esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Centre of Excellence in Simulation of Weather and Climate in Europe
Phase 2

Summary of Container Hackathon Experiences

Deliverable D2.8

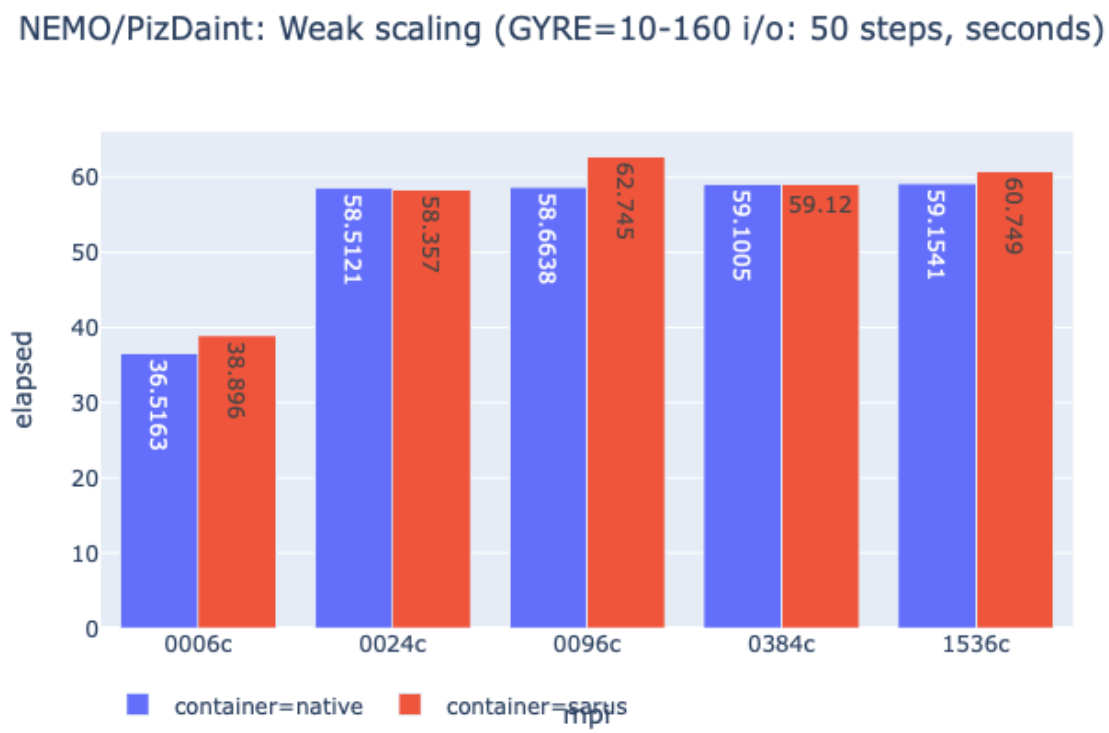

[1]

The project Centre of Excellence in Simulation of Weather and Climate in Europe Phase 2 (ESiWACE2) has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 823988

[1] Weak scaling results from the NEMO team indicate that running containers within the SARUS framework on Piz Daint introduces little overhead compared to running natively without containers.

About this document

Work package in charge: WP2 Establish, evaluate and watch new technologies for the community

Actual delivery date for this deliverable: 28 Feb. 2020
Dissemination level: PU (for public use)

Lead authors
Swiss National Supercomputing Centre, ETH Zürich (ETHZ): William Sawyer, Lucas Benedicic

Other contributing authors

All authors of project reports (see appendix A.1)

Contact details
Project Office: esiwace@dkrz.de
Visit us on: www.esiwace.eu

Access our documents in Zenodo:
https://zenodo.org/communities/esiwace

Follow us on Twitter: @esiwace

Table of Contents

1. Abstract / publishable summary

Atmosphere and ocean models are characterised by complex dependencies, external configurations, and performance requirements. The objective to containerise such software stacks helps to provide a consistent environment to ensure security, portability and performance. Since the container is built only once, but then can be deployed on multiple platforms, productivity is increased.

ETH Zurich (ETHZ) has organised and carried out a "hackathon" (programming session) to help ESiWACE community scientists create containers for their Earth-system models. Seven teams joined the hackathon, and each team was assigned a mentor with extensive container experience. The teams were given the challenge to get as far as they could in containerising their model, and, if possible, to analyse performance at scale on the Piz Daint computing platform at the Swiss National Supercomputing Centre (CSCS).

The teams made faster progress on their containers than even we organisers anticipated. Although many of the participants had never created a container previously, every team managed to complete and a run containerised version of its model, and some teams were able to benchmark on hundreds of computing nodes on Piz Daint. This report summarises the experiences made by the individual teams at the hackathon and draws some conclusions about the viability of 'kick-starting' containerisation efforts with such an intense programming event.

2. Conclusion & Results

Preparation for this event took place in 2019Q4 with mentors from ETHZ/CSCS contacting individual teams to ready their applications for subsequent containerisation. The 3-day hackathon took place Dec. 3-5, 2019 at CSCS in Lugano, Switzerland. The central achievements were:

- Seven teams from the ESiWACE-2 community participated, 15 people in total
- Each team successfully containerised its model and was able to deploy it on at least one target platform (e.g., a laptop)
- Several teams managed to deploy their containers at scale on the ETHZ/CSCS Piz Daint high-performance computing platform
- Each team delivered a project report (see Appendix A.1) either during the hackathon or immediately thereafter.
- A post-event survey was taken, and the feedback (see Appendix A.2) was favourable.

For many of the teams this event represented their first exposure to containers, which makes the results even more impressive. All the teams left with the knowledge and ability to create their own containers and to pursue their own modelling objectives. All of the teams were offered the possibility to continue their efforts to deploy and benchmark their containers on the Piz Daint HPC platform.

3. Project objectives

This deliverable contributes directly and indirectly to the achievement of all the macro-objectives and specific goals indicated in section 1.1 of the Description of the Action:

Macro-objectives	Contribution of this deliverable?
(1) Enable leading European weather and climate models to leverage the available performance of pre-exascale systems with regard to both compute and data capacity in 2021.	Yes
(2) Prepare the weather and climate community to be able to make use of exascale systems when they become available.	Yes

Specific goals in the workplan	Contribution of this deliverable?
Boost European climate and weather models to operate in world-leading quality on existing supercomputing and future pre-exascale platforms	Yes
Establish new technologies for weather and climate modelling	Yes
Enhance HPC capacity of the weather and climate community	Yes
Improve the toolchain to manage data from climate and weather simulations at scale	Yes
Strengthen the interaction with the European HPC ecosystem	Yes
Foster co-design between model developers, HPC manufacturers and HPC centres	Yes

4. Detailed report on the deliverable

The Container Hackathon was organised by the following ETHZ personnel:

- William Sawyer
- Lucas Benedicic
- Theofilos Manitaras

Just after the ESiWACE-2 kick-off meeting, invitations were sent to ESiWACE-2 team members to propose their models within WP2 for potential containerisation and then to participate at the hackathon. There was considerable interest, and ultimately seven models were accepted:

- **ICON** (Icosahedral Grid Non-hydrostatic): atmospheric model from the Max Planck Institute for Meteorology (MPI-M) and the German Weather Service (DWD).
- **NEMO** (Nucleus for European Modelling of the Ocean): Ocean Model created by a European Consortium with participating institutions in the UK, France and Italy.
- **EC-Earth**: A European community Earth-System Model, developed by a consortium of 27 research institutes in 10 European countries.
- **LFRic** (implementing concepts from researchers Lewis, Fry and Richardson a century ago) atmospheric model developed at the United Kingdom Meteorological Office (UKMet), among other institutions.
- **NorESM** (Norwegian Earth System Model): Earth System Model developed by a consortium of Norwegian institutions, which is based on the CESM (Community Earth System Model), developed in the USA.

- **OpenIFS** (Open Integrated Forecasting System): an open-source version of the IFS model developed at the European Centre for Medium-range Weather Forecasts (ECMWF).
- **WFPPDL** (Weather Forecast Parallel Prediction Deep Learning): Deep Learning prediction of 2-metre temperatures from the Jülich Supercomputing Centre.

The following teams were put together, including one mentor from ETHZ and the team participants:

Team	Mentor	Participants
EC-Earth	Alberto Madonna	Pablo Echevaria (BSC) Julian Berlin (BSC) Uwe Fladrich (SMHI)
ICON	Rafael Sarmiento	Remo Dietlicher (MeteoSuisse) Andre Walser (MeteoSuisse) Sergey Kosukhin (MPI-M) William Sawyer (CSCS)
LFRic	Luca Marsella	Iva Kavcic (UKMet) Simon Wilson (UKMet)
NEMO	Jean-Guillaume Piccinali	Italo Epicoco (CMCC) Giorgio Micaletto (CMCC)
NorESM/CESM	Kean Mariotti	Anne Claire Fouilloux (Univ. Oslo)
OpenIFS	Theofilos Manitaras	Marcus Koehler (ECMWF)
WFPPDL	Tomas Aliaga	Amirpasha Mozaffari (JSC) Bing Gong (JSC) Jan Vogelsang (JSC)

The mentors started contacting the team members in Oct. 2019 to help the team prepare its application for the hackathon. Each team was also encouraged to install Docker on its target platform, whether it be a laptop or a larger system at their home institution. Although this step required root privileges, each team was able to prepare a target platform. On the other hand, not every team was able to invest time learning about containers before the event, and several participants came without any initial exposure to containers.

ETHZ provided resources through a GitHub repository[2], into which the teams checked in their work. A SLACK[3] channel was created for intra- and inter-team communication. The Hackathon agenda contained only a presentation of the above-mentioned hackathon objectives, as well as a brief introduction to Docker, followed by 2-1/2 days of teamwork with an occasional scrum. Working alongside the mentors the participants made quick progress, and most achieved the objectives set out for each day:

1. After the first day, a sane build environment is working, and a "Hello, World" type of container exists.
2. By the end of the second day, the code should be containerised, and the test case selected by the team should be passing.

[2] https://github.com/eth-cscs/ContainerHackathon
[3] `containerhack-t3m5315.slack.com`

3. The third day was reserved for launching the containers at scale on Piz Daint, and preparing the final report.

The participants were encouraged to write their reports as a way of documenting the container procedures for themselves, but also for other teams who intended to follow their steps. The reports give detailed information on the activities of each team, as well as conclusions and comments about the effectiveness of the hackathon as a preparatory event for containerisation efforts.

5. References *(Bibliography)*

- ***Get Docker Engine - Community for Ubuntu***[4]
- ***Install Docker Desktop on Mac***[5]

6. Changes made and/or difficulties encountered, if any

As we realised the quick progress of the participants we encouraged them to pursue benchmarks on Piz Daint and to focus on their final reports.

7. How this deliverable contributes to the European strategies for HPC

Containers are a crucial technology to allow wide-scale deployment of applications to multiple computing environments. In the past, the cross-platform performance of the application has not been a key consideration. However, with the advent of Docker-compatible frameworks, such as Singularity and SARUS[6], which offer similar performance on HPC platforms of containerised applications to the application running on "bare-metal", containers are now viable for HPC. SARUS, for example, is an OCI-compliant, publicly available (BSD-3 license) container framework, which is used extensively on CSCS Piz Daint, among other systems. Thus, this deliverable illustrates that containers are now an integral part of the EU HPC strategy.

8. Sustainability

Deliverable D2.8 is the basis for a much larger effort with WP2 Task 2.3 to evaluate containers as a technique to port Earth system models to new architectures. The teams participating in the hackathon will take the knowledge acquired back to their institutions in order to augment and extend their model containers, and then evaluate them on a wide range of new architectures, in particularly with respect to their performance on HPC platforms. Their collective experiences with these containers will form Deliverable D2.8.

In the wider context, we see key links to WP3 *(HPC services to prepare the community for the pre-exascale)*, in particular Task 3.2 *(Model portability)* and Task 3.4 *(Weather and climate*

[4] `https://docs.docker.com/install/linux/docker-ce/ubuntu/`
[5] `https://docs.docker.com/docker-for-mac/install/`
[6] `https://user.cscs.ch/tools/containers/sarus/`

benchmarking). We already have contacts with some of the team involved in those tasks and will promote the idea of using containers to port models to new platforms and to benchmark on HPC platforms without extensive loss in performance. In addition, ETHZ plans to teach Containers as a topic for the ESiWACE-2 summer schools in 2020 and 2022.

Finally, ETHZ has contacts in the Climate and Weather Prediction communities outside of ESiWACE-2 who could potentially profit from containerising their own models.

The experiences from the hackathon have been resoundingly positive, and convince us that this is a viable technique to "kick-start" the containerisation of Earth system models.

9. Dissemination, Engagement and Uptake of Results

9.1 Target audience

As indicated in the Description of the Action, the audience for this deliverable is:

X	The general public (PU)
	The project partners, including the Commission services (PP)
	A group specified by the consortium, including the Commission services (RE)
	This reports is confidential, only for members of the consortium, including the Commission services (CO)

This is how we are going to ensure the uptake of the deliverables by the targeted audience:

This deliverable was always foreseen as open to the general public, and the teams were instructed to prepare their reports so that they could be for public consumption. Thus this deliverable provides numerous actual containerisation case studies that can act as a guide for new container efforts. The ETHZ will make this deliverable available to teams planning their own container efforts. Already teams who did not make it to the hackathon, e.g., non-ESiWACE-2 members, have requested this deliverable. We will also encourage the ESiWACE-2 community to propagate this useful result of the project.

9.2 Record of dissemination/engagement activities linked to this deliverable

Type of dissemination and communication activities	Details	Date and location of the event	Type of audience	Zenodo Link	Estimated number of persons reached
Hackathon	As explained in this document	Dec. 3-5, 2019 at CSCS Lugano, Switzerland	ESiWACE-2 community		15 through 3 days (45 person days in total).

9.3 Publications in preparation OR submitted

No publications are foreseen beyond this deliverable. The modelling teams may consider publication of their individual containerisation efforts.

9.4 Intellectual property rights resulting from this deliverable

None.

Appendix A.1: Modelling Team Reports

The modelling reports (PDF) are attached to this deliverable (the LFRic team delivered reports on two separate container frameworks). All teams managed to containerize their models. All teams but two managed to run their code at scale (at least 32 nodes) on Piz Daint using SARUS.

Team Reports	Authors
EC-Earth Containerization	A. Madonna, P. Echevaria. J. Berlin, U. Fladrich
Report on the Containerization of ICON	R. Sarmiento, R. Dietlicher. A. Walser, S. Kosukhin, W. Sawyer
Containerisation of LFRic with Docker	I. Kavcic
Containerisation of LFRic with Singularity	S. Wilson
NEMO Report	J.-G. Piccinali, I. Epicoco, G. Micaletto
NorESM/CESM Report	K. Mariotti, A. C. Fouilloux
The ECMWF OpenIFS atmospheric model in a Docker Container	T. Manitaras, M. Koehler
WFPPDL: Container Hackathon for Modellers	T. Aliaga, A. Mozaffari, B. Gong, J. Vogelsang

Appendix A.2: Course Evaluation Summary

Workshop surveys are collected after all ETHZ/CSCS educational events, and the Container Hackathon was no exception. The survey result (PDF) is attached. Key results of the survey were:

- None of the participants felt they had extensive experience with containers at the start of the hackathon; most felt they were 'somewhat' experienced or 'inexperienced'.
- At the end of the hackathon, the majority felt they were 'very comfortable' with containers, while the all the rest were 'somewhat comfortable'.
- All teams managed to create a docker container with their simulation code.
- The participants rated the overall experience with 4.91 out of a possible 5.

One quotation sums up the success of the hackathon for many of the teams:

> *All experiences were excellent. If I want to pinpoint the best thing, it would be the full engagement of the project mentors in all three days with us. It helped us to be able to deliver a complete product at the end of the third day.*