



E E S S I


EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

Lara Peeters - Ghent University (Belgium)

lara.peeters@ugent.be

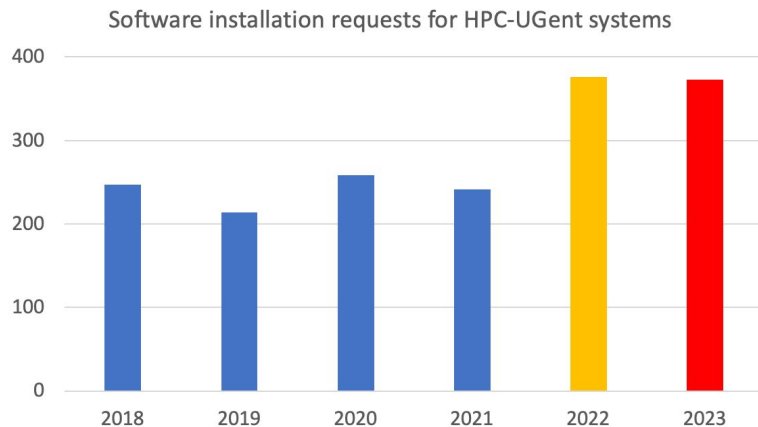
Tue 3 Dec 2024 - Dnevi hackathon in Ljubljana

Multi-scale




The changing landscape of scientific computing

- **Explosion of available scientific software** applications (bioinformatics, AI boom, ...)
- Increasing interest in **cloud** for scientific computing (flexibility!)
- **Increasing variety in processor (micro)architectures** beyond Intel & AMD: Arm is coming already here (see [Fugaku](#), [JUPITER](#), ...), RISC-V is coming (soon?)
- In strong contrast: available (wo)manpower **in HPC support teams is (still) limited...**



Optimized scientific software installations

- Software should be optimized for the system it will run on (keep the P in HPC!)
- Impact on performance is often significant for scientific software!
- Example: GROMACS 2020.1 (PRACE benchmark, Test Case B)
- Metric: (simulated) ns/day, higher is better
- Test system: dual-socket Intel Xeon Gold 6420 (Cascade Lake, 2x18 cores)
- Performance of different GROMACS binaries, on exact same hardware/OS



*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 in a nutshell

- European Environment for Scientific Software Installations (EESSI)
- **Shared repository of (optimized!) scientific software installations**
- Uniform way of providing software to users, regardless of the system they use!
- Should work on any Linux OS (+ WSL, macOS via Lima) and system architecture
- From laptops and personal workstations to HPC clusters and cloud
- Support for different CPU (micro)architectures, interconnects, GPUs, etc.
- **Focus on performance, automation, testing, collaboration**



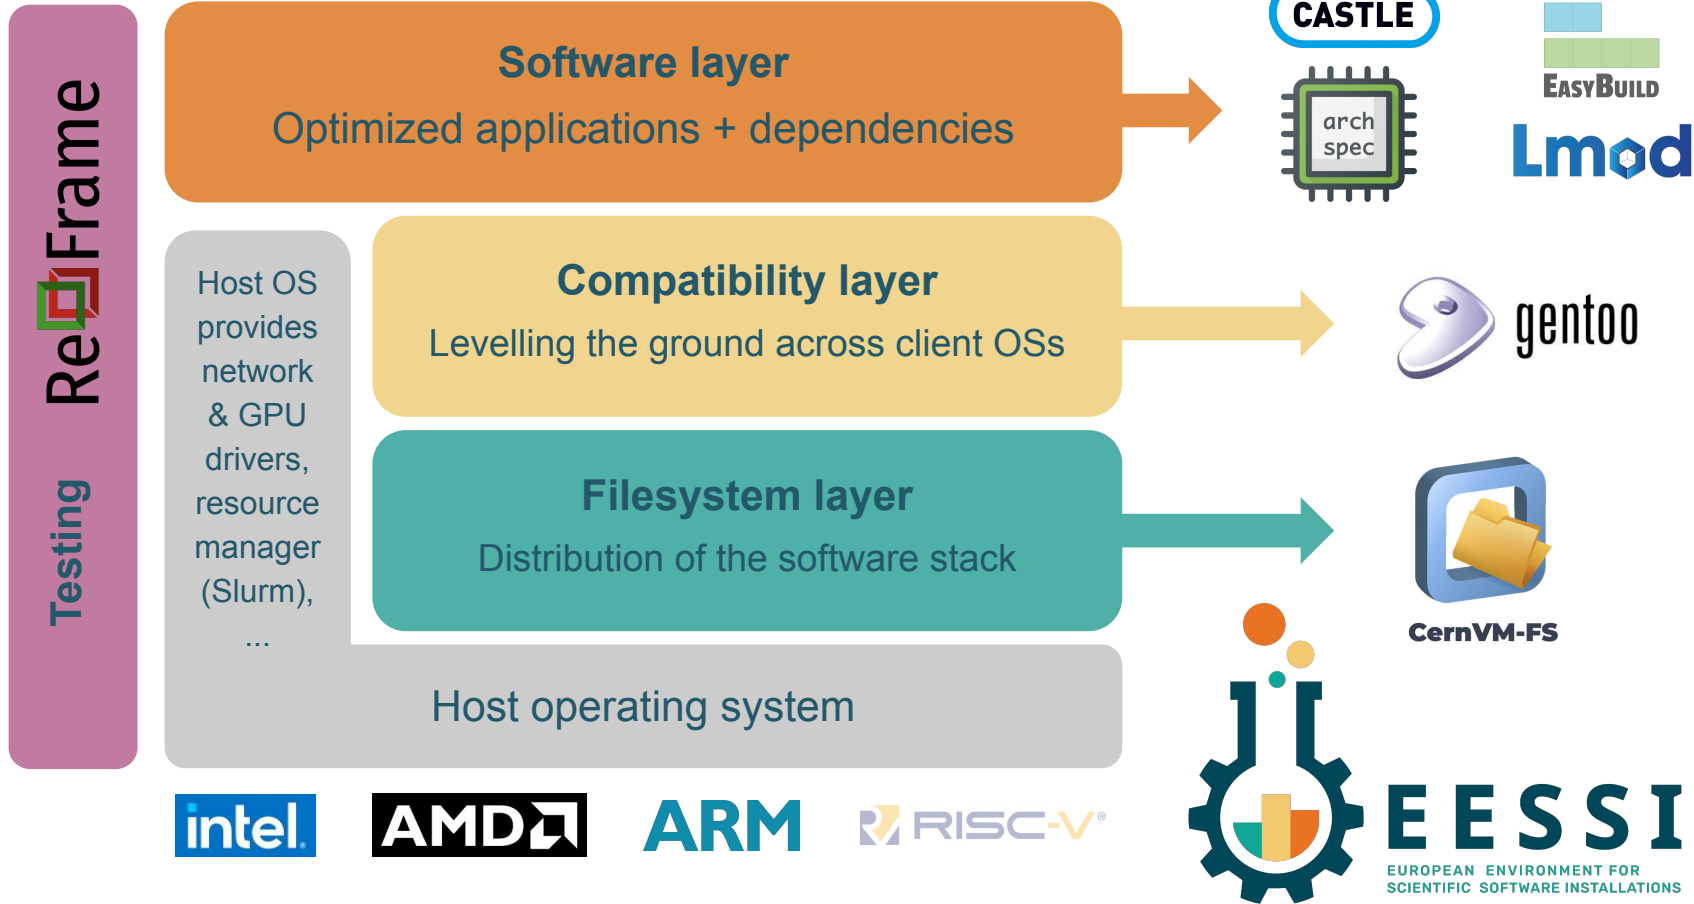
E E S S I

EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

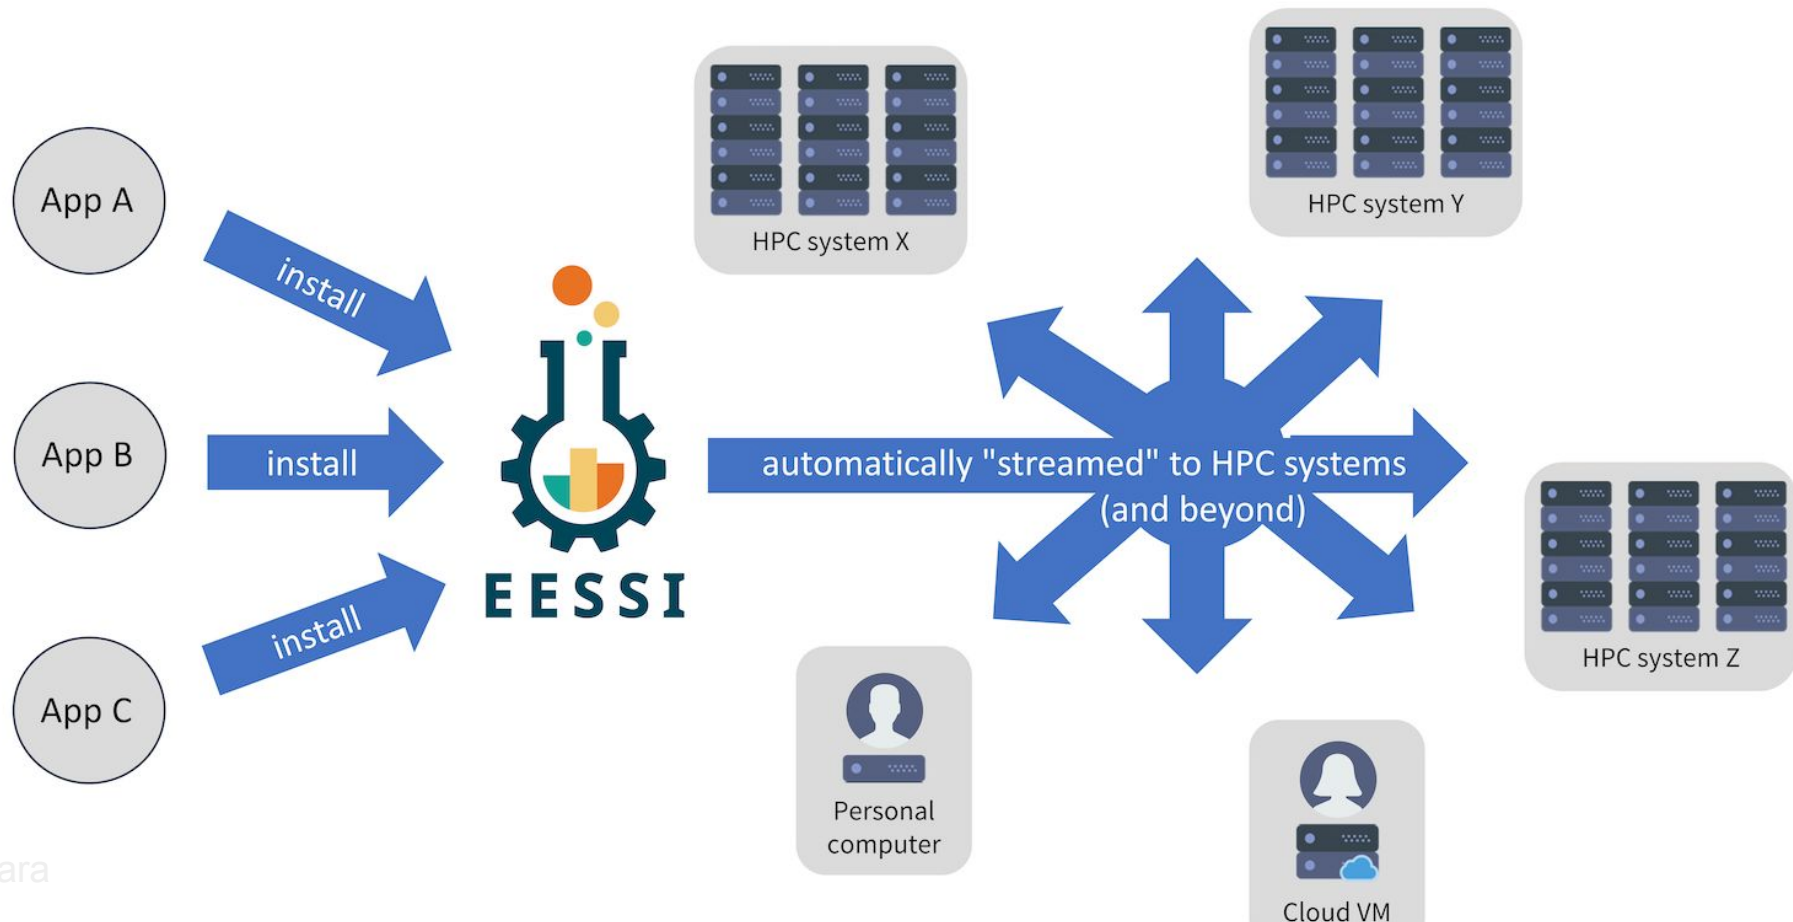
<https://eessi.io>

<https://eessi.io/docs>

High-level overview of EESSI

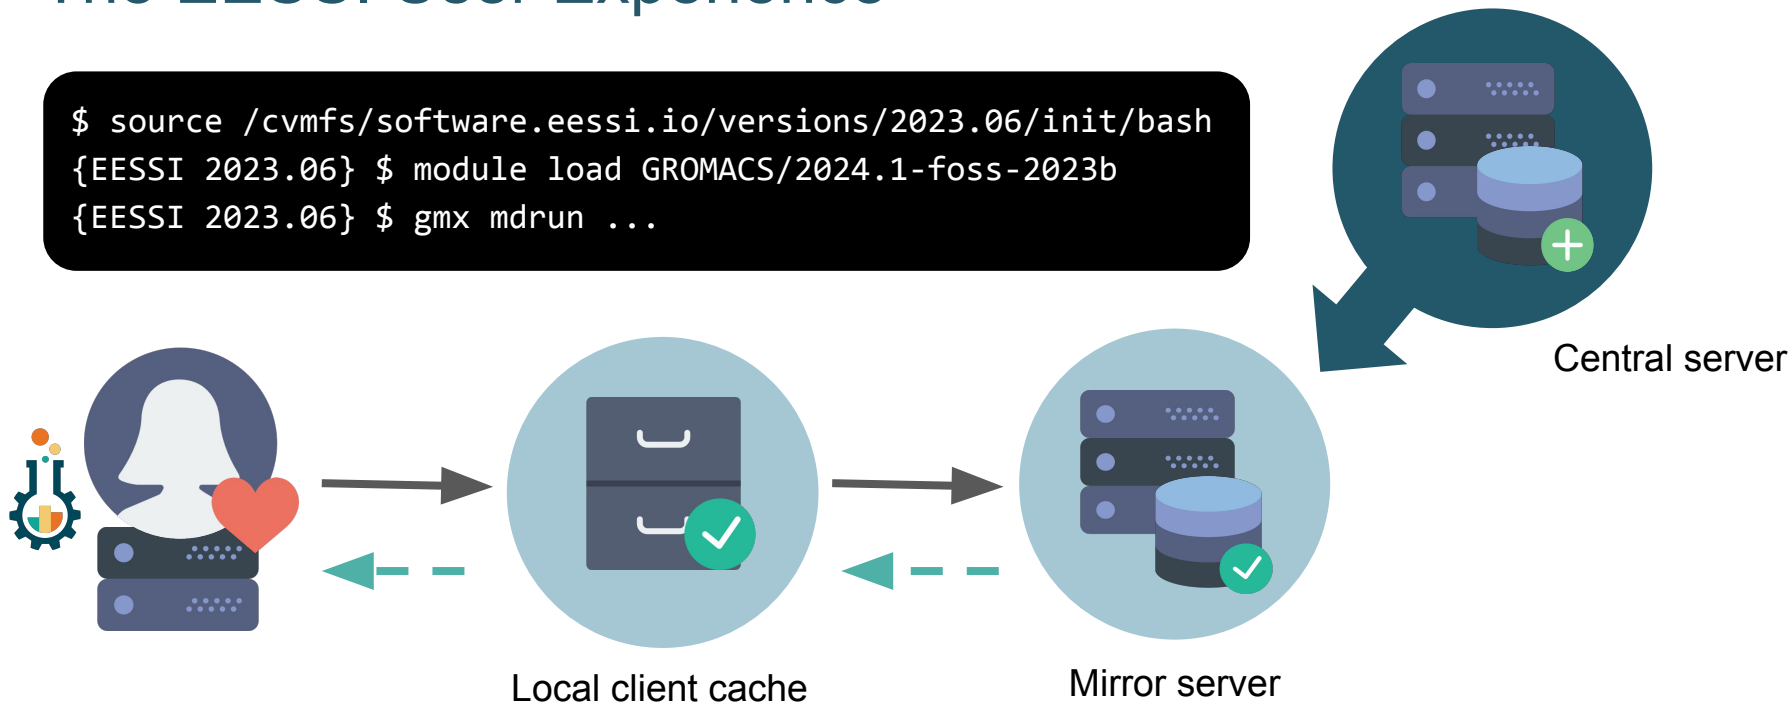


EESSI as a shared software stack



The EESSI User Experience

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
{EESSI 2023.06} $ module load GROMACS/2024.1-foss-2023b
{EESSI 2023.06} $ gmx mdrun ...
```



EESSI provides **on-demand streaming**
of (scientific) software (like music, TV-series, ...)

Demo: Accessing EESSI via CernVM-FS on Vega

```
[{user}@localhost ~]$ ssh {user}@login.vega.izum.si
# Go through two-factor authentication
# Once successful you should be logged in
# Start interactive session on Vega
[{user}@vglogin0006 ~]$ srun --partition=dev --nodes=1 --ntasks=4 --time=0:30:00 --pty
bash -i
[{user}@{partition} ~]$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
Found EESSI repo @ /cvmfs/software.eessi.io/versions/2023.06!
archdetect says x86_64/amd/zen2
Using x86_64/amd/zen2 as software subdirectory.
Using
/cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/amd/zen2/modules
/all as the directory to be added to MODULEPATH.
Using
...
Environment set up to use EESSI (2023.06), have fun!
```



Alternative ways of accessing EESSI are available, via a container image, via cvmfsexec, ...
eessi.io/docs/getting_access/native_installation - eessi.io/docs/getting_access/eessi_container

Native installation of CernVM-FS

```
# Native installation
# Installation commands for RHEL-based distros
# like CentOS, Rocky Linux, AlmaLinux, Fedora, ...

# install CernVM-FS

sudo yum install -y
https://ecsft.cern.ch/dist/cvmfs/cvmfs-release/cvmfs-release-latest.noarch.rpm
sudo yum install -y cvmfs

# create client configuration file for CernVM-FS
# (no proxy, 10GB local CernVM-FS client cache)
sudo bash -c "echo 'CVMFS_CLIENT_PROFILE='single'' > /etc/cvmfs/default.local"
sudo bash -c "echo 'CVMFS_QUOTA_LIMIT=10000' >> /etc/cvmfs/default.local"

# Make sure that EESSI CernVM-FS repository is accessible
sudo cvmfs_config setup
```



See docs for alternative ways of installing CernVM-FS natively, via a VM on a personal computer
eessi.io/docs/getting_access/eessi_wsl/ - eessi.io/docs/getting_access/eessi_limactl/

Demo: Using EESSI

eessi.io/docs/using_eessi/eessi_demos



```
/cvmfs/software.eessi.io/versions/2023.06/software
`-- linux
   |-- aarch64
   |   |-- generic
   |   |-- neoverse_n1
   |   `-- neoverse_v1
   `-- x86_64
       |-- amd
       |   |-- zen2
       |   `-- zen3
       |-- generic
       `-- intel
           |-- haswell
           `-- skylake_avx512
               |-- modules
               `-- software
```

```
$ source /cvmfs/software.eessi.io/versions/2023.06/init/bash
Found EESSI pilot repo @
/cvmfs/software.eessi.io/versions/2023.06!
```

```
archdetect says x86_64/amd/zen3
Using x86_64/amd/zen3 as software subdirectory
```

```
...
Environment set up to use EESSI pilot software stack, have fun!
```

```
{EESSI 2023.06} $ module load R/4.3.2-gfbb-2023a
```

```
{EESSI 2023.06} $ which R
/cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/
amd/zen3/software/R/4.3.2-gfbb-2023a/bin/R
```

```
{EESSI 2023.06} $ R --version
R version 4.3.2
```

Demo: Running LAMMPS on Vega

```
#!/bin/bash
# EESSI_demo.sh script
#SBATCH --job-name="EESSI_Demo_LAMMPS_lj"
#SBATCH --ntasks=4
#SBATCH --ntasks-per-node=4
#SBATCH --cpus-per-task=1
#SBATCH --output=EESSI_demo.out
#SBATCH --error=EESSI_demo.err
#SBATCH --time=0:30:0
#SBATCH --partition=dev
```

```
source /cvmfs/software.eessi.io/versions/2023.06/init/bash
module load LAMMPS/29Aug2024-foss-2023b-kokkos
mkdir /tmp/$USER && cd /tmp/$USER
curl -o in.lj https://raw.githubusercontent.com/lammps/lammps/refs/heads/develop/bench/in.lj
export OMP_NUM_THREADS=1
mpirun -np 4 lmp -in in.lj
rm -r /tmp/$USER
```



```
$ ssh {user}@login.vega.izum.si
# Run the script in the slide
$ sbatch EESSI_demo.sh
$ queue -u $USER
$ cat EESSI_demo.out
```

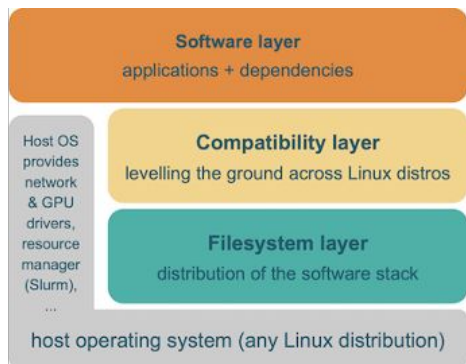


CernVM-FS

How does EESSI 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




Overview of available software



Currently ~920 software software installations available
per CPU target via software.eessi.io CernVM-FS repository;
increasing every day

- Over 465 different software packages
- Excl. extensions: Python packages, R libraries
- Including ESPResSo, GROMACS, LAMMPS, OpenFOAM, PyTorch, R, QuantumESPRESSO, TensorFlow, waLBerla, WRF, ...
- eessi.io/docs/available_software/overview
- Using recent compiler toolchains: currently focusing on `foss/2023a` and `foss/2023b`



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)
 - For now, only software installations for AMD Rome (Zen2) + NVIDIA A100 are available
- **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**

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 ([see blog post](#))
 - Native installation on **MeluXina, Deucalion, MareNostrum5** 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.
- EESSI has been **integrated in Azure HPC OnDemand Platform** (<https://azure.github.io/az-hop>)
- **Overview of (known) systems that have EESSI available at eessi.io/docs/systems**

Software testing is an important part of EESSI



- EESSI test suite: eessi.io/docs/test-suite
 - Collection of *portable* tests for software available in EESSI
- Example: failing tests in GROMACS test suite when installing it in EESSI
 - See <https://gitlab.com/eessi/support/-/issues/47>
 - Filesystem race in GROMACS test suite when running tests concurrently
 - **Bug in Arm SVE support**, leading to (very) wrong results for several tests
 - See <https://gitlab.com/gromacs/gromacs/-/issues/5057>
 - Works fine on A64FX (512-bit SVE), but problem on Graviton 3 + NVIDIA Grace!

Leveraging EESSI in CI environments



We have an EESSI GitHub Action that provides EESSI+di renv:

See it in action in the `github-eessi-action` repository:

github.com/EESSI/github-action-eessi

github.com/EESSI/github-action-eessi/blob/main/.github/workflows/tensorflow-usage.yml

```
name: ubuntu_tensorflow
on: [push, pull_request]
jobs:
```

```
  build:
```

```
    runs-on: ubuntu-latest
```

```
    steps:
```

```
      - uses: actions/checkout@v3
```

```
      - uses: eessi/github-action-eessi@v3
```

```
      with:
```

```
        eessi_stack_version: '2023.06'
```

```
      - name: Test EESSI
```

```
        shell: bash
```

```
        run: |
```

```
          module load TensorFlow
```

```
          python -c 'import tensorflow; print(tensorflow.__version__)'
```



Deployment of test-release of Scientific Software with EESSI



- **Dev.eessi.io**
 - Available on Vega
 - <https://eessi.io/docs/repositories/dev.eessi.io>
 - More information Coming soon

```
/cvmfs/dev.eessi.io/versions/2023.06/example/software
`-- linux
  `-- x86_64
    |-- amd
    |  |-- zen2
    |     |-- modules
    |        |-- all
    |           |-- Espresso
    |              |-- 4.2.2-foss-2023a-2ba17de6096933275abec0550981d9122e4e5f28.lua
    |           |-- LAMMPS
    |              |-- 0cb72423b8ed2fdf138831c145a3bfb6ea42394e-foss-2023a-kokkos-dev_OBMD.lua
    |           |-- ...
$ module use
/cvmfs/dev.eessi.io/versions/2023.06/example/software/linux/x86_64/amd/zen2/modules/all
```

Support for installing, using, contributing to EESSI



eessi.io/docs/support

- Via GitLab, or via email: support@eessi.io
- Report problems
- Ask questions
- Request additional software
- Get help with contributing to EESSI
- Suggest enhancements, additional features, ...
- Confidential tickets possible (security issues, ...)

Project

Search or go to...

EESSI / EESSI support portal

Project

EESSI support portal

Manage >

Plan >

Code >

Build >

Deploy >

Operate >



Monitor >

Analyze >

Help

README.md

EESSI support portal

MultiXscale  

Thanks to the [MultiXscale EuroHPC project](#) we are able to provide support to the u

Contact

Create an issue with you GitLab account

If you have a GitLab account or create one you can create and manage your issue also use one of our issue templates.

Contact us via E-mail


If you do not have a GitLab account you can also ask for support via E-mail.

Dedicated support team, thanks to EuroHPC Centre-of-Excellence



EESSI won an HPCWire Reader's Choice award!





Website: eessi.io

GitHub: github.com/eessi

Documentation: eessi.io/docs

Blog: eessi.io/docs/blog

[Join](#) the EESSI Slack

YouTube channel: youtube.com/@eessi_community

Paper (open access): doi.org/10.1002/spe.3075

EESSI support portal: gitlab.com/eessi/support

[Bi-monthly online meetings](#) (1st Thu, odd months, 2pm CE(S)T)

MultiXscale

Web page: multixscale.eu

Facebook: [MultiXscale](https://www.facebook.com/MultiXscale)

Twitter: [@MultiXscale](https://twitter.com/MultiXscale)

LinkedIn: [MultiXscale](https://www.linkedin.com/company/multixscale)



Co-funded by
the European Union



EuroHPC
Joint Undertaking



UNIVERSITAT DE
BARCELONA



Universität
Stuttgart



SORBONNE
UNIVERSITÉ



Université
de Toulouse



Consiglio Nazionale
delle Ricerche



MAX-PLANCK-GESSELLSCHAFT

