



Funded by the European Union

Co-ordinated by  ECMWF



Energy efficient
Scalable
Algorithms for
weather
Prediction at
Exascale

Peter Bauer



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627



Forecast production workflow: what is the challenge?

Funded by the European Union



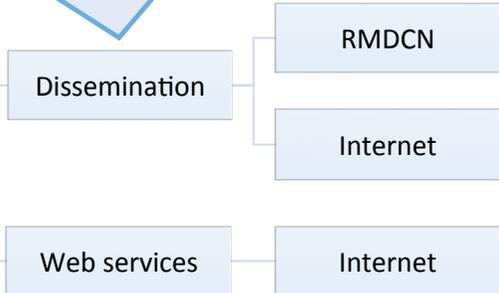
Data acquisition



Forecast run

Product generation

today: 10 TB
tomorrow: 300 TB
→ f=30 per day in critical path



today: 60 million
tomorrow: 600 million
→ f=10 per day

today: 10 million
100 vert
10 progr
50 enser

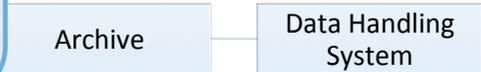
tomorrow: 500 milli
200 vert
100 progr
100 enser

today: 20 TB written
20 million fields
85 million products

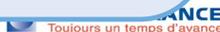
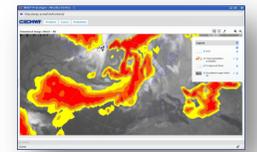
tomorrow: 500 TB written
500 million fields
2 billion products

→ f=25 per day in critical path

today: 100 TB
tomorrow: 10 PB
→ f=100 per day



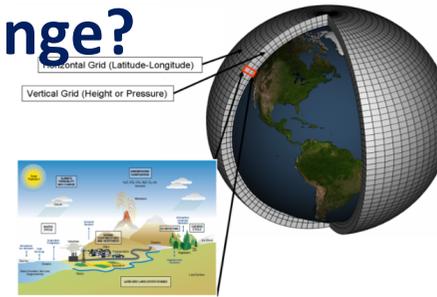
→ f=2000 per time step





What is the challenge?

Funded by the European Union



Today:

	Observations	Models
Volume	20 million = 2×10^7	5 million grid points 100 levels 10 prognostic variables = 5×10^9
Type	98% from 60 different satellite instruments	physical parameters of atmosphere, waves, ocean

Tomorrow:

	Observations	Models
Volume	200 million = 2×10^8	500 million grid points 200 levels 100 prognostic variables = 1×10^{13}
Type	98% from 80 different satellite instruments	physical and chemical parameters of atmosphere, waves, ocean, ice, vegetation

→ **Factor 10 per day**

→ **Factor 2000 per time step**

(10-day forecast today = 1440 time steps, but more time steps with increased resolution)





Funded by the European Union



Traditional science workflow

Mathematical description

$$\text{Wind } \rho \dot{\mathbf{v}} = -\nabla p + \rho \mathbf{g} - 2\boldsymbol{\Omega} \times (\rho \mathbf{v}) + \mathbf{F}$$

$$\text{Pressure } \dot{p} = -(c_{pd}/c_{vd}) p \nabla \cdot \mathbf{v} + (c_{pd}/c_{vd} - 1) Q_h$$

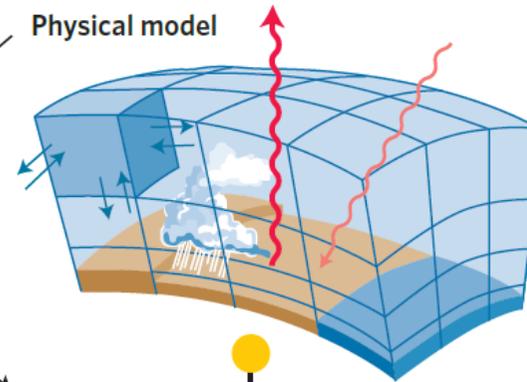
$$\text{Temperature } \rho c_{pd} \dot{T} = \dot{p} + Q_h$$

$$\text{Water } \rho \dot{q}^v = -\nabla \cdot \mathbf{F}^v - (I^l + I^f)$$

$$\rho \dot{q}^{lf} = \nabla \cdot (\mathbf{P}^{lf} + \mathbf{F}^{lf}) + I^{lf}$$

$$\text{Density } \rho = p [R_d (1 + (R_v/R_d - 1) q^v - q^l - q^f) T]^{-1}$$

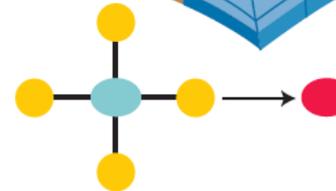
Physical model



Domain science and applied mathematics

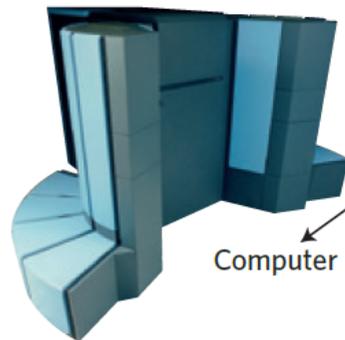
```
lap(i, j, k) = -4.0 * data(i, j, k) +
              data(i+1, j, k) + data(i-1, j, k) +
              data(i, j+1, k) + data(i, j-1, k);
```

Algorithmic description



Imperative code

Computer engineering



Compilation

Computer

[Schulthess 2015]

NATURE PHYSICS





Future science workflow

Funded by the European Union



Mathematical description

Wind $\rho \dot{\mathbf{v}} = -\nabla p + \rho \mathbf{g} - 2\Omega \times (\rho \mathbf{v}) + \mathbf{F}$

Pressure $\dot{p} = -(c_{pd}/c_{vd}) \rho \nabla \cdot \mathbf{v} + (c_{pd}/c_{vd}-1) Q_h$

Temperature $\rho c_{pd} \dot{T} = \dot{p} + Q_h$

Water $\rho \dot{q}^v = -\nabla \cdot \mathbf{F}^v - (I^l + I^f)$

$\rho \dot{q}^{lf} = \nabla \cdot (\mathbf{P}^{lf} + \mathbf{F}^{lf}) + I^{lf}$

Density $\rho = \rho [R_d (1 + (R_v/R_d - 1) q^v - q^l - q^f) T]^{-1}$

Physical model

Science applications using a descriptive and dynamic developer environment

Algorithmic description

Multidisciplinary co-design of tools, libraries, programming environment

Imperative code

Compiler front-end

Optimization / low-level libraries / runtime

Architecture specific back-ends

Architecture 1

Architecture 2

Architecture N

↑ science specific code
↓ generic code

**Energy efficient
SCalable Algorithms for
weather Prediction at
Exascale**

→ www.hpc-escape.eu

This project is funded by the European Union

[Schulthess 2015]

NATURE PHYSICS

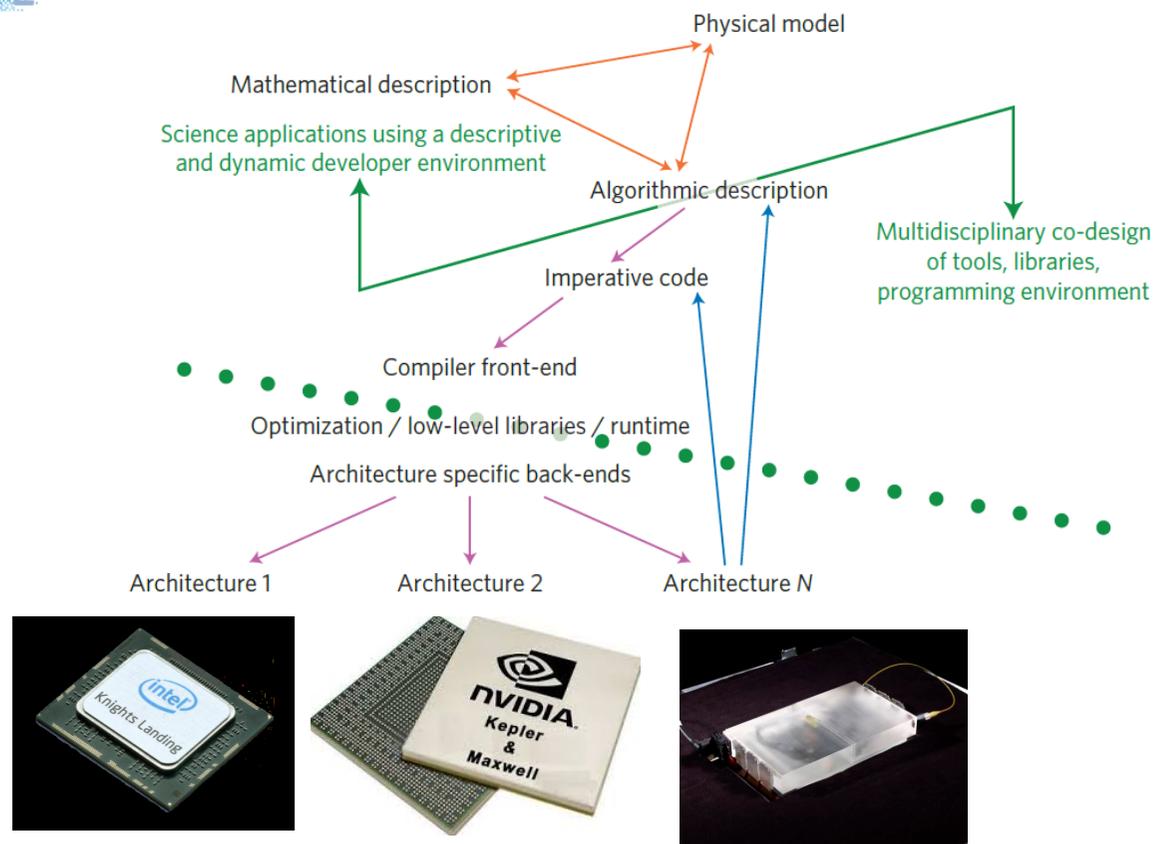




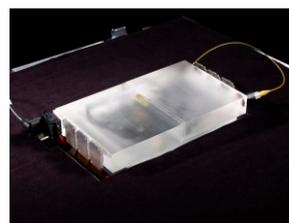
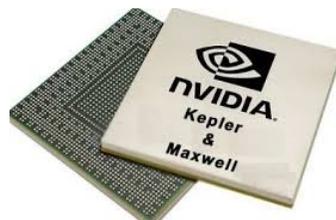
Funded by the European Union



Energy efficiency



.. aiming at minimizing Watts per forecast





ESCAPE key objectives

Funded by the
European Union



- Define fundamental **algorithm building blocks** (“*Weather & Climate Dwarfs*”) to co-design, advance, benchmark and efficiently run the next generation of NWP and climate models on energy-efficient, heterogeneous HPC architectures.
- Combine frontier research on **algorithm development** and extreme-scale, high-performance computing applications with **novel hardware technology**, to create a flexible and sustainable weather and climate prediction system.
- Foster the **future design of Earth-system models** and commercialisation of weather-dependent innovative products and services in Europe through enabling open-source technology.
- Pairing **world-leading NWP** with **innovative HPC solutions**.





Funded by the European Union



ESCAPE

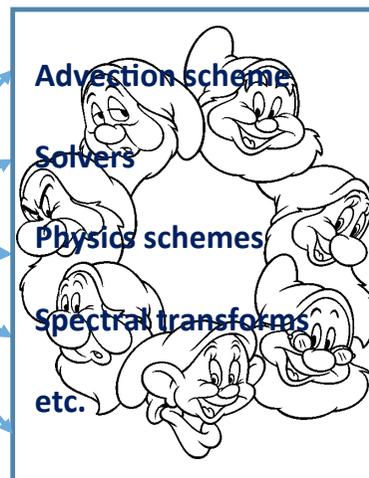
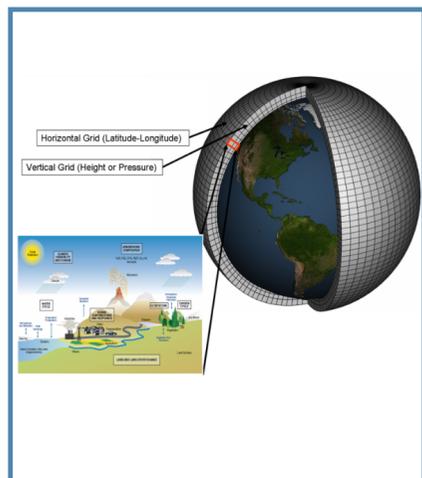
Energy efficient Scalable Algorithms for weather Prediction at Exascale

Disassemble ...
NWP models

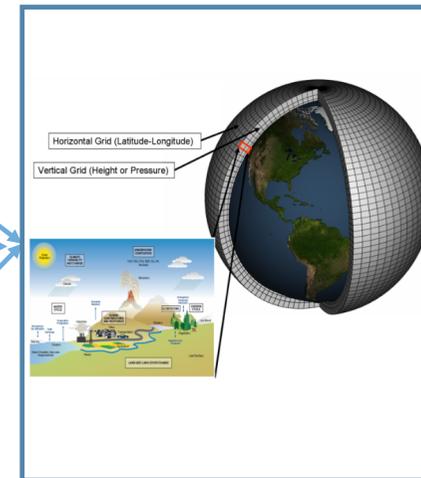
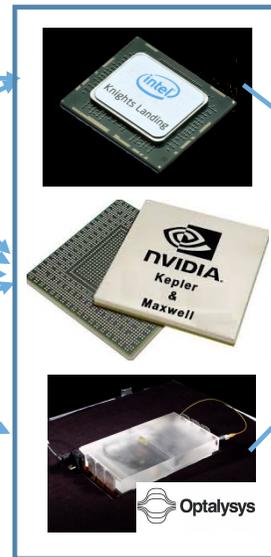
Extract, redesign...
key components

Optimize for energy...
efficiency on new hardware

... Reassemble
NWP models



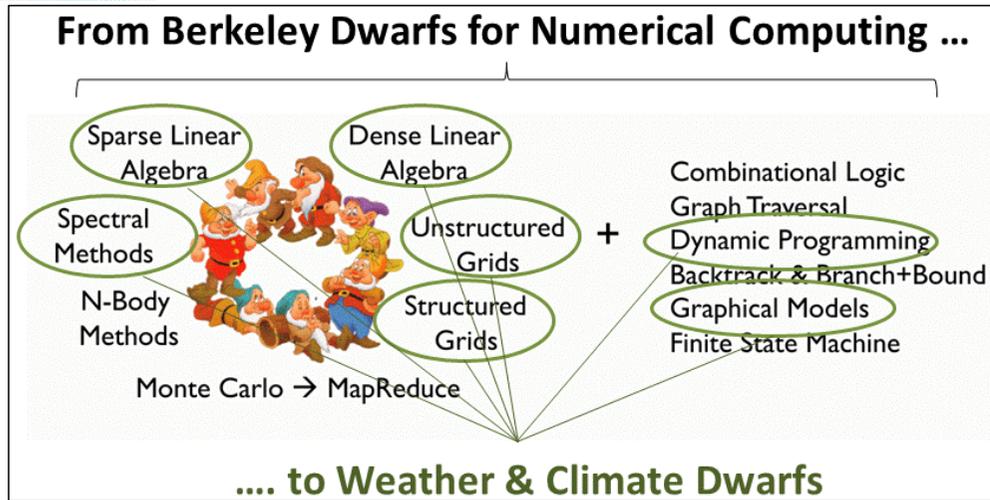
(memory, communication,
compute intensity)





What is a dwarf?

Funded by the European Union



A dwarf encapsulates a relevant characteristic or required functionality of an weather/climate prediction model and it is meant to be a **runnable** and **verifiable mini-application**

Candidates:

- Spectral transforms (FT/LT and bi-FT): very memory and communication bandwidth intensive, possibly limited scalability
- 2 & 3-dimensional elliptic solver: new, compute and communication latency intensive, possibly limited scalability
- Semi-Lagrangian advection: communication intensive, possibly limited scalability
- Cloud physics parameterization: expensive computation, scalable
- Radiation parameterization: expensive computation, scalable

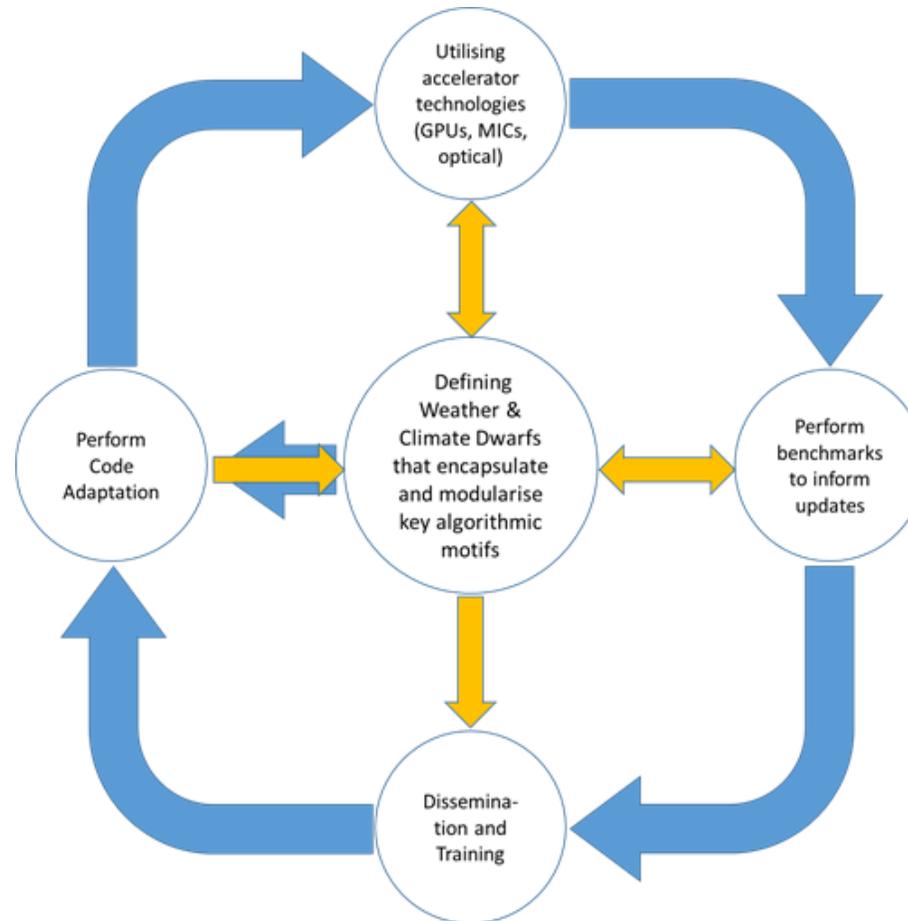
... more to follow





ESCAPE work flow

Funded by the European Union





ESCAPE European impact map

Funded by the European Union



34 countries

- Member States
- Co-operating States
- Under negotiation

7 countries

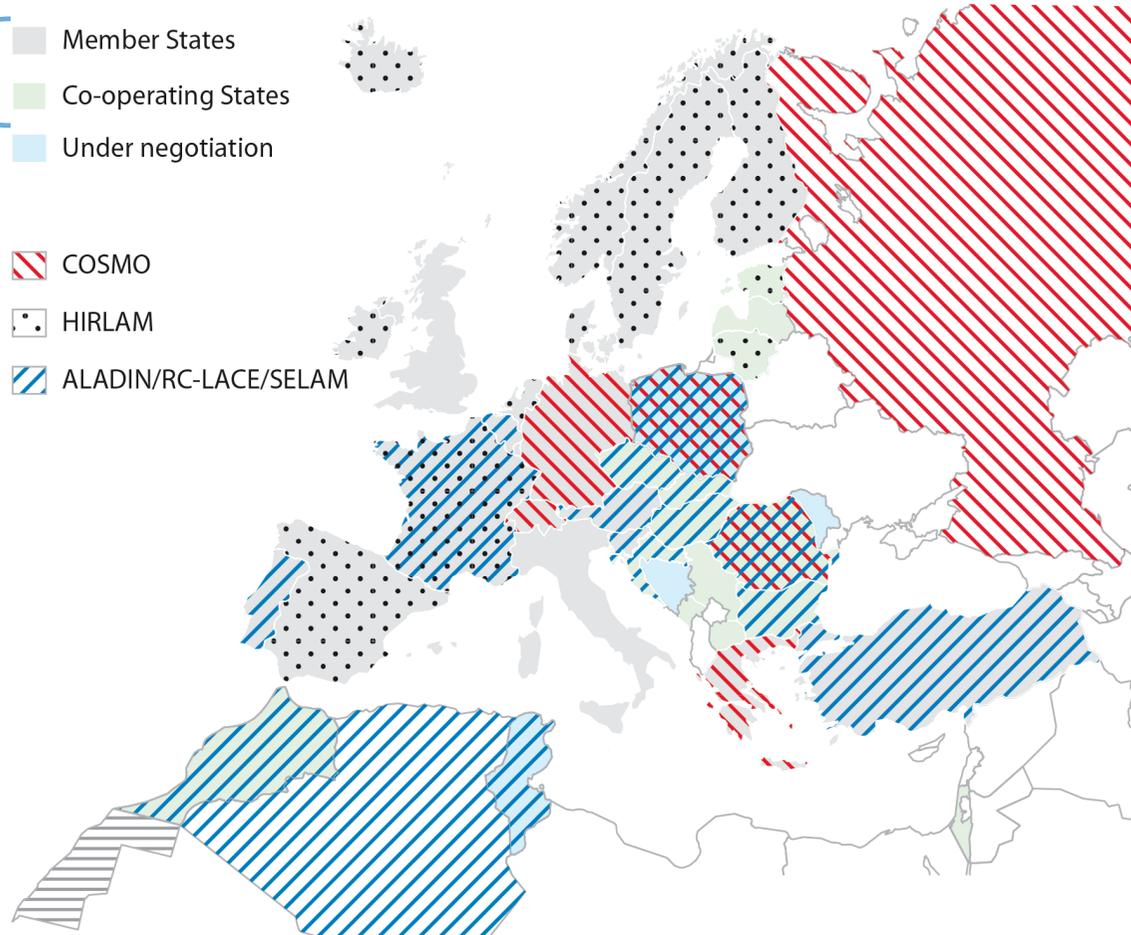
COSMO

11 countries

HIRLAM

16 countries

ALADIN/RC-LACE/SELAM

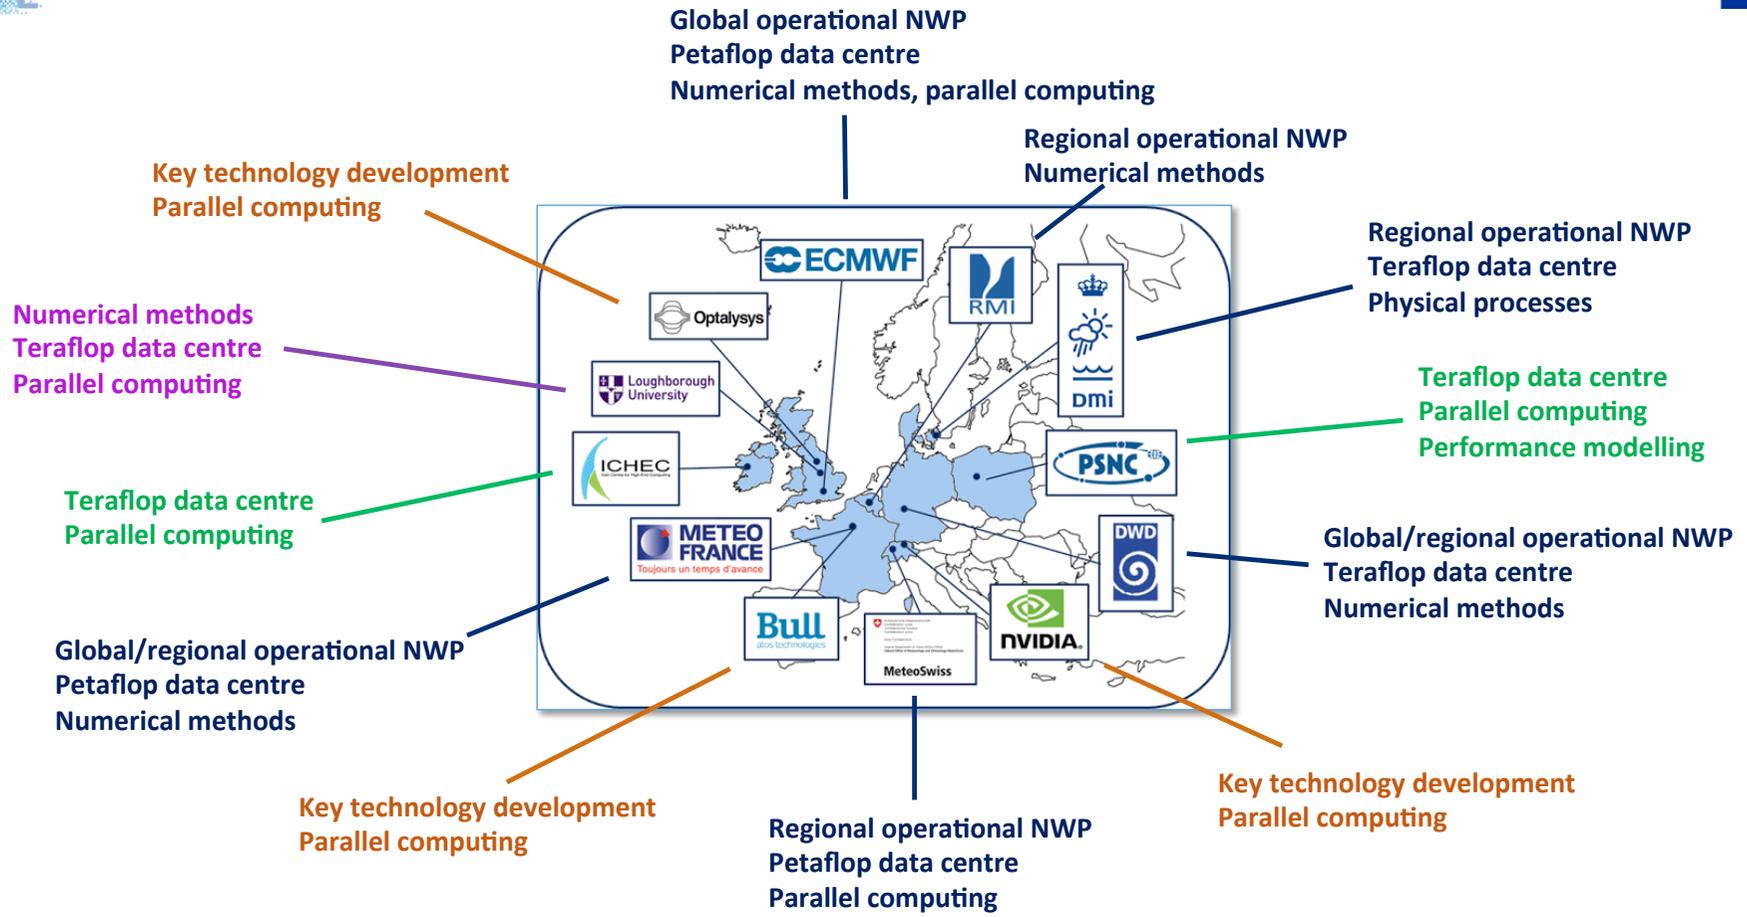




Funded by the European Union



ESCAPE partners & expertise





http://www.hpc-escape.eu

Funded by the European Union




The screenshot shows the ESCAPE website homepage. At the top, there is a navigation bar with a search box and a 'Search this site' button. The main header features the word 'ESCAPE' in large, stylized letters, with a background image of a city skyline and a weather map. Below the header, there is a 'NAVIGATION' menu on the left, including links for 'ESCAPE', 'PROJECT BACKGROUND', 'PROJECT OBJECTIVES', 'PROJECT IMPACT', 'PROJECT STRUCTURE', 'PROJECT PARTNERS', 'INTERACTIVE PORTAL', 'MEDIA HUB', 'ESCAPE NEWS', 'ESCAPE EVENTS', 'ESCAPE PUBLICATIONS', 'CONTACT', and 'SITEMAP'. The main content area is titled 'ESCAPE' and contains a paragraph describing the project's goals: '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. ESCAPE addresses the ETP4HPC Strategic Research Agenda 'Energy and resiliency' priority topic, developing a holistic understanding of energy-efficiency for extreme-scale applications using heterogeneous architectures, accelerators and special compute units. The three key reasons why this project will provide the necessary means to take a huge step forward in weather and climate modelling as well as interdisciplinary research on energy-efficient high-performance computing are:

- Defining and encapsulating the fundamental algorithmic building blocks ('Weather & Climate Dwarfs') underlying weather and climate services. This is the prerequisite for any subsequent co-design, optimization, and adaptation efforts.
- Combining ground-breaking frontier research on algorithm development for use in extreme-scale, high-performance computing applications, minimizing time- and cost-to-solution.
- Synthesizing the complementary skills of all project partners. ECMWF and leading European regional forecasting consortia are teaming up with excellent university research and experienced high-performance computing centres, two world-leading hardware companies, and one European start-up SME, providing entirely new knowledge and technology to the field.

ECMWF's partners in the project are Denmark's Meteorologiske Institut; Deutscher Wetterdienst; Institut Royal Météorologique de Belgique; Météo-France; MeteoSchweiz; Instytut Chemii Bioorganicznej Polskiej Akademii Nauk; Loughborough University; National University of Ireland, Galway; Bull SAS; NVIDIA Corporation; and Optalysys Ltd.

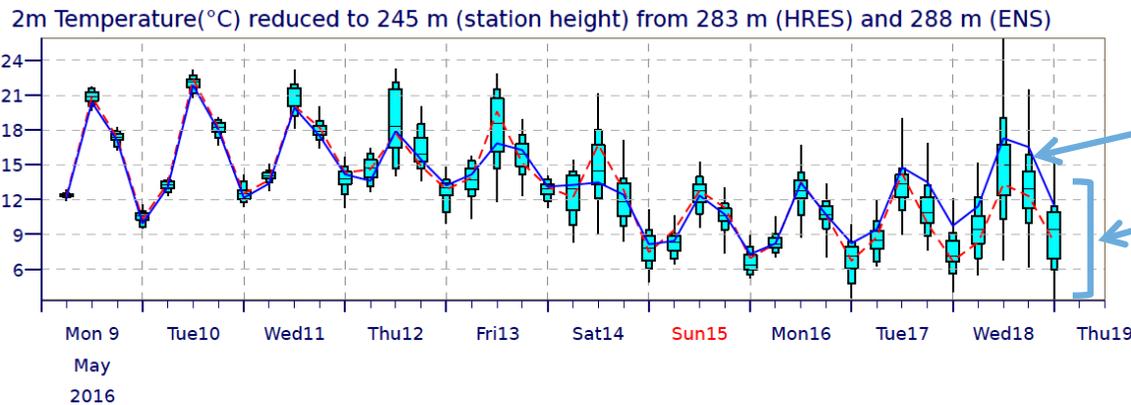
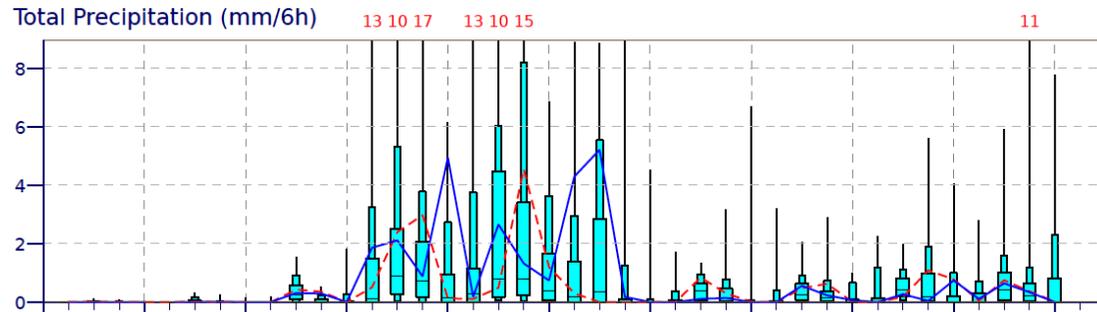
On the right side of the page, there is a 'News and Events' section. It features a 'NEW FORECAST MODEL AT ECMWF' article dated March 2016, which mentions that ECMWF introduced a new version of its operational model that provides much enhanced horizontal resolution in both forecasts and analyses. Below this is an 'INVITATION AND CALL FOR CONTRIBUTIONS FOR 1ST ESCAPE DISSEMINATION WORKSHOP' article, which states that the ESCAPE project is organizing its 1st Dissemination and Training Workshop from the 18th to the 20th October in Elsinore, Denmark.





And finally, the weather forecast

Funded by the European Union



Operational forecasts:

- 1x 10km forecast run on 352 nodes in 60'
- 51x 20km forecasts run on 1632 nodes in 74'
- Post-processing run on 500 nodes in 30'

Total about 2,500 nodes
(24 Ivybridge cores per node on Cray XC-30)





Software collaboration platform

Funded by the European Union



ESCAPE Public Website
http://www.hpc-escape.eu

<https://software.ecmwf.int/wiki/display/ESCAPE>

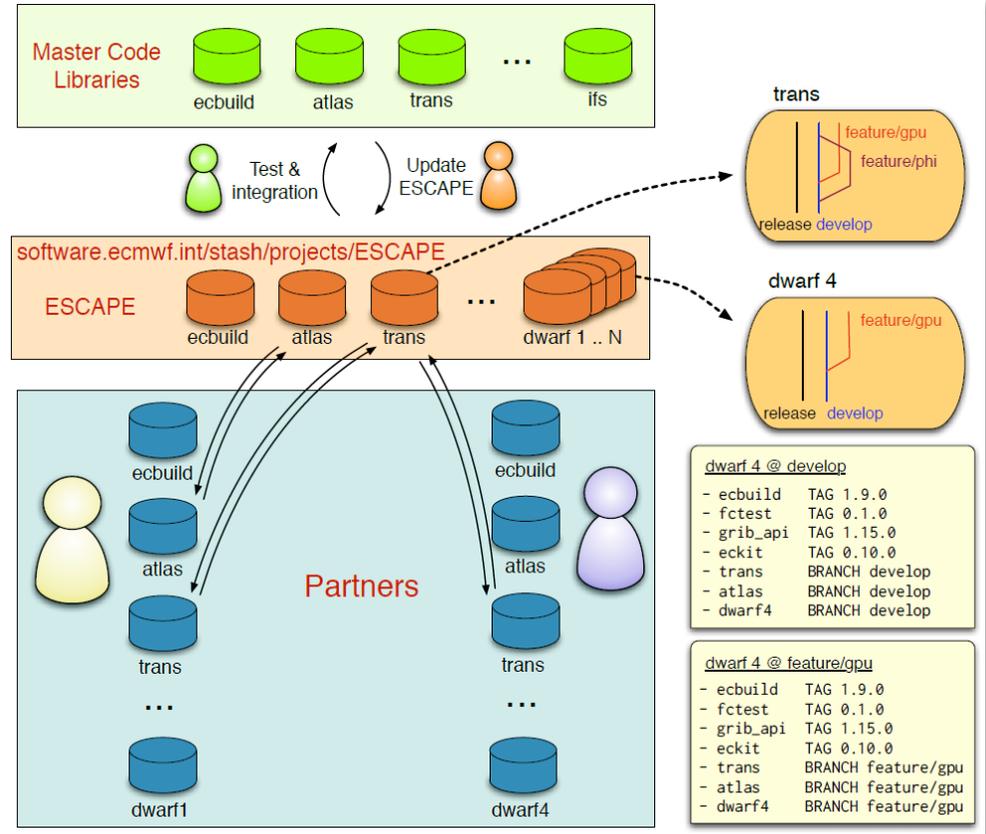
Confluence
- Wiki
- Howto
- Blog

Stash
- Software repositories
- git hosting

JIRA
- Bug submission
- Issue tracking
- Feature requests

Bamboo
- Compilation with several compilers
- Unit testing (small)
- Continuous integration

ESCAPE





Software stack

Funded by the European Union

