer Springer Briefs in Earth System Sciences , doi:10.1007/978-3-642-23830-7
5. Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S., Osuna, C., Porter, A., Serradell, K., Valcke, S., Wedi, N., and Wilson, S.: Crossing the Chasm: How to develop weather and climate models for next generation computers?, Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-186, in review, 2017.

Recent Projects (up to 5)

- PRISM - EVR1-CT-2001-40012 – to share the development, maintenance and support of a comprehensive Earth System Modelling software environment
- ISENES1- 228203 –
    - The integration of the European climate and Earth system modelling community;
    - The development of Earth System Models for the understanding of climate change;
    - High-end simulations enabling better understanding and prediction of future climate change;
    - The application of Earth system model simulations to better predict and understand future climate change impacts.
- Continued in ISENES2 – 312979
    - Foster the integration of the European Climate and Earth system modeling community
    - Enhance the development of Earth System Models for the understanding of climate variability and change
    - Support high-end simulations enabling us to better understand and predict climate variations and change
    - Facilitate the application of Earth system model simulations to better predict and understand climate change impacts on society

- ESiWACE – 675191 – ESiWACE substantially improves 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 High Performance Computing environments

Major Hardware/ Infrastructure available

Easy access to computing and support resources of DKRZ

Large programming department for software infrastructure development tasks

To satisfy the extensive computational needs of the MPI-M, the institute built strategic partnership with the German Climate Computing Centre (DKRZ), the German collaborator of the Earth System Grid Federation (ESGF). MPG is also the major shareholder of the DKRZ.

The MPI-M has for a long time benefited from a close integration of its systems with those of DKRZ, offering the MPI-M seamless access to the DKRZ-managed machines such as the Bull machine "mistral", a supercomputer with a peak performance of 3.14 Petaflops consisting of approx. 3,000 compute nodes, 100,000 compute cores, 240 Terabytes of memory, and 54 Petabytes of disk. “Mistral” will also be used for CMIP6 simulations.

04 - Federal Institute of Meteorology and Climatology MeteoSwiss (MSwiss)

About the Organisation

The Federal Office of Meteorology and Climatology MeteoSwiss is by federal mandate the national provider for weather and climate services in Switzerland. In this role, it serves the general public, authorities, research and industry. MeteoSwiss monitors the atmosphere over Switzerland and operates the corresponding networks, it issues weather forecasts, warns the authorities and the general public of dangerous weather conditions and also monitors the Swiss climate. The legal duties include the provision of climate information and climatological services for the benefit of the general public. MeteoSwiss provides generic and tailor-made datasets and services for customers, and conducts research on themes from weather and climate to high-performance computing. Weather and climate in the Alpine region is one of its core competences. MeteoSwiss hosts the national GCOS office and is the official representative of Switzerland in various international organisations (WMO, ECMWF, EUMETSAT, EUMETNET etc.) and member of the Swiss Centre for Climate Systems Modelling (C2SM). In its research MeteoSwiss collaborates with academia (e.g. ETH Zurich and Swiss National Supercomputing Centre CSCS), with other governmental offices (e.g. hydrology) and the private sector (e.g. reinsurance). In the framework of the Swiss HP2C and PASC Initiatives, MeteoSwiss has led the adaption of the regional weather and climate model (COSMO) to hybrid high-performance computing systems and has spearheaded the application of domain-specific languages in operational atmospheric codes. On the basis of its numerical weather forecasting models, MeteoSwiss has issued weather and climate forecasts to commercial customers for over ten years now and has also a profound experience in the communication of such forecasts to the public and media.

Role in this Project

Partner, Co-lead of WP-2

MeteoSwiss will lead the development of the DSL frontend and toolchain.

Expertise with Relevance to this Project and Role

MeteoSwiss has significant expertise in the domain of developing DSLs for weather and climate. It has lead the development of the STELLA DSL (Gysi et al. 2015) which is running in the operational implementation of the regional numerical weather prediction model COSMO (Fuhrer et al. 2014; Lapillonne and Fuhrer, 2014). Since 3 years, and together with the Swiss National Supercomputing Centre (CSCS) and the computer science department of the ETH Zurich, MeteoSwiss is leading the development of the GridTools DSL, which will supersede and extend STELLA with new functionality. The GridTools DSL has also been the basis for the DSL work that MeteoSwiss has done in collaboration with ECMWF, where the DSL has been extended for unstructured grids used in global atmospheric models. Within the Swiss Platform of Advanced Scientific Computing (pasc.org) this development is still on-going. Further, MeteoSwiss has developed the CLAW compiler, a directives based source-to-source translator that can be applied to the physical parametrizations. Within two new projects that have recently been awarded funding (ENIAC and PASCHA), MeteoSwiss will continue this work on the GridTools and CLAW DSLs, which will directly benefit the work proposed in the ESCAPE-2 project.

Key Persons

Oliver Fuhrer (Male) – holds a PhD in Physics and has many years of experience in developing and deploying weather and climate models. Currently, he is leading the model development group at MeteoSwiss and a lecturer at ETH Zurich. His group is responsible for both improvements in model quality as well as the high-performance computing aspects of the model development. Over the past 7 years, he has attracted the funding for multiple research projects focusing on the adaption of weather and climate codes to emerging and future hardware architectures. A main thrust has been the development of novel programming models and DSLs. His team at MeteoSwiss has been awarded the Swiss ICT Award in 2016 for having implemented the first operational weather model fully running on a Cray CS Storm cluster with fat GPU nodes.

Carlos Osuna (Male) – has a PhD in Particle Physics which he obtained in 2009. Since then he has been working in the field of high-performance computing with large scientific codes. In 2011 he joined, as a Post-Doctoral at ETH, the HP2C project to develop a new GPU capable version of the numerical weather and climate model prediction COSMO, which significantly improve the performance and energy efficiency of the model used in production at MeteoSwiss and ETH institutions. Since then he has been involved in the development of novel domain specific languages (STELLA and GridTools) for numerical methods used in atmospheric sciences. Since 2017 he is co-PI of the PASC PASCHA project that aims at developing a universal toolchain for high-level DSLs that increase the productivity of model developments in heterogeneous architectures.

Stefan Moosbrugger (Male) – is a scientific software developer working at the Scalable Parallel Computing Laboratory (ETH Zurich) as one of the core developers of the PASCHA project. He is mainly involved in implementation of core features, optimization, software architecture and design. Before working at ETH Zurich he worked as a computer science researcher at the Distributed and Parallel systems group (University of Innsbruck) in the Insieme (C/C++) source-to-source compiler project. He holds a master's degree in computer science from the University of Innsbruck. Beside the working experience in complex projects he has strong knowledge in C++, high performance computing, compiler development, and software engineering.

Xavier Lapillonne (Male) – is a senior scientist in the model development group at the Federal Office of Meteorology and Climatology MeteoSwiss, Zurich since 2015. He is working in the field of high performance computing, in particular on hybrid system, and is currently the project leader of the Performance On Massively Parallel Architectures project for the Consortium for Small-scale Modeling (COSMO). He received a PhD in physics from EPF Lausanne in 2010, where he worked on massively parallel codes to simulate turbulence in hot plasmas. In 2010-2015 he was Post-doctoral researcher at ETH Zurich working on the development of a new GPU capable version of the numerical weather prediction code COSMO used at MeteoSwiss for weather prediction and at several universities for climate research. He was the lead developer for porting several parts of the model using OpenACC compiler directives, including physical parameterization and data assimilation. He has experience in developing and running scientific models on many supercomputers, including hybrid architectures (IBM blue gene, Bull NovaScale, Cray XT-4/5, XE6, XK7, XC30, CS-Storm)

Relevant Publications (up to 5)

1. Fuhrer, O., Chadha, T., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, Ch., Schulthess, T. C., Vogt, H., Höfler, T., Kwasniewski, G.: Near-global climate simulation at 1 km resolution on a GPU-accelerated supercomputer: establishing a performance baseline with COSMO 5.0, Geosci. Model Dev., submitted
2. Fuhrer, O., Osuna, C., Lapillonne, X., Gysi, T., Cumming, B., Arteaga, A., and Schulthess, T. C.: Towards a performance portable, architecture agnostic implementation strategy for weather and climate models, Supercomp. Front. Innov., 1, doi:10.14529/jsfi140103, 2014.
3. Gysi, T., Osuna, C., Fuhrer, O., Bianco, M., and Schulthess, T. C.: STELLA: A Domain-specific Tool for Structured Grid Methods in Weather and Climate Models, in: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC ’15, 41:1–41:12, doi:10.1145/2807591.2807627, 2015.
4. Lapillonne, X. and Fuhrer, O.: Using Compiler Directives to Port Large Scientific Applications to GPUs: An Example from Atmospheric Science, Parallel Process. Lett., 24, 1450003, doi:10.1142/S0129626414500030, 2014.

5. Leutwyler, D., Fuhrer, O., Lapillonne, X., Lüthi, D., and Schär, C.: Towards European-scale convection-resolving climate simulations with GPUs: a study with COSMO 4.19, Geosci. Model Dev., 9, 3393-3412, https://doi.org/10.5194/gmd-9-3393-2016, 2016.

Recent Projects (up to 5)

- H2020 ESCAPE – extension of GridTools DSL to unstructured grids used in global models and inter-operability with the Atlas data storage library
- HP2C COSMO – development of a version of COSMO fully capable of running on GPU-accelerated hybrid HPC architectures based on the STELLA DSL
- PASC GridTools – development of the GridTools DSL, which will supersede the STELLA DSL
- ETH CLAW – development of the CLAW DSL and source-to-source translator
- PASC ENIAC – application of the GridTools DSL to parts of the ICON model

Major Hardware/ Infrastructure available

Access to all systems at the Swiss National Supercomputing Centre CSCS (through research grants).

05 - Barcelona Supercomputing Center – Centro Nacional de Supercomputación (BSC)

About the Organisation

The Barcelona Supercomputing Center (BSC) was established in 2005 and is a key element of and coordinates the Spanish Supercomputing Network, which is the main framework for granting competitive HPC time to Spanish research institutions. Furthermore, BSC-CNS is one of six hosting nodes in France, Germany, Italy and Spain that form the core of the Partnership for Advanced Computing in Europe (PRACE) network. PRACE provides competitive computing time on world-class supercomputers to researchers in the 25 European member countries.

The Center houses Mare Nostrum, one of the most powerful supercomputers in Europe with 48,128 cores and 1.1 Pflops capacity. The mission of BSC is to research, develop and manage information technologies in order to facilitate scientific progress. BSC combines HPC service provision, and R&D into both computer and computational science (life, earth and engineering sciences) under one roof and currently has over 450 staff from 44 countries. BSC has collaborated with industry since its creation, and participates in various bilateral joint research centers with companies such as IBM, Microsoft, Intel, NVIDIA and Spanish oil company Repsol. The Center has been extremely active in the EC Framework Programs and has participated in over 100 projects funded by it. BSC is a founding member of HiPEAC, the ETP4HPC and other international fora.

The ES-BSC activities with the focus on global climate modelling and prediction are based on research, development and predictions with the EC-Earth climate forecast system. EC-Earth is the state-of-the art coupled climate model that is being developed and used for climate predictions and projections by the European consortium of more than 20 research and operational institutions from European Centre for Mid-range weather Forecasts (ECMWF is provider of the atmospheric and land components) to ES-BSC. Beside contributing to the 5th phase of the Coupled Model Intercomparison Project (CIMP5) critical for the UN IPCC Fifth Assessment Report (AR5), global climate research activities at ES-BSC enable provision of various historical reconstructions and initial conditions to the EC-Earth community for analysis of climate dynamics and for seasonal to decadal climate predictions. The ES-BSC is a contributor to the IS- ENES FP7 European project fostering the integration of the European climate modelling community and the development of Earth System Models (ESM) for advancing the understanding and predictions of climate variability and change. The ES-BSC is already active in the planning and design of the future coupled climate model intercomparison project, CIMP6, and is preparing to make key contributions including the groundbreaking high-resolution climate simulations with EC-Earth.

Role in this Project

BSC-CNS will participate in the project in the following:

- Defining and analyzing benchmarks and assessing their performance

- Collaborating in the work package devoted to VVUQ framework with CEA to adapt URANIE to Earth System models.

Expertise with Relevance to this Project and Role

Experience in performance analysis earth system models running in different HPC platforms. This knowledge includes deploy and configure EC-Earth coupled climate model.

Key Persons

Kim Serradell (Male) - Kim Serradell is the manager of the Computational Earth Science (CES) group at the Earth Sciences department in the Barcelona Supercomputing Center (BSC-CNS). The CES group is a multidisciplinary team of 19 members with different IT profiles that interacts closely with all the other groups of the Earth Sciences Dept. The group has among its tasks providing help and guidance to the scientists with the technical issues related to their work and developing a framework for the most efficient use of HPC resources.

Mario Acosta (Male) - Mario C. Acosta is a post-Doctoral fellow in the Computational Group of the Earth Sciences Department at the Barcelona Supercomputing Center. He received his PhD in Computer Science (related to High Performance Computing applied to Earth System Modelling) from University of Granada in 2015. His research interests and expertise include wide knowledge in numerical models (governing equation, numerical algorithms and computational implementation), performance analysis to highlight the main bottlenecks of the models and how to adapt and optimize them efficiently to actual and new High Performance Platforms.

Alicia Sanchez (Female) - Alicia Sanchez holds a PhD in "High-Precision Gamma and X-Ray Spectroscopy" from the Johannes Gutenberg University, Mainz, (Germany). After working at the Helmholtz Institute Mainz (Germany) in the Nuclear Spectroscopy Physics Division as a research manager, she joined the Barcelona Supercomputing Center Earth-Sciences co-leading the computational Earth Sciences group. She has a long experience in developing computational techniques (Geant4, ROOT) as well as data analysis algorithms to deal with huge amount of data (Big Data) in the framework of nuclear research facilities. She has participated on several international Scientific Collaborations worldwide and has helped with her interdisciplinary knowledge to consolidate important European research projects within the framework of High Precision Energy technologies. Currently, her interest focuses on developing Big Data and Machine Learning techniques in the field of earth sciences.

Relevant Publications (up to 5)

1. Yepes-Arbós, X., M. C. Acosta, K. Serradell, A. Sanchez Lorente, F.J. Doblas-Reyes (2017). Simulation-based performance analysis of EC-Earth 3.2.0 using Dimemas. BSC-CES Technical Memorandum 2017-001, 30 pp.
2. Acosta, M.C., X. Yepes-Arbós, S. Valcke, E. Maisonnave, K. Serradell, O. Mula-Valls and F.J. Doblas-Reyes (2016). Performance analysis of EC-Earth 3.2: Coupling BSC-CES Technical Memorandum 2016-006, 38 pp.
3. Yepes-Arbós, X., M.C. Acosta, K. Serradell, O. Mula-Valls, F.J. Doblas-Reyes (2016). Scalability and performance analysis of EC-Earth 3.2.0 using a new metric approach (Part I) BSC-CES Technical Memorandum 2016-001, 28 pp.
4. Yepes-Arbós, X., M.C. Acosta, K. Serradell, O. Mula-Valls, F.J. Doblas-Reyes (2016). Scalability and performance analysis of EC-Earth 3.2.0 using a new metric approach (Part II) BSC-CES Technical Memorandum 2016-004, 56 pp.

Recent Projects (up to 5)

- FP7 project Infrastructure for the European Network for Earth System modelling phase 2 (IS-ENES2-312979) IS-ENES2 is the second phase project of the distributed e-infrastructure of models, model data and metadata of the European Network for Earth System Modelling (ENES). This network gathers together the European modelling community working on understanding and predicting climate variability and change. ENES organizes and supports European contributions to international experiments used in assessments of the Intergovernmental Panel on Climate Change. This activity provides the predictions on which EU mitigation and adaptation policies are built.
- H2020 project ESiWACE Excellence in SImulation of Weather and Climate in Europe (GA 675191): 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.

- H2020 project: PRocess-based climate sIMulation: AdVances in high resolution modelling and European climate Risk Assessment (PRIMAVERA-641727) the main objective is to develop a new generation of advanced and well-evaluated high-resolution global climate models, capable of simulating and predicting regional climate with unprecedented fidelity, for the benefit of governments, business and society in general
- H2020 project Seasonal-to-decadal climate Prediction for the improvement of European Climate Services (SPECS-308378) will undertake research and dissemination activities to deliver a new generation of European climate forecast systems, with improved forecast quality and efficient regionalization tools to produce reliable, local climate information over land at seasonal-to-decadal time scales, and provide an enhanced communication protocol and services to satisfy the climate information needs of a wide range of public and private stakeholders.

Major Hardware/ Infrastructure available

BSC-CNS is the National Supercomputing Facility of Spain and hosts a range of high-performance computing (HPC) systems including Mare Nostrum IV the new supercomputer, will be 12.4 times more powerful than the current Mare Nostrum 3 that will have a performance capacity of 13.7 Petaflop/s. The general-purpose element, will have 48 racks with more than 3,400 nodes with next generation Intel Xeon processors and a central memory of 390 Terabytes. The second element of Mare Nostrum 4 will be formed of clusters of three different technologies that will be added and updated as they become available. These are technologies currently being developed in the US and Japan to accelerate the arrival of the new generation of pre-exascale supercomputers.

The BSC-CNS is a key element of and coordinates the Spanish Supercomputing Network, which is the main framework for granting competitive HPC time to Spanish research institutions. Furthermore, BSC-CNS is one of six hosting nodes in France, Germany, Italy and Spain that form the core of the Partnership for Advanced Computing in Europe (PRACE) network. PRACE provides competitive computing time on world-class supercomputers to researchers in the 25 European member countries.

06 - Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA)

About the Organisation

The French Alternative Energies and Atomic Energy Commission (CEA) is a key player in research, development and innovation in four main areas:

- nuclear and renewable energies,
- defence and security,
- technological research for industry,
- fundamental research in the physical sciences and life sciences.

Drawing on its widely acknowledged expertise, the CEA actively participates in collaborative projects with a large number of academic and industrial partners.

The CEA is established in nine centers spread throughout France. It works in partnership with many other research bodies, local authorities and universities. Within this context, the CEA is a stakeholder in a series of national alliances set up to coordinate French research in energy (ANCRE), life sciences and health (AVIESAN), digital science and technology (ALLISTENE), environmental sciences (AllEnvi) and human and social sciences (ATHENA).

Role in this Project

Work package leader (WP4: Validation, Verification and Uncertainty).

Scientific contributor.

Expertise with Relevance to this Project and Role

The LGLS laboratory (Software Engineering for Numerical Simulation), within the Nuclear Energy Department of CEA, has developed over the last 15 years a strong expertise within the area of VVUQ (Verification, Validation and Uncertainty Quantification). Initially dedicated to nuclear reactor physics, this expertise has been extended and capitalized into the VVUQ open-source platform URANIE.

URANIE provides a generic framework for the VVUQ analysis of any physical simulation code, regardless of the underlying physic and numeric. Thanks to its various modules, URANIE is able to perform sensitivity analysis, design of experiment, optimisation and surrogate model building, all of it in a generic fashion, without having to deal with the specific implementation of the analysed code.

Key Persons

A. Bruneton (Male) – graduated from Telecom ParisTech, one of the top French public institutions of higher education and research. After having worked for 5 years in the banking sector in London (market risk modelling and management), I went back to a research activity in Germany, at Fraunhofer ILT (Aachen), focused on the design and simulation of freeform optics. Since 2013 I have been working at CEA in the LGLS laboratory (Software Engineering for Numerical Simulation) which is in charge of developing three major open-source platforms: TRUST (CFD simulation), SALOME (pre and post processing toolbox for numerical simulation), and URANIE (generic VVUQ platform).

J-B. Blanchard (Male) – graduated in fundamental physics from Paris-Sud University in 2008, I defended my ph.D in particle physics on "Charge asymmetry and angular distributions of W bosons in the ATLAS experiment at the LHC" within CNRS, at Paris-sud university. I have kept on working on W boson mass related measurements for four years at IRFU (CEA), leading to the best single-measurement up-to now. Since late 2015, I have joined the LGLS laboratory (Software Engineering for Numerical Simulation), developing the URANIE platform used to handle uncertainty-related analysis within numerical simulation frameworks.

Relevant Publications (up to 5)

1. BOULORÉ, A., STRUZIK, C., et GAUDIER, F. Uncertainty and sensitivity analysis of the nuclear fuel thermal behavior. Nuclear Engineering and Design, 2012, vol. 253, p. 200-210.
2. REUX, C., DI GALLO, L., IMBEAUX, F., et al. DEMO reactor design using the new modular system code SYCOMORE. Nuclear Fusion, 2015, vol. 55, no 7, p. 073011.
3. LE TELLIER, R., SAAS, L., et BAJARD, S. Transient stratification modelling of a corium pool in a LWR vessel lower head. Nuclear Engineering and Design, 2015, vol. 287, p. 68-77.
4. JABOULAY, J.-C., PUMA, A. Li, et ARROYO, J. Martínez. Neutronic predesign tool for fusion power reactors system assessment. Fusion Engineering and Design, 2013, vol. 88, no 9, p. 2336-2342.
5. GHIONE, Alberto, NOEL, Brigitte, VINAI, Paolo, et al. Assessment of thermal–hydraulic correlations for narrow rectangular channels with high heat flux and coolant velocity. International Journal of Heat and Mass Transfer, 2016, vol. 99, p. 344-356.

Recent Projects (up to 5)

- NURESAFE ("NUCLEAR REACTOR SAFETY SIMULATION PLATFORM") -- Grant agreement no: 323263: The objective of NURESAFE was to deliver to European stakeholders a reliable software capacity usable for safety analysis needs and to develop a high level of expertise in the proper use of the most recent simulation tools. The NURESIM platform developed in the project includes the S&U (Sensitivity and Uncertainties analysis) software URANIE. S&U methods of URANIE were used to optimize and validate the applications, the calculation routes and the coupling schemes developped in the framework of NURESAFE project.
- EUROfusion Consortium - DEMO fusion reactor -- Grant agreement no: 633053: A roadmap to the realization of fusion energy was adopted by the EFDA system at the end of 2012. The roadmap aims at achieving all the necessary know-how to start the construction of a demonstration power plant (DEMO) by 2030, in order to reach the goal of fusion electricity in the grid by 2050. A coupling method has been applied to SYCOMORE (SYstem COde for MOdeling tokamak REactor) which is a system code developed in form of a modular workflow for designing magnetic fusion reactors. The coupling of SYCOMORE with the optimization platform URANIE enables design optimization along various figures of merit and constraints.

Major Hardware/ Infrastructure available

N/A

07 - Loughborough University (LU)

About the Organisation

The Wolfson School of Mechanical, Electrical and Manufacturing Engineering is one of the biggest engineering schools in the UK. The school is rated top 3 in the UK for overall quality of research in aeronautical, mechanical, chemical , and manufacturing engineering, with 95% of the school's research rated as world leading or of international standing in the latest UK Government Research Assessment Exercise (REF UoA B12 based on GPA x volume). In addition to holding two prestigious research Queen's Anniversary Prizes, the school is home to two EPSRC (Engineering and Physical Sciences Research Council) National Centres. The £52m research portfolio is undertaken within the school's diverse range research programmes with new research grants and contracts that generate annual research funds in excess of £10m/year. Among other disciplinary areas, the school is renowned for its world-class expertise in thermo-fluids, additive manufacturing, optical engineering, regenerative medicine, materials and structures. This research effort includes work of an established numerical methods development group, working on novel numerical methods and advanced numerical scientific applications to multi-disciplinary problems. Loughborough University hosts the Midland+ HPC facility, which is a top 500 ranked machine with more than 14000 Intel Xeon cores, 64TB of memory and a petabyte of disc storage; in addition to the University's HPC 2460-core cluster. The Wolfson School is also a part of the rewarded in 2014 Natural Environmental Research Council (NERC) Doctoral Training Partnership for Central England.

Role in this Project

LU will co-lead WP1, lead task 1.3 and co-lead task 1.5.

LU will contribute to the development of multigrid and multilevel technologies for the finite volume module of IFS.

LU will lead task 5.2 and will contribute to formulating and executing a plan for gender balance.

Expertise with Relevance to this Project and Role

Development of novel numerical methods for atmospheric models such as unstructured mesh based MPDATA, Krylov solvers, and non-oscillatory forward-in-time schemes. Expertise in mesh generation and adaption, FV, FE and FD discretisations, parallel processing, and multigrid – including custom built techniques for IFS-FVM. The proposed work will build on expertise and developments contributed to the ESCAPE project.

Key Persons

Dr Joanna Szmelter (Female) - is a reader in computational fluid and solid dynamics at Wolfson School. She joined Loughborough University from Cranfield University in October 2006, where she was a senior lecturer in the Ballistics and CFD Centre. Prior to this she was a principal engineer and a Head of the Aerodynamic Technology Group at BAe Airbus Ltd responsible for the development and implementation of all aspects of numerical methods and methodologies for application to the aerodynamic design of civil transport aircraft. She gained her Ph.D. and worked as a Research Associate at Swansea University. Among other appointments, she served as a member of the UK Advisory Group for Computational Fluid Mechanics (1994-96) and the BAe representative in the Airbus Industries Unsteady Aerodynamics Working Group at Toulouse (1993-95). Between 2005- 2012 she was awarded annually a visiting scientist appointment by the National Center for Atmospheric Research (NCAR), Boulder, Colorado and in 2012 the Newton Institute Fellowship from Cambridge University. Currently, she is a member of the EPSRC Peer Review College. Her work has been funded by industry (Airbus, Roxel, BAe Systems, Royal Ordnance), the MoD, Dstl, the EU, the British Council, The Royal Society and the EPSRC. Her edge-based framework and mesh generation developments formed a base for unstructured mesh NFT MPDATA atmospheric models that she originated jointly with Dr PK Smolarkiewicz, and which are now integrated into the Finite Volume Module for IFS. She was a co-leader of WP1 in the ESCAPE project.

Dr. Mike Gillard (Male) - is a research associate at the Wolfson School of Mechanical, Electrical and Manufacturing Engineering, working on elliptic operator preconditioning for bespoke Krylov solvers to solve atmospheric flows at exascale. Mike has an MSci in Mathematical Physic (2005) from the University of

Nottingham, and an MSc in Elementary Particle Theory (2007) and a Ph.D. in Mathematics (2011) from Durham University. He has experience in scientific computing, programming, massively parallel systems, algorithm design, analysis, and has worked on several large sized software codebases. He has worked in both a commercial and academic environments, including both developing and maintaining a very large software suite for Dstl while working as a consultant for BAE Systems Detica (2012-2013) which has been in active use and development for over 10 years. He also developed algorithms to solve non-linear PDE's permitting Topological solitons as a researcher in the School of Mathematics, University of Leeds (2013-2016), which formed part of a large software suite designed for use on massively parallel systems. This work was supported by EPSRC under the project Geometry, holography and Skyrmions - linking fundamental theory and experimental physics through a geometric description of nuclear matter. Mike joined Loughborough University (20016-present) as part of the ESCAPE project, to work on the development of a class of scalable, energy-efficient Krylov solvers bespoke for NWP. Particular focus was given to using multigrid techniques to solve a 3D Helmholtz problem through effective preconditioning, in the framework of the Finite Volume Module for IFS and the Atlas library.

Relevant Publications (up to 5)

1. Smolarkiewicz, P.K., J. Szmelter, and F. Xiao, 2016: Simulation of all-scale atmospheric dynamics on unstructured meshes. J. Comp. Phys., 322, 267-287.
2. Smolarkiewicz, P.K., J. Szmelter, and Z. Zhang, 2015: An unstructured mesh atmospheric model for nonhydrostatic dynamics: Towards optimal mesh resolution. J. Comp. Phys., 294, 363-381.
3. Smolarkiewicz, P.K., J. Szmelter, and A.A. Wyszogrodzki, 2013: An unstructured-mesh atmospheric model for nonhydrostatic dynamics. J. Comp. Phys., 254, 184-199.
4. Szmelter, J. and P.K. Smolarkiewicz, 2011: An edge-based unstructured mesh framework for atmospheric flows. Comput. Fluids, 46, 455-460.
5. Szmelter, J. and P.K. Smolarkiewicz, 2010: An edge-based unstructured mesh discretisation in geospherical framework. J. Comput. Phys., 229, 4980-4995.

Recent Projects (up to 5)

- ESCAPE, (2016-2018) H2020-FETHCP-2014.
- Evolving multimaterial interface simulations, (2013-2017) AWE PLC. This project provides high order MPDATA based novel remapping schemes.
- Unstructured Adaptive Mesh Model for Stratified Turbulence in Atmospheric Flows, (2008-2013) NERC, NE/G004358/1.

Major Hardware/ Infrastructure available

Wolfson school HPC provision includes two HPC facilities Hydra (LU cluster) Athena HPC midlands+ (for use by a consortium of 7 institutions in the Midlands). Midlands+ is a new entry to the Top500. The facilities are detailed as follows:

Hydra is a 2460-core 64-bit Intel Xeon cluster supplied by Bull, 103TB storage, with 185 total compute nodes.

Athena: 1 PB file store, supplied by ClusterVision and Huawei, with Power8 large memory nodes. 512 compute nodes in Huawei X6000 quad-node chassis, each with 28 cores, in the form of two Intel Xeon E5-2680v4, and 128 GB of memory - giving a total of 14,336 cores and 64 TB of memory. 1 Petabyte of disk storage, delivered via GPFS from four storage nodes. Dedicated 10 TB SSD GPFS filestore for prestaging files. An EDR Infiniband high-performance internal network. 5 OpenPOWER compute nodes, each with 20 cores and 1 TB of memory. One of the OpenPOWER nodes also has two Nvidia GP100 GPGPU cards, connected via NVlink.

Remote access to ECMWF HPC facilities will be used for final implementation and validation of the proposed developments to IFS-FVM.

08 - Royal Meteorological Institute (RMI)

About the Organisation

The Royal Meteorological Institute is the Belgian National Meteorological Service (NMS). Its mission is to provide weather and climate services and to carry out research. The RMI is a Member State of the ECMWF and uses the model output of the ECMWF model IFS for its operational weather forecasts. The RMI is also a Partner of the ALADIN NWP consortium.

Role in this Project

The RMI will contribute to various work packages and will co-lead WP3. Specifically,

- it will contribute with the delivery of a LAM EPS called RMI EPS and its radiation scheme, and
- it will carry out research and development on the verification and validation methods.

Expertise with Relevance to this Project and Role

Piet Termonia is leading the ALADIN consortium as Program Manager. The model configurations of the ALADIN System are based on a code that is shared with the one of the IFS of ECMWF. The research and development of the ESCAPE-2 project will be aligned with the ALADIN work plan, which is currently written jointly with the HIRLAM consortium. The Belgian ALADIN team has actively contributed to the developments of the ALADIN model.

Key Persons

Dr. Piet Termonia (Male) - is the head of the department of meteorological and climatological research of the Royal Meteorological Institute (RMI). Since 2004 he is associated to Ghent University as a guest professor, where he is teaching dynamic meteorology and climatology and where he supervises currently 4 PhD students. He obtained his PhD in theoretical physics in 1997. After a 2-year postdoc at the Institute for Theoretical Physics of the University of Torino, he switched to meteorology in the field of numerical weather prediction (NWP). At the RMI he established a new team that is currently active in regional climate modelling producing state-of-the-art RCP simulations for the CORDEX project based on the ALARO model. His team has two collaborators specialized in numerics and dynamics. He co-authored 50 A1 papers on a wide range of topics in atmospheric model ranging from dynamics, data assimilation, validation studies and regional climate modelling. He is a member of the board of the ECMWF Scalability Programme.

Dr. D. Degrauwe (Male) - is researcher at the Royal Meteorological Institute (RMI). In 2007, he obtained his PhD at KULeuven with a thesis on how to deal with uncertain parameters in numerical models that are used in the design of structures. After a 1-year postdoc at KULeuven, he joined the RMI in 2008, where he has since worked on several aspects of the ALADIN NWP model, including numerics, dynamics, thermodynamics, physics parameterizations and code efficiency. Since 2016, he holds the position of Code Architect within the ALADIN consortium. He teaches a course on numerical techniques in NWP at Ghent University.

Dr. G. Smet (Male) - is researcher at the Royal Meteorological Institute (RMI), in the department of meteorological and climatological research. He obtained a PhD in theoretical physics at KULeuven in October 2006. Afterwards, he switched to meteorology in the field of numerical weather prediction (NWP). He has been working at the RMI since December 2008, focusing on ensemble prediction, forecast verification, and economical value of probabilistic weather forecasts.

Relevant Publications (up to 5)

1. Termonia, P., F. Voitus, D. Degrauwe, S. Caluwaerts and R. Hamdi, 2012: Application of Boyd's periodization and relaxation method in a spectral atmospheric limited-area model – Part I : implementation and reproducibility tests. Mon. Wea. Rev., 140, 3137 – 3148. doi: http://dx.doi.org/10.1175/MWR-D-12-00033.1

2. De Troch, R., R. Hamdi, H. Van De Vyver, J.-F. Geleyn, P. Termonia 2013: Multiscale performance of the ALARO-0 model for simulating extreme summer precipitation climatology in Belgium, J. Climate, 26, 8895-8915.
3. Caluwaerts S., D. Degrauwe, P. Termonia, F. Voitus, P. Bénard and J-F. Geleyn, 2015: Importance of temporal symmetry in spatial discretization for geostrophic adjustment in semi-implicit Zgrid schemes. Q. J. R. Meterorol. Soc., 141 128138. doi: 10.1002/qj.2344.
4. Degrauwe, D., Seity, Y., Bouyssel, F., and Termonia, P.: Generalization and application of the flux-conservative thermodynamic equations in the AROME model of the ALADIN system, Geosci. Model Dev., 9, 2129-2142, doi:10.5194/gmd-9-2129-2016, 2016.
5. Smet, G., Termonia, P., Deckmyn, A. (2012). Added economic value of limited area multi-EPS weather forecasting applications. Tellus A, 64, 18901, 20 pages. Available online at http://dx.doi.org/10.3402/tellusa.v64i0.18901.

Recent Projects (up to 5)

- The ALADIN program (2015-2020). This is a consortium of 16 National Meteorological Services (NMSs) with an annual reported man power of about 90 FTE. Its aim is to maintain and develop a state-of-the-art Limited-Area Model NWP system. The ALADIN consortium is working closely together with the European HIRLAM consortium and these two consortia are now sharing its model code.
- CORDEX.be (2015-2017). This is a Belgian project funded by the Belgian Science Policy grouping all of the Belgian climate modelling groups. The aim of the project was to compute climate scenarios for Belgium following the guidelines to the CORDEX programme and to perform simulations with model configurations at higher-resolutions than the ones of the CORDEX. The RMI was the coordinator of this project.
- ESCAPE (2016-2018), H2020-FETHPC-2014. The RMI is a partner in the ESCAPE project and co-leading its work package 4.
- ERA4CS, H2020-SC5-2015. This is a European Research Area (ERA) project to boost European climate services. The RMI is Partner in following the three funded projects: URCLIM, INDECIS and MEDSCOPE, that will be carried out within the frame of ERA4CS.

Major Hardware/ Infrastructure available

The RMI uses for its NWP and climate applications an SGI Cluster with:

- 112 compute nodes with each 2 Xeon E5-2680V3 processors.
- 24 compute cores per node: 2688 cores in total;
- 88 compute nodes with 64GB Ram, 24 compute nodes with 256GB Ram;
- Infiniband FDR (56 Gbit/sec) interconnect;
- 2 scratch storage systems with each 60 TB of storage capacity.

09 - Politecnico di Milano - MOX (POLIMI)

About the Organisation

Politecnico di Milano (POLIMI) is the largest school of engineering, architecture and design and engineering in Italy. POLIMI counts about 40,000 students, roughly 10% of which from foreign countries. It provides a large number of Bachelor of Science, Master of Science and PhD programmes, of which 19 MS programmes, 3 BS programmes and 24 PhD programmes taught in English. As of 2013, POLIMI has participated in 211 projects financed by EU (FP7). It has started 24 Spin-offs since 2000 (20 still active) with a consolidated turnover of 13.7 millions Euro in 2013. Researchers of POLIMI produced a total of 337 inventions and 603 patents as of February 2013. According to QS World University Rankings 2015, POLIMI is ranked 24th in the world and 1st in Italy in the general category in the category “Engineering & Technology”. Strategic research is carried out mainly in the fields of energy, planning, management, design, applied mathematics, ICT, built environment and cultural heritage, with more than 250 laboratories, among which a Wind Tunnel and a Crash Test Centre. The Department of Mathematics at POLIMI is the second largest in Italy by number of full time staff and has been ranked first in Italy for the quality of research in recent assessments. The Laboratory for Modeling and Scientific Computing (MOX) at the Department of Mathematics of Politecnico di Milano started its activity in 2002. Its goal is coordinating and developing the expertise in mathematical modeling, numerical and statistical methods available in the Department of Mathematics of Politecnico di Milano. MOX collaborates with several research laboratories, from Politecnico and from other institutions belonging to industrial, clinical and academic world.

Role in this Project

POLIMI-MOX will contribute as co-coordinator of WP1 to the assessment and development in domain specific language framework of the advanced computational techniques that have been identified as main scientific goals of the project. POLIMI-MOX will also contribute specifically with the development of a semi-implicit, semi-Lagrangian DG model and with contributions to the study and application of other advanced mathematical techniques in the other workpackages.

Expertise with Relevance to this Project and Role

The key person has devoted most of his research activity to advanced numerical methods for NWP. Over the last 10 years, he has demonstrated in various the feasibility of the semi-implicit, semi-Lagrangian DG techniques and has supervised the development of a novel semi-implicit, semi-Lagrangian DG dynamical core ongoing at ICTP Trieste.

Key Persons

Luca Bonaventura (Male) – Luca Bonaventura graduated cum laude in Mathematics from University Sapienza of Rome in 1989 and received his PhD in Mathematics from University of Trento in 1994. He was recipient of a TMR-Marie Curie grant in 1996 and post-doctoral fellow at GKSS Forschungszentrum Geesthacht. He has been Assistant Professor in Numerical Analysis at the University of Trento and at Politecnico di Milano. He has been Scientist at Max Planck Institute for Meteorology, where he was in charge of the Numerics working group of the ICON project in 2002-2005. He is presently Associate Professor of Numerical Analysis at Politecnico di Milano. He has authored or co-authored 42 peer reviewed papers, mostly devoted to methods for numerical weather prediction and climate models. He has been a visiting scientist at Deutscher Wetterdienst, Institute for Pure and Applied Mathematics – UCLA, Keio University, ICTP Trieste and Euler visiting lecturer at ETH Zurich supported by ERCOFTAC. He has supervised 10 doctoral theses in applied mathematics and scientific computing.

Relevant Publications (up to 5)

1. Tumolo, G., & Bonaventura, L. (2015). A semi-implicit, semi-Lagrangian discontinuous Galerkin framework for adaptive numerical weather prediction. Quarterly Journal of the Royal Meteorological Society, 141(692), 2582-2601.
2. Bonaventura, L., & Ferretti, R. (2014). Semi-Lagrangian methods for parabolic problems in divergence form. SIAM Journal on Scientific Computing, 36(5), A2458-A2477.
3. Garcia, F., Bonaventura, L., Net, M., & Sánchez, J. (2014). Exponential versus IMEX high-order time integrators for thermal convection in rotating spherical shells. Journal of Computational Physics, 264, 41-54.
4. Wan, H., Giorgetta, M. A., Zängl, G., Restelli, M., Majewski, D., Bonaventura, L., et al. (2013). The ICON-1.2 hydrostatic atmospheric dynamical core on triangular grids, Part I: formulation and performance of the baseline version. Geoscientific Model Development, 6, 735-763.
5. Tumolo, G., Bonaventura, L., & Restelli, M. (2013). A semi-implicit, semi-Lagrangian, p-adaptive discontinuous Galerkin method for the shallow water equations. Journal of Computational Physics, 232(1), 46-67.

Recent Projects (up to 5)

- ERC Advanced Grant iHeart (2016-2022)
- INDAM project “Semi-Lagrangian methods for hyperbolic systems”, 2015-16
- Italian Ministry of Research PRIN2009 project “Numerical models for scientific computing and advanced applications” 2011-2013
- ERC Advanced Grant N.227058 “MATHCARD” (2009-2014)
- RTN project HaeMOdel (2002-2006)

Major Hardware/ Infrastructure available

The main scientific computing cluster at the Department of Mathematics of Politecnico di Milano, Gigat, is currently equipped as follows: Cluster HP Proliant DL560 Gen8 5 nodes equipped with 4 Xeon® E5-4610 v2, 2.3 GHz CPUs each, a total of 32 cores per node or 160 total cores, 256 GB RAM per node, a total of 1.25TB RAM, 3x300GB HDD SAS per node. In addition, two smaller clusters, IDRA with 128 intel XEON cores and 24 GB RAM per node and CERBERO with 24 intel Xeon cores and 16 GB RAM per node. The Department of Mathematics is also being allocated by CINECA, the Italian Supercomputing Centre, a relevant number of core hours on its Tier-0 system Marconi.

10 - Danish Meteorological Institute (DMI)

About the Organisation

DMI provides meteorological services in the Commonwealth of the Realm of Denmark, the Faroe Islands, Greenland, and surrounding waters and airspace.

Meteorological services include forecasting, warnings and monitoring of weather, climate and related environmental conditions in the atmosphere, on land and at sea. DMI also collects and processes meteorological, climatological and oceanographic measurements/observations throughout the Realm.

DMI has solid and long-lasting experience in developing and operating limited area weather models (LAMs) and regional climate models (RCMs). The LAMs are developed in the framework of the international ALADIN-HIRLAM consortium and share many of the model components with the ECMWF Integrated Forecast System.

Role in this Project

DMI will focus on developing and training neural networks for demanding critical physics computation. Here radiation is the primary focus.

Expertise with Relevance to this Project and Role

DMI has expertise with radiation physics in atmospheric models. A PhD candidate (Peter Ukkonen) will be employed, who has experience with using neural networks for weather models. Furthermore, the SME: “Geminali AS” will be consulted in the project. Geminali has worked on neural networks for atmospheric radiation for more than 10 years.

Key Persons

Peter Ukkonen (Male) – Peter Ukkonen is interested in the use of machine learning in atmospheric science, and has several years of experience in this field. Most recently he has worked with predicting convective precipitation in a regional climate model using neural networks.

Kristian Pagh Nielsen (Male) – Kristian Pagh Nielsen has worked with radiative transfer for almost 20 years and with radiative transfer in weather models specifically for almost 10 years. His main focus areas are clouds, aerosols, and solar energy resource assessment and forecasting.

Relevant Publications (up to 5)

1. Ukkonen, P., A. Manzato, and A. Mäkelä (2017): Evaluation of Thunderstorm Predictors for Finland Using Reanalyses and Neural Networks, doi:10.1175/JAMC-D-16-0361.1, J. Appl. Meteorol. Clim., 56, 2335-2352.
2. Nielsen, K. P., E. Gleeson, and L. Rontu (2014): Radiation sensitivity tests of the HARMONIE 37h1 NWP model, doi:10.5194/gmd-7-1433-2014, Geosci. Model Dev., 7(4), 1433-1449.
3. Fan, L., W. Li, A. Dahlback, J. J. Stamnes, S. Stamnes, and K. Stamnes (2014): New neural-network-based method to infer total ozone column amounts and cloud effects from multi-channel, moderate bandwidth filter instruments, doi:10.1364/OE.22.019595, Opt. Expr., 22, 19595-19609.
4. Nielsen, K. P., L. Zhao, G. A. Ryzhikov, M. S. Biryulina, E. R. Sommersten, J. J. Stamnes, K. Stamnes, and J. Moan (2008): Retrieval of the physiological state of human skin from UV-VIS reflectance spectra: A feasibility study, doi: 10.1016/j.jphotobiol.2008.06.010, J. Photochem. Photobiol. B: Biol., 93, 23-31.
5. Hestenes, K., K. P. Nielsen, L. Zhao, J. J. Stamnes, and K. Stamnes (2007): Monte Carlo and discrete-ordinate simulations of spectral radiances in a coupled air-tissue system, doi:10.1364/AO.46.002333, Appl. Opt. 46, 2333-2350.

Recent Projects (up to 5)

- HIRLAM C – HIRLAM C is a collaborative project of several national weather services running from 2016 to 2020. HIRLAM is an abbreviation for "High Resolution Limited Area weather Modelling". The main focus is integrated short-range ensemble forecasting systems especially suited for extreme weather.
- ESCAPE – H2020 (GA 671627) – In the ESCAPE project world-class, extreme-scale computing capabilities for European operational numerical weather prediction (NWP) and future climate models is being developed. Existing extreme-scale application software of weather and climate services is ill-equipped to adapt to the rapidly evolving hardware. ESCAPE is redressing this imbalance through innovative actions that fundamentally reform Earth system modelling.
- MarcoPolo – FP7 (GA 606953) – Due to the strong economic growth in the China in the past decade, air pollution has become a serious issue in many parts of the country. In this project, which ran from 2013 to 2017, the focus was on high resolution up-to-date regional air pollution information. DMI focussed particularly on aerosol feedback effects on weather forecasts.
- BeyondTMY – (IEA SolarPACES) – Proper knowledge on the availability and variability of solar irradiance – in particularly solar direct normal irradiance – is the single most important thing for estimating the performance of planned concentrating solar power systems. Developing best practice documents for making and utilizing such data sets is the focus of this project that runs from 2015 to 2018.

Major Hardware/ Infrastructure available

DMI has a Cray XC30 with 280 high performance compute nodes and 16 general purpose compute nodes. This is due to be upgraded with better processors in 2018 to reach of computer performance of ~700 TFlops. Half of the compute nodes are used for operational weather models, while the other half is used for research and development projects.

11 - CMCC Foundation Euro-Mediterranean Center on Climate Change (CMCC)

About the Organisation

The Euro-Mediterranean Center on Climate Change Foundation (CMCC Foundation) is a non-profit research institution. CMCC's mission is to investigate and model our climate system and its interactions with the environment and the society to provide rigorous, and timely scientific results in support to sustainable development and science driven adaptation and mitigation policies in a changing climate. Within CMCC, experienced scientists, economists, and technicians, work together to provide multidisciplinary analyses of climate impacts on various systems such as agriculture, ecosystems, coasts, water resources, and economics. CMCC also supports policymakers in setting and assessing costs, mitigation, and adaptation policies.

CMCC benefits from the extensive applied research experience of its members and institutional partners: Istituto Nazionale di Geofisica e Vulcanologia (INGV); Centro Italiano di Ricerche Aerospaziali (CIRA S.c.p.a.); Politecnico di Milano (PoliMi); Università Ca' Foscari Venezia; Università del Salento; Università di Sassari; Università della Tuscia.

CMCC acquired portfolio of research projects includes 225 funded projects: 37 in FP6 and FP7, 33 in H2020 and 155 funded projects under other EU and international research grants. In about a half of the implemented projects, CMCC was the project coordinator.

CMCC Foundation is also member of the European Network for Earth System Modelling (ENES), partner of the NEMO Consortium, contributing to its System Team, and the Earth System Grid Federation (ESGF), providing access to 100TB CMIP5 data through its data node deployed at the CMCC Supercomputing Center.

During 2016, the number of people who worked at CMCC is equal to 94 employees. People carrying out scientific and technical activities prevail (83%), while 17% of the staff performs administrative roles and carries out communication activities. The research activities are distributed among eight research divisions that share different knowledge and skills in the field of climate science: Advanced Scientific Computing (ASC) Division; Climate Simulation and Prediction (CSP) Division; Economic analysis of Climate Impacts and Policy (ECIP) Division; Impacts on Agriculture, Forests and Ecosystem Services (IAFES) Division; Ocean modeling and Data Assimilation

(ODA) Division; Ocean Predictions and Applications (OPA) Division; Risk Assessment and Adaptation Strategies (RAAS) Division; Regional Models and geo-Hydrological Impacts (REHMI) Division.

The Advanced Scientific Computing (ASC) Division of CMCC Foundation carries out Research & Development activities on Computational Sciences applied to Climate Change, focusing on the optimization of numerical models on HPC architectures and the management of large volumes of scientific data looking forward at exascale scenarios. In particular, a research unit of ASC Division works on: (I) the optimization and the parallelization of the numerical models for climate change simulations (both climate and impacts models); (II) the evaluation of how new emergent technologies impact on the implementation and the design of the current climate models; (III) the innovation of the main computational kernels by means of a re-design at algorithmic and software level.

The Ocean Modelling and Data Assimilation (ODA) Division of CMCC Foundation focuses on the development and improvement of the CMCC Earth System Model components with a particular emphasis on the ocean physical and biogeochemical models. In the last year, the ODA Division has tackled the challenge of simulating the high-resolution ocean dynamics at global scales, developing an eddying ocean/sea ice system, currently the highest resolution implemented in a global NEMO ocean-model domain.

Role in this Project

CMCC will contribute to the following activities:

WP1 - Extraction and providing NEMO dwarf

WP2 – Providing requirements for high-level domain specific language specification and demonstrating the application of the domain specific language toolchain to NEMO dwarf

WP3 - implementation of a NEMO reference global configuration at eddying resolution for the HPCW v0-benchmarks, integration, validation and benchmark of the novel DSL NEMO dwarf in the tier-1 benchmark suite

Expertise with Relevance to this Project and Role

Expertise on the analysis, optimization and parallelization of numerical models for climate change simulations, and on the development, implementation and evaluation of model physics of the CMCC Earth System Model components.

Key Persons

Silvia Mocavero (Female) Ph.D. - is Scientist at the "Advanced Scientific Computing" (ASC) Division of the CMCC Foundation, where is leading the "High End Computing" research group. Her skills include parallel programming on hybrid architectures, distributed environments, and deep experience in several parallel programming models: message passing, shared memory, many-threads programming with accelerators. She is exploring new computing issues, such as exascale computing. She works on performance analysis and optimisation of climate models with a particular focus on the NEMO ocean framework. Since 2012 she is member of the NEMO System Team and, since 2014, she is member of the HPC group of the NEMO Consortium. She has been strongly involved into several EU projects such as GridLab, CoreGRID, IS-ENES1, IS-ENES2 and national projects like FIRB Grid.it and TESSA. Currently she is involved in ESIWACE CoE working on models scalability. She is co-author of more than 25 papers in journals/proceedings on high-performance, distributed and grid computing.

Giovanni Aloisio (Male) - is full professor of Information Processing Systems at the Dept. of Innovation Engineering of the University of Salento, Lecce, Italy, where is leading the HPC laboratory. Former director of the Scientific Computing and Operations (SCO) Division at CMCC, he is now the Director of the CMCC Supercomputing Center and member of the Strategic Council. His expertise concerns high performance computing, grid & cloud computing and distributed data management. He was strongly involved in EU grid projects such as GridLab, EGEE, IS-ENES1, EUBrazilCC, IS-ENES2, CLIP-C and the G8 ExArch. As CMCC, he was also the coordinator of the OFIDIA (Operational FIre Danger preventIon plAtform) project, in the context of the European Territorial Cooperation Program Greece-Italy 2007-2013 and, currently, he is coordinating CMCC activities in the ESiWACE project. He has been the responsible for ENES of the EU-FP7 EESI (European Exascale Software Initiative) project and for University of Salento (as PRACE Third Party) of the EU-FP7 EESI 2 project, also chairing in both cases the WCES (Weather, Climate and solid Earth Sciences) European Working Group. He is member of the ENES HPC Task Force and one of the key experts of the IESP project (International Exascale Software Project), whose main goal is the definition of the roadmap for a common, open source software infrastructure for scientific computing at exascale. He is the author of more than 100 papers in referred journals on high performance computing, grid computing and distributed data management.

Sandro Fiore (Male) Ph.D. - is the Director of the Advanced Scientific Computing (ASC) Division of the Euro-Mediterranean Centre on Climate Change. His research activities focus on parallel, distributed, grid and cloud computing, in particular on distributed data management, data analytics/mining and high performance database management. He is Visiting Scientist at Lawrence Livermore National Laboratory (LLNL) working at PCMDI in the context of the Earth System Grid Federation (ESGF). Since 2004, he has been involved into several national and international projects like: EGEE (the 3 cycles), EGI-InSPIRE, IS-ENES1 and IS-ENES2, EUBRAZILCC, ExArch, ORIENTGATE, TESSA, OFIDIA, CLIP-C, INDIGO-DataCloud, ESIWACE working on data management topics. Since 2010, he is the Principal Investigator of the Ophidia project, a research project on high performance data analytics and mining for eScience. He is author and co-author of more than 50 papers in refereed books/journals/proceedings on distributed and grid computing and holds a patent on data management topics. He is editor of the book “Grid and Cloud Database Management” (Springer, 2011). He is ACM Member.

Doroteaciro Iovino (Female) - holds a Ph. D. in physical oceanography from the Geophysical Institute, University of Bergen (Norway). She expanded her knowledge and experience in oceanography and numerical modeling during the post-doctorate at the Laboratoire d'Océanographie et du Climat: Expérimentations et Approches Numériques (LOCEAN) in Paris (France), and is currently a researcher at the Euro-Mediterranean Centre on Climate Change (CMCC), in Bologna, where she coordinates the research activities of Ocean and Sea-Ice Modeling group within the Ocean Modeling and Data Assimilation Division at CMCC since 2015. She has been involved into several national and international projects, with more than 10 years of experience working on ocean and sea ice modeling, both on the technical and scientific aspects, with particular interest in the high-resolution ocean dynamics. She was involved in high-resolution modeling study within ENS4OCEAN, a PRACE funded projects. Since 2017, she is member of the CLIVAR/CliC Northern Oceans Regional Panel (NORP). As NEMO Officer, she leads the CMCC effort within the NEMO System Team in developing the model system, member of the NEMO sea ice Working Group since 2016. Since 2012, she teaches in Ph.D. programme in Science and Management of Climate Change at Ca’ Foscari University in Venice.

Relevant Publications (up to 5)

1. Balaji V., Maisonnave E., Zadeh N., Lawrence B. N., Biercamp J., Fladrich U., Aloisio G., Benson R., Caubel A., Durachta J., Foujols M-A., Lister G., Mocavero S., Underwood S., Wright G. (2017) CPMIP: measurements of real computational performance of Earth system models in CMIP6, Geoscientific Model Development, 10, 19-34, DOI: 10.5194/gmd-10-19-2017
2. Epicoco I., Mocavero S., Porter A. R., Pickles S. M., Ashworth M., Aloisio G. (2017). Hybridisation strategies and data structures for the NEMO ocean model, International Journal of High Performance Computing Applications, SAGE Publications, ISSN 1094-3420, Online ISSN: 1741-2846, DOI: 10.1177/1094342016684930
3. Epicoco I., Mocavero S., Macchia F., Vichi M., Lovato T., Masina S., Aloisio G. (2016). Performance and results of the high-resolution biogeochemical model PELAGOS025 v1.0
4. J. Dongarra, P. Beckman, T. Moore, P. Aerts, G. Aloisio, J. C. Andre, D. Barkai, J. Y. Berthou, T. Boku, B. Braunschweig, F. Cappello, B. M. Chapman, X. Chi, A. N. Choudhary, S. S. Dosanjh, T. H. Dunning, S. Fiore, A. Geist, B. Gropp, R. J. Harrison, M. Hereld, M. A. Heroux, A. Hoisie, K. Hotta, Z. Jin, Y. Ishikawa, F. Johnson, S. Kale, R. Kenway, D. E. Keyes, B. Kramer, J. Labarta, A. Lichnewsky, T. Lippert, B. Lucas, B. Maccabe, S. Matsuoka, P. Messina, P. Michielse, B. Mohr, M. S. Mueller, W. E. Nagel, H. Nakashima, M. E. Papka, D. A. Reed, M. Sato, E. Seidel, J. Shalf, D. Skinner, M. Snir, T. L. Sterling, R. Stevens, F. Streitz, B. Sugar, S. Sumimoto, W. Tang, J. Taylor, R. Thakur, A. E. Trefethen, M. Valero, A. van der Steen, J. S. Vetter, P. Williams, R. Wisniewski, K. A. Yelick (2011). The International Exascale Software Project roadmap, International Journal of High Performance Computing Applications, 25(1): 3-60 (2011), ISSN 1094-3420, doi: 10.1177/1094342010391989
5. Iovino, D., S. Masina, A. Storto, A. Cipollone, and V. N. Stepanov, 2016: A 1/16 eddying simulation of the global NEMO sea-ice–ocean system. Geosci. Model Dev., 9, 2665-2684, doi:10.5194/gmd-9-2665-2016

Recent Projects (up to 5)

- EU H2020 Project ESIWACE (Centre of Excellence in Simulation of Weather and Climate in Europe; 2015-2019). ESIWACE is a user-driven Centre of Excellence in Simulation of Climate and Weather in Europe. In particular, CMCC Foundation participates to several activities including: enhancing community capacity in

HPC; scheduling and workflow capabilities; implementation of new storage layouts for Earth System Data to be evaluated on fast (in-memory) data analytics frameworks for scientific data management.

- EU FP7 Project IS-ENES (2009-2013) and its follow on IS-ENES2 (2013-2017). The goal of IS-ENES has been the development of a common climate and Earth system modelling distributed research infrastructure in Europe. IS-ENES2 further integrated the European climate modelling community, stimulated common developments of software for models and their environments, fostered the execution and exploitation of high-end simulations and supported the dissemination of model results to the climate research and impact communities. CMCC worked on the use of exascale to develop next generation climate models and on the data node monitoring infrastructure.
- EU H2020 Project PRIMAVERA (PRocess-based climate sIMulation: AdVances in high resolution modelling and European climate Risk Assessment; 2015-2019). The goal of PRIMAVERA is to deliver novel, advanced and well-evaluated high-resolution global climate models (GCMs), capable of simulating and predicting regional climate with unprecedented fidelity, out to 2050. The ASC division of the CMCC Foundation will work in WP9 addressing HPC and Data management challenges.
- EU FP7 Project EESI1 (2010-2011). The goal of the project is the production of the European Roadmap on Exascale. CMCC chaired the WP3/Task3.2 on Weather, Climatology and solid Earth Sciences.
- Tender 2015/S009-011211 Copernicus Marine environment Service (CMS) Mediterranean Monitoring and Forecasting Center (Med-MFC; 2015-2018). The Consortium aims to provide regular and systematic information about the physical state of the ocean and marine ecosystems for the Mediterranean Sea, starting from the pre-operational system developed during the Myocean projects. The CMS products will include analyses, 10 days forecasts, specific and targeted products and reanalysis. The CMS Med-MFC products will describe waves, currents, temperature, salinity, sea level and pelagic biogeochemistry.

Major Hardware/ Infrastructure available

The HPC infrastructure managed at the CMCC Supercomputing Centre is composed of a 8000 cores iDataPlex cluster equipped with Intel Xeon Sandy Bridge processors (peak performance 160TFlops). Athena (its name) is integrated with two DDN SFA10000 storage subsystems capable to offer a storage capacity of about 840TB in total and an I/O performance of 6GBytes/sec per disk array. Additionally, CMCC runs a big data analytics cluster with 5 fat nodes, 100 cores, 1.3TB RAM, 56TB GFS. An upgrade to a 2PFlops HPC facility is planned for 2018. Part of these resources is expected to be committed to this project

12 - Bull SAS (BULL)

About the Organisation

Bull (Bull SAS) is one of the newest member of the Atos family. In 2016, Atos SE is a leader in digital services with pro forma annual revenue of 5 e.6 billion and 96,000 employees in 72 countries. Serving a global client base, the Group provides Consulting & Systems Integration services, Managed Services & BPO, Cloud operations, Big Data & Cyber-security solutions, as well as transactional services. With its deep technology expertise and industry knowledge, the Group works with clients across different business sectors: Defense, Financial Services, Health, Manufacturing, Media, Utilities, Public sector, Retail, Telecommunications, and Transportation. With 80+ years of technology innovation expertise, the new « Big Data & Cyber-security » Service Line (ATOS BDS) gathers the expertise in Big Data, Security and Critical Systems brought by Bull acquisition and the ones already developed by Atos in this domain. The Service Line is structured into 3 complementary activities: Big Data, Cyber-security and Critical Systems.

- Big Data: the expertise of extreme performance to unleash the value of data (detailed below)
- Cyber-security: the expertise of extreme security for business trust
- Critical Systems: the expertise of extreme safety for mission-critical activities.

In recent years, the Bull R&D labs have developed many major products that are recognized for their originality and quality. These include the Sequana supercomputer which concretes the first results of the “Bull Exascale Program” announced during SuperComputing 2014, bullion servers for the private Clouds and Big Data, the Shadow intelligent jamming system designed to counter RCIEDs, the libertp tool for modernization of legacy applications and hoox, the first European smartphone featuring native security. To explore new areas and develop tomorrow's solutions, today, Bull R&D is investing heavily in customers – with whom it has forged many successful technological partnerships – as well as in institutional collaborative programs (such as competitiveness clusters and European projects) and in partnerships with industry (Open Source, consortiums). Bull is involved with the strategy towards HPC in Europe with the active leadership of ETP4HPC and contribution to the Strategic Research Agenda.

Role in this Project

European industrial HPC vendor, Bull will mainly contribute to ESCAPE-2 by helping at defining, optimizing, installing and testing a representative set of application benchmarks (the High Performance Climate and Weather (HPCW) Benchmark Suite) to provide valuable performance and energy metrics through its knowledge and long standing experience in procurements and acceptance testing of high-end HPC systems, including for Numerical Weather and Climate Prediction codes.

Moreover, during the project, Bull will bring access to various state-of-the-art HPC systems, technologies and tools that will be used to test ESCAPE-2 tools and to provide reference performance of the HPCW Benchmark Suite on a wide variety of systems.

Finally, Bull will support, through its contributions to the ESCAPE-2 project, the development of excellence in European HPC technology and thus strengthen its competitiveness and growth, especially for the weather and climate community.

Bull will be mainly involved in WP3 (T3.1, T3.2, T3.3, T3.5) which it co-lead.

Smaller participation is also planned in the WP5 for dissemination and communication.

Expertise with Relevance to this Project and Role

Long standing experience in procurements and acceptance testing of HPC systems including for weather and climate centres (DKRZ, Meteo-France, KNMI, AEMET) with access to a wide variety of state-of-the-art HPC systems.

Bull already took part in several European projects and Centre of Excellence in which it led work packages, tasks, and especially the industrial contributions and dissemination (ESCAPE, ESiWACE, etc.).

Key Persons

Xavier Vigouroux (Male) - after a Ph.D. from Ecole Normale Supérieure de Lyon in Distributed computing, worked for several major companies in different positions. He has now been working for bull for more than 10 years. He led the HPC benchmarking team for the first five years, then in charge of the "Education and Research" market for HPC at Bull, he is now managing the "Center for Excellence in Parallel Programming" of Bull.

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

Dr Erwan Raffin (Male) - obtained his PhD degree in computer science in 2011 from the University of Rennes1, France in the framework of collaboration between Technicolor and the INRIA/IRISA laboratory. During his PhD, he has focused his research on compilation and synthesis aspects of multimedia application acceleration on coarse-grained reconfigurable architectures using constraint programming. From 2011 to 2014, he was a research and development engineer at CAPS entreprise, which developed and commercialized innovative software for high performance application tuning in the domains of HPC. He then worked at INSA Rennes, IETR, UMR CNRS 6164, UEB in the Image Group as a research engineer on low power video decoder. He joined the Center for Excellence in Parallel Programming (CEPP) of Bull in 2016 focusing on Weather and Climate community and its HPC application domain.

Relevant Publications (up to 5)

As an industrial, Bull rarely submits publication in academic journals or to conferences. However, in association with academic organizations Bull participated to some publications, and Bull is present publicly with innovative technologies presentations at HPC major events such as ISC or SuperComputing.

As an expert in delivering ultra-high performance, Bull is one of the world leaders in Extreme Computing. As an IT manufacturer, Bull has a strong presence in the list of the world's top supercomputers. With more HPC specialists than any other player in Europe, Bull is recognized for the technological excellence of its bullx range, its HPC applications expertise and its ability to manage large-scale projects. Bull has developed a complete vertical HPC offer containing Hardware and Software technologies dedicated to scalable HPC systems:

- Bull Sequana: The open exascale-class supercomputer. With the new Bull sequana range of supercomputers, Atos confirms its strategic commitment to the development of innovative high performance computing systems, the systems needed to meet the major challenges of the 21st century. Designed by the Bull R&D in close cooperation with major customers, the Bull sequana X1000 supercomputer leverages the latest technological advances, so as to guarantee maximum performance for a minimized operation cost.
- bullx S6000 series: This scalable server has been specially developed to meet today's challenges of mass data processing, high availability, virtualization and eco-efficiency. The bullx S6130, the first available model of the bullx 6000 series, is fully scalable up to 16 CPUs/240 cores/24 TB. It supports full memory protection and features hot-swappable memory and I/O capabilities. The bullx S6130 is the HPC version of the bullion S16 enterprise server, currently recognized as the fastest x86 server on the market (SPECint_rate2006 for the 16 socket configuration).
- Bull eXascale Interconnect (BXI): This is a new interconnection network developed by Bull for High Performance Computing (HPC). BXI is highly flexible and it scales from small configurations to Exascale with up to 128k nodes. To boost HPC applications communications, besides its impressive links bandwidth (200Gb/s), BXI features hardware acceleration for message passing, one-way messages and collective operations. Finally, thanks to optimized remote direct memory access (RDMA) capabilities BXI can be for used data storage as well. The BXI network uses two types of ASIC, a Network Interface Connector (NIC) Lutetia and a switch Divio. It is managed out-of-band and can be configured with many topologies including fat-tree (non-blocking or pruned), flattened butterfly, dragonfly, hypercubes, n-D torus … BXI improves communications efficiency with a congestion control mechanism and adaptive routing. To balance priorities of different workflows (e.g. communications and IO), Quality of Service (QoS) is implemented through virtual networks and channels. Data integrity is maintained on the BXI network with Cyclic Redundancy Check (CRC) both from-end-to-end and for each individual link.
- bullx supercomputer suite (scs) is a comprehensive, powerful and robust software solution that meets the requirements of even the most challenging high performance computing needs. It is the result of Bull's long experience in deploying HPC software to the strictest specifications and major investments and continued efforts in R&D

Recent Projects (up to 5)

Bull is involved with the strategy toward HPC in Europe with the active leadership of ETP4HPC and contribution to the Strategic Research Agenda.
Bull participated to the following cooperative projects connected to the subject of this proposal:

- Mont-Blanc 3 :
  - The main target of the H2020 Mont-Blanc 3 project is the creation of a new high-end HPC platform that is able to deliver a new level of performance / energy ratio whilst executing real applications. The project builds upon the previous Mont-Blanc & Mont-Blanc 2 FP7 projects, with ARM, BSC & Bull, it adopts a co-design approach to ensure that hardware and system innovations are readily translated into benefits for HPC applications. This encompasses the three following objectives:
  - To design a well-balanced architecture and to deliver the design for an ARM based SoC (System-on-a-Chip) or SoP (System-on-Package) capable of providing pre-exascale performance when implemented in the 2019-2020 time frame.
  - To maximise the benefit for HPC applications with new high-performance ARM processors and throughput-oriented compute accelerators designed to work together.

- To develop the necessary software ecosystem for the future SoC. This additional objective is key to ensure that the project successfully translates into an industrial offer for the HPC market.

- SAGE: The SAGE project is working on building an Extreme scale data centric computing capable storage system. The project is building a storage system with in-storage compute capability, consisting of multiple tiers of storage devices. The use cases consist of extreme scale data generators (eg: scientific experiments) as well as extreme scale I/O HPC use cases.
- CompBioMed: CompBioMed is a user-driven Centre of Excellence in Computational Biomedecine, to nurture and promote the uptake and exploitation of high performance computing within the biomedical modelling community. Our user communities come from academia, industry and clinical practice.
- ESCAPE: ESCAPE stands for Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. The project will develop world-class, extreme-scale computing capabilities for European operational numerical weather prediction (NWP) and future climate models. The biggest challenge for state-of-the-art NWP arises from the need to simulate complex physical phenomena within tight production schedules. Existing extreme-scale application software of weather and climate services is ill-equipped to adapt to the rapidly evolving hardware. This is exacerbated by other drivers for hardware development, with processor arrangements not necessarily optimal for weather and climate simulations. ESCAPE will redress this imbalance through innovation actions that fundamentally reform Earth system modelling.
- ESIWACE: A main goal of ESiWACE is to 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. Besides, with regard to the upcoming exascale era, ESiWACE will establish demonstrator simulations, which will be run at highest affordable resolutions (target 1km). This will yield insights into the computability of configurations that will be sufficient to address key scientific challenges in weather and climate prediction.

Major Hardware/ Infrastructure available

The Atos CEPP has access to two computing clusters hosted in Bull France facilities. Both are continuously upgraded to the latest technologies (network, CPU, memory…).

The 1st one, called “genji”, is a medium scale system dedicated to experimentation bed. It is heterogeneous and its resources are evolving quickly in order to test early upcoming technologies. Thus it is equipped with standard x86 CPUs, Nvidia GPUs, Intel Xeon Phis, as well es with ARM processors. One large in-memory server MESCA-2 is also mounted to genji. This platform can be tuned. So depending on the needs its fabric could be either EDR IB-2, OPA or BXI, and different types of parallel filesystems may be available.

The 2nd one, called “manny”, is a large scaled system based on systems in delivery testing. It is designed for performance and is continuously equipped with several hundreds of nodes embedding last generations of CPUs and possibly accelerators.

With these platforms partners can bring their data and then install, test and validate their developments throughout the project whether they are technology providers or applications developers.

4.2 Third Parties involved in the project

Partner 10 - DMI

Does the participant plan to subcontract certain tasks (please note that core tasks of the project should not be sub-contracted)	Y
DMI will subcontract Geminali AS to deliver the radiative transfer tool AccuRT and provide the following services: *- Establishing a synthetic database using AccuRT* *- Training of forward-inverse model using Neural Networks* *This will support DMI in achieving the goals of T1.4 Training & Verification of Artificial Neural Networks (ANN). Jakob Stamnes and Knut Stamne of Geminali AS have worked on ANNs for radiation for more than 10 years in the fields of biomedical, atmospheric and ocean optics.*	

The value of the subcontract is EUR 32,215.	
Does the participant envisage that part of its work is performed by linked third parties[1]	N
n/a	
Does the participant envisage the use of contributions in kind provided by third parties (Articles 11 and 12 of the General Model Grant Agreement)	N
n/a	

No other third parties involved

5 Ethics and Security

5.1 Ethics

No ethical issues have been identified.

5.2 Security

- activities or results raising security issues: NO
- 'EU-classified information' as background or results: NO

[1] A third party that is an affiliated entity or has a legal link to a participant implying a collaboration not limited to the action. (Article 14 of the Model Grant Agreement).

6 Support Letters

6.1 Swiss National Supercomputing Centre - CSCS

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Lugano, September 25, 2017

Support for ESCAPE 2 project proposal to FET-HPC-02-2017

To Whom It May Concern:

I write this letter in support of the ESCAPE 2 proposal submitted to the FET-HPC-02-2017 call by Dr. Peter Bauer and Dr. Nils Wedi (both ECMWF). The proposed project will be conducted by a consortium of weather and climate prediction centres that operate at global, as well as European regional scales, along with universities performing research on numerical methods, HPC centres and a leading European HPC vendor. The proposal addresses challenges of energy-efficient high-performance computing and will provide key technologies for weather and climate simulations at exascale.

I have stated numerous times in public that weather and climate simulation constitute one of the most promising problems for the development of exascale computing platforms in Europe. Unlike petascale and earlier systems that were designed for general purpose, the early exascale systems will have to be designed to solve challenging (exascale) problems in specific domains. These problems have to represent a value that justifies multi-billion dollar investments and need to be well-articulated in terms of a factor ~100 increase over the current state of the art at petascale. Weather and climate fulfils these criteria in a domain where Europe is the recognized leader today. This means Europe owns the baseline and is in an optimal position to win the race to exascale.

Furthermore, I am convinced that systems designed to solve exascale problems in a specific domain like weather and climate will result in supercomputing platforms that are immediately useable for other domains/purposes as well. They will be much more productive at the outset than systems developed to perform on the high-performance LINPACK benchmark of the Top500 list.

ESCAPE 2 will be a key contributor for exascale technologies in weather and climate. It is thus completely aligned with the European HPC strategy and represents a sound investment for FET-HPC-02-2017 to make. For these reasons, I would like to express my strongest support and will certainly do everything within my means at CSCS to help this project succeed.

I would be pleased to provide any additional information upon request.

Yours sincerely,

Thomas C. Schulthess
Professor of Computational Physics at ETH Zurich and Director of the Swiss National Supercomputing Centre

Via Trevano 131
6900 Lugano
Switzerland

Tel +41 (0)91 610 82 11
Fax +41 (0)91 610 82 82
www.cscs.ch

6.2 ETHZ

ETH zürich

Institute for Atmospheric and Climate Science
ETH Zürich
Universitätsstr. 16
8092 Zürich, Switzerland

Prof. Christoph Schär
Tel. +41 44 632 8199
Fax +41 44 632 1311
schaer@env.ethz.ch
http://www.iac.ethz.ch/people/schaer

To Whom It May Concern

Letter of support for the proposal "Energy-efficient Scalable Algorithms for weather and climate Prediction at Exascale" (ESCAPE-2).

Zürich, August 30. 2017

Currently large efforts are underway to refine the horizontal resolution of global and regional weather and climate models to about 1 km, with the intent to represent critical processes such as convective clouds explicitly, rather than using semi-empirical parameterizations. This refinement will move the governing equations closer to first principles and is expected to reduce the uncertainties of weather and climate models. High resolution is particularly attractive in order to better represent critical cloud feedback processes (e.g. related to global climate sensitivity and extratropical summer convection) and extreme events (heavy precipitation events, floods, and hurricanes).

ESCAPE-2 is an exciting project of leading weather and climate prediction centers. Its predecessor project (ESCAPE) has already delivered rapid progress regarding numerical algorithms and the implementation of atmospheric and ocean models on emerging high-performance compute architectures.

The project is of great significance far beyond the immediate project participants. The increase in computational resolution of weather and climate models will be beneficial in a number of areas. It will strengthen and enable:

- research and applications related to extreme weather events, including socio-economic impacts and applications related to national authorities and the re-insurance sector.
- interdisciplinary research at regional to local scales, for instance related to hydrology, agriculture and ecosystem services. This research requires information at regional to local scales that is not available from conventional models.
- research on climate change adaptation. Climate change has its origins at global scales, but its impacts will show up regionally. Climate change adaptation has been identified on high levels (e.g. IPCC, WMO, EU, many national governments) as a key element of our response to this challenge.
- collaborations with university institutions, for instance related to limited-area high-resolution numerical modeling (convection-resolving and large-eddy-simulation

models) at horizontal resolutions between 10 m and 10 km. Such models can be used to pioneer the targeted resolutions at a time when they are not yet available in global models.

My own group has been working in this area and has developed a continental-scale climate modeling capability at km-scale resolution[1]. Currently we are starting up a (externally funded) PhD project that will address the role of topography in the climate system using km-scale continental-scale simulations. The project includes collaborations with ECMWF and other ESCAPE partners.

I am excited about ESCAPE-2: it is an important strategic and scientific element towards the development of the next generation weather and climate models.

Based on the foregoing, it is a pleasure to strongly recommend this project for funding.

Yours sincerely

[1] Project "Convection-resolving climate modeling on future supercomputing platforms (crCLIM)" funded by the Swiss National Science Foundation. See http://www.c2sm.ethz.ch/research/crCLIM.html

6.3 DWD

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Deutscher Wetterdienst - Postfach 10 04 65 - 63004 Offenbach

European Commission

Geschäftsbereich Forschung und Entwicklung

Ansprechpartner:
Prof. Dr. Sarah Jones
Telefon:
069/8062 5820
E-Mail:
Ruud.Dirksen@dwd.de

Geschäftszeichen:
FE/ESCAPE-2
Fax:
069 / 8062 - 2731

UST-ID: DE221793973

Offenbach, 3 August 2017

Letter of support for ESCAPE-2

Deutscher Wetterdienst (DWD), the national meteorological service of Germany, whole-heartedly supports the new H2020 proposal ESCAPE-2 led by the European Centre for Medium-Range Weather Forecasts (ECMWF).

ESCAPE-2 builds on the successful ESCAPE (Energy-Efficient Scalable Algorithms for Weather Prediction at Exascale) project (2015 – 2018) which combines expertise from world-leading global and regional numerical weather prediction centres, academia, high-performance computing centres and hardware vendors. In ESCAPE, so-called 'Weather& Climate dwarfs' are being developed, e.g. spectral transform with spherical harmonics, spectral transform with bi-Fourier transform, 3D elliptic solver, semi-Lagrangian advection scheme, a cloud scheme, and a radiation scheme.

ESCAPE-2 extends this expertise by bringing leading European weather and climate prediction centres together in a seamless modelling approach aiming at simulation time scales from minutes to decades. The open source ATLAS library of ECMWF will contain all ESCAPE-2 developments and will provide grid generators for structured and unstructured model grids as well as complete parallelization approaches and discretizations of different numerical operators like grad, div and curl. Because of this ATLAS will be a very attractive framework of future global and regional model developments for all European national meteorological services and research institutes.

Therefore the research and development planned in ESCAPE-2 will also provide an important basis for DWD's global and high-resolution regional numerical weather prediction (NWP) strategy until 2022 which foresees a migration of the regional NWP component from the COSMO-model to the regional mode of ICON, the introduction of a regional hourly rapid update cycle of analyses and 12-h forecasts, and a more than 10-fold increase in operational data production from 17 to 200 TByte per day.

DWD will closely collaborate with ESCAPE-2 by providing the ICON Modelling Framework (jointly developed by DWD and the Max-Planck-Institute for Meteorology in Hamburg) as one testbed of the methodologies in preparation for exascale computing.

Yours faithfully,

Deutscher Wetterdienst
Geschäftsbereich
Forschung und Entwicklung
Postfach 10 04 65
63004 Offenbach

Prof. Dr. Sarah Jones
Head Research and Development
Member of the Board of Directors of DWD

6.4 SMHI/ EC-EARTH

Letter of Support

To: Dr. Peter Bauer
ECMWF

16 Aug 2017

Dear Dr. Bauer,

We hereby state our full support to the "ESCAPE-2" proposal.

ESCAPE-2 aims at adressing key engineering challenges for weather and climate models of the future exascale age. Climate and Earth System Models such as EC-Earth will greatly benefit from those envisioned developments. EC-Earth is building on the atmosphere component IFS of ECMWFs prediction system. Thus, developments towards novel mathematical and algorithmic approaches, multi-grid methods, higher-order discretizations, machine learning techniques, new pathways for increased arithmetic intensity, enhanced robustness and fault resilience, and towards substantially enhanced code efficiency, are most useful and welcomed by the EC-Earth consortium.

Furthermore we welcome basic structural enhancements such as open and user-accessible programming models through domain-specific language concepts, reducing the hardware dependence of the science code and enhancing performance portability. Those measures are expected to increase accessibility of complex codes by climate scientists.

The EC-Earth consortium is looking forward to collaborate with ESCAPE-2, especially in the field of adaption of new methods for climate modelling, prediction and projection.

Sincerely,

Dr. Ralf Döscher, chair of the Steering Commitee of EC-Earth, science coordinator of the Rossby Centre at SMHI

6.5 University of Copenhagen

UNIVERSITY OF COPENHAGEN

Peter Bauer and Nils Wedi
ECMWF, Shinfield Park,
Reading, RG2 9AX, England

Letter of support and commitment.

30 AUGUST 2017

PROFESSOR EIGIL KAAS
NIELS BOHR INSTITUTE
JULIANE MARIES VEJ 30
DK-2100 COPENHAGEN
DENMARK

TLF. DIR. +45 35 32 05 14

kaas@gfy.ku.dk

REF: Eigil Kaas

Web: http://www.gfy.ku.dk/~kaas

Dear Peter and Nils

I am writing to express my strong support to the ESCAPE-2 proposal. ESCAPE-2 targets some of the most important obstacles in the next generations of numerical weather prediction systems, and in Earth system modelling in general.
The scientific methods outlined in the application are novel and represent the future state of art. ESCAPE-2 is an ambitious project addressing some of the most important challenges we are facing, and it does so via a relevant range of scientific, technical and (future) user oriented approaches.
I should like to offer my commitment to collaborate with ESCAPE-2 in two ways:

1) I will act as supervisor for a PhD student working on machine learning within ESCAPE-2. The focus will mainly be on (parts of) radiation schemes.
2) I am currently working on new mathematical methods enabling long stepping without the need for global elliptic solvers. In this way one can achieve high numerical stability for fast atmospheric modes using a numerical domain of dependence, which is dictated solely by the phase speed of these waves. This should improve scalability significantly.

I look forward very much to a fruitful collaboration.

Sincerely

Eigil Kaas,
Niels Bohr Institute,
University of Copenhagen

6.6 Freie Universitaet Berlin

Freie Universität Berlin

Fachbereich Mathematik
und Informatik
Institut für Mathematik

Professor Dr.-Ing. Rupert Klein
Arnimallee 6
14159 Berlin

Phone +49 (0) 30 - 838 75414
Fax +49 (0) 30 - 838 75412
E-Mail rupert.klein@math.fu-berlin.de
Internet www.math.fu-berlin.de/groups/fluid-dyn/

Freie Universität Berlin, Fachbereich Mathematik + Informatik
Institut für Mathematik, Arnimallee 6, 14195 Berlin

Dr. Peter Bauer
Deputy Director
Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX
UK

September 21, 2017

Letter of support for the H2020-FETHPC-2016-2017 project ESCAPE2

Dear Dr. Bauer (dear Peter),

thanks for informing me about the follow-up proposal to the H2020 project ESCAPE! All the project's main objectives find my full support. The projects will clearly spearhead the development of exascale-ready weather forecasting and climate research codes at the international level. As an applied mathematician, I welcome especially the project's ambitious goal of incorporating very recent top research results in computational science and numerical mathematics in accurate, practically useful high-performance code components, of course.

A particularly appealing key aspect of the project is, in my view, its generalization of the concept of modularity from its standard meaning in the IT and programming world to modularity at the levels of both scientific expertise and interdisciplinary cooperation. Let me be a bit more specific and provide a personal perspective as a university professor leading a relatively small applied mathematics research team: The continuing increase in complexity of modern high-performance computing (HPC) architectures has had drastic, often negative, ramifications for scientific knowledge transfer. Both the development time and the breadth of expertise needed to transfer fundamental research results in applied mathematical modelling and numerical methods to full-fledged production codes for real-life applications have drastically increased as a consequence of increasing hardware complexity. Obviously, this is not a desirable development, but it is also hard to counter, as modern science does rely increasingly on the available HPC power for a number of very good reasons.

The ESCAPE2 project, lifting ESCAPE to the next level, comes at exactly the right time in this respect as it suggests and implements a systematic way out of this dilemma in the context of weather and climate research, thereby providing a prototypical pathway also for other areas of science. The systematic separation of concerns that will be achieved by realization of the envisioned hardware-independent programming and code development paradigms will allow experts to refocus on their core competences while nevertheless contributing directly to highly complex HPC-ready modelling systems. In this way, successful completion of the ESCAPE projects would very substantially boost scientific progress in the weather and climate area while it will also positively stimulate and possibly jump-start related developments in other fields.

That stated, I truly look forward to continuing interactions with the ESCAPE/ESCAPE2 team both through my ECMWF fellowship and through cooperations with the Collaborative Research Center "Scaling Cascades in Complex Systems" which is funded by Deutsche Forschungsgemeinschaft, operated by Freie Universität Berlin, and of which I am the scientific coordinator (see http://sfb1114.mi.fu-berlin.de).

Berlin, September 21, 2017

Prof. Dr.-Ing. R. Klein

6.7 Alfred Wegener Institute

ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR-
UND MEERESFORSCHUNG

Alfred-Wegener-Institut, Postfach 12 01 61, 27515 Bremerhaven

To whom it may concern

Prof. Dr. Thomas Jung
Head of Climate Dynamics
Alfred Wegener Institute
Bussestraße 24
D-27570 Bremerhaven
Germany
T +49 471 4831 1761
F +49 471 4831 1797
thomas.jung@awi.de

21 September 2017
Letter of support for the ESCAPE-2 project

Dear Dr. Bauer,

it is my pleasure to express my *full* support for your ESCAPE-2 proposal in response to the call "Transition to Exascale Computing" (FETHPC-02-2017) withing the Horizon 2020 framework.

The transition to exascale computing constitutes a major challenge for the weather, climate and Earth system modelling communities that will require fundamentally new approaches in modelling along with novel cooperations across discplinies including partners from industry (e.g. HPC providers). Building on the success of the ESCAPE project, ESCAPE-2 promises major progress in this critical area by focussing on very promising "Mathematics and algorithms for extreme scales HPC system" and "High-productivity programming environments for exascale".

Given the scale of the challenges ahead in the emergent field of exascale modelling, international collaboration—as envisaged in ESCAPE-2—is absolutely key.

I hereby confirm my intention to participate in the ESCAPE-2 project through attendance of relevant project meetings, and by ensuring clustering with two major European and national projects—namely APPLICATE and Earth System Modelling project of Germany's Helmholtz Association, respectively—that will benefit substantially from the outcomes generated by ESCAPE-2.

Yours sincerely,

(Thomas Jung)

Alfred-Wegener-Institut
Helmholtz-Zentrum für
Polar- und Meeresforschung

BREMERHAVEN

Am Handelshafen 12
27570 Bremerhaven
Telefon 0471 4831-0
Telefax 0471 4831-1149
www.awi.de

Stiftung des öffentlichen Rechts

Sitz der Stiftung:
Am Handelshafen 12
27570 Bremerhaven
Telefon 0471 4831-0
Telefax 0471 4831-1149
www.awi.de

HELMHOLTZ
GEMEINSCHAFT

6.8 University of Oxford

Department of Physics

Atmospheric, Oceanic and Planetary Physics
Clarendon Laboratory, Parks Road, Oxford OX1 3PU

UNIVERSITY OF OXFORD

Tel: +44-(0)1865-272901 Fax: +44-(0)1865-272924 www.physics.ox.ac.uk

---

Dr Peter Bauer
Deputy Director, ECMWF
Shinfield Park
Reading, Berkshire
RG2 9AX

Ref: 001 2017 0831

31 August 2017

To whom it may concern

Dear Sir/Madam

I am writing in strong support of the Horizon 2020 proposal ESCAPE-2 (Energy-efficient Scalable Algorithms for Weather Prediction at Exascale) submitted in response to the Commission's call on Future and Emerging Technologies.

I am writing this as Hurricane Harvey devastates the Gulf Coasts of the USA, with enormous societal and global business implications. Two key questions have emerged: should the citizens of Houston have been evacuated before the Hurricane hit, and will these events become more commonplace in the future with climate change. Both of these questions require the development of a new generation of weather\climate models which can predict not only the development of the hurricane itself and associated rainfall bands, but also how the storm interacts with the larger-scale circulation patterns. In the case of Harvey, the key issue that led to devastation was the fact that these circulation patterns did not cause it to be swept inland after landfall – the normal situation. Instead it "stalled". Current generation climate models simply do not have enough resolution to be able to determine whether these "stalled" hurricanes will become more commonplace in the future. The new generation of weather/climate model, a key focus for ESCAPE-2, must approach convective-scale resolution and this will necessarily require extremely efficient integration on next-generation exascale computers.

Running such models efficiently on exascale computers present significant algorithmic challenges. The European Centre for Medium-Range Weather Forecasts (ECMWF) is at the vanguard of such research, and has successfully led a consortium composed of numerical weather prediction centres operating at both global and European regional scales, in ESCAPE. Here, this research will continue in ESCAPE-II, leading, by the end of the project, to the first generation of weather/climate models fit for purpose for the dawn of the era of the exascale. This work will be a vital element of an emerging FET Flagship Proposal, which seeks funding for dedicated exascale computing in order to predict extremes of weather and climate, like Hurricane Harvey, and the European heat waves this summer, with unprecedented accuracy and reliability. ESCAPE-II is a technical project whose societal and business implications are simply enormous, and touch almost every aspect of human life.

Yours faithfully

T.N.Palmer
Royal Society Research Professor in Climate Physics Professorial Fellow, Jesus College, Oxford.

6.9 CNRS/IPSL

Sylvie Joussaume
CNRS/Institut Pierre Simon Laplace
Orme des Merisiers Bâtiment 712
CE Saclay
91191 Gif sur Yvette cedex, FRANCE
Sylvie.joussaume@lsce.ipsl.fr

Letter of Support to ESCAPE-2

To: Peter Bauer, ECMWF

Saclay, September 22, 2017

Dear Dr Peter Bauer,

I hereby state the support of the climate modelling community gathered under the European Network for Earth System modelling (ENES) to the "*Energy-efficient Scalable Algorithms for weather and climate Prediction at Exascale*" (ESCAPE-2) proposal, aiming at preparing efficient weather and climate models for extreme computing.

The proposed work is highly relevant to ENES. As emphasized in the ENES infrastructure strategy (https://enes.org), preparing efficient climate models for the future exascale architectures is challenging, but necessary to be able to run models at very high resolution, as well as to address large ensembles and complexity. Moreover, the focus of ESCAPE-2 on the IFS, ICON and NEMO codes is fully relevant to the ENES community. The EC-Earth community in several European countries uses the IFS code, as well as Météo-France under the name ARPEGE. A large community in Germany and Switzerland shares the ICON code, and NEMO is the European ocean platform shared by 5 of the 7 European climate models. The ENES network will help spread the developments and experience gained from ESCAPE-2, in particular thanks to its HPC Task Force and its engagement in the ESIWACE Centre of Excellence.

The ENES community therefore strongly supports your ESCAPE-2 proposal and we look forward to working with you and your team in the near future.

Sincerely,

Sylvie JOUSSAUME
Chair of the ENES Board

6.10 Netherlands eScience Centre

netherlands eScience center

Date
19 September 2017

Our reference
NLeSC C 17.0157

Name of contact
Prof. Dr. Wilco Hazeleger

E-mail
w.hazeleger@esciencecenter.nl

Telephone
+31 20 460 4770

Subject
Support Letter

ECMWF
Attn. Peter Bauer

To whom this may concern,

With this letter I would like express my full support to the ESCAPE-2 project. There is an increasing need to bring computer and applied sciences together. Societal and scientific needs require highly accurate, detailed and skillful predictions of weather and climate. Novel computer architectures and innovative algorithmic developments allow major steps in weather and climate capabilities to be made. To do so, it is essential to bring these communities together to advance both the field of meteorology and climate research and to interact closely with the computer sciences. The programming and algorithmic challenges are shared among climate and weather research. This calls for a seamless approach in the technology development.

The Scalability Programme at ECMWF, which I fully support as well, aims to do this at ECMWF. The broader embedding of the Scalability Programme in European computer sciences and applications thereof will offer a tremendous opportunity for ECMWF and the European community. ECMWF is therefore ideally positioned to coordinate this project.

Best regards,

Prof. Dr. Wilco Hazeleger
Director/CEO Netherlands eScience Center
Chair Climate Dynamics Wageningen University

6.11 University of Helsinki

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Support letter for the ESCAPE-2 proposal

Dynamic meteorology at the University of Helsinki represents a small user group of high-performance computing in atmospheric sciences. Projects, like ESCAPE, are vital for us since our challenges are highly parallel with those expressed by ESCAPE:

> *... the need to simulate complex physical phenomena within tight production schedules ... existing extreme-scale application software ... is ill-equipped to adapt to the rapidly evolving hardware.*

These translate easily to constrains of Academic research: we need to be able to make best use of existing super-computing infrastructures we have little leverage to affect, and need to do so with limited technical and human resources, upon which our deliverables - MSc and PhD degrees - are completely dependent on. We see projects, like ESCAPE, vital because they help us to effectively use HPC resources, and produce useful and usable knowledge and software that is pervasive across the chain form universities and academia to weather services and world-lead weather and climate centres. These benefits have also been demonstrated through the ESCAPE Young Scientist Summer Schools (YSSS) – three PhD students from the Doctoral School for Atmospheric Sciences, University of Helsinki, participated YSSS at the University of Copenhagen, Denmark, 7 - 12 August 2017 with great enthusiasm.

I very much support the preparation of your ESCAPE-2 proposal, and hope the best of success for its evaluation process.

Helsinki, 18 September 2017

Heikki Järvinen
Professor of Meteorology
University of Helsinki

Fysiikan laitos
Matemaattis-luonnontieteellinen tiedekunta
PL 64 (Gustaf Hällströmin katu 2), 00014 Helsingin yliopisto
Puhelin (09) 191 50600, faksi (09) 191 50610, www.physics.helsinki.fi/

Institutionen för fysik
Matematisk-naturvetenskapliga fakulteten
PB 64 (Gustaf Hällströms gata 2), FI-00014 Helsingfors universitet
Telefon +358 9 191 50600, fax +358 9 191 50610, www.physics.helsinki.fi/

Department of Physics
Faculty of Science
P.O. Box 64 (Gustaf Hällströmin katu 2), FI-00014 University of Helsinki
Telephone +358 9 191 50600, fax +358 9 191 50610, www.physics.helsinki.fi/

6.12 Naval Postgraduate School

NAVAL POSTGRADUATE SCHOOL
Department of Applied Mathematics
833 Dyer Road, Bldg 232, SP-250
Monterey, CA 93943-5216

August 22, 2017

To the European Commission on Future and Emerging Technologies,

I am writing this letter to strongly support the ESCAPE-2 project being proposed by the European Center for Medium-Range Weather Forecasting (ECMWF). Specifically, I can offer my expert opinion regarding the first two goals of the ESCAPE-2 project; namely, (1) the use of high-order and efficient mathematical algorithms for the governing equations of motion and (2) the use of computer languages that are hardware independent.

Before I do so allow me to list my areas of expertise. I am an applied mathematician that has been working on the development of accurate and efficient mathematical algorithms (numerical methods) for solving systems of nonlinear partial differential equations. Although my area of expertise is in developing new numerical methods, for the past 20 years I have been applying these methods to weather prediction models as well as to ocean models for climate. I work in close collaboration with the U.S. Navy on their next-generation global and mesoscale weather prediction system (called NEPTUNE) that is derived from the numerical libraries we developed in my group (called GNuMe which houses our NUMA[1] atmospheric solver). In addition, I work with the U.S. Department of Energy on an ocean modeling component for possible use in the ACME[2] climate system. This ocean component we call NUMO[3] and is housed within the same GNuMe[4] software library. Let me now offer my opinion on why the first two goals of the ESCAPE-2 project are vitally important to the weather and climate communities worldwide.

(1) <u>Use of High-Order Numerics</u>
ECMWF is proposing to explore various numerical methods including high-order discontinuous Galerkin (DG) methods and semi-Lagrangian time-integrations methods.

---

[1] Information on the NUMA model can be found at http://frankgiraldo.wixsite.com/mysite/numa

[2] Information on ACME can be found at https://climatemodeling.science.energy.gov/projects/accelerated-climate-modeling-energy

[3] Information on NUMO can be found at http://frankgiraldo.wixsite.com/mysite/numo

[4] Information on GNuMe can be found at http://frankgiraldo.wixsite.com/mysite/gnume

ECMWF has led the global numerical weather prediction (NWP) community in constructing the most accurate and efficient weather models in the world. One of the key reasons for this is due to their use of high-order numeric (spectral transforms) and their semi-Lagrangian (SL) time-integration. The attraction of exploring methods such as DG is that while they are high-order, they only require a small communication stencil thereby making them more local in nature and, hence, more amenable for use in emerging hardware such as accelerators. Since I also have worked in both DG and SL methods, I welcome future collaborations with the members of the ESCAPE-2 project because any work developed in this project will benefit the global NWP community. DG methods offer a means to represent the spatial differential operators with high-fidelity, robustly, and with the potential to be exascale performant; however, they tend to be relatively expensive (from a floating point count perspective). In addition, SL methods offer the means to use extremely large time steps that may translate to a faster time-to-solution of the model. Constructing SL methods at scale and that can preserve the conservation properties of DG methods will be extremely important to tackle.

(2) Use of Hardware Independent Computer Languages

I sit on the U.S. Earth Systems Prediction Capability High-Performance Computing panel[5] and many in this group are worried about how we, as a community, will harness the promised capabilities of exascale systems that are due to arrive in the U.S. and Europe in the early 2020s. I also sit on the U.S. Department of Energy's ACME review panel and this is a concern that arises at all of our meetings. Therefore, we must tackle this issue as a community and the ESCAPE-2 project offers a way forward. That is, rather than worrying about what the hardware will look like, why not use computer languages that are hardware independent. Examples of such languages being considered in the U.S. include: Kokkos[6] and OCCA[7]. The key point here is to reconfigure our computer models just once regardless of how the hardware will change in the next 30 years. For example, if we write a computer program that is able to run on both Nvidia and AMD GPUs we could achieve this by either (i) writing two versions of the same code, one in CUDA and the other in OpenCL (of course OpenCL would work on both but the performance hit would be substantial). Or we could write (ii) one version of the code in a hardware independent language that we can then translate to either CUDA or OpenCL. If you take this concept further, then we could construct back-ends to any type of hardware. My group has worked on this idea[8] and I would be very happy to collaborate with the ESCAPE-2 members to offer my advice and to help them pursue the optimal path forward.

---

[5] Our position paper on the future of HPC can be found at https://repository.library.noaa.gov/view/noaa/14319

[6] Information on Kokkos can be found at https://github.com/kokkos

[7] Information on OCCA can be found at http://libocca.org/

[8] For results on porting our NUMA atmospheric model to hardware independent languages see https://arxiv.org/abs/1702.04316

In closing, I welcome the opportunity to collaborate with the ESCAPE-2 members. We have already begun a collaboration of sorts through a research article we have been writing on time-integration methods for numerical weather prediction. In addition, we are planning my visit to ECMWF for a few months. Our interests in both of the ESCAPE-2 goals are clearly aligned and I look forward to seeing this project funded. One of the most impressive aspects of ECMWF is that they work on difficult basic research problems but with the goal of making them practical so that they can improve both the accuracy and efficiency of their weather models – this is the perfect applied mathematics problem.

Sincerely,

Frank Giraldo 13" Macbook

Digitally signed by Frank Giraldo 13" Macbook
DN: cn=Frank Giraldo 13" Macbook, o=Naval Postgraduate School, ou=Department of Applied Mathematics, email=fxgirald@nps.edu, c=US
Date: 2017.08.24 07:53:36 -07'00'

Francis X. Giraldo
Professor of Applied Mathematics
Department of Applied Mathematics
Naval Postgraduate School
Monterey, CA 93943-5216
Voice Mail: 1-831-656-2293
Email: fxgirald@nps.edu
http://frankgiraldo.wix.com/mysite

6.13 Argonne National Laboratory

Paul C. Messina, Ph.D.

Director of DOE's Exascale Computing Project
Argonne Distinguished Fellow
Advisor of the ALD and Lab on Exascale
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240 Rm 4136
Argonne, IL 60439
Office: 630-252-3045
messina@anl.gov

Letter of Support

September 5, 2017

To: Dr. Peter Bauer
Deputy Director, Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, UK
Email: peter.bauer@ecmwf.int

Dear Dr. Bauer,

I hereby state my full support to the "ESCAPE-2" proposal, whose main goal is to facilitate Europe's transition to Exascale. By addressing the challenge of energy-efficient high-performance computing and exposing new mathematical and algorithmic approaches from a unified interface, ESCAPE-2 can help to address the needs of the climate & weather community. Therefore, I see great value in developing the proposed technology.

The proposed work is highly relevant; this assessment is based on my experience in High Performance Computing over the last four decades as well as my experience during the last two years as Director of the U.S. Exascale Computing Project, which includes development of exascale applications in the same domain that is the focus of your proposal.

My full support for the proposal is based on the ways in which you plan to address the challenge. The proposal identifies key challenges for an extremely important application -- climate prediction -- and ways to meet them on exascale systems. For example, it recognizes the importance of using mathematical methods and algorithms that will produce high-fidelity predictions (such as higher-order methods) as well as implementing them in ways that take into account the characteristics of exascale systems architectures. Portability of the software and increased productivity are also addressed through the use of domain specific languages, both of which are important considerations given the rapid changes in computer architectures. Additional valuable products that this project will yield are workflows that enable Uncertainty Quantification, benchmarks and proxy applications. By providing an open source software and data structure framework your project promises to have an extremely valuable impact on earth-system modelling beyond the life of this project).

If it would be of interest to you, I would be pleased to serve on your advisory board.

In summary, I strongly support your "ESCAPE-2" proposal and look forward to working with you and your team in the near future.

Sincerely,

Paul Messina

Paul Messina

A U.S. Department of Energy laboratory managed by UChicago Argonne, LLC

6.14 University of Tennessee

THE UNIVERSITY OF
TENNESSEE
KNOXVILLE

JACK DONGARRA
UNIVERSITY DISTINGUISHED PROFESSOR
865-974-8295
dongarra@icl.utk.edu

August 11, 2017

Dr. Peter Bauer
Deputy Director
Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, UK

Dear Peter,

I am a University Distinguished Professor of Computer Science in the Electrical Engineering and Computer Science Department at the University of Tennessee. I hold several leadership positions at international centers of learning, including Distinguished Researcher in the Computer Science and Mathematics Division at Oak Ridge National Laboratory (ORNL); Turing Fellow at Manchester University; and Adjunct Professor in the Computer Science Department at Rice University. I specialize in numerical algorithms for linear algebra, parallel scientific computing, advanced computer architectures, programming methodology, and tools for large scale supercomputers. I have made foundational contributions to several high-quality software packages that form the cornerstone of scalable parallel computing across academia and industry. These include: EISPACK, LINPACK, BLAS, LAPACK, ScaLAPACK, Netlib, PVM, MPI, NetSolve, Top500, HPCG, ATLAS, and PAPI. For these contributions, I have been awarded the IEEE Sidney Fernbach Award; the first IEEE Medal of Excellence in Scalable Computing; the SIAM Career Achievement Award; the IEEE Charles Babbage Award; and the ACM/IEEE Ken Kennedy Award. I am a Fellow of the AAAS, ACM, IEEE, and SIAM, and a foreign member of the Russian Academy of Science, and a member of the US National Academy of Engineering.

I have been involved in the development of mathematical software for the last 40 years and have a strong believe that when one analyzes what will be required to enable numerical software to achieve high performance on extreme computing systems, it becomes clear that creating them will take the coordinated efforts of a team that can simultaneously tackle research and design problems on three main fronts — algorithms, software, and management of parallelism. The weather and climate prediction community is facing the challenge that significant predictive skill enhancements are required to support decision making in response to climate change effects that will dramatically change our society.

1122 Volunteer Blvd., Suite 203 INNOVATIVE COMPUTING LABORATORY Knoxville, TN 37996

BIG ORANGE. BIG IDEAS.

These skill enhancements can only come from a substantial investment in computational science so that future prediction models can be run efficiently on future HPC hardware technologies that are currently undergoing disruptive change.

Under ECMWF's leadership, the ESCAPE project has already initiated steps towards a fundamental revision of numerical algorithm design focusing on energy efficient code adaptation to today's technologies. The ESCAPE concept of 'weather and climate dwarfs' has offered the community a way to disentangle individual sources of performance bottlenecks following the functional design of a forecast model. These dwarfs also serve as universal domain specific mini-application benchmarks – an approach that has been pioneered by myself with the High Performance Conjugate Gradient (HPCG) Benchmark. ESCAPE-2 will set new standards for the weather and climate prediction community as it aims at further elaborating the dwarf concept to enhanced levels of complexity with a strong focus on benchmarking readiness for future exascale systems with yet unknown architecture. Targeting those mathematical and numerical methods that offer the largest potential for parallelization and that maximize data locality is at the core of ESCAPE-2. The project's focus on the leading European models and its aim to design a weather and climate domain specific language appear to be the only ways to ensure code portability and maximize code performance at the same time.

In my view, a European investment in ESCAPE-2 is required to equip European's excellence in forecasting with the necessary expertise in computing to prepare for exascale systems. ECMWF's global recognition as the world-leader in numerical weather prediction and its ambition to enable the European community to exploit exascale facilities needs unprecedented support.

I have no hesitation in strongly supporting the ESCAPE-2 effort and hope to help with its success.

With Best Regards,

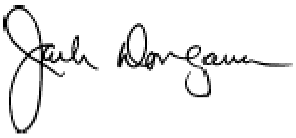

Jack Dongarra
University Distinguished Professor of Computer Science, University of Tennessee

6.15 Princeton University

Princeton University

CICS
Cooperative Institute
for Climate Science

Princeton University
NOAA/GFDL
201 Forrestal Road
Princeton NJ 08540
USA

30 August 2017

To: Dr. Peter Bauer
Deputy Director
Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, UK

Dear Dr. Bauer,

Your proposal entitled "ESCAPE-2: Energy-efficient Scalable Algorithms for weather and climate Prediction at Exascale" which you and Dr. Nils Wedi are submitting to the EU's *FETHPC-02-2017: Transition to Exascale* funding call is of great interest to us. This project will address many issues of common interest for exascale architectures, including exploration of domain-specific languages and novel methods of probabilistic predictive modeling and uncertainty quantification.

Here at NOAA/GFDL and Princeton University, we are undertaking work along similar lines. ECMWF and NOAA/GFDL have previously collaborated on common scientific and performance benchmarks under the NGGPS Phase I and NGGPS Phase II efforts to define the US's Next Generation Global Prediction System.

Should ESCAPE-2 be funded, which would be of great benefit to the entire community, Princeton University and NOAA/GFDL would like to continue common efforts to refine and contribute results to a set of common scientific and performance benchmarks for exascale-class weather and climate prediction models. We would also be willing serve on common steering committees or scientific advisory bodies, if requested. We look forward to being in touch with you at the commencement of your project.

V. Balaji
balaji@princeton.edu

Please do not hesitate to get in touch if you have any questions.

Yours sincerely,

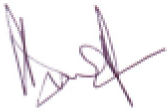

V. Balaji
Head, Modeling Systems Group
Cooperative Institute for Climate Science
Princeton University and NOAA/GFDL
Tel: +1-609-452-6516

6.16 NCAR

Computational & Information Systems Laboratory

P.O. Box 3000 ~ Boulder, CO 80307-3000 ~ www.cisl.ucar.edu

August 29, 2017

Dr. Peter Bauer
Deputy Director
Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, United Kingdom

Dear Dr. Bauer,

The purpose of this letter is to describe general areas of collaboration between the Computational and Information Systems Laboratory (CISL) at the National Center for Atmospheric Research (NCAR) located in Boulder, Colorado and the European Centre for Medium Range Weather Forecasts (ECMWF), pursuant to the proposal entitled Energy-efficient Scalable Algorithms for Weather Prediction at Exascale, Phase 2 (ESCAPE2).

After reviewing the ESCAPE-2 project summary, we are convinced the project has a high potential for advancing weather and climate prediction, given its focus on mathematics and algorithms, we have identified two potential areas of collaboration between NCAR/CISL and the proposed ESCAPE-2 project:

Machine learning NCAR is interested in working with ESCAPE-2 in exploring areas where machine learning techniques can replace conventional parameterization techniques. Historically, atmospheric physics parameterizations can be complex, laden with math intrinsics and branchy, and therefore difficult to optimize. However, computer architectures are now emerging that are designed to efficiently run deep learning algorithms as autoencoders that can substitute for these parameterizations.

Benchmarking NCAR would be happy to contribute reference benchmarks to ESCAPE-2. In particular, NCAR could provide two scalable atmospheric dynamical cores representative of global weather and climate modeling: the dry, non-hydrostatic dynamical core from the Model for Prediction Across Scales - Atmosphere (MPAS-A) based on finite volumes, and the High-Order Methods Modeling Environment (HOMME) hydrostatic, spectral element dynamical core, used in the atmospheric component of the Community Earth System Model (CESM). These dynamical cores are widely-used, highly scalable and efficient, and portable.

We understand that ECMWF is coordinating the ESCAPE-2 effort with a broad, international group of organizations, HPC centers and vendors. NCAR also recognizes ECMWF's leadership and indeed preeminence in numerical weather prediction. This is why we are delighted to be working with this elite team on the exascale challenges described above.

NCAR The National Center for Atmospheric Research is operated by the University Corporation for Atmospheric Research under sponsorship of the National Science Foundation.

Computational & Information Systems Laboratory

P.O. Box 3000 ~ Boulder, CO 80307-3000 ~ www.cisl.ucar.edu

We have had early engagement with the goals and objectives of exascale weather and climate modeling and have independently begun working along similar dimensions as ESCAPE-2. We look forward to engaging with you in these aspects of your proposal and participating in the discoveries and progress that will ensue from your work.

Sincerely,

Anke Kamrath

CISL Director
Associate NCAR Director

6.17 UCAR

CPAESS™ | Cooperative Programs for the Advancement of Earth System Science | University Corporation for Atmospheric Research (UCAR)

September 6, 2017

Dr. Peter Bauer
Deputy Director, Research Department
Dr. Nils Wedi
Head, Earth System Modeling Section
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, UK

Dear Drs. Bauer and Wedi,

I am writing to express enthusiastic support for your proposal to the European Union's Horizon 2020 research and innovation program: Energy-Efficient Scalable Algorithms for Weather and Climate Prediction at Exascale (ESCAPE-2), which builds on the successes of ECMWF and its partners under the current ESCAPE project.

As a UCAR/CPAESS visiting scientist with the U.S. Naval Research Laboratory (NRL) Marine Meteorology Division (MMD), I lead efforts to address HPC performance on advanced architectures for the development of NEPTUNE, the U.S. Navy's next-generation global weather prediction model. Before this, I was principal software architect of the widely used Weather Research and Forecast (WRF) model. My background and experience with HPC spans three decades at the Argonne National Laboratory, National Center for Atmospheric Research, National Renewable Energy Laboratory, and NOAA National Centers for Environmental Research. I am a co-author of the report from an ECMWF-lead workshop that formed the basis for ECMWF's Scalability Program. I represent NRL/MMD on the HPC working group of the Earth System Prediction Capability (ESPC) project, a multi-agency effort in the U.S. that shares ESCAPE-2's objective of engaging in co-design activities with the HPC industry.

We know that as HPC verges both towards exascale (10^{18} floating point operations per second) performance of our current weather and climate applications will bottleneck even more severely on requirements for massive parallelism, high data locality, minimal memory footprint and balanced memory system bandwidth. In light of these challenges, NRL is developing the NEPTUNE model that features an entirely reformulated dynamical core based on spectral elements to be more scalable than its spectral-transform predecessor and more computationally intense than a purely finite volume formulation. We expect NEPTUNE will provide a more accurate higher order solution while running more efficiently on next-generation computer architectures where data movement is a key bottleneck.

Your new ESCAPE-2 proposal correctly identifies co-design and the need for international cooperation in a concerted effort to characterize and refactor our applications while collectively enunciating our requirements for exascale architectures to HPC research and industry. The reformulation of NEPTUNE's numerics combined with our ongoing efforts to understand and improve computational, memory system, and I/O performance of both isolated kernels and overall application performance is directly in line with the goals and activities of the current

P.O. Box 3000 | Boulder, Colorado 80307 | Phone 303.497.8666 | cpaess.ucar.edu

CPAESS™
Cooperative Programs for the Advancement of Earth System Science
University Corporation for Atmospheric Research (UCAR)

ESCAPE project and provides an ideal framework for cooperative engagement with your ESCAPE-2 proposal.

As part of my work at NRL/MMD and the U.S. ESPC, I feel there is a compelling need to establish strong links to international efforts such as ESCAPE-2. I wholeheartedly endorse and support the proposal and look forward to finding avenues for substantive cooperation and collaboration.

Best regards,

John Michalakes
UCAR/CPAESS Visiting Scientist
NRL Marine Meteorology Division

P.O. Box 3000 | Boulder, Colorado 80307 | Phone 303.497.8666 | cpaess.ucar.edu

6.18 WMO

Ref.: 32996/2017-1.3 RES-ARE

WMO OMM

World Meteorological Organization
Organisation météorologique mondiale
Organización Meteorológica Mundial
Всемирная метеорологическая организация
المنظمة العالمية للأرصاد الجوية
世界气象组织

Secrétariat
7 bis, avenue de la Paix – Case postale 2300
CH 1211 Genève 2 – Suisse
Tél.: +41 (0) 22 730 81 11
Fax: +41 (0) 22 730 81 81
wmo@wmo.int – public.wmo.int

FAX

From:	Paolo Ruti	To:	Dr Peter Bauer Deputy Director of Research Department European Centre for Medium-Range Weather Forecasts
Date:	5 September 2017	City:	Reading
Our ref.:	32996/2017-RES/ARE	Country:	United Kingdom
No pages:	1	E-mail:	peter.bauer@ecmwf.int

Subject: "**E**nergy-efficient **Sc**albale **A**lgorithms for weather and climate **P**rediction at **E**xascale" (ESCAPE-2)

Dear Dr Bauer,

I am writing on behalf of the World Weather Research Programme (WWRP) in support of the proposal "**E**nergy-efficient **Sc**albale **A**lgorithms for weather and climate **P**rediction at **E**xascale" (ESCAPE-2) submitted by the European Centre for Medium-Range Weather Forecasts (ECMWF).

WWRP is very pleased to endorse this ESCAPE-2 project because its objectives and expected outcomes align very nicely with activities conducted at the intersection between weather and climate science priorities, infrastructures and communities. ESCAPE-2 is expected to make a substantial contribution to the seamless prediction framework in further developing the numerical background and the infrastructure to support Earth System modelling for the benefit of both weather and climate communities.

WMO expects substantial progress and benefit from ESCAPE-2 on his research initiatives under WWRP (polar prediction, high-impact weather prediction, sub-seasonal to seasonal prediction) and on future development of downstream services for all his members. ESCAPE-2 proposes an innovative approach to the challenge of energy-efficient high-performance computing and will provide the key components for weather and climate model benchmarks to be deployed on extreme-scale demonstrators and to facilitate the transition to Exascale.

Looking forward to our fruitful collaboration.

Yours sincerely,

(P. Ruti)
Chief
World Weather Research Division

6.19 Global Scientific Information and Computing Center

GSIC TOKYO INSTITUTE OF TECHNOLOGY
GLOBAL SCIENTIFIC INFORMATION AND COMPUTING CENTER

Letter of Support

To: Dr. Peter Bauer
Deputy Director
Research Department
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading RG2 9AX, UK
Email: peter.bauer@ecmwf.int

Sep. 1, 2017

Dear Dr. Bauer,

I hereby state my full support to the "ESCAPE-2" proposal, whose main goal is to facilitate Europe's transition to Exascale,
By addressing the challenge of energy-efficient high-performance computing and exposing new mathematical and algorithmic approaches from a unified interface, ESCAPE-2 can help to address the needs of the climate & weather community. Therefore, I see great value in developing the proposed technology for Europe and worldwide.

The proposed work is highly relevant, as accurate, high-resolution weather prediction is tantamount to our future societal stability to prevent disasters, as well as to foster economical growth. Natural disasters through adverse weather conditions still are posing significant problems as witnessed by recent hurricanes that devastated Texas in the US. In Japan advanced weather modeling and prediction towards exascale is one of the 9 strategic application areas for the Post-K next generation national flagship supercomputer, and our research group as well as our TSUBAME supercomputers have been playing collaborator roles in various projects with regards to the objective. Based on such experiences I truly feel that ESCAPE-2 to be one of the most important numerical weather prediction project to date, in that it will attempt to effectively unify the European efforts for the exascale objective, involving all the key institutions.

In this regard I would also be happy to serve on your advisory board if it would meet your interest.

I strongly support your "ESCAPE-2" proposal and I look forward to working with you and your team in the near future.

Best Regards,
Satoshi Matsuoka
Professor / Division Director of HPC Infrastructure
Global Scientific Information and Computing Center / Dept. Mathematical and Computing Sciences
Tokyo Institute of Technology
Fellow, AI Research Center, Japan Advanced Institute for Science and Technology (AIST).

European Commission

PARTICIPANT PORTAL E-RECEIPT

Digitally signed by sealingservice.grants.ec.europa.eu
Date: 2017.09.26 17:48:08 CEST
Reason: Acknowledgment of Receipt

This electronic receipt is a digitally signed version of the document submitted by your organisation. Both the content of the document and a set of metadata have been digitally sealed.

This digital signature mechanism, using a public-private key pair mechanism, uniquely binds this eReceipt to the modules of the Participant Portal of the European Commission, to the transaction for which it was generated and ensures its full integrity. Therefore a complete digitally signed trail of the transaction is available both for your organisation and for the issuer of the eReceipt.

Any attempt to modify the content will lead to a break of the integrity of the electronic signature, which can be verified at any time by clicking on the eReceipt validation symbol.

More info about eReceipts can be found in the FAQ page of the Participant Portal. (http://ec.europa.eu/research/participants/portal/page/faq)