HPC ALLIANCE FOR APPLICATIONS AND SUPERCOMPUTING INNOVATION: THE EUROPE – JAPAN COLLABORATION

Funded by the European Union

EuroHPC Joint Undertaking

This project received funding from the European High Performance Computing Joint Undertaking (EuroHPC JU) under the European Union's Horizon Europe framework program for research and innovation and Grant Agreement No.101136269. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

DELIVERABLE D4.3

Report on detailed work plan for extension of HPCW

Funded by the European
EuroHPC Joint Undertaking

Project Title	Hpc AlliaNce for Applications and supercoMputing Innovation: the Europe - Japan collaboration
Project Ref	EuroHPC International Cooperation (HORIZON-EUROHPC-JU-2022-INCO-04)
Project Acronym	HANAMI
Project Number	101139786
Type of Action	HORIZON JU Research and Innovation Actions
Topic	HORIZON JU Research and Innovation Actions
Starting Date of Project	2024-03-01
Ending Date of Project	2028-02-28
Duration of the Project	36 months
Website	http://hanami-project.com/

Work Package	4
Task	4.2
Lead Authors	Samuel Hatfield (ECMWF)
Contributors	Niclas Schroeter (DKRZ)
Peer Reviewers	Mario Acosta (BSC), Joachim Biercamp (DKRZ)
Version	1.0
Due Date	31/08/24
Submission Date	29/08/24

Dissemination Level

X	PU: Public
	SEN: Sensitive – limited under the conditions of the Grant Agreement
	EU-RES. Classified Information: RESTREINT UE (Commission Decision 2005/444/EC)
	EU-CON. Classified Information: CONFIDENTIEL UE (Commission Decision 2005/444/EC)
	EU-SEC. Classified Information: SECRET UE (Commission Decision 2005/444/EC)

Disclaimer

Funded by the European
EuroHPC Joint Undertaking

Version History

Revision	Date	Editors	Comments
1.0	30/07/2024	Samuel Hatfield (ECMWF)	Compiling the comments from reviewers.
1.1	29/08/2024	France Boillod-Cerneux (CEA)	Compile the document for the submission.

Glossary of Terms

Item	Description
HPCW	High-Performance Climate and Weather benchmarking suite
HPC	High-performance computing
NICAM-DC	Nonhydrostatic ICosahedral Atmospheric Model (Dynamical Core)
NEMO	Nucleus for European Modelling of the Ocean
ICON	ICOsahedral Nonhydrostatic (Model)
IFS	Integrated Forecasting System
CI	Continuous Integration
EPI	European Processor Initiative

Funded by the European
EuroHPC Joint Undertaking

Executive Summary

The High-Performance Climate and Weather (HPCW) benchmarking suite brings together some of the key European models and model components used for simulating the Earth system on large computing systems. This suite allows us to quickly and easily evaluate a new computing platform from the perspective of weather and climate codes, providing a much more representative benchmark than those typically used to evaluate supercomputers, such as High-Performance Linpack. In this deliverable, we outline our plans for extending HPCW with the support of HANAMI, focusing on continued development of existing European models and model components, the inclusion of the first Japanese model (NICAM), and supporting HPCW on EuroHPC and Japanese (notably, Fugaku) high-performance computers.

Funded by the European
EuroHPC Joint Undertaking

Table of Contents

Funded by the European

1. Introduction

The High-Performance Climate and Weather suite (HPCW) is a software framework for preparing, building, and benchmarking a range of weather- and climate-related codes, both models and model components, in a unified manner. Taking advantage of the versatility of CMake, HPCW allows one to carry out all of these actions with just a few commands. This is in sharp contrast to the traditional benchmarking paradigm, in which every model and model component has its own method for preparing source code, preparing input data, building, and running. As such, HPCW is intended to be a complement to more traditional benchmarks such as the High-Performance Linpack library used for ranking supercomputers on the Top500 list and High-Performance Conjugate Gradient which is intended to be a more realistic benchmark of a modern memory-bound application. HPCW was developed through earlier EU-funded projects, namely ESCAPE-2, ESiWACE2, and is being maintained and governed by the EuroHPC funded Center of Excellence (CoE) ESiWACE3. As a result, it focuses entirely on European codes at present. This includes the ICON (ICOsahedral Nonhydrostatic) global atmospheric model, the NEMO (Nucleus for European Modelling of the Ocean) ocean model, and various open-source subcomponents of the ECMWF atmospheric model IFS (Integrated Forecasting System; the remainder of the IFS is still closed-source). We intend to develop HPCW further through HANAMI, adding the first Japanese (and indeed, first non-European) components to HPCW and deploying all the components on EuroHPC and Japanese machines.

2. State of HPCW prior to the HANAMI project

Prior to the start of HANAMI, HPCW was developed over the course of multiple projects and phases, namely ESCAPE-2, ESiWACE2 and ESiWACE3. This work culminated in the most recent release of the benchmark suite, HPCW v2.0. The work carried out during HANAMI will contribute towards another release of HPCW.

At the time of writing (July 2024) we are - under the governance of ESiWACE3 - in the process of consolidating the various components of HPCW. The components that we refer to as the HPCW-core are the following open-source codes:

- The HPCW framework itself, which handles the downloading of source code, the build process, the downloading of input files, and the running of benchmarks themselves.
- ecTrans: A module developed by ECMWF for performing efficient spectral transforms of global meteorological fields in a distributed memory context. It consists of compute- and communication-heavy routines and is therefore an excellent stress-test for any machine.

Funded by the European
EuroHPC Joint Undertaking

- CLOUDSC: A cloud microphysics parameterization scheme developed by ECMWF. CLOUDSC is considered an archetypal physical parameterization scheme, operating on a globe of atmospheric columns with no horizontal dependency, and is therefore a common target for optimization and porting efforts intended for parameterization schemes as a whole. CLOUDSC contains no communication routines and is therefore a useful test-bed for CPU or GPU-local optimisation approaches such as vectorization.
- ecRad: An atmospheric radiation scheme designed for computing profiles of solar and thermal infrared irradiances.

While these standalone model components (also known as “dwarfs”) are less complex than the complete models, they represent certain characteristics of said models.

The HPCW-core is further supplemented by two complete models. The first is ICON, an ocean and atmosphere model that is used in the context of both weather prediction and climate research. The second model is the NEMO ocean model, which is used for co-design in central European projects such as EPI (European Processor Initiative). Prior to the start of HANAMI, the version of ICON required by HPCW was not openly available and required a non-disclosure agreement from the user. In the context of ESiWACE3, plans were made to consolidate the ICON benchmark within HPCW by making ICON available under an open-source license. For the second model, NEMO, explicit funding to support it is currently not available. However, it is still one of the proclaimed goals to maintain NEMO as a part of HPCW.

The latest release of HPCW v2.0 also contains other components that are not supposed to receive further support and development in the context of this benchmark suite. The components include the following:

- IFS-RAPS: This code has a restricted license model and cannot be freely distributed. It could remain in HPCW as a reference due to its history and because it represents many motifs of weather codes.
- The KRONOS workload simulator and the ACRANEB2 radiation code: They will no longer be considered for HPCW due to the need to restrict the number of codes and wanting to concentrate on open-source codes.
- IFS-FVM dwarf: It is dependent on tools developed by Swiss institutions and thus cannot be supported by EuroHPC JU funded activities.

These components will be removed before the next release of HPCW.

One of the additions made during ESiWACE2 was the support for the package manager Spack, which is widely used in the current high-performance computing (HPC) landscape, easing the build process of many applications and their dependencies. All of the major HPCW components and certain dependencies received Spack recipes in this context.

The development version of HPCW is currently hosted on the DKRZ GitLab, where it is tightly integrated with a continuous integration (CI) system, used to test both the compilation and the execution of the individual components.

Over the course of ESiWACE2, HPCW has been deployed on many different HPC systems, highlighting the portability of the benchmark suite. Included in this are:

Funded by the European Union EuroHPC Joint Undertaking

- Several Eviden in-house systems
- Levante, the HPC system at DKRZ, Atos BullSequana XH2000
- the HPC system at ECMWF, also an Atos BullSequana XH2000
- The Swiss ALPS system of the Swiss National Supercomputing Centre, an HPE Cray EX
- LUMI of CSC, Finland
- Marenostrum IV of the Barcelona Supercomputing Centre

In the Appendix we have listed all extant systems which we already officially support for HPCW or will officially support during the course of HANAMI. Notably, ALPS is not mentioned here as we don't have an official way to maintain access to this system going forward and MareNostrum IV is not present in this table as the machine has been decommissioned.

3. Plans for extension of HPCW

3.1 Developing ICON

A recent ESiWACE3 deliverable provided recommendations for the future of ICON in the context of HPCW. In said deliverable, it was suggested to incorporate an open-source version of ICON into HPCW to circumvent the licensing issues of past releases. We intend to act on this plan in the context of HANAMI. Early in 2024, a first open-source version of ICON was released. We are going to replace the existing ICON version in HANAMI with this open-source release. Furthermore, we intend to incorporate every following update to the open-source release of ICON into HPCW in a timely manner. This change in versions also requires the replacement of all existing test cases for ICON. We intend to replace them with a different range of test cases, once again providing small, medium and large test cases, for different model and hardware setups, which includes GPU test cases. After this change, the open-source version of ICON can be considered a part of the HPCW-core as well.

3.2 A first open-source release

After incorporating the update to the open-source release of ICON into HPCW, we intend to make a first open-source release of HPCW. Currently, we are targeting a release window in the second half of 2024. This release enables other external users to gather first experiences with HPCW, while also not requiring any more non-disclosure agreements to access the complete benchmark suite. This first release is vital for the future of HPCW in many aspects. In order to foster the adoption of HPCW by the wider community, it is important to enable external users to utilize the entire benchmark suite without encumbering it through licensing issues. Users can then deploy HPCW on whichever machine they choose and gather first experiences with the benchmark suite. In an optimal scenario, they could provide feedback afterwards (for example via opening an issue on the HPCW GitLab repository), which in turn leads to further improvements of HPCW. As such, a viable option to consider, alongside the open-source release, is the incorporation of open-development, engaging the

Funded by the European Union
EuroHPC Joint Undertaking

community as a whole to foster the continued development of HPCW and make it sustainable for the future.

3.3 Continuous integration

In its current iteration, HPCW benefits from a CI system that is running on Levante. The concept of CI is very important to ensure that changes made to HPCW do not break any existing functionality, while also providing reassurance that new features work as intended. With CI, we could ease the porting process to different HPC systems as well. As such, we adjusted the CI system to utilize GitLab CI/CD. This CI system should enable an easy integration of other HPC systems into the tests, as said HPC systems need only to provide a GitLab runner. Integrating, for example, LUMI into the CI with this system should require comparatively little effort, which in turn leads to valuable feedback for any porting efforts to this particular system. Furthermore, any changes to HPCW components afterwards could be validated automatically across different systems, providing even greater assurance that all is working as intended.

3.4 Adding NICAM-DC to HPCW

So far, HPCW has been comprised exclusively of European models. For HANAMI, we will be adding the first subcomponents developed in Japan. The global atmospheric model NICAM (Nonhydrostatic ICosahedral Atmospheric Model) will be added to HPCW such that this can be built and benchmarked alongside existing components. More specifically, the “NICAM-DC” package will be added, which consists of the NICAM dynamical core only. This refers to the part of the model responsible for the solution of the large-scale fluid equations across the globe and doesn’t include small-scale, unresolved processes which are handled by “physical parametrisations”. We decided upon NICAM-DC instead of NICAM as the former has a permissive, open-source license, which allows us to include it in a future open-source HPCW release.

As of the time of writing (July 2024), initial work to include NICAM-DC in HPCW has just concluded. It is now possible, through HPCW, to automatically download the NICAM-DC source code, build the model, and execute one of NICAM-DC’s low resolution built-in benchmark suites. The latter consists of a standard atmospheric “test case” known as the Jablonowski Baroclinic Wave case. Development took place on the ECMWF HPC system but it was trivial to extend support to the DKRZ Levante system. Going forward, we will be working with colleagues in Japan to outline a full suite of NICAM-DC benchmarks at different resolutions.

3.5 Supporting Japanese machines (Fugaku)

Currently only a handful of European machines support HPCW, notably DKRZ Levante and the ECMWF HPC system. One of our main goals in HANAMI is to extend support of HPCW to machines in Japan. So far, only Fugaku has been identified as a machine of interest. This is the fastest supercomputer in Japan, and took the #1 place

Funded by the European U EuroHPC Joint Undertaking

on the Top500 list from June 2020 to May 2022. We anticipate two challenges in porting HPCW and its components to Fugaku:

1. Compilation. Fugaku uses the Fujitsu compiler for which we have relatively little experience. Each compiler has its own quirks, so it is unlikely that the full suite will work out of the box. However, drawing on ECMWF's experience on similar systems such as CEA Irene with the EUPEX project, we do not anticipate any fundamental barriers to building the full HPCW suite on Fugaku.
2. Hardware limitations. Fugaku is unusual in that it is comprised of a huge number of homogeneous compute nodes based on ARM CPU technology. Each node of Fugaku has one Fujitsu A64fx CPU. Again, based on earlier experiences, we do not anticipate any barriers on paper with running HPCW codes on this CPU. A potential barrier, however, is the limited memory capacity of each node of Fugaku: only around 30 GB. This may limit the test cases we are able to officially support on Fugaku.

We will tackle both of these technical hurdles with support from our colleagues in Japan once we have successfully negotiated accounts and a computational budget on Fugaku.

3.6 Extending support of EuroHPC machines

Of the three pre-exascale EuroHPC machines, LUMI, Leonardo, and MareNostrum 5, only LUMI has been considered for HPCW support. Furthermore, only the NEMO component has been tested on the LUMI-C (CPU-based) partition, so support there is very much limited. During HANAMI we aim to add support for execution of as many HPCW components as possible on both LUMI and MareNostrum 5.

For the ECMWF-specific components (ecTrans, CLOUDSC, ecRad) we do not anticipate difficulties. Both LUMI and MareNostrum 5 are targeted for deployment of the Destination Earth digital twins, one of which is based on ECMWF's forecasting model IFS. Therefore we can benefit from these developments. We expect to have support for GPU execution for at least ecTrans and CLOUDSC, as these two have official and well-developed plans for GPU portability.

For the remaining components, ICON and NICAM-DC, we also do not expect great difficulties. ICON has already been ported to LUMI, supporting execution on both the CPU and GPU partitions, and Mare Nostrum 5 provides a hardware and software stack that is similar to existing systems that already support ICON. As for NICAM-DC, we will only be able to run on CPUs due to limitations of the code, and we do not expect any major difficulties in compiling and running the model on these machines.

3.7 Application and Synergies

It is worth recalling that there is a strong scientific collaboration between Japan and Europe based on several components of HPCW and NICAM. The DYAMOND study (https://www.esiwace.eu/the-project/past-phases/dyamond-initiative) - which is comparing the scientific performance of high resolution climate models - was initiated by scientific leaders of the NICAM and ICON teams and includes also IFS. It is further worth mentioning that the NICAM team has been pioneering very high resolution

Funded by the European Union EuroHPC Joint Undertaking

atmospheric models for more than ten years.

In the context of HANAMI and using HPCW as a benchmark we will be able to complement the scientific comparison by an assessment of the computational performance of the different implementations on different platforms. This will lead to recommendation for co-design and ultimately improve the performance portability of our models.

4. Conclusions

HPCW is already a useful framework for assessing a new computing facility for performance from the perspective of weather and climate codes. However, it is far from realising its full potential. With support from HANAMI, we will extend HPCW to cover a fuller set of codes, including full models such as ICON, and make it available to the wider community through an open-source license. We will follow good modern software development strategies throughout, including a pivotal role for continuous integration suites. We will implement NICAM-DC, the first Japanese (and indeed, non-European) component of HPCW. We will extend HPCW to support more European HPC machines as well as Fugaku. Through this exercise, we will learn better about how to make the system-specific aspects of HPCW as generic as possible to ease the process of supporting a new machine. Finally we will use HPCW to compare computational implementations of the different models on different platforms aiming at improving the models and by way of co-design identify the best HPC platform for high-resolution climate and weather models.

Appendix: Systems targeted for HPCW

Name	Institution	Status
Levante	DKRZ	Full support
HPC2020	ECMWF	HPCW-core
(unnamed internal cluster)	Eviden	Full support
LUMI	CSC	NEMO tested on LUMI-C
MareNostrum 5	BSC	Not yet supported
Fugaku	R-CCS	Not yet supported

Funded by the European U EuroHPC Joint Undertaking