Newsletter 07/2019

To subscribe to/unsubscribe from this newsletter, please visit:
https://lists.dkrz.de/mailman/listinfo/esiwace-external
You find further details on ESiWACE at www.esiwace.eu.
All newsletters are also available at www.esiwace.eu/newsletter.
Deliverables and milestone documents referred to in the newsletters can be downloaded from Zenodo: (https://zenodo.org/communities/esiwace).

Contents

Dissemination and Upcoming Events: Meet Us!

ESiWACE/ESiWACE2 will participate in

- Euro-Par 2019, 26-30 Aug 2019, Göttingen/DE, https://www.europar.org/
- Workshop on Machine Learning for Weather and Climate Models, 2-5 Sep 2019, Oxford/UK, http://users.ox.ac.uk/~phys0895/mlwc2019/index.html
- ParCo Symposium: Tools and Infrastructure for Reproducibility in Data-Intensive Applications, 10 Sep 2019, Prague/CZ, https://sites.google.com/cmcc.it/parco2019-reprod-symposium/home
- International Computing for Atmospheric Sciences Symposium (iCAS), 8-12 Sep 2019, Stresa/Italy, https://www2.cisl.ucar.edu/events/conferences/icas/2019
- eScience 2019, 24-27 Sep 2019, San Diego/USA, https://escience2019.sdsc.edu
- ECMWF Workflows Workshop, 14-16 Oct 2019, Reading/UK, https://www.ecmwf.int/en/learning/workshops/building-reproducible-workflows
- ESCAPE-2 1st Dissemination Workshop, 21-22 Oct 2019, Reading/UK, http://www.hpc-escape2.eu/site/hpcescape2/media-hub/escape2-events/1st-dissemination-workshops
- Docker Container Hackathon for ESiWACE2 models, 3-5 Dec 2019, Lugano/CH

ESiWACE @ Teratec, PASC and ISC HPC

Figure 1: Rupert Ford talking about ESiWACE2 work at PASC

As usual, summer started with several computational science and HPC conferences in June.

First in a row, the Teratec forum was held 11-12 June in Palaiseau with ESiWACE participating in the European Research Café. The PASC conference took place in Zurich, 12-14 June. Five minisymposia were organized under ESiWACE participation and featured talks on domain-specific languages, the data challenge of weather and climate modelling at exascale and on global high-resolution modelling. Worthy a particular mention, the PASC Best Paper Award was given to Sam Hatfield and colleagues (including ESiWACE contributors) for their work on "Accelerating High-Resolution Weather Models with Deep-Learning Hardware", in which the use of half-, single- and double-precision is investigated and sheds light on the potential of emerging AI hardware—such as the NVIDIA Tensor Core—for the modelling community. More information on the work can be found here – congratulations!

At ISC HPC which took place 16-21 June in Frankfurt, ESiWACE was represented in several ways. Being an elementary building block of this year's DKRZ booth, visitors were informed about ESiWACE and global high-resolution developments. This was complemented by talks on weather and climate predictions at exascale and ESiWACE at the booth of the Gauss Allianz and in a session on challenges and visions for future exascale systems, as well as by a ESiWACE poster in the project poster track.

Figure 2: Left: ESiWACE at the DKRZ booth. Middle: DKRZ booth sign. Right: Philipp Neumann giving an overview of ESiWACE

2nd DYAMOND-ESiWACE Hackathon

A second hackathon to explore and experiment with the high-resolution data from the DYAMOND runs took place in Mainz, 18-21 June. The first day introduced the fifty international participants—coming from Europe, USA and Japan—to DYAMOND and the models involved as well as to the supporting ESiWACE infrastructure project and related activities. Afterwards groups were formed according to the participants' interests; the following 24h were devoted to hacking and, thus, carrying out a multitude of investigations on clouds and convection, tropical cyclones, extreme precipitation, the vertical structure of the general circulation and spectra of vertical velocities and the underlying topography. Several of the

results achieved in this short period of time will feed into publications in the future, amongst others into contributions for a special issue on the DYAMOND runs; for the corresponding call, click here.
Besides, a continuation of the DYAMOND project was discussed and plans for the next phase of DYAMOND have evolved. Updates on the next phase of DYAMOND will be published soon on the project website.

Figure 3: Participants of the second DYAMOND-ESiWACE Hackathon

An Update on the Earth-System Data Middleware

We had been improving the Earth-System Data Middleware (ESDM) over this year by integrating a prototype for NetCDF and MPI calls. The evaluation of the performance for large benchmark runs was conducted. These runs demonstrate the capability to enable executions of a high-resolution model which is the primary focus of ESiWACE.

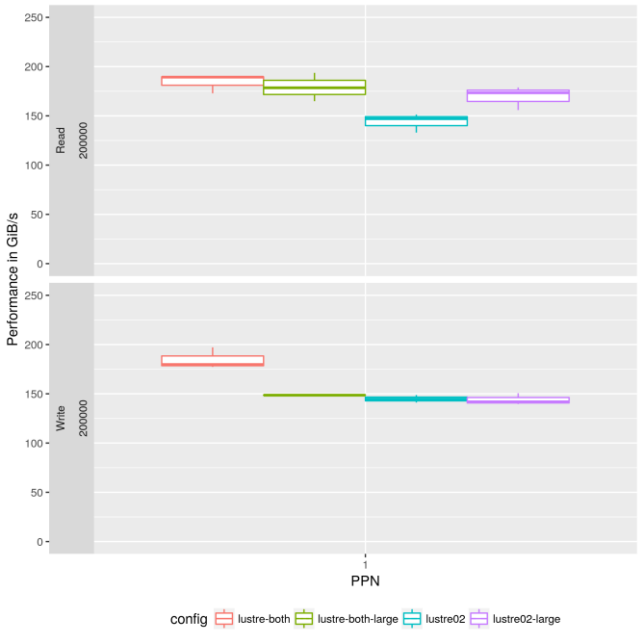

Figure 4: Performance of the current ESDM version

Figure 4 shows the performance of the current ESDM version running on Mistral (DKRZ) with 500 nodes, one process per node, and ten time steps of 300 GB data each. Data are clustered in boxplot summaries for *read* and *write* processes considering three runs. Mistral has two file systems (Lustre01 and Lustre02) and four configurations were tested: storing data only in Lustre02, settings where data are stored on both Lustre file systems concurrently (both), and an environment where data are fragmented into large chunks (both-large).
By utilising the two file systems resembling a heterogeneous environment effectively, we can improve the performance further by 20%. Note that the performance achieved on a single file system is similar to the best-case performance of file system benchmarks achieved with optimal settings (not shown in Fig. 4).
ESDM builds upon a data model similar to NetCDF and utilises a self-describing on-disk data format for storing structured data. We aim to deliver the NetCDF integrated version by the end of the first phase of

ESiWACE. This improvement can then be used as a drop-in replacement for typical use cases without changing anything from the application perspective. While our current version utilises the manual configuration by data-center experts, the ultimate long-term goal is to employ machine learning to automatise the decision making and reduce the burden for users and experts.

Calls for ESiWACE2 Support Services in Preparation

An essential building block of the ESiWACE2 infrastructure is given by services that are meant to support the weather and climate communities at pushing their models and software towards exascale. A call for proposals for corresponding service projects on model portability and refactoring will be opened on September 4. It will be focused towards porting existing model codes to accelerator hardware and it will be open to both ESiWACE partners and external groups, where the latter will be given preference. Two selected proposals will then be granted 6PMs in kind contribution within 2020. A call for services on the XIOS IO server and on the OASIS3-MCT coupler offering 2 PMs of dedicated support for each tool will also be published in September. This call together with the call for projects on model portability and refactoring will be further advertised soon. Calls for 2021 and 2022 will follow.
Are you interested? Do you have questions? Feel free to contact:

Model portability and refactoring:	Gijs van den Oord, g.vandenoord@esciencecenter.nl
	Ben van Werkhoven, b.vanwerkhoven@esciencecenter.nl
XIOS/OASIS support:	Sophie Valcke, sophie.valcke@cerfacs.fr

Save the Date: 6th ENES HPC Workshop

The next ENES HPC workshop, being the sixth in a series, will take place 25-26 May 2020 in Hamburg at Elbpanorama, Bernhard-Nocht-Straße 113, 20359 Hamburg. Save the date!

Snapshots

DKRZ is currently pushing the developed in-situ visualization techniques towards production in the ICON framework and contributes to the archival of the 1.2 PB of DYAMOND data.

ECMWF is investigating the scaling efficiency of the IFS extreme scale demonstrator at 1.45 km resolution on the fastest supercomputer of the world – Summit at Oak Ridge; click here for more information.

The Barcelona Supercomputing Center (in collaboration with the Netherlands eScience Center and ECMWF) has completed the developments to port XIOS into IFS/OpenIFS.

Arm has been working on porting LFRic to their current HPC architecture along with performance investigation for NEMO and the BENCH test case.

Bull has finalised the technical report on NEMO optimisations as a first step to submit a result paper soon. Bull has also represented the ESiWACE CoE at the Teratec forum, the international meeting for simulation and high-performance computing near Paris, with a booth at the European Research Café. Our booth

showed two posters and a high resolution forecast simulation video. The distribution of ESiWACE-branded USB sticks was a real success!

The Met Office is currently focused on the development of the next major releases of Cylc and Rose. This includes porting the code to Python 3 and introducing a new architecture which will support a web based GUI to replace the existing GTK based GUIs.

CERFACS continued to improve OASIS3-MCT. To support surface models including multiple types of surface, fractional masks are now supported by OASIS3-MCT for the global conservation operation.

The University of Reading is currently finalising the ESDM plugin for NetCDF. Over the last month, we switched to a new metadata scheme and improved the MPI support library; performance results from Mistral yield 200 GB/s showing that ESDM is able to extract available bandwidth from both Lustre file systems concurrently.

Seagate is working on providing an Object based Storage Backend for the ESDM developed by the consortium.

CMCC is finalising a simple performance model for assessing the I/O throughput and latency of the ESDM component. Besides, a multi-model analytics workflow on 17 CMIP5 models has been implemented on the Athena HPC cluster. To this end, HPC scheduling, deployment and parallel file system integration aspects with the analytics workflow manager of Ophidia have been properly addressed.

ETH Zurich / CSCS has fixed the deadline for the Docker Container hackathon as Dec. 3-5, 2019 at CSCS headquarters in Lugano. ESiWACE2 stakeholders have committed to bringing the following models to the hackathon: OpenIFS/EC-Earth, Unified Model, NEMO, ICON, COSMO, and LFRIC.

UKRI/STFC are developing PSyclone to be able to transform NEMO code to run on GPUs.

The projects ESiWACE and ESiWACE2 have received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreements No 675191 and 823988.