Co-funded by
the European Union

Multi scale

EuroHPC
Joint Undertaking

EESSI
EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

EuroHPC JU Centre of Excellence

HPC ecosystem tools session

EuroHPC User Day 2024 @ Amsterdam

Tue 22 Oct 2024

Kenneth Hoste (Ghent University)

kenneth.hoste@ugent.be

cecam
Centre Européen de Calcul Atomique et Moléculaire

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and countries participating in the project. Neither the European Union nor the granting authority can be held responsible for them.

The MultiXscale CoE in a nutshell

- Collaboration between **EESSI** and **CECAM** (total of 16 partners)
  - EESSI primarily addresses technical aspects
  - CECAM network provides scientific expertise
- Scientific target: **multiscale simulations** with 3 key use cases
  - Helicopter design and certification for civil transport
  - Battery applications to support the sustainable energy transition
  - Ultrasound for non-invasive diagnostics and biomedical applications
- More info: multixscale.eu

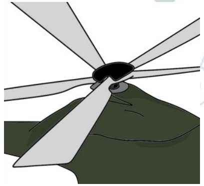

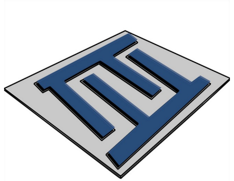

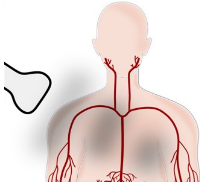

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Today’s challenges for scientific researchers:

- Growing **diversity** in system architectures of (EuroHPC JU) systems
- More and more different **software** used by researchers (cfr. AI boom)
- Lack of sufficient **manpower** in user support teams to help everyone
- Lack of a **uniform software environment** across EuroHPC JU systems
- No **central CI/CD service** across EuroHPC systems (testing & deploying)

This project has received funding from the European
High Performance Computing Joint Undertaking under
grant agreement No. 101093169

What if you no longer have to install

a broad range of scientific software

from scratch on every laptop, HPC cluster,

or cloud instance you use or maintain,

without compromising on **performance**?

EESSI

EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

EuroHPC Joint Undertaking

European Environment for Scientific Software Installations

- **Shared repository <u>optimized</u> scientific software installations**
- **Avoid duplicate work** by collaborating on a shared software stack
- **Uniform** way of providing software to users
- Should work on **any system architecture** (laptop, HPC, cloud)
- Focus on **performance**, automation, testing, collaboration

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

EESSI as a **uniform software environment** for EuroHPC systems (and beyond)

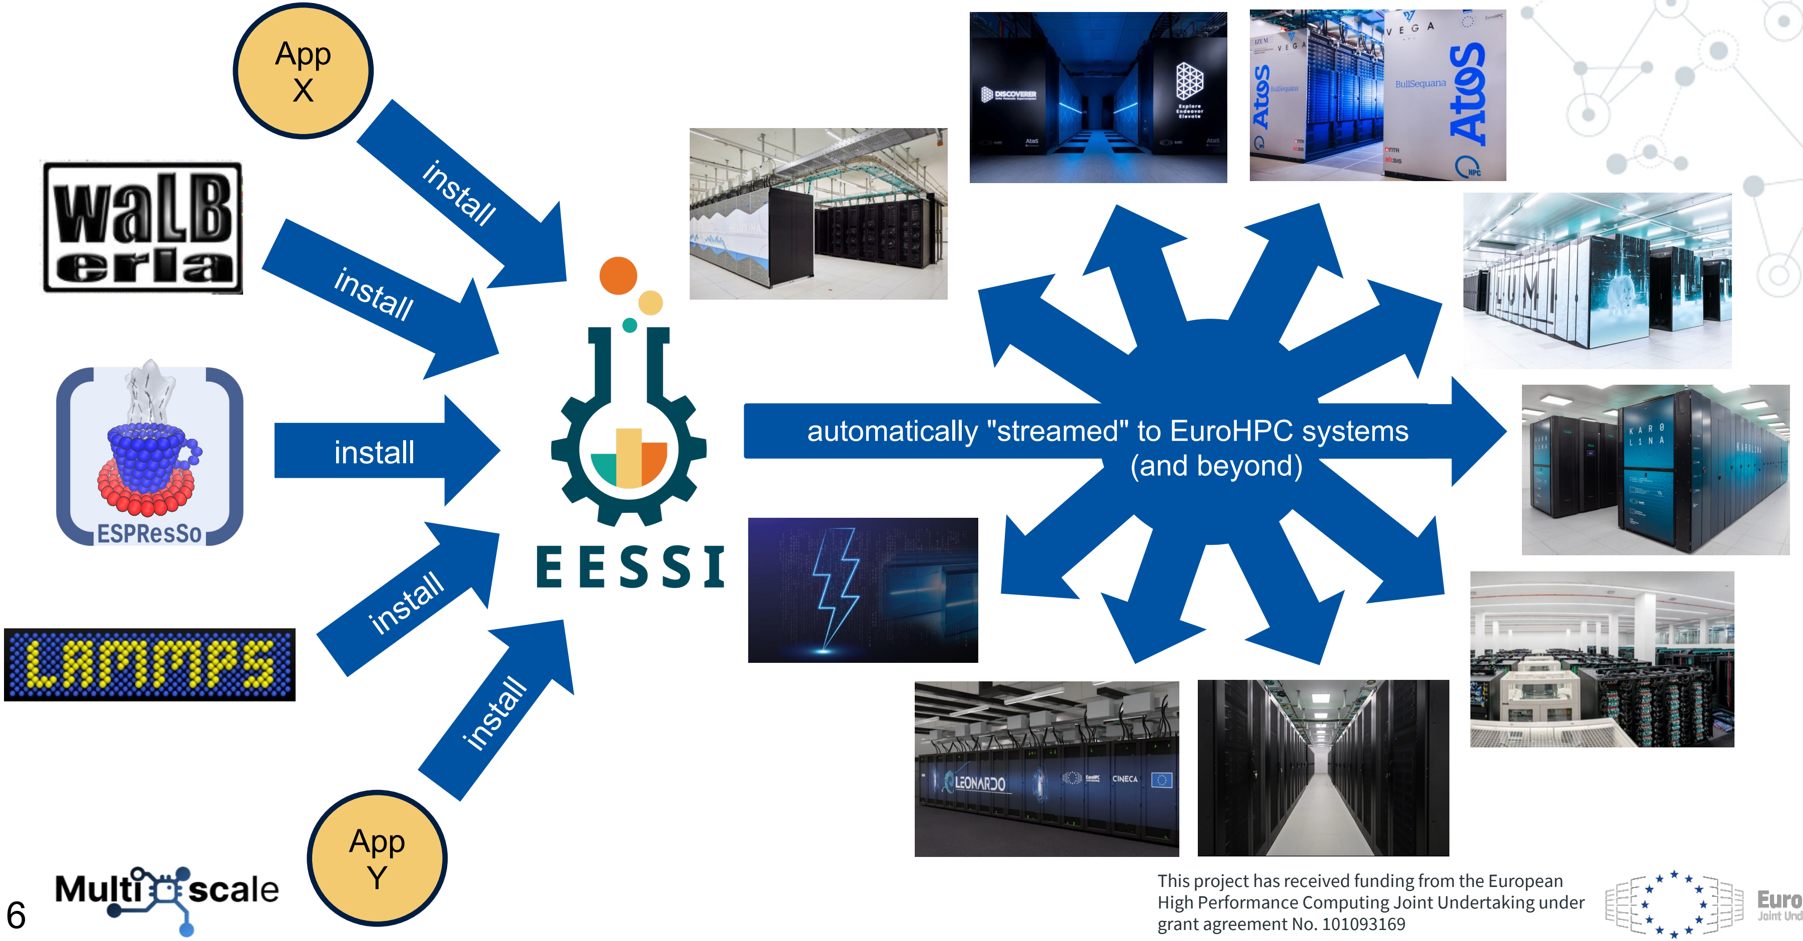

App X

install

install

install

install

install

App Y

EESSI

automatically "streamed" to EuroHPC systems (and beyond)

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

End-user experience: initialize, load modules, run! **(DEMO)**

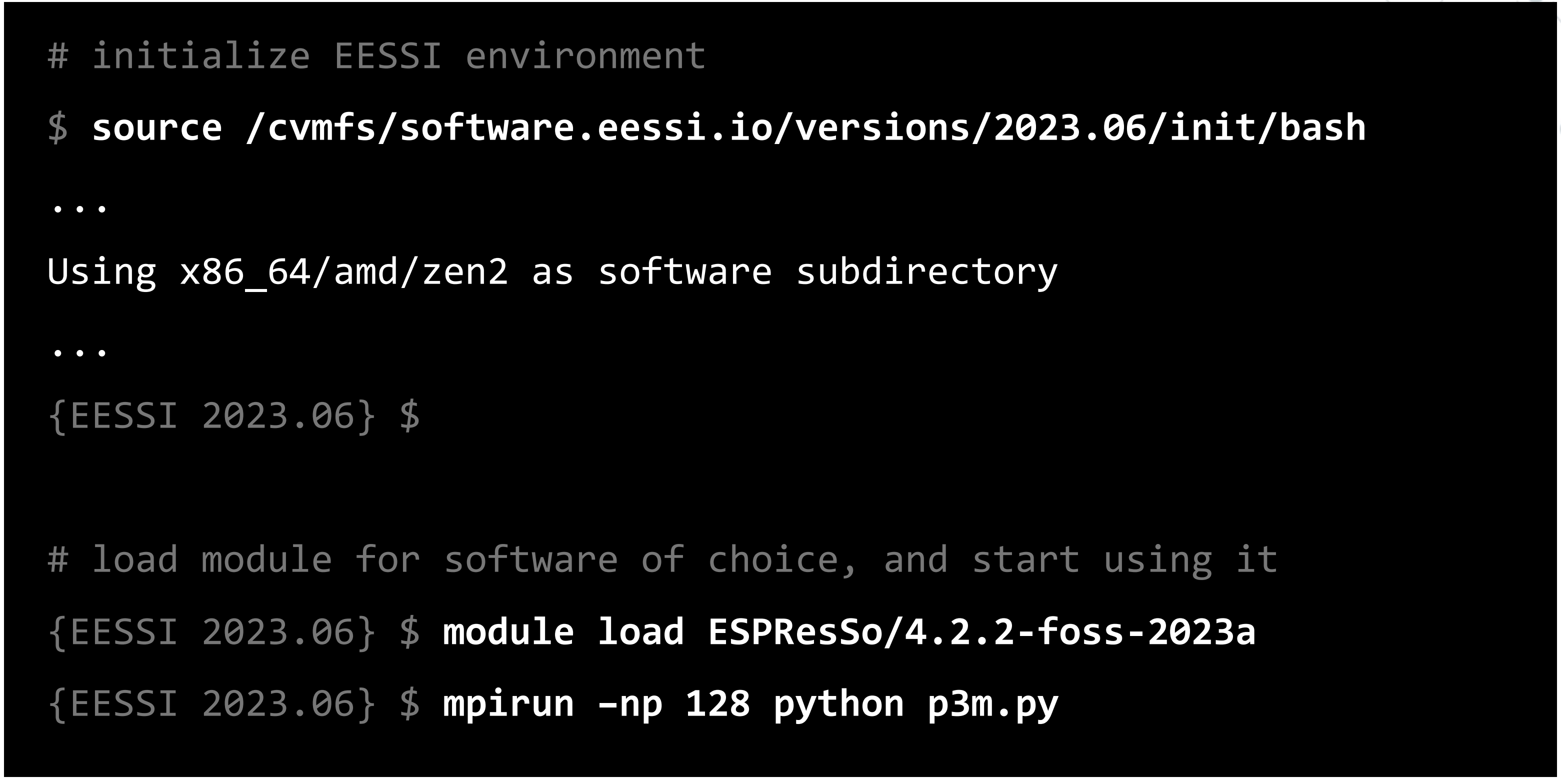

```
# initialize EESSI environment
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
...
Using x86_64/amd/zen2 as software subdirectory
...
{EESSI 2023.06} $


# load module for software of choice, and start using it
{EESSI 2023.06} $ module load ESPResSo/4.2.2-foss-2023a
{EESSI 2023.06} $ mpirun -np 128 python p3m.py
```

eessi.io/docs/using_eessi/eessi_demos

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

End-user experience: initialize, load modules, run! **(DEMO)**

*On an **NVIDIA A100 GPU** node of the Vega EuroHPC JU system:*

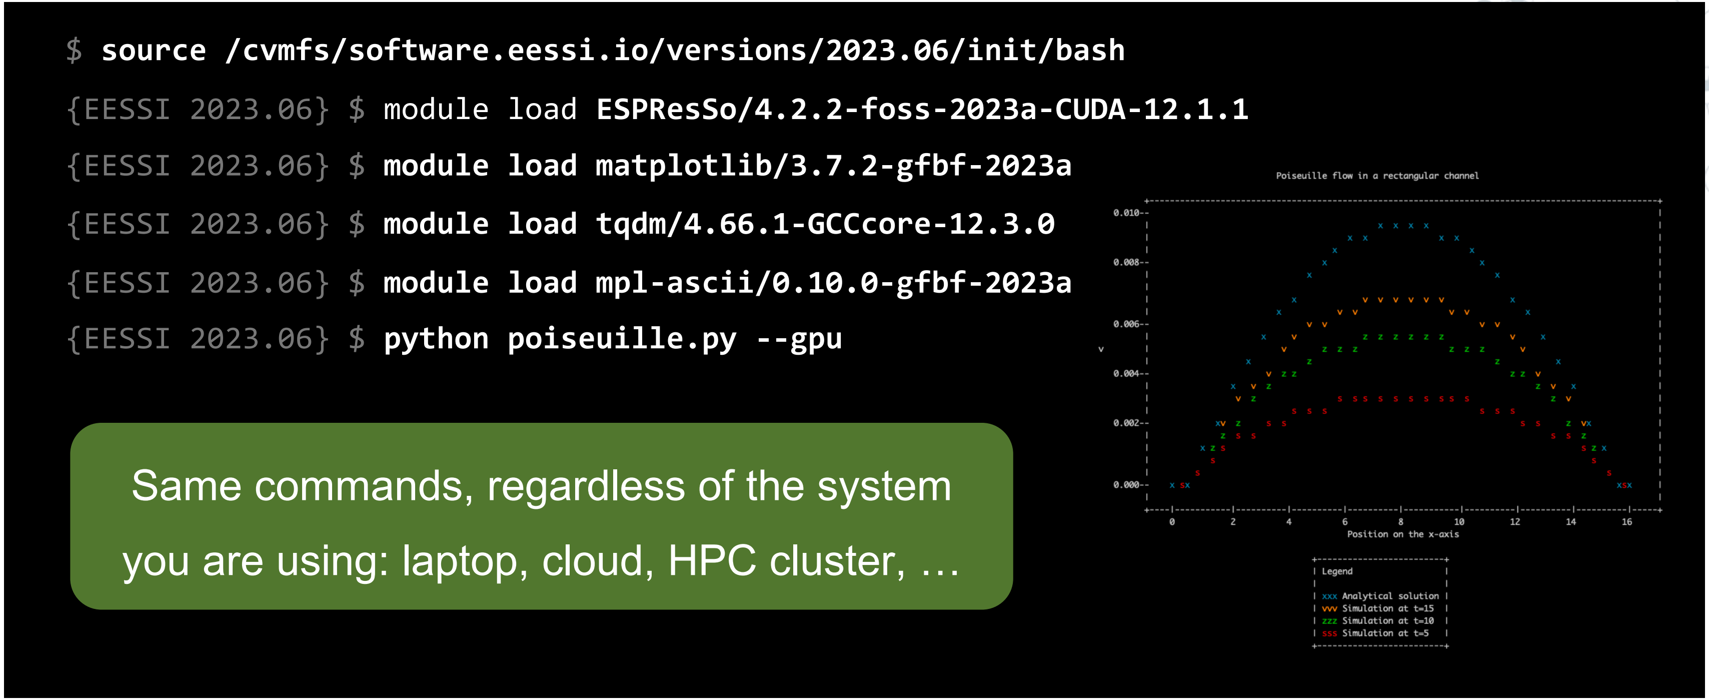

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
{EESSI 2023.06} $ module load ESPResSo/4.2.2-foss-2023a-CUDA-12.1.1
{EESSI 2023.06} $ module load matplotlib/3.7.2-gfbf-2023a
{EESSI 2023.06} $ module load tqdm/4.66.1-GCCcore-12.3.0
{EESSI 2023.06} $ module load mpl-ascii/0.10.0-gfbf-2023a
{EESSI 2023.06} $ python poiseuille.py --gpu
```

Same commands, regardless of the system you are using: laptop, cloud, HPC cluster, …

eessi.io/docs/using_eessi/eessi_demos

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

How does it work?

- Software installations included in EESSI are:
  - Automatically **“streamed in” on demand** (via CernVM-FS)
  - Built to be **independent of the host operating system**
    - *“Containers without the containing”*
  - **Optimized** for specific CPU generations + specific GPU types
- Initialization script **auto-detects** CPU + GPU of the system

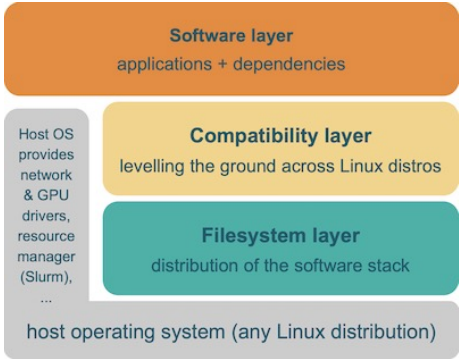

Software layer
applications + dependencies
Compatibility layer
levelling the ground across Linux distros
Filesystem layer
distribution of the software stack
Host OS provides network & GPU drivers, resource manager (Slurm), ...
host operating system (any Linux distribution)

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Getting access

- EESSI is **publicly accessible** (for free) from anywhere in the world
- **System-wide native installation** via CernVM-FS is recommended
  - See our *“Best Practices for CernVM-FS in HPC” tutorial*

    multixscale.github.io/cvmfs-tutorial-hpc-best-practices
  - Must be done by system administrators…
- Various alternative ways to access EESSI (as a regular user):
  - Using Singularity or Apptainer + EESSI client **container image** (incl. CernVM-FS)
  - Using **`cvmfsexec`** tool (requires support for user namespaces)
  - Via a pre-created **squashfs** image for EESSI (for offline workernodes)
  - See eessi.io/docs

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Supported system architectures

- Different generations of **x86_64** (Intel, AMD) and **Arm 64-bit** CPUs; **RISC-V** is WIP
  - Including A64FX (Deucalion, *WIP*) & NVIDIA Grace (JUPITER, *coming soon*)
  - Also works on laptops, in virtual machines in the cloud, on Raspberry Pi boards, etc.
- Different accelerators: **NVIDIA GPUs** (today) + **AMD GPUs** (soon)
- **Various interconnects** like Infiniband, via “fat” MPI libraries
  - Support for injecting a vendor-provided MPI library is available
- Goal is to support system architecture of **all** (current & future) **EuroHPC systems**

This project has received funding from the European
High Performance Computing Joint Undertaking under
grant agreement No. 101093169

On which systems is EESSI already available?

EuroHPC JU systems:

- Native installation (via CernVM-FS) on **Vega** + **Karolina**
- EESSI can be used via `cvmfsexec` tool on Deucalion, Discoverer, MeluXina
- Native installation on **MeluXina** and **Deucalion** is a work-in-progress
- JSC has expressed significant interest to make EESSI available on **JUPITER**

EESSI is already available on various other European systems (and beyond)

- Snellius @ SURF, EMBL, Univ. of Stuttgart, VSC sites in Belgium, Sigma2 in Norway, etc.

Overview of (known) systems that have EESSI available at eessi.io/docs/systems

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Over 450 different software already available

- Including: ALL, CP2K, **ESPResSo**, **LAMMPS**, OpenFOAM, ParaView, PETSc, Python + various PyPI packages incl. numpy/scipy/pandas/…, QuantumESPRESSO, ROOT, R + packages from CRAN and Bioconductor, **waLBerla**, WRF, …

- CUDA software installations for NVIDIA GPUs: **ESPResSo**, **LAMMPS** *(more coming soon: GROMACS, PyTorch, TensorFlow, AlphaFold, …)*

- See also eessi.io/docs/available_software/overview

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Support for NVIDIA GPUs

- Initial support for software built to run on **NVIDIA GPUs** is available
- Currently including installations of ESPResSo, LAMMPS (+ testing tools)
- Availability of NVIDIA GPUs is **auto-detected**, along with CPUs
- **GPU drivers must be exposed** to EESSI (by system administrators)
  - $LD_PRELOAD can be used as a (temporary) workaround if necessary
- Support for AMD GPUs is a work-in-progress (ETA summer 2025)
- More information: eessi.io/docs/site_specific_config/gpu

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Continuous Integration & Deployment (CI/CD)

- EESSI can be used in **CI environments** like GitHub Actions and GitLab CI
  - To provide required (build) dependencies for your software
  - EESSI makes setting up CI environment a lot easier
- Uniform software environment across systems **and CI**
- Deploy your software *once* via EESSI, run *everywhere*
- Already used in the wild, see for example pyMBE's GitHub repository
- See also eessi.io/docs/using_eessi/eessi_in_ci

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

Building and deploying pre-release software versions

- Our production repository `software.eessi.io` is for software *releases*
- Software developers sometimes also want to deploy **pre-release builds**
  - To test improvements, evaluate performance enhancements, etc.
- We are currently setting up a procedure for this in the scope of MultiXscale
- Developers will be able to easily **build & deploy development versions** of their code via a separate dedicated repository (`dev.eessi.io`)
- Keep an eye on our documentation for more information: eessi.io/docs

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

EPICURE

EESSI

- Introduction to EESSI, the European Environment for Scientific Software Installations
- **Friday 15 Nov 2024, 14:00-15:30 CEST**
- Free webinar, online
- Registration via https://epicure-hpc.eu/2024/10/17/webinar-streaming-optimised-scientific-software-an-introduction-to-eessi

Multiscale

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

EuroHPC Joint Undertaking

Walk-in networking session at EuroHPC User Day 2024

- Wednesday 23 Oct 2024, **10:30-12:00 in room Waterfront**
- Walk-in networking session *“Application Support & Skills”*
- Come visit the **MultiXscale stand**, and **experience EESSI hands-on**!
- Join our prize draw and **win a Raspberry Pi 5** starter pack!

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No. 101093169

EuroHPC Joint Undertaking

Multi scale

Web page: multixscale.eu

Facebook: MultiXscale

X: @MultiXscale

LinkedIn: multixscale

YouTube: @MultiXscale

Co-funded by the European Union

EuroHPC Joint Undertaking

KEMIJSKI INŠTITUT

UNIVERSITAT DE BARCELONA

Universität Stuttgart

UNIVERSITAS BERGENSIS

SORBONNE UNIVERSITÉ

Université de Toulouse

Consiglio Nazionale delle Ricerche

MAX-PLANCK-GESELLSCHAFT

HPC Now!

university of groningen

SURF

JÜLICH Forschungszentrum

GHENT UNIVERSITY

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

LEONARDO

iit ISTITUTO ITALIANO DI TECNOLOGIA

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and countries participating in the project under grant agreement No 101093169.