MAchinE Learning for Scalable meTeoROlogy and climate

MAELSTROM

Deliverable 3.3: Benchmarking and Test Infrastructure

Daniele Gregori

www.maelstrom-eurohpc.eu

MAELSTROM

D3.3 Benchmarking and test infrastructure

Author(s):	Daniele Gregori (E4)
Dissemination Level:	Public
Date:	13/03/2022
Version:	0.4
Contractual Delivery Date:	30/03/2022
Work Package/ Task:	WP3/ T3.3
Document Owner:	E4
Contributors:	E4, FZJ
Status:	Final

MAELSTROM

Machine Learning for Scalable Meteorology and Climate

Research and Innovation Action (RIA)
H2020-JTI-EuroHPC-2019-1: Towards Extreme Scale Technologies and Applications

Project Coordinator: Dr Peter Dueben (ECMWF)
Project Start Date: 01/04/2021
Project Duration: 36 months

Published by the MAELSTROM Consortium

Contact:
ECMWF, Shinfield Park, Reading, RG2 9AX, United Kingdom
Peter.Dueben@ecmwf.int

The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway

Contents

1 EXECUTIVE SUMMARY 6

2 INTRODUCTION 7

2.1 ABOUT MAELSTROM 7

2.2 SCOPE OF THIS DELIVERABLE 7

2.2.1 OBJECTIVES OF THIS DELIVERABLE 7

2.2.2 WORK PERFORMED IN THIS DELIVERABLE 7

2.2.3 DEVIATIONS AND COUNTER MEASURES 8

3 CURRENT COMPUTING INFRASTRUCTURES 9

3.1 E4 9

3.1.1 ARM THUNDERX-2 9

3.1.2 AMD 10

3.1.3 INTEL 11

3.1.4 POWER METER 12

3.2 FZJ 13

3.2.1 JUWELS CLUSTER (2018) 13

3.2.2 JUWELS BOOSTER (2020) 14

4 INFRASTRUCTURES AVAILABLE IN THE NEAR FUTURE 16

4.1 E4 16

ALL OF THE ABOVE IS DEPENDENT ON THE CAVEATS REPORTED IN 2.2.3 DEVIATIONS AND COUNTER MEASURES. 16

4.2 FZJ 16

4.2.1 JURECA DC 16

4.2.2 JURECA DC EVALUATION NODES 17

5 CONCLUSION 19

6 REFERENCES 20

1 Executive Summary

In the co-design approach adopted within the MAELSTROM project, an aspect of fundamental importance is the availability of different computing systems to guide development towards the top performing hardware configurations.

These configurations can be distinguished by the type of CPU architecture (for example from Intel, AMD, ARM), by the accelerators made available (for example from NVIDIA, AMD, Intel, Graphcore or Xilinx), by internal configuration of the servers in terms of memory quantity and type of storage (e.g. HDD, SSD, NVMe, distributed or parallel filesystem), or the type of network (e.g. InfiniBand or Ethernet).

In the first chapter of this deliverable we describe the computing infrastructures that are currently available at Jülich Supercomputing Center (FZJ) and at E4 laboratories (E4) for the development and benchmarking of the MAELSTROM applications developed within WP1.

The subsequent chapter gives an outlook about the systems that will be made available over the next few months.

2 Introduction

2.1 About MAELSTROM

To develop Europe’s computer architecture of the future, MAELSTROM will co-design bespoke compute system designs for optimal application performance and energy efficiency, a software framework to optimise usability and training efficiency for machine learning at scale, and large-scale machine learning applications for the domain of weather and climate science.

The MAELSTROM compute system designs will benchmark the applications across a range of computing systems regarding energy consumption, time-to-solution, numerical precision and solution accuracy. Customised compute systems will be designed that are optimised for application needs to strengthen Europe’s high-performance computing portfolio and to pull recent hardware developments, driven by general machine learning applications, toward needs of weather and climate applications.

The MAELSTROM software framework will enable scientists to apply and compare machine learning tools and libraries efficiently across a wide range of computer systems. A user interface will link application developers with compute system designers, and automated benchmarking and error detection of machine learning solutions will be performed during the development phase. Tools will be published as open source.

The MAELSTROM machine learning applications will cover all important components of the workflow of weather and climate predictions including the processing of observations, the assimilation of observations to generate initial and reference conditions, model simulations, as well as post-processing of model data and the development of forecast products. For each application, benchmark datasets with up to 10 terabytes of data will be published online for training and machine learning tool-developments at the scale of the fastest supercomputers in the world. MAELSTROM machine learning solutions will serve as blueprint for a wide range of machine learning applications on supercomputers in the future.

2.2 Scope of this deliverable

2.2.1 Objectives of this deliverable

The purpose of this document is to describe the computing systems that were made available to developers in the first year of operation of MAELSTROM by E4 and FZJ. An anticipation of the systems that will be made available in the coming months is also provided.

2.2.2 Work performed in this deliverable

The work carried out was to provide computing systems with heterogeneous architectures, equipped with computing accelerators.

The work is characterized by the preparation of hardware systems, their installation, the configuration of user management services, the configuration of shared filesystems that contain user data and applications and the configuration of compilers, mathematical libraries and the queue manager to submit jobs.

In this way, developers of weather forecasting applications can compare the different performances in terms of time-to-solution and energy-to-solution to make a codesign analysis.

2.2.3 Deviations and counter measures

The current situation of shortage of components is well-known. Up to this report, FZJ and E4 have been providing systems and infrastructures to the project featuring components usually already available in their inventories. Moving forward, and according to the currently un-predictability of the availability of components, the following actions will be taken to source components to continue to provide additional systems and infrastructures:

- Look for inventories in FZJ's and E4's supply chain
- Look for refurbished spare parts, possibly assembling a functioning part out of several malfunctioning parts

However, it should be highlighted that the study of the MAELSTROM applications could also be continued on the current systems and infrastructure.

3 Current Computing Infrastructures

3.1 E4

E4 [1] is an Italian SME that has a laboratory that allows to integrate and test HPC computing systems, including both commercial servers or prototypes. E4 is therefore able to supply "small" computing systems with very heterogeneous hardware, particularly suitable for the codesign approach taken in MAELSTROM.

For this first round of benchmarks, clusters of three different architectures are currently available at the laboratory: ARM, AMD and Intel. Each of these has different hardware configurations that are explained below.

Users share the following directories on all compute nodes: /home, /opt/share and /data/maelstrom via the NFS filesystem.

In particular, a NAS Infortrend GS 1012R2 with 10 HDD x 14 TB was made available to share the storage area that contains the datasets (/data/maelstrom) and home directories (/home).

Simultaneous Multithreading (*hyperthreading*) is turned off on all systems to maximize per-core-performance.

3.1.1 ARM ThunderX-2

ARM architecture systems can be reached via the tlnode01 login node. On all nodes the Operating System is Red Hat Enterprise Linux release 8.0 (Ootpa), with Linux kernel 4.18.0-80.

3.1.1.1 Tcnode0[1-7]

All nodes are dual socket, 32 cores per socket with Cavium ThunderX2 processors, 256 GB RAM, connected with InfiniBand 100 Gb/s network.

3.1.1.2 Tcnode0[5-6]

These two nodes offer access to one NVIDIA V100 GPU per node.

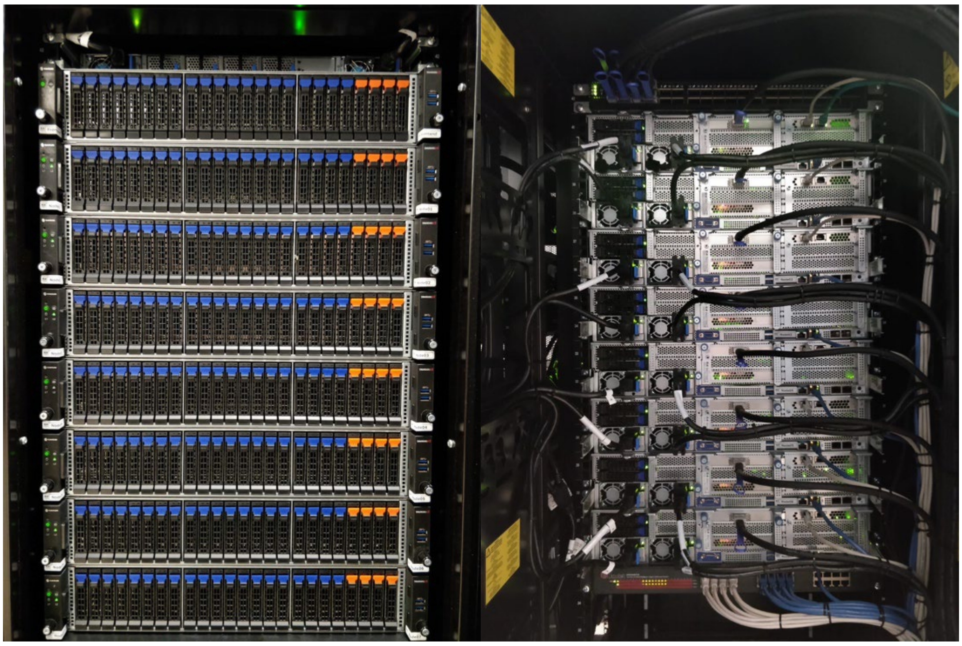

3.1.2 AMD

AMD architecture systems can be reached via the alnode01 login node. On all nodes the Operating System is Red Hat Enterprise Linux release 8.5 (Ootpa), with Linux kernel 4.18.0-305.

3.1.2.1 Alnode01

It is the Virtual Machine node on AMD Hypervisor used as login node.

3.1.2.2 Awnode0[1-4]

The overall system is a supermicro twinsquare:
https://www.supermicro.com/en/Aplus/system/2U/2124/AS-2124BT-HNTR.cfm

Each node is a dual socket with 2x AMD EPYC 7313 16-core processor, 256 GB RAM, connected via InfiniBand 100 Gb/s network.

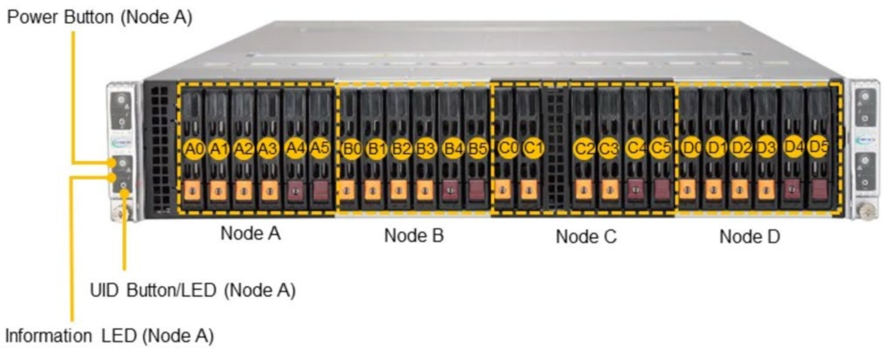

3.1.2.3 Acnode0[1-2]

Each node is a dual socket with 2x AMD EPYC 7313 16-core processor, 512 GB RAM, connected via 1 Gb/s Ethernet network (due to InfiniBand card shortage) and equipped with 2x AMD Mi100 GPU per node.

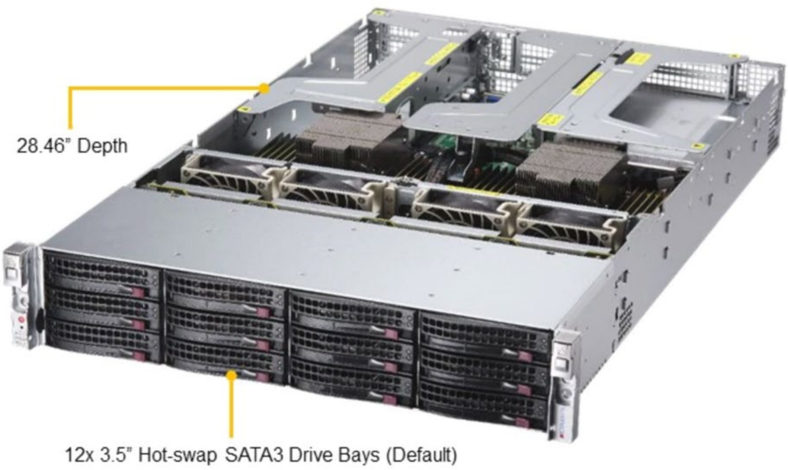

3.1.3 Intel

Intel architecture systems can be reached via the ilnode01 login node.

3.1.3.1 Ilnode01

It is the Virtual Machine node on Intel Hypervisor used as login node.

3.1.3.2 Iwnode0[1-4]

The overall system is a supermicro twinsquare:
https://www.supermicro.com/en/products/system/2U/2029/SYS-2029BT-HNTR.cfm

On all nodes the Operating System is Red Hat Enterprise Linux release 8.4 (Ootpa), with Linux kernel 4.18.0-305.25.

Each node is a dual socket with 2x Intel Xeon Gold 6226R CPU @ 2.90GHz, 16-core processor, 192 GB RAM, connected via InfiniBand 100 Gb/s network.

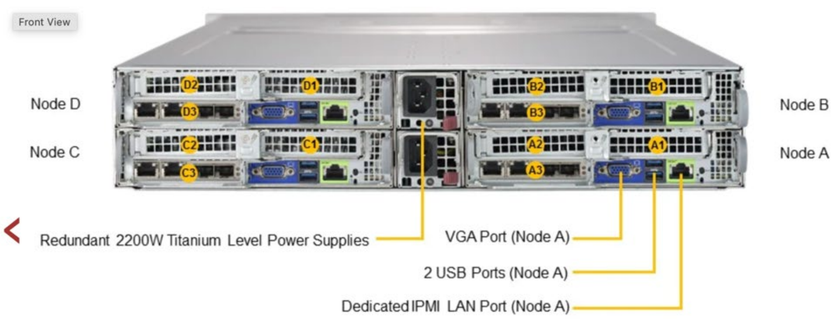

3.1.3.3 Icnode0[1-2]

On all nodes the Operating System is Red Hat Enterprise Linux release 8.5 (Ootpa), with Linux kernel 4.18.0-305.25.

Each node is a dual socket with 2x Intel Xeon Gold 6326 CPU @ 2.90GHz, 16-core processor, 512 GB RAM, connected via InfiniBand 100 Gb/s network.

Each node is equipped with a NVIDIA A100 GPU.

3.1.4 Power Meter

A peculiarity of the E4 laboratory is the availability of a Zimmer LMG95 power meter that enables to measure the power consumption at the plug of a power supply.

The product is described in the following link:

https://www.zes.com/en/Products/Discontinued-Products/Energy-and-Power-Meters/LMG95

This power meter therefore allows to perform measurements of the overall consumption of a server to determine the consumption value in idle and at full load. This value can be compared with the consumption of GPUs alone, the values of which can be obtained with the "nvidia-smi" utility for NVIDIA GPUs or "rocmsmi" utility for AMD GPUs.

The meter is available upon request and can be connected to one server at a time.

The connection diagram is shown in the following figure:

E4 "Single Server" Power Meter System

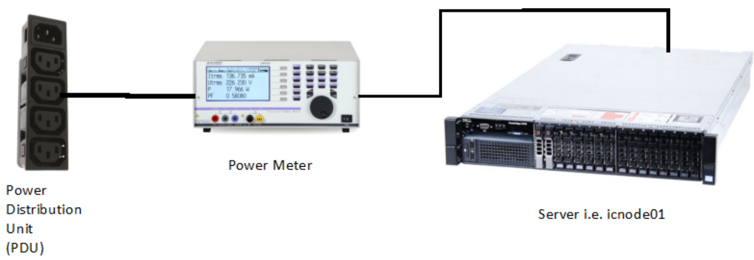

The LMG95 allows to measure the active power with a sampling time of 10 seconds and an accuracy of 0.025%.

To allow the simultaneous measurement of more servers, we are evaluating to use a metered PDU, which is able to perform a measurement to the plug of several servers connected to it.

3.2 FZJ

Forschungszentrum Jülich ("Jülich Research Centre") [2] is one of the largest interdisciplinary research centres in Europe. Jülich Supercomputing Centre hosts the systems available within MAELSTROM.
The systems supplied are partitions of computing clusters among the largest and top performing in Europe. In particular, the JUWELS Booster HPC system was #7 in the Top500 list of November 2020 and #1 in the Green500 list (counting large-scale installations). These computing systems are based on the Bull Sequana platform and are liquid-cooled. A storage system based on the parallel GPFS file system is connected with all computing systems (called JUDAC).

3.2.1 JUWELS Cluster (2018)

The general data of the cluster are the following:

- 2271 standard compute nodes
    - 2× Intel Xeon Platinum 8168 CPU, 2× 24 cores, 2.7 GHz
    - 96 (12× 8) GB DDR4, 2666 MHz
    - InfiniBand EDR (Connect-X4)
    - Intel Hyperthreading Technology (Simultaneous Multithreading)
    - diskless
- 240 large memory compute nodes
    - 2x Intel Xeon Platinum 8168 CPU, 2× 24 cores, 2.7 GHz
    - 192 (12× 16) GB DDR4, 2666 MHz
    - InfiniBand EDR (Connect-X4)
    - Intel Hyperthreading Technology (Simultaneous Multithreading)
    - diskless
- 56 accelerated compute nodes
    - 2× Intel Xeon Gold 6148 CPU, 2× 20 cores, 2.4 GHz

    - 192 (12× 16) GB DDR4, 2666MHz
    - 2× InfiniBand EDR (Connect-X4)
    - Intel Hyperthreading Technology (Simultaneous Multithreading)
    - 4× NVIDIA V100 GPU, 16 GB HBM
    - diskless
- 12 login nodes
    - 2× Intel Xeon Gold 6148 CPU, 2× 20 cores, 2.4 GHz
    - 12× 64 GB DDR4, 2666MHz
    - InfiniBand EDR (Connect-X5)
    - Intel Hyperthreading Technology (Simultaneous Multithreading)
    - 100 GigE
    - 2× 1TB HDD (RAID 1)
- 4 visualisation nodes
    - 2× Intel Xeon Gold 6148 CPU, 2× 20 cores, 2,4 GHz
    - 768 (12× 64) GB DDR4, 2666 MHz
    - InfiniBand EDR (Connect-X5)
    - Intel Hyperthreading Technology (Simultaneous Multithreading)
    - 100 GigE
    - 2x 1TB HDD (RAID 1)
    - 1× NVIDIA Pascal P100
- 122,768 CPU cores
- 10.6 (CPU) + 1.7 (GPU) Petaflop per second peak performance
- Mellanox InfiniBand EDR fat-tree network with 2:1 pruning at leaf level and top-level HDR switches
    - 40 Tb/s connection to Booster
- 250 GB/s network connection for storage access

3.2.2 JUWELS Booster (2020)

The general data of the cluster are the following:

- 936 compute nodes

- 2× AMD EPYC Rome 7402 CPU, 2× 24 cores, 2.8 GHz
  - 512 GB DDR4, 3200 MHz
  - 4× NVIDIA A100 GPU, 4× 40 GB HBM2e
  - 4× InfiniBand HDR (Connect-X6)
  - diskless
- 4 login nodes
  - 2x 24 cores, 2.7 GHz
  - 12x 16 GB, 2666 MHz
  - EDR-Infiniband (Connect-X4)
  - Intel Hyperthreading Technology (Simultaneous Multithreading)
  - diskless
- 3,744 GPUs
- 73 Petaflop per second peak performance
- Mellanox InfiniBand HDR DragonFly+ topology with 20 cells - 40 Tb/s connection to Cluster
- 350 GB/s network connection for storage access

The configuration of JUWELS Booster compute nodes is the following

- CPU: AMD EPYC 7402 processor; 2 sockets, 24 cores per socket, SMT-2 (total: 2×24×2 = 96 threads) in NPS-4 configuration
- Memory: 512 GB DDR4-3200 RAM (of which at least 20 GB is taken by the system software stack, including the file system); 256 GB per socket; 8 memory channels per socket (2 channels per NUMA domain)
- GPU: 4 × NVIDIA A100 Tensor Core GPU with 40 GB; connected via NVLink3 to each other
- Network: 4 × Mellanox HDR200 InfiniBand ConnectX 6 (200 Gbit/s each)
- Periphery: CPU, GPU, and network adapter are connected via 2 PCIe Gen 4 switches with 16 PCIe lanes going to each device (CPU socket: 2×16 lanes). PCIe switches are configured in synthetic mode.

The InfiniBand network of JUWELS Booster is implemented as a DragonFly+ network.

4 Infrastructures Available in the Near Future

In the near future, new architectures and accelerators will be available on the market and for this reason they will be made available to the developers of the project. At present, given the uncertainties due to current contingencies (war, pandemic, chip shortage), the release dates of new products and their procurement may be delayed. Below we will try to outline what we believe to be a feasible roadmap.

4.1 E4

During 2022, the x86 servers will be updated with new architectures, in particular Intel Sapphire Rapids and AMD Genoa processors.

The ARM servers will be replaced with the new AMPERE Altra CPUs.

The accelerators we intend to mount on these servers are NVIDIA A100, A100x, A30, AMD Mi 200. We estimate that the Intel Pontevecchio accelerator will be available to developers at the beginning of 2023.

Furthermore, there will be the possibility of providing Xilinx FPGA accelerators if the developers show interest in the test.

All of the above is dependent on the caveats reported in 2.2.3 Deviations and counter measures.

4.2 FZJ

4.2.1 JURECA DC

The JURECA DC supercomputer is already installed at FZJ and can be made available upon request to MAELSTROM developers. The sysem design is similar to JUWELS Booster but not identical. The system contains the following parts:

- 480 standard compute nodes
    - 2× AMD EPYC 7742, 2× 64 cores, 2.25 GHz
    - 512 (16× 32) GB DDR4, 3200 MHz

  - InfiniBand HDR100 (NVIDIA Mellanox Connect-X6)
  - diskless
- 96 large-memory compute nodes
  - 2× AMD EPYC 7742, 2× 64 cores, 2.25 GHz
  - 1024 (16× 64) GB DDR4, 3200 MHz
  - InfiniBand HDR100 (NVIDIA Mellanox Connect-X6)
  - diskless
- 192 accelerated compute nodes
  - 2× AMD EPYC 7742, 2× 64 cores, 2.25 GHz
  - 512 (16× 32) GB DDR4, 3200 MHz
  - 4× NVIDIA A100 GPU, 4× 40 GB HBM2e
  - 2× InfiniBand HDR (NVIDIA Mellanox Connect-X6)
  - diskless
- 12 login nodes
  - 2× AMD EPYC 7742, 2× 64 cores, 2.25 GHz
  - 1024 (16× 64) GB DDR4, 3200 MHz
  - 2× NVIDIA Quadro RTX8000
  - InfiniBand HDR100 (NVIDIA Mellanox Connect-X6)
  - 100 Gigabit Ethernet external connection
- 3.54 (CPU) + 14.98 (GPU) Petaflops per second peak performance
- 98,304 CPU cores, 768 GPUs
- Mellanox InfiniBand HDR (HDR100/HDR) DragonFly+ network
  - Ca. 15 Tb/s connection to Booster via gateway nodes
- 350 GB/s network connection for storage access

4.2.2 JURECA DC Evaluation Nodes

To investigate new and upcoming technology, different evaluation nodes are currently being installed into JURECA DC as dedicated partitions.

4.2.2.1 AMD MI250 Nodes

Two Gigbyte G262-ZO0 nodes with 4 AMD MI250 OAM GPUs each are installed. The nodes feature 512 GB RAM and two sockets with an AMD EPYC 7443 24-core processors each. More details can be found at https://www.gigabyte.com/Enterprise/GPU-Server/G262-ZO0-rev-A00.

The nodes are available through the dedicated “dc-mi200” Slurm partition.

4.2.2.2 NVIDIA ARM HPC Dev Kit Nodes

Two Gigabyte G242-P32 nodes with 2 NVIDIA A100 (PCIe, 40 GB) each are installed. The nodes feature 512 GB RAM and a single-socket Ampere Altra 80-core processor. More details can be found at https://www.gigabyte.com/Enterprise/GPU-Server/G242-P32-rev-100.

4.2.2.3 Graphcore Node

A Graphcore IPU-POD4 system is installed, featuring an access server (with two sockets with AMD EPYC 7413 24-core processors, 512 GB RAM) and the core Graphcore IPU-M2000 accelerator node.

More details about the IPU-M2000 can be found at:
https://docs.graphcore.ai/projects/graphcore-ipu-m2000-datasheet/en/latest/

and details about the access server can be found at:
https://docs.graphcore.ai/projects/graphcore-approved-server-list/en/latest/bullsequana-x400-series.html

5 Conclusion

In this first year of MAELSTROM, the computing resources provided by E4 and FZJ have been characterized by being heterogeneous in terms of processors and accelerators.

This is the necessary condition to be able to perform codesign work in search of the best solution for pre-exascale systems.

Over the next year, further additions will be made, making an increasing number of systems available to developers.

In the first round of benchmarks, not all architectures and accelerators made available were used. The lesson learned is that computational resources must be made available well in advance and this will be even more true for the new accelerators that will require more work to port the application code.

6 References

[1] https://www.e4company.com/

[2] https://www.fz-juelich.de/portal/DE/Home/home_node.html

Document History

Version	Author(s)	Date	Changes
0.0	Daniele Gregori (E4)	14/03/2022	First draft
0.1	Daniele Gregori (E4)	04/04/2022	Added informations

Internal Review History

Internal Reviewers	Date	Comments
Peter Dueben (ECMWF)	19/04/2022	Minor corrections and comments
Andreas Herten (FZJ)	20/04/2022	Minor corrections, additions relating to FZJ systems
Fabian Emmerich	26/04/2022	Minor corrections

Estimated Effort Contribution per Partner

Partner	Effort
E4	0.5 PM
Total	**0.5 PM**