Multiscale

EPICURE

EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

Installing software in ARM can be EESSI

Kenneth Hoste (Ghent University, Belgium)

Wed 5 Feb 2025 @ EPICURE HPC in ARM Architecture Hackathon

MultiXscale Centre-of-Excellence in a nutshell

- 4-year project (started in Jan 2023), ~€6M budget
- 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: https://multixscale.eu

Alan

Landscape of scientific computing is changing

- **Explosion of available scientific software** applications (bioinformatics, AI, …)
- Increasing interest in **cloud** for scientific computing (flexibility!)
- **Increasing variety in processor (micro)architectures** beyond Intel & AMD:
  Arm is ~~coming~~ already here (Fugaku, Deucalion, JUPITER, AWS Graviton, NVIDIA Grace, …),
  RISC-V is coming (soon?)
- Broader adoption of **accelerated computing**, beyond NVIDIA GPUs (AMD, Intel, …)
- In strong contrast: available (wo)manpower **in HPC support teams is (still) limited**…

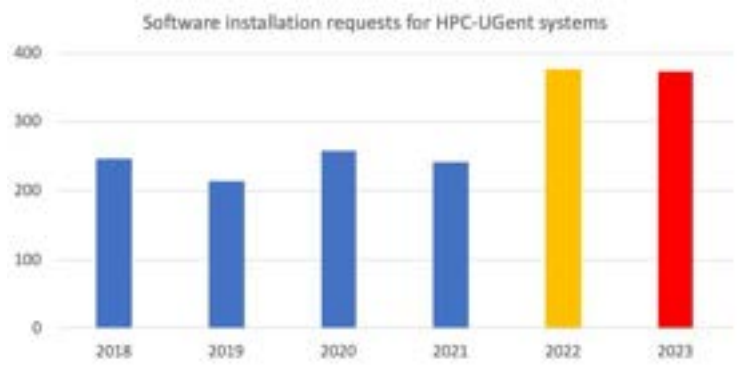

Optimized scientific software installations

- 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**

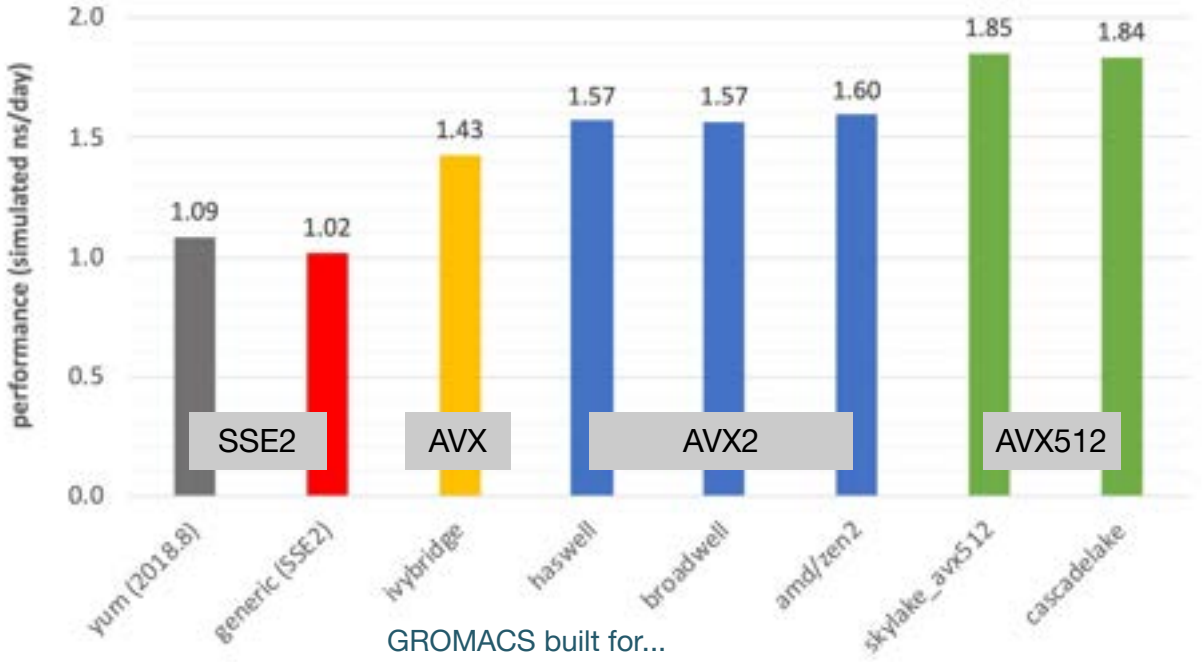

Alan

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?

Alan

European Environment for Scientific Software Installations

- **Public repository of (optimized!) scientific software *installations***
- **Avoid duplicate work** by collaborating on a shared software stack
- **Uniform way of providing software** to users, regardless of the system they use!
- **Should work on any Linux OS** (incl. WSL & macOS) **and system architecture**
  - From laptops and personal workstations to HPC clusters and cloud
  - Support for different CPUs (AMD, Intel, Arm, RISC-V), interconnects, GPUs, etc.
- Focus on **performance, automation, testing, collaboration**

EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

https://eessi.io

https://eessi.io/docs

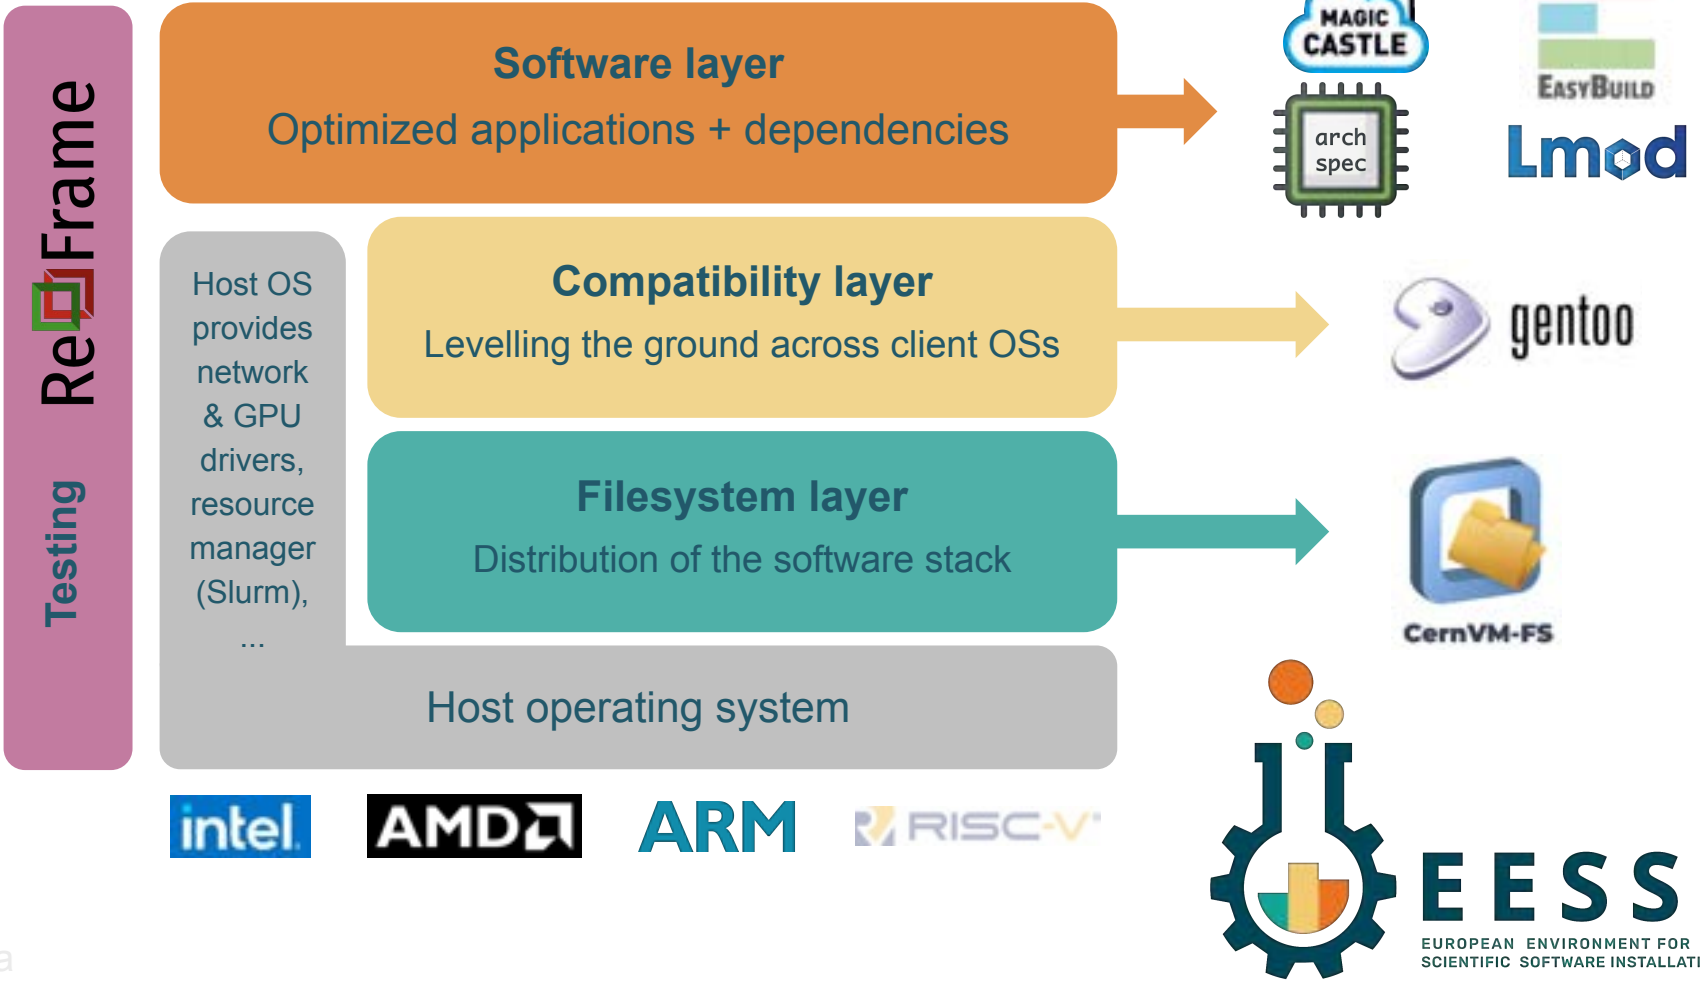
ReFrame
Testing
Software layer
Optimized applications + dependencies
MAGIC CASTLE
arch spec
EasyBuild
Lmod
Host OS provides network & GPU drivers, resource manager (Slurm), ...
Compatibility layer
Levelling the ground across client OSs
gentoo
Filesystem layer
Distribution of the software stack
CernVM-FS
Host operating system
intel
AMD
ARM
RISC-V
EESSI
EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

EESSI as a shared software stack in EuroHPC

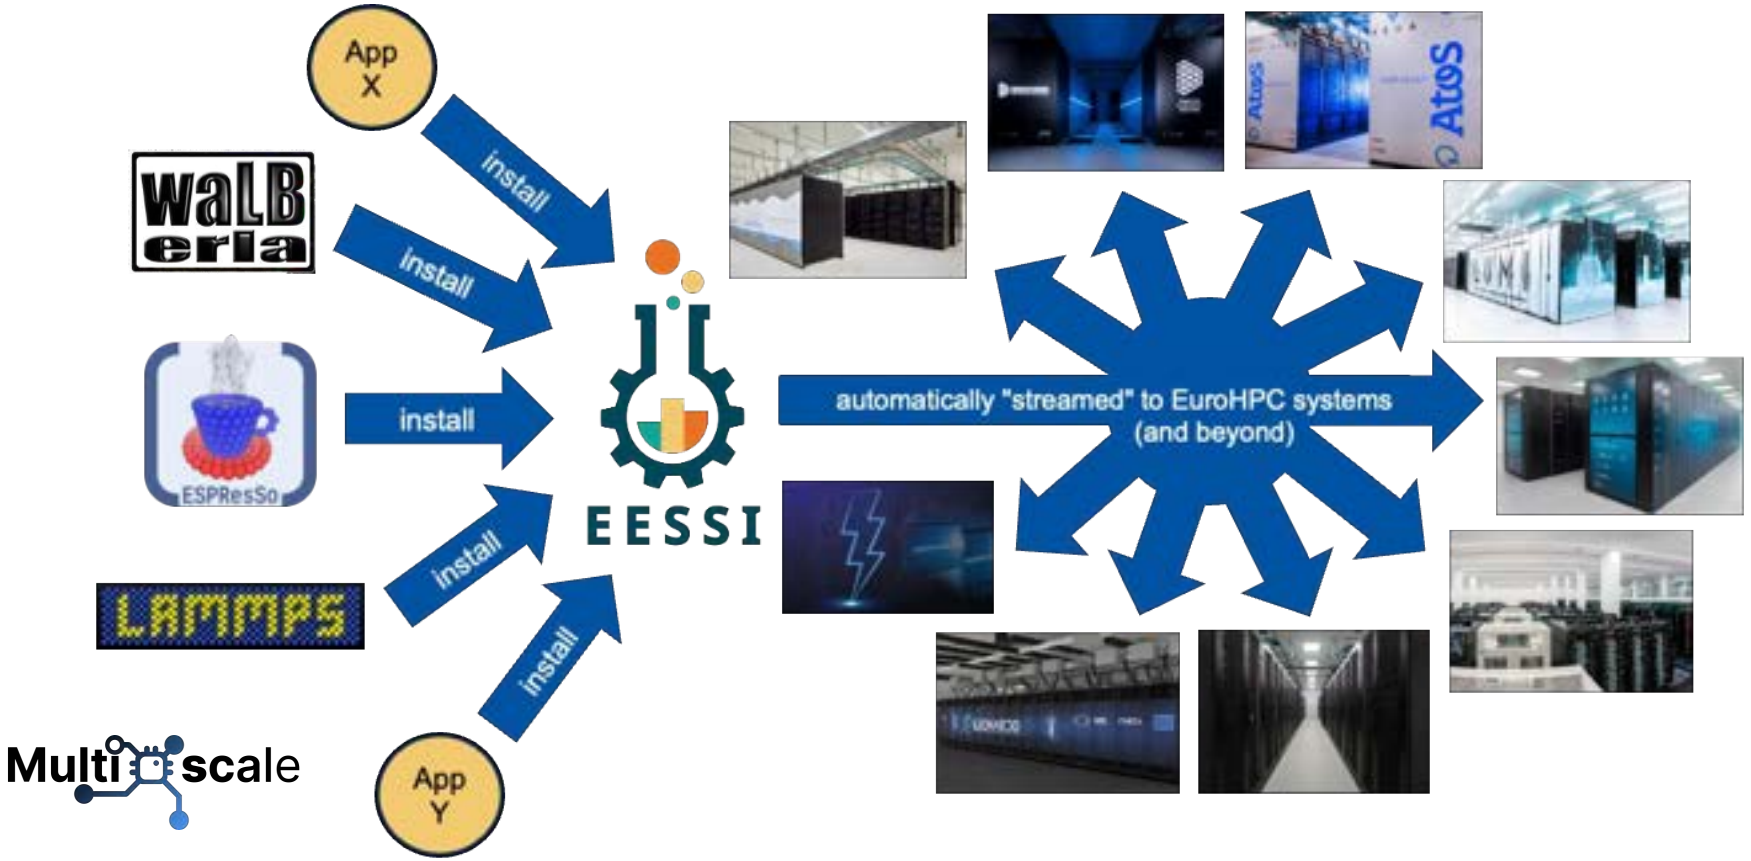

EESSI as a shared software stack (more general)

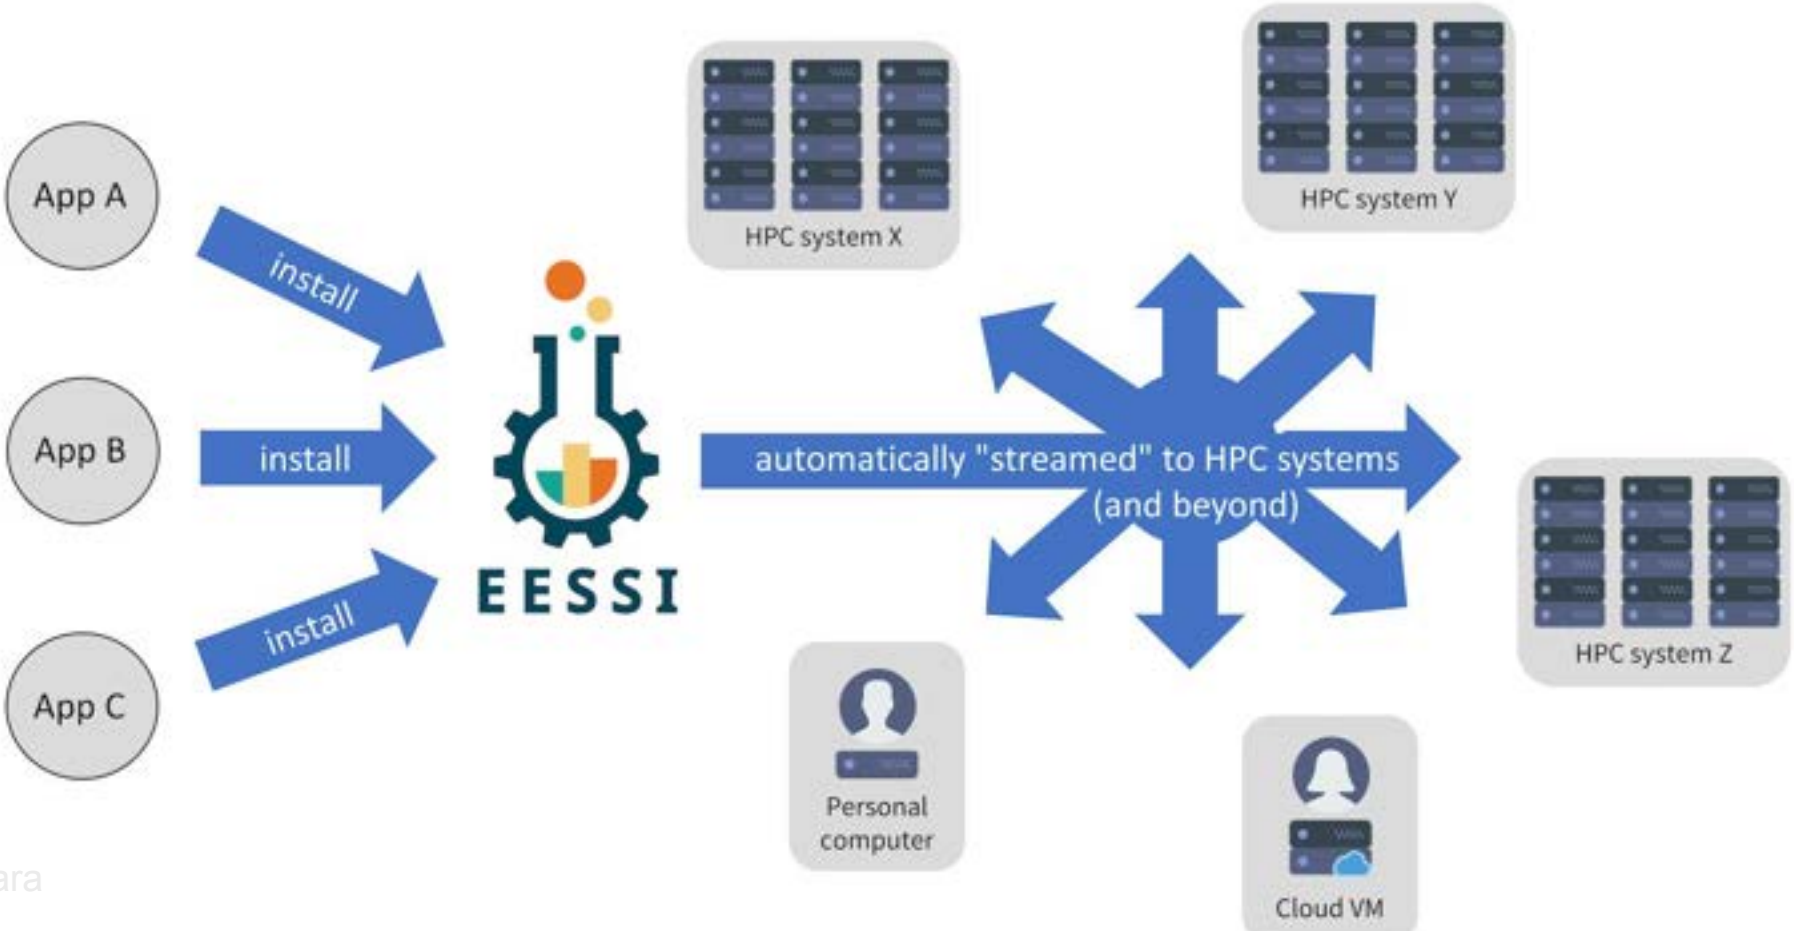

Getting access to EESSI

- Native installation of CernVM-FS *(requires admin privileges)*
  eessi.io/docs/getting_access/native_installation
- Using a container (via Apptainer)
  eessi.io/docs/getting_access/eessi_container
- Via cvmfsexec
  github.com/cvmfs/cvmfsexec

To check whether you have access to EESSI:

```
ls /cvmfs/software.eessi.io
```

Getting access EESSI via CernVM-FS (demo)

```
# 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
```

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

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 ...
```

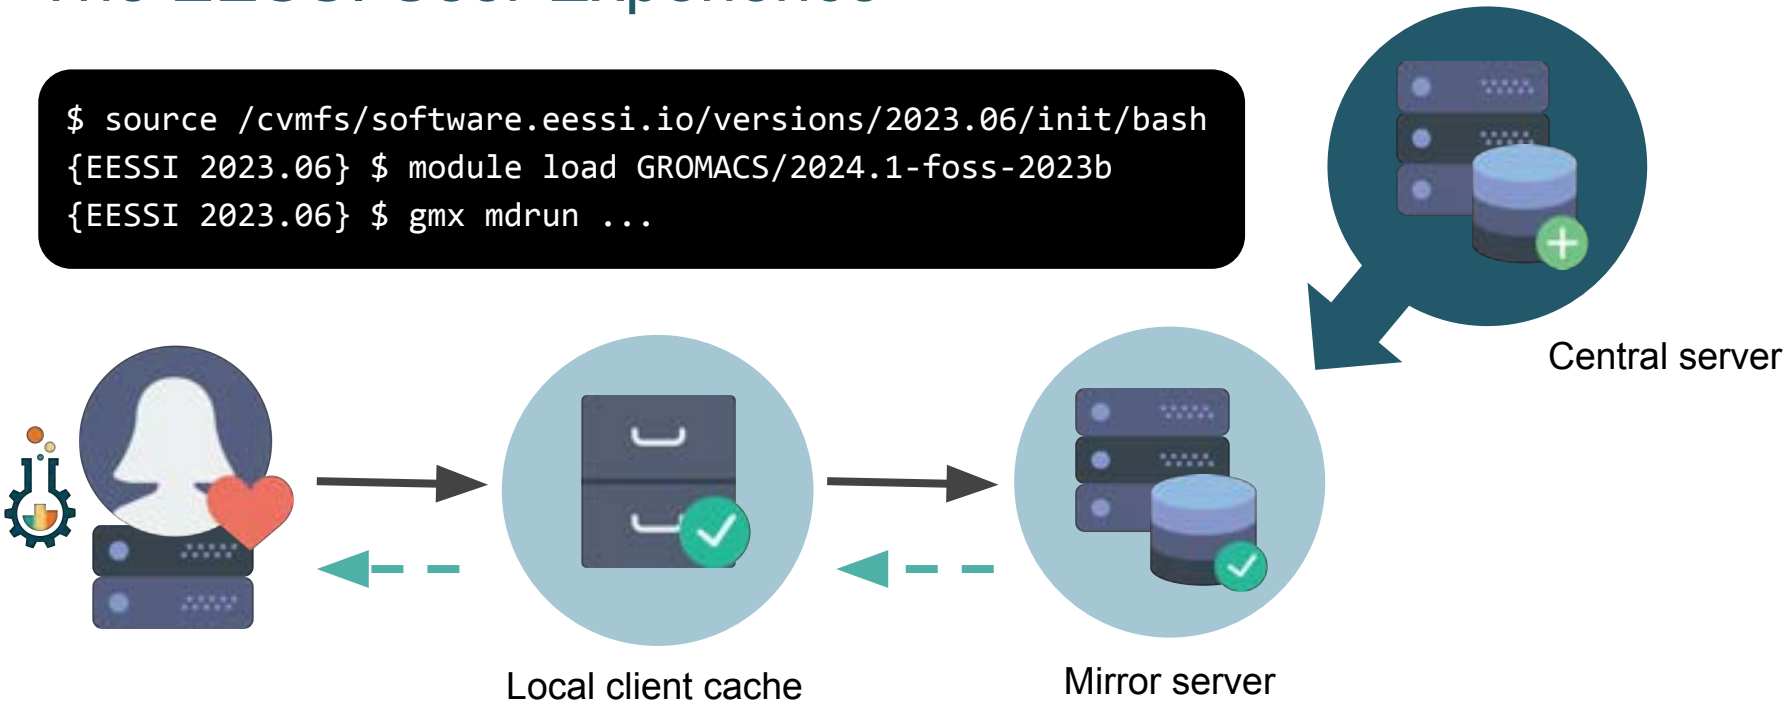

EESSI provides **on-demand streaming**

of (scientific) software (like music, TV-series, …)

Using EESSI (demo)

eessi.io/docs/using_eessi/eessi_demos

```
/cvmfs/software.eessi.io/versions/2023.06/software
`-- linux
    |-- aarch64
    |   |-- a64fx
    |   |-- generic
    |   |-- neoverse_n1
    |   `-- neoverse_v1
    `-- x86_64
        |-- amd
        |   |-- zen2
        |   |-- zen3
        |   `-- zen4
        |-- generic
        `-- intel
            |-- haswell
            |-- sapphire_rapids
            `-- 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 aarch64/a64fx
Using aarch64/a64fx as software subdirectory
...
Environment set up to use EESSI pilot software stack, have fun!

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

{EESSI 2023.06} $ which R
/cvmfs/software.eessi.io/versions/2023.06/software/linux/aarch64
/a64fx/software/R/4.3.2-gfbf-2023a/bin/R

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

Demo: Running LAMMPS


```
#!/bin/bash

source /cvmfs/software.eessi.io/versions/2023.06/init/bash

module load LAMMPS/2Aug2023_update2-foss-2023a-kokkos

if [ ! -f in.lj ]; then
    curl -OL https://github.com/lammps/lammps/raw/refs/tags/stable_2Aug2023/bench/in.lj
fi

export OMP_NUM_THREADS=1

mpirun -np 4 lmp -in in.lj
```

How does EESSI work?

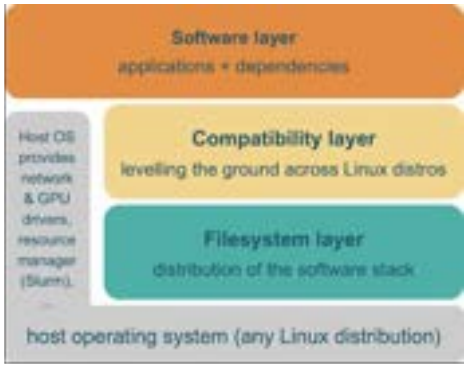

- Software installations included in EESSI are:
  - Automatically **“streamed in” on demand** (via CernVM-FS)
  - Built to be **independent of the host operating system**

    “Containers without the containing”
  - **Optimized** for specific CPU generations + specific GPU types
- Initialization script **auto-detects** CPU + GPU of the system

Lara

CernVM-FS network for EESSI

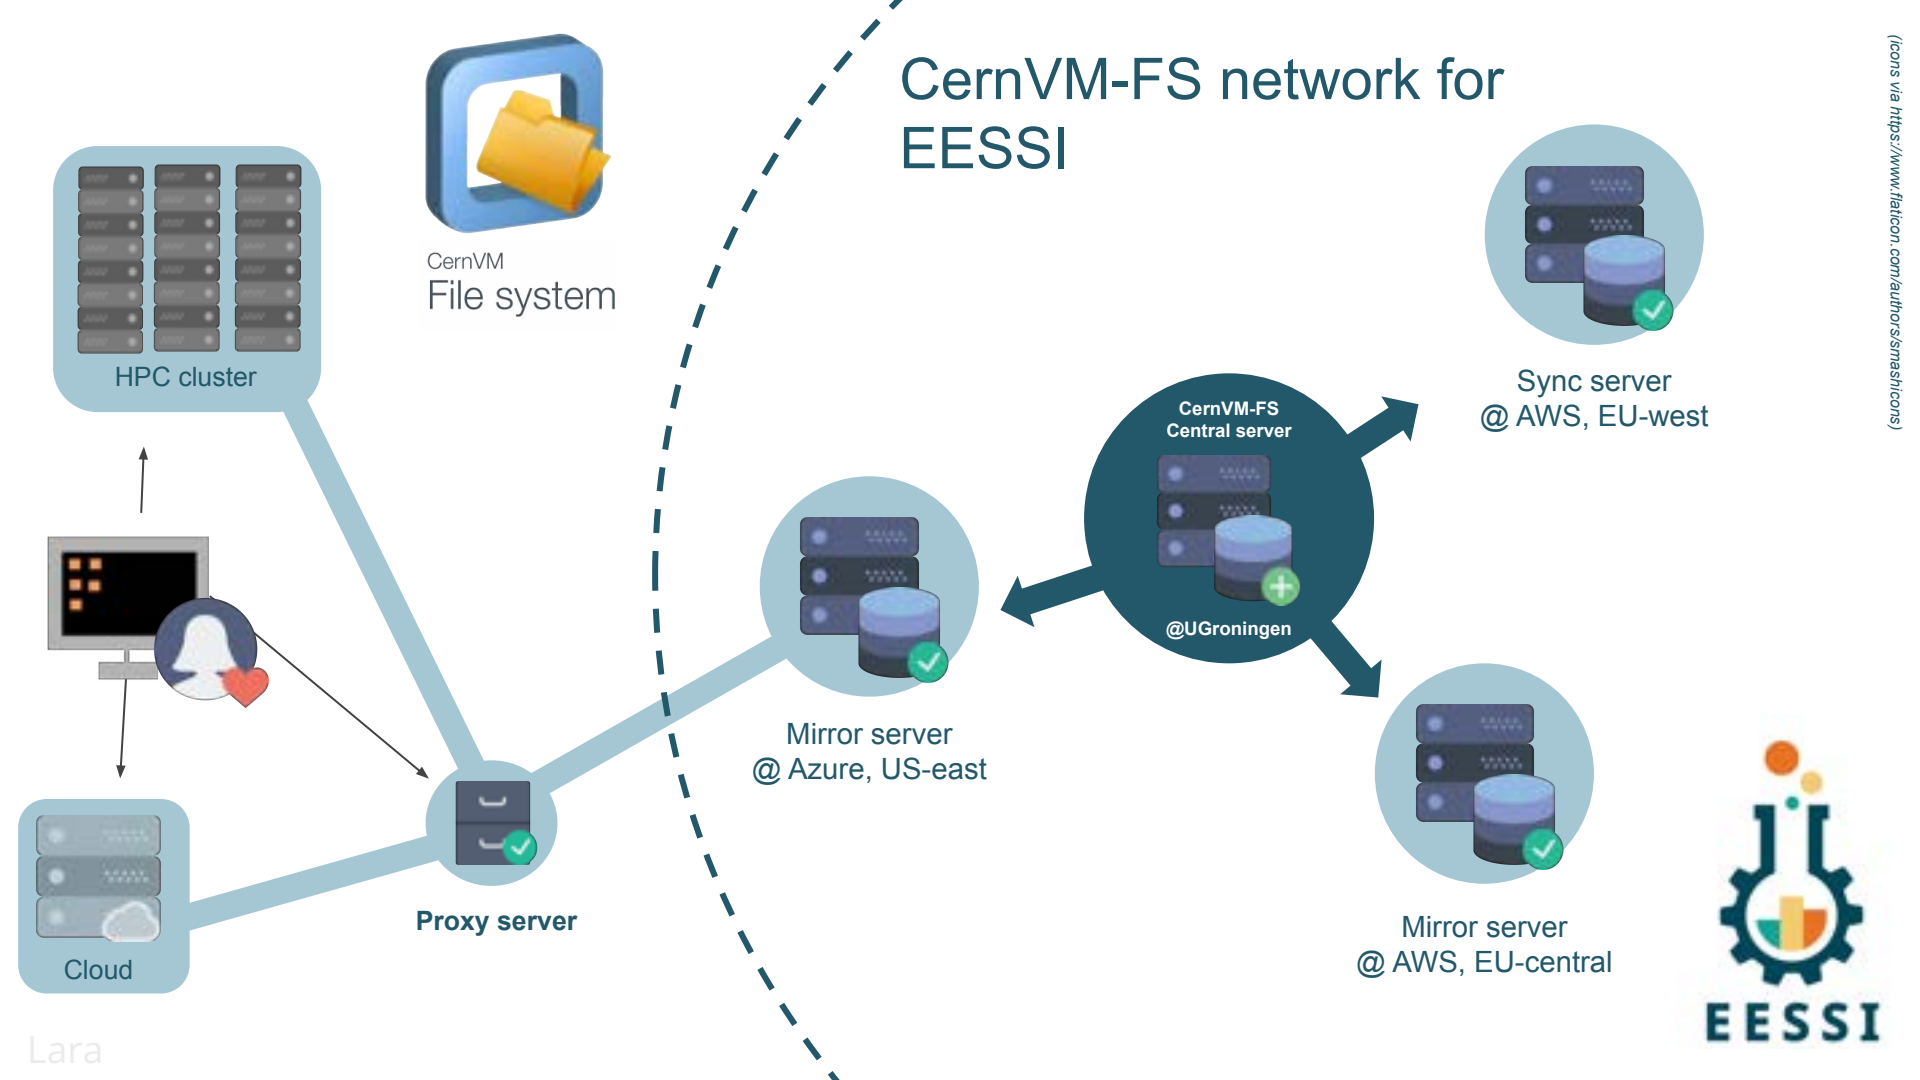

(icons via https://www.flaticon.com/authors/smashicons)

Lara

Tutorial “Best Practices for CernVM-FS in HPC”

- multixscale.github.io/cvmfs-tutorial-hpc-best-practices
- Held online on 4 Dec 2023 (~3 hours), **recorded & available on YouTube**
- Over 200 registrations, ~125 attending the meeting
- Lecture + hands-on demos
- Topics:
  - Introduction to CernVM-FS + EESSI
  - Configuring CernVM-FS: client, Stratum 1 mirror server, proxy server
  - Troubleshooting problems
  - Benchmarking of start-up performance w/ TensorFlow

CernVM
File system

Lara

EESSI compatibility layer

github.com/EESSI/compatibility-layer

- *“Containers without the containing”*
- **Minimal collection of tools and libraries** (incl. glibc, bash, Python, Lmod, …)
- **Built from source per CPU family** (`x86_64`, `aarch64`, …) with Gentoo Prefix
- Installations included in software layer **only link to compat layer** (RPATH)
- Ensures **compatibility** with any client system running Linux

```
$ ls /cvmfs/software.eessi.io/versions/2023.06/compat/linux/aarch64/
bin  etc  lib  lib64  opt  reprod  run  sbin  stage1.log  stage2.log
stage3.log  startprefix  tmp  usr  var
```

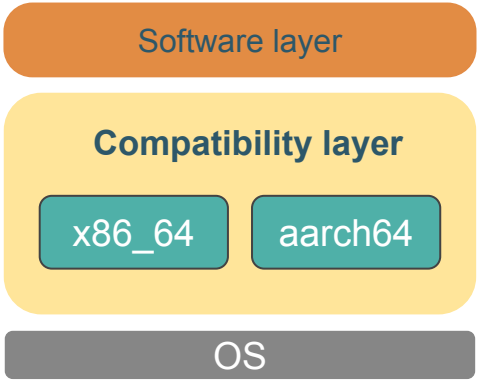

Lara

Software layer

github.com/EESSI/software-layer

- Installations of scientific software applications
- **Optimized for specific CPU targets**
- Works on any client system running Linux, since we only link to libraries in compat layer
- Built using EasyBuild
- Environment modules as user interface (via Lmod)
- Detection of host CPU via archspec (Python) or archdetect (bash)
- Magic Castle to manage our (build) clusters in the cloud (AWS, Azure)
- **Best subset of software installations for host CPU is automatically selected**

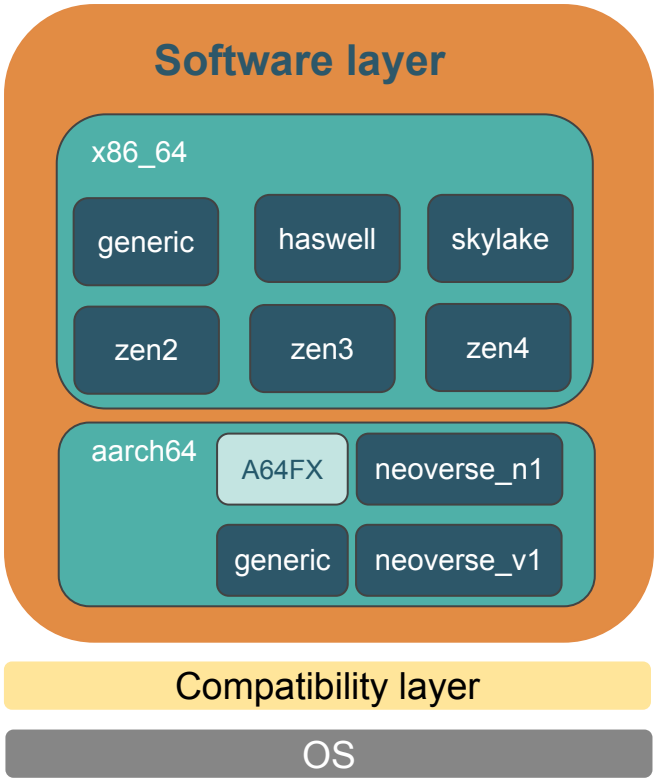

Lara

Overview of installed software

Currently ~933 software software installations available per CPU target via software.eessi.io CernVM-FS repository; increasing every day

- Over 470 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`
- **Partial coverage for A64FX, work-in-progress**

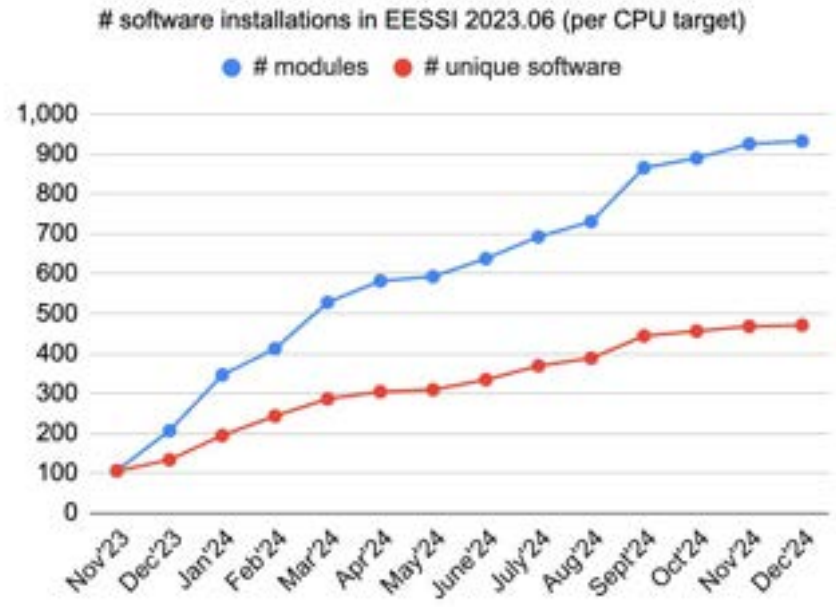

Semi-automated workflow for adding software to EESSI

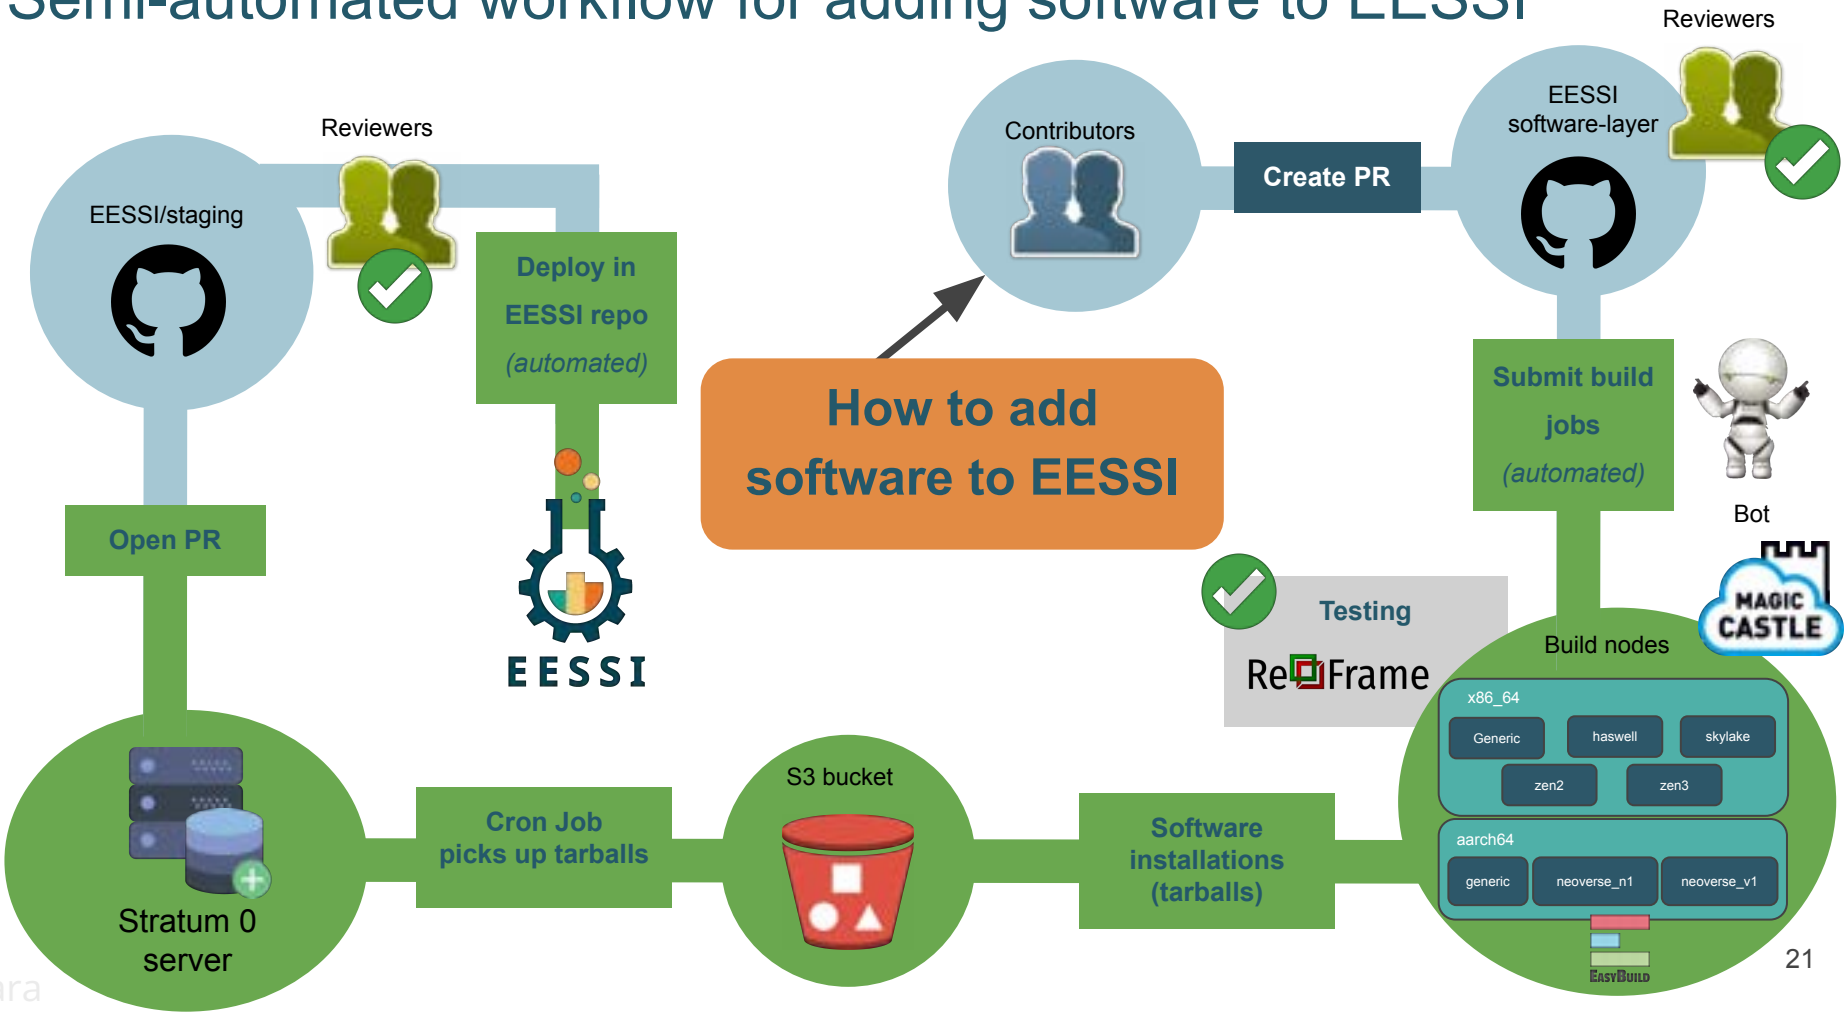

Software testing is an important part of EESSI

- 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 (now fixed)
  - Works fine on A64FX (512-bit SVE), but problem on Graviton 3 + NVIDIA Grace!

Leveraging EESSI in CI environment

Using EESSI in GitHub Actions is trivial (and works *really* well):

```
name: ubuntu_gromacs
on: [push, pull_request]
jobs:
 build:
   runs-on: ubuntu-latest
   steps:
   - uses: actions/checkout@v2
   - uses: eessi/github-action-eessi@v3
     with:
       eessi_stack_version: '2023.06'
   - name: Test EESSI
     run: |
       module load TensorFlow/2.13.0-foss-2023a
       python test_with_tensorflow.py
     shell: bash
```

github.com/EESSI/github-action-eessi

EESSI in a nutshell

- **On-demand streaming** of **optimized** scientific software installations
- **Works on any Linux distribution** thanks to EESSI compat layer
- **Uniform software stack** across various systems: laptop, HPC, cloud, …
- Community-oriented: **let's tackle the challenges we see together!**

EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

Getting support for EESSI

eessi.io/docs/support

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

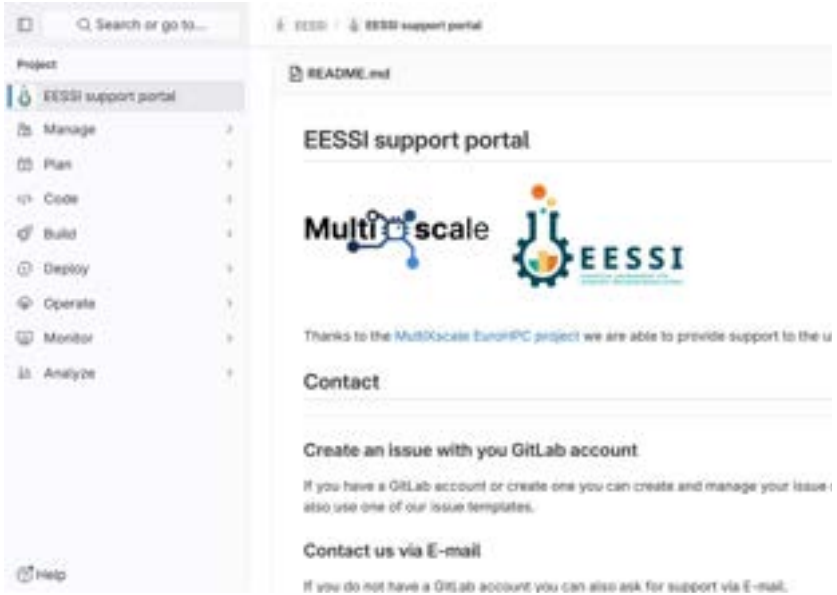

Dedicated support team, thanks to EuroHPC Centre-of-Excellence

HPCwire 2024 Readers' Choice Awards

HPCwire readers voted EESSI as
Best HPC Programming Tool or Technology 🎉

- Award was presented at the Supercomputing 2024 conference in Atlanta (Nov'24)
- Link to blog post with more details & pictures

Integration of EESSI in EuroHPC Federation Platform

EuroHPC
Joint Undertaking

- Consortium was selected to implement the EuroHPC Federation Platform (EFP), a “one-stop shop” to facilitate usage of EuroHPC systems
- 5-year project (2025-2029)
- Consortium led by CSC, Ghent University as one of the partners
- https://eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en
- Includes federated account, web portal (Open OnDemand), project management (Waldur), workflow service (LEXIS), …
- **EESSI will be integrated into EFP as part of the Federated Software Stack component !**
  - Important to note: EESSI itself (and its development) is **not** funded by EFP
- Talk at FOSDEM’25 on EFP in HPC devroom on Sun 2 Feb 2025 (recorded + live-streamed): https://fosdem.org/2025/schedule/event/fosdem-2025-6718-eurohpc-fp-a-federated-platform-for-hpc-infrastructure-in-europe-built-with-open-source-software/

[Kenneth, Lara]

Acknowledgements

- 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.

- Thanks to Amazon Web Services (AWS) and Microsoft Azure for generously sponsoring the EESSI project with cloud credits, feedback, and guidance.

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

Co-funded by the European Union

EuroHPC Joint Undertaking

Web page: multixscale.eu

Facebook: MultiXscale

Twitter: @MultiXscale

LinkedIn: MultiXscale

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.