EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

MultiScale

Introduction to EESSI

9th EasyBuild User, Umeå (2024-04-25)

Lara Peeters (HPC-UGent)

lara.peeters@ugent.be

GHENT UNIVERSITY

VLAAMS SUPERCOMPUTER CENTRUM

Who am I ?

GHENT UNIVERSITY

VLAAMS SUPERCOMPUTER CENTRUM

ugent.be/hpc

vscentrum.be

Lara Peeters: Digital Art Historian from Belgium

- Hired on the MultiXscale project at Ghent University (Belgium) since May 2023
- Active contributor to **EasyBuild** & **EESSI**, in the context of MultiXscale EuroHPC CoE
- Just getting started, still figuring out the “art” of software packaging

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**…

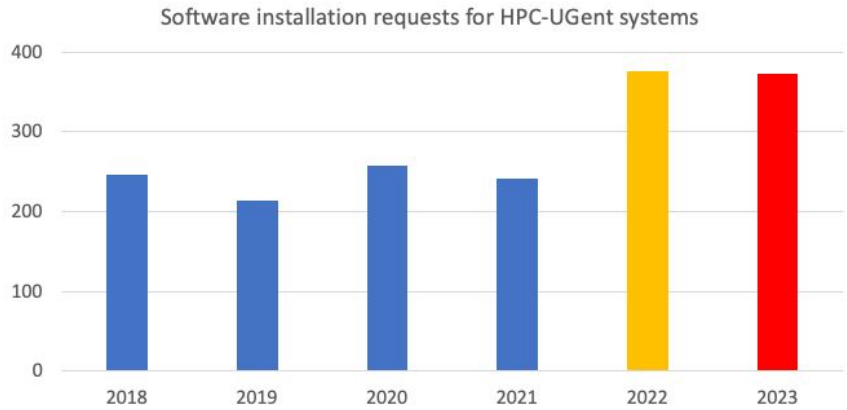

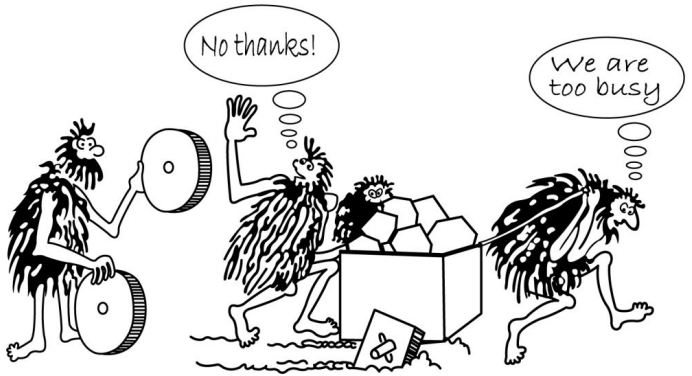

We need to collaborate more!

- **Too much software** for a single support team to handle
- Different systems (CPU, OS, …) => different problems
- **EasyBuild is not sufficient anymore…**
- **Duplicate work** across HPC sites and scientists
- **Diverse software stacks** across different platforms

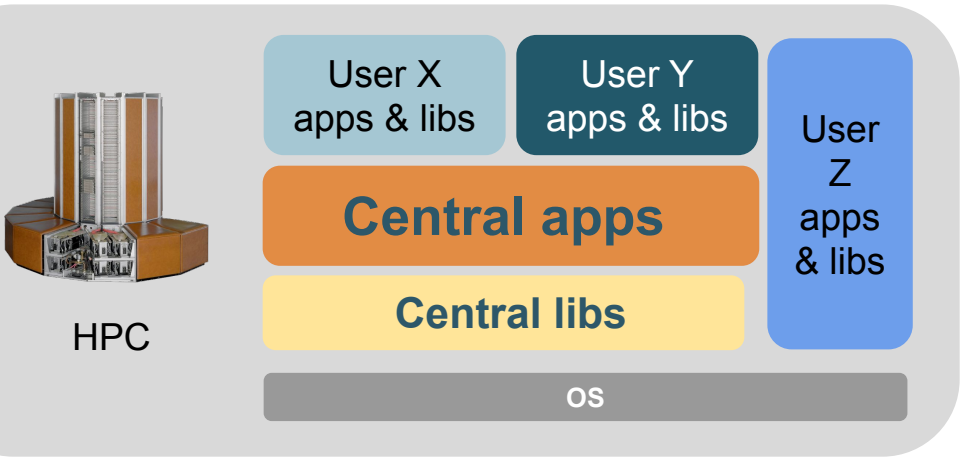

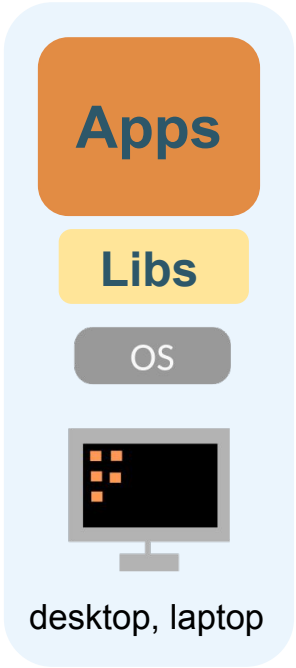

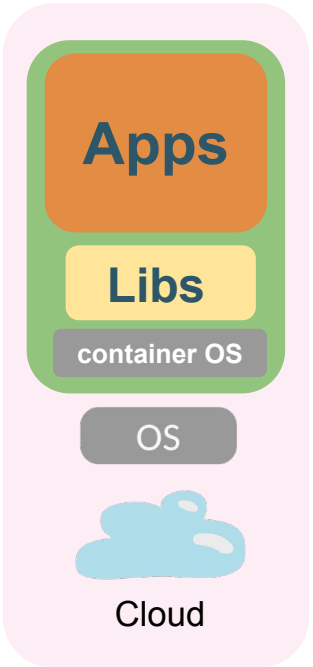

EESSI to the rescue

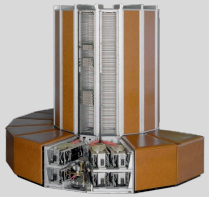

HPC

User X | User Y | User Z

EESSI

compat

OS

CPU

Software layer

x86_64

generic | haswell | skylake | zen2 | zen3

aarch64

generic | neoverse_n1 | neoverse_v1

Compatibility layer

x86_64 | aarch64

OS

Shared repository of **(optimized)** scientific software **installations**

Same software stack everywhere!

EESSI

compat

OS

CPU

desktop, laptop

EESSI

compat

OS

CPU

cloud

EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

eessi.io

eessi.io/docs (try out the pilot setup!)

eessi.io/docs/support

EESSI ingredients

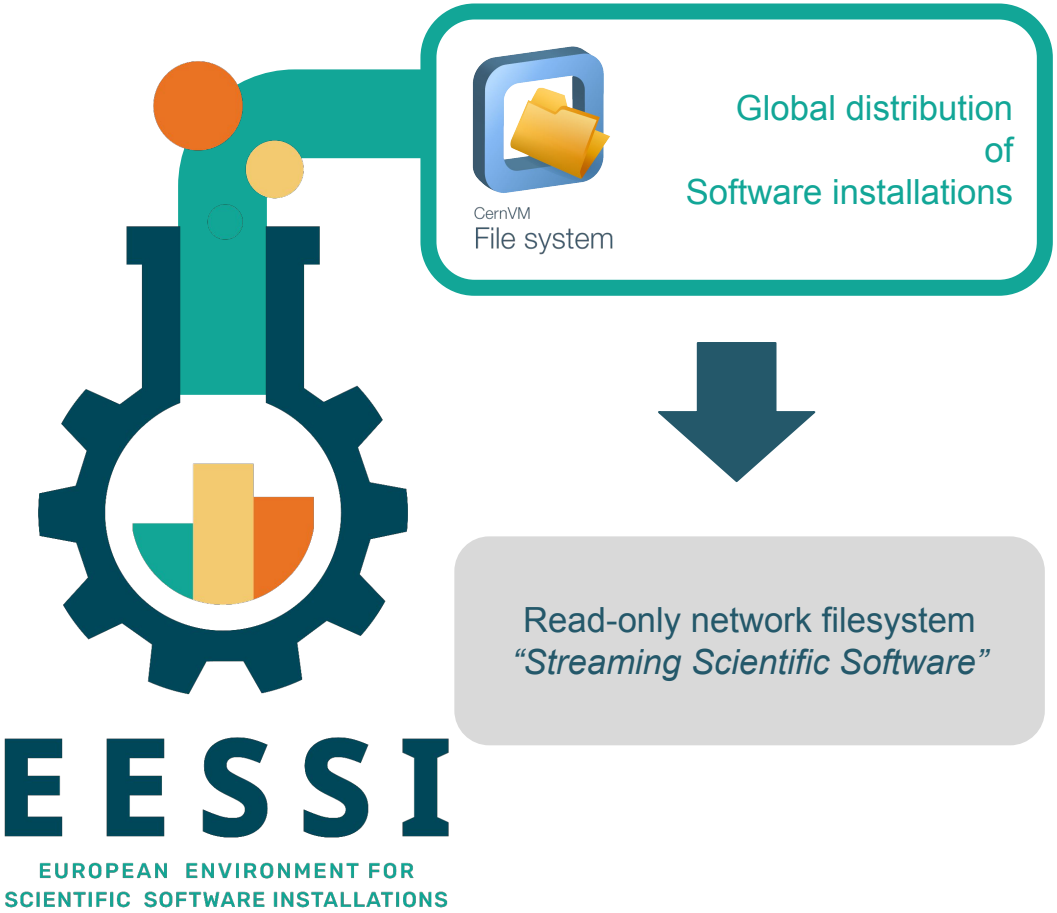

CernVM-FS network for EESSI

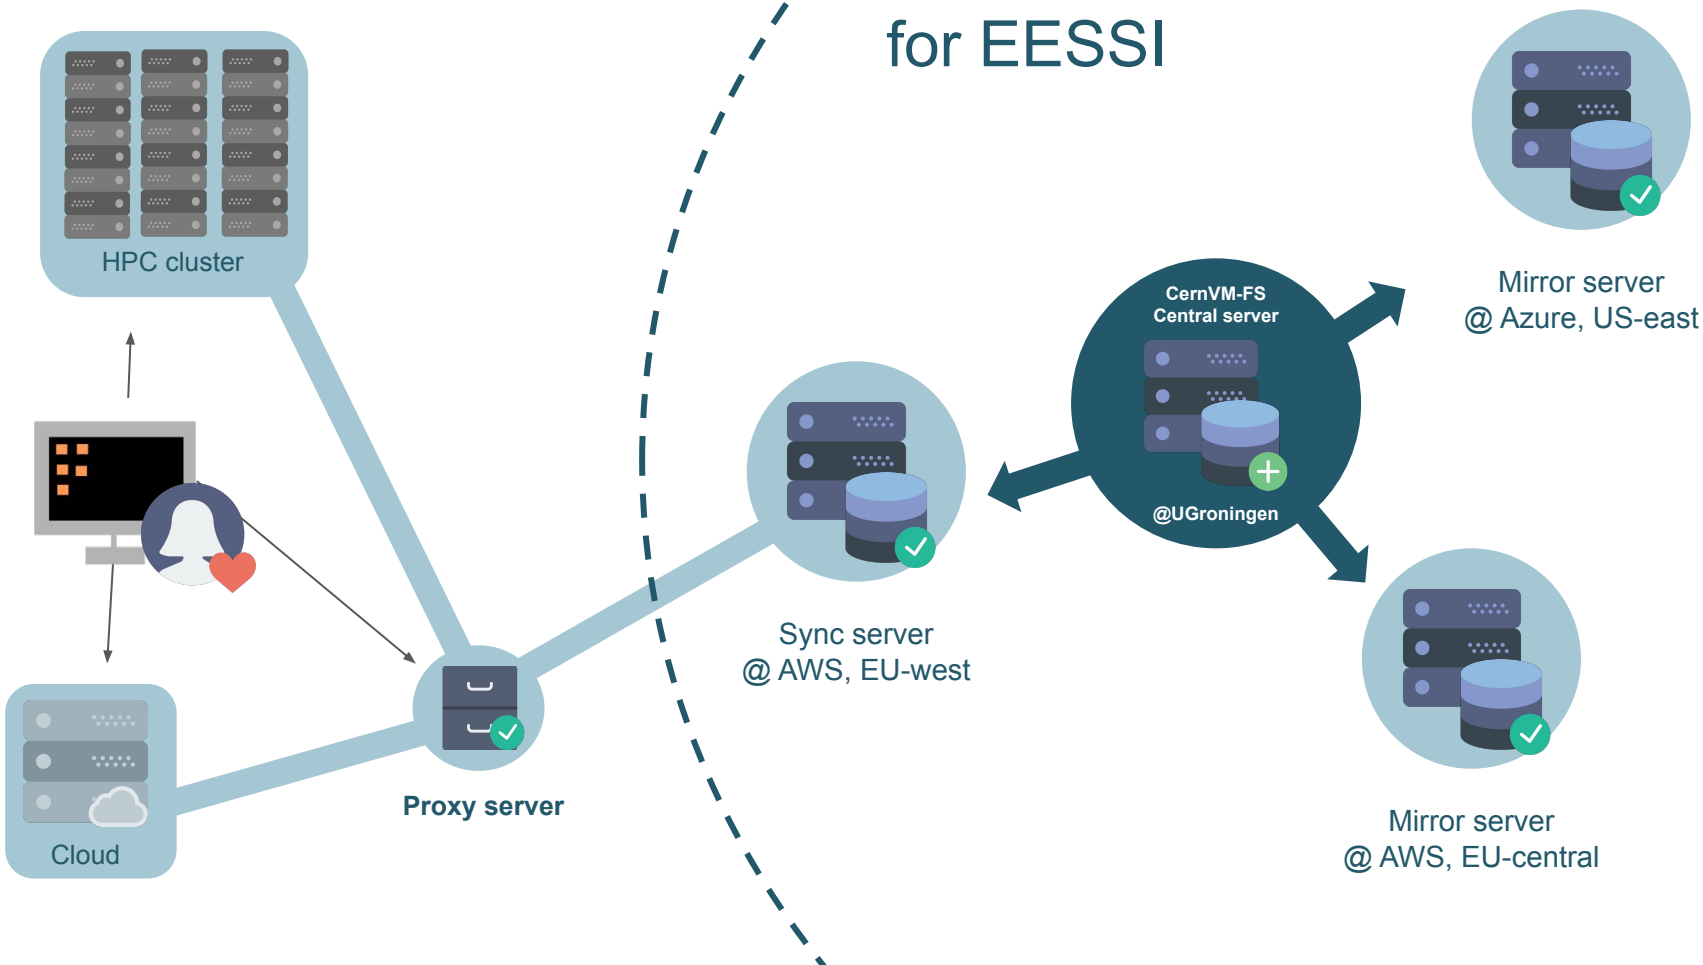

(icons via https://www.flaticon.com/authors/smashicons)

Best Practice for CernVM-FS in HPC

- Online tutorial (~3h15min)
- Tutorial website: https://multixscale.github.io/cvmfs-tutorial-hpc-best-practices
- YouTube video: https://youtu.be/L0Mmy7NBXDU

EESSI ingredients

Global distribution
of
software installations

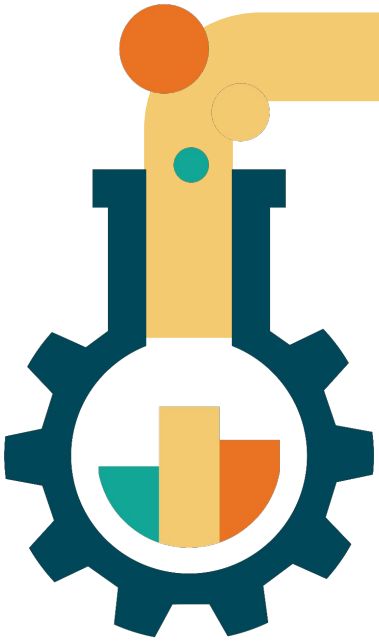

EESSI

EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

Abstraction away from
The host OS

Similar to container images
"Containers without the Containing"

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
```

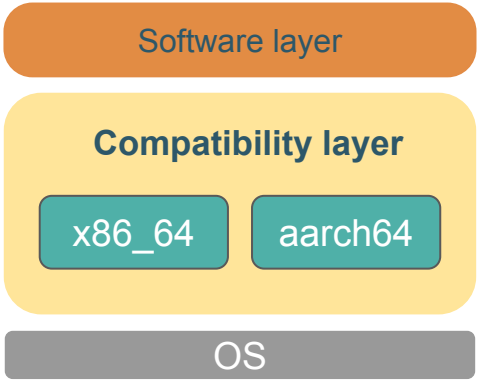

EESSI ingredients

Abstraction away from the host OS

Global distribution of software installations

EESSI

EUROPEAN ENVIRONMENT FOR SCIENTIFIC SOFTWARE INSTALLATIONS

Optimized software Installations for specific CPU microarchitectures

Intuitive user interface: module avail, module load, …

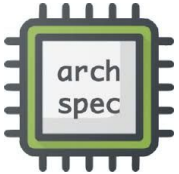

Automatic selection of Best suited part of Software stack for CPU microarchitectires

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 archdetect or archspec
- **Best subset of software installations for host CPU Is automatically selected**

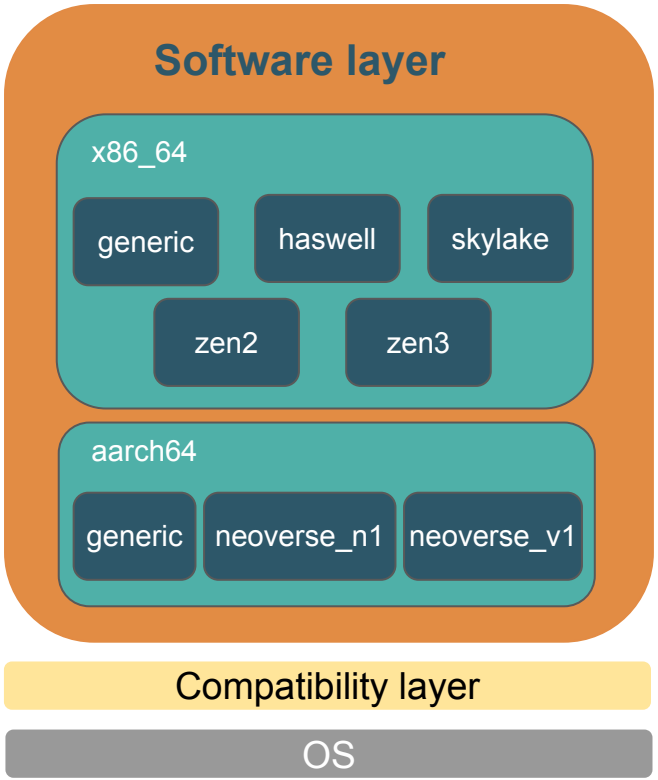

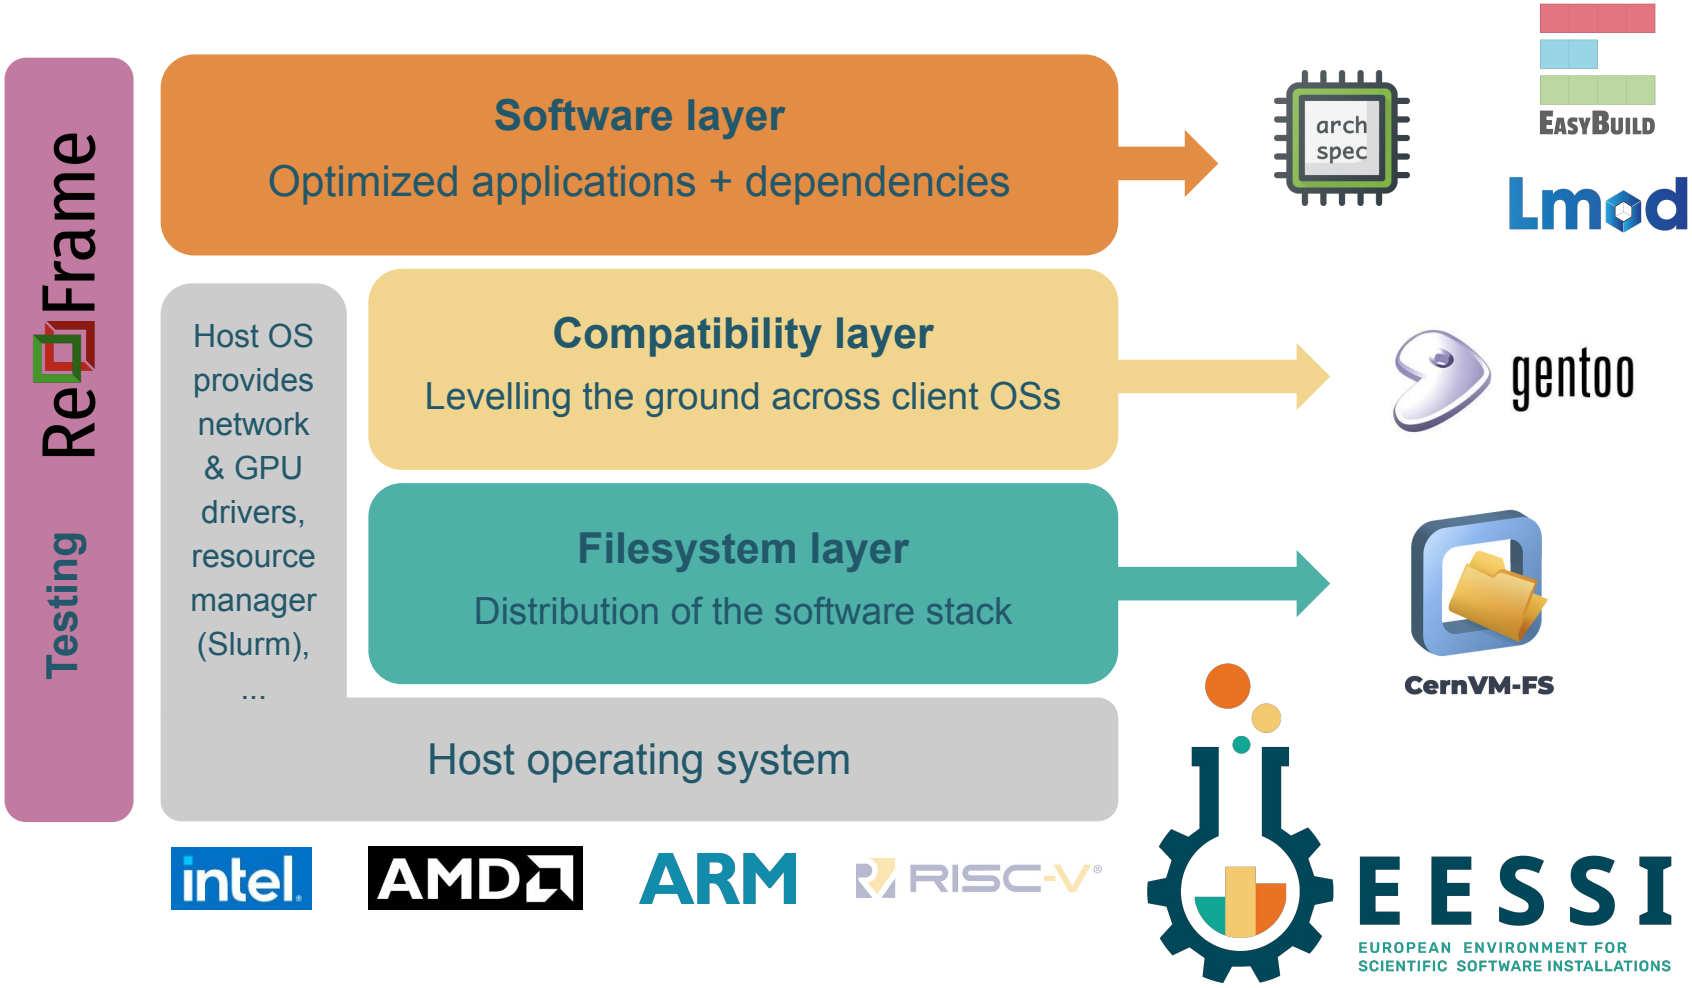
ReFrame
Testing
Software layer
Optimized applications + dependencies
arch spec
EasyBuild
Lmod
Host OS provides network & GPU drivers, resource manager (Slurm), ...
Compatibility layer
Levelling the ground across client OSs
gentoo
Filesystem layer
Distribution of the software stack
CernVM-FS
Host operating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

EESSI paper (open access, Feb'22)

doi.org/10.1002/spe.3075

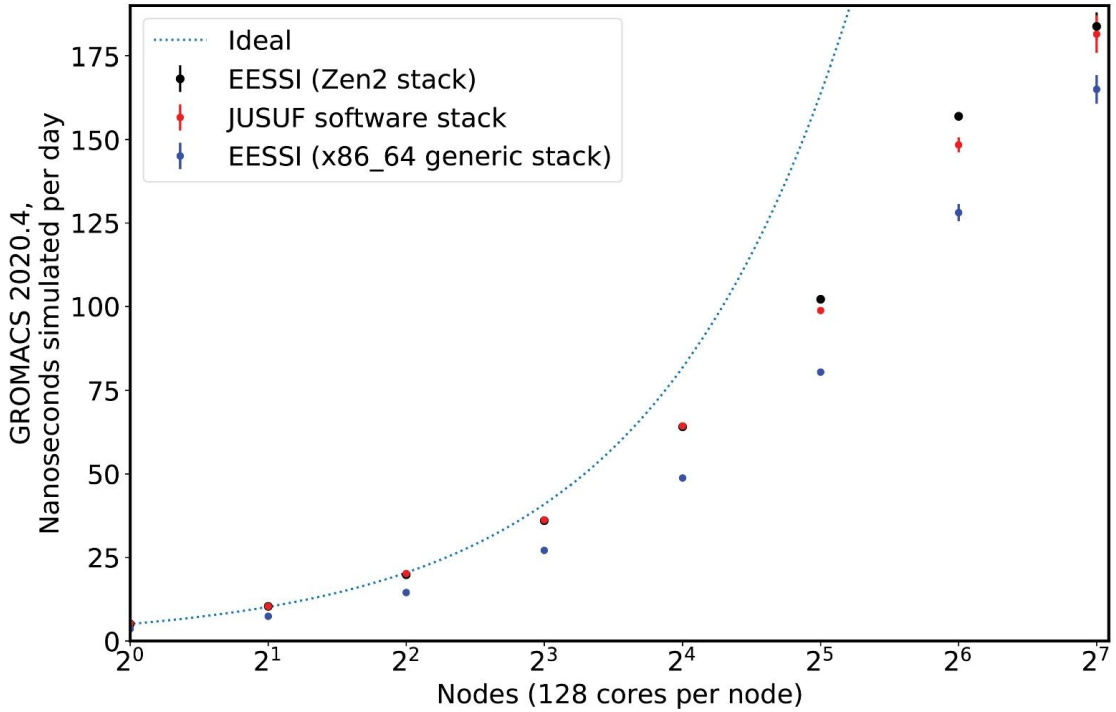

Paper includes **proof-of-concept performance evaluation** compared to system software stack, performed at JUSUF @ JSC using GROMACS 2020.4, up to 16,384 cores (CPU-only)

www.multixscale.eu

- EuroHPC Centre of Excellence: 4 year project (2023-2026), €6M budget (50% for EESSI)
- Collaboration between EESSI and CECAM: total of 16 partners (academic + industry)
- EESSI focuses on technical aspects: providing a shared stack of scientific software
- 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 

Lara

Support portal

GitLab

gitlab.com/eessi/support

- Via GitLab gitlab.com/eessi/support or email

  `support (@) eessi.io`
- Report issues
- Request software
- Suggest features
- Confidential tickets possible

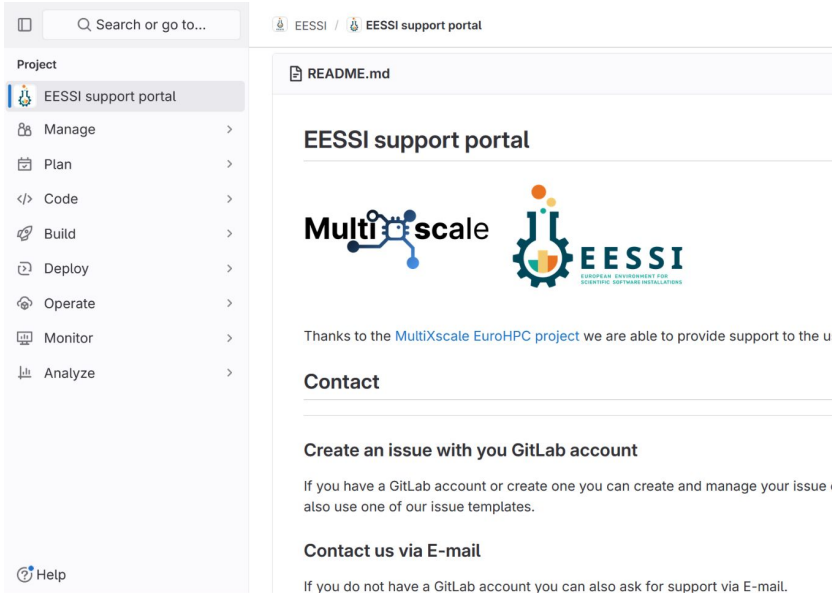

The EESSI User Experience

```
$ source
/cvmfs/software.eessi.io/versions/2023.06/init/bash
{EESSI 2023.06} $ module load TensorFlow
{EESSI 2023.06} $ python -c ‘import tensorflow’
```

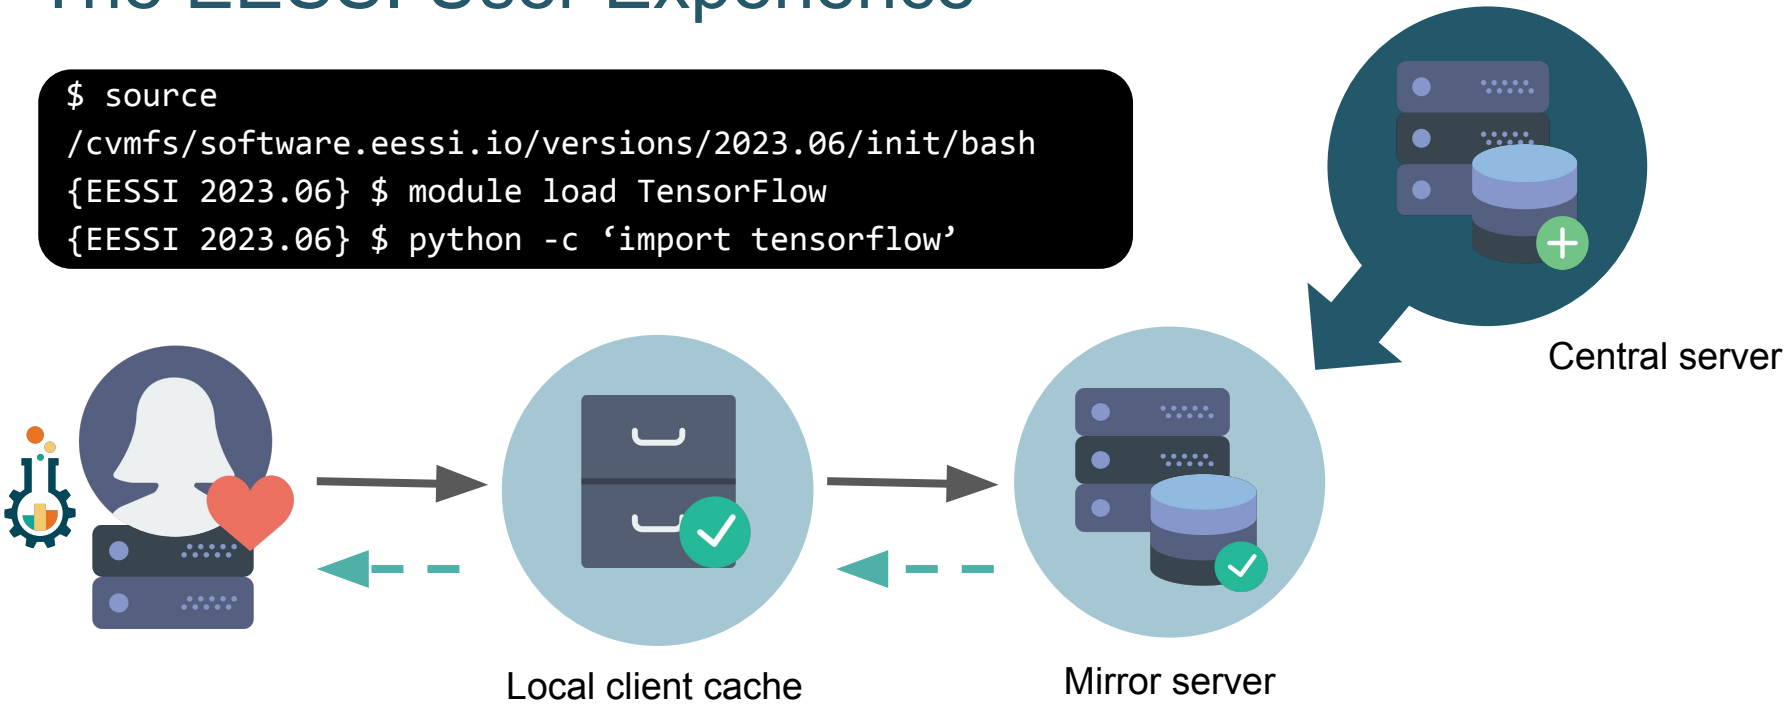

EESSI provides **on-demand streaming**

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

Accessing EESSI via 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 squid 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 container image, cvmfsexec, …
eessi.io/docs/getting_access/native_installation - eessi.io/docs/getting_access/eessi_container

Software layer

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/intel/skylake_avx512
Using x86_64/intel/skylake_avx512 as software subdirectory
...
Environment set up to use EESSI pilot software stack, have fun!

$ module load R

$ which R
/cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/
intel/skylake_avx512/software/R/4.3.2-gfbf-2023a/bin/R

$ R --version
R version 4.3.2
```

Adding software to EESSI

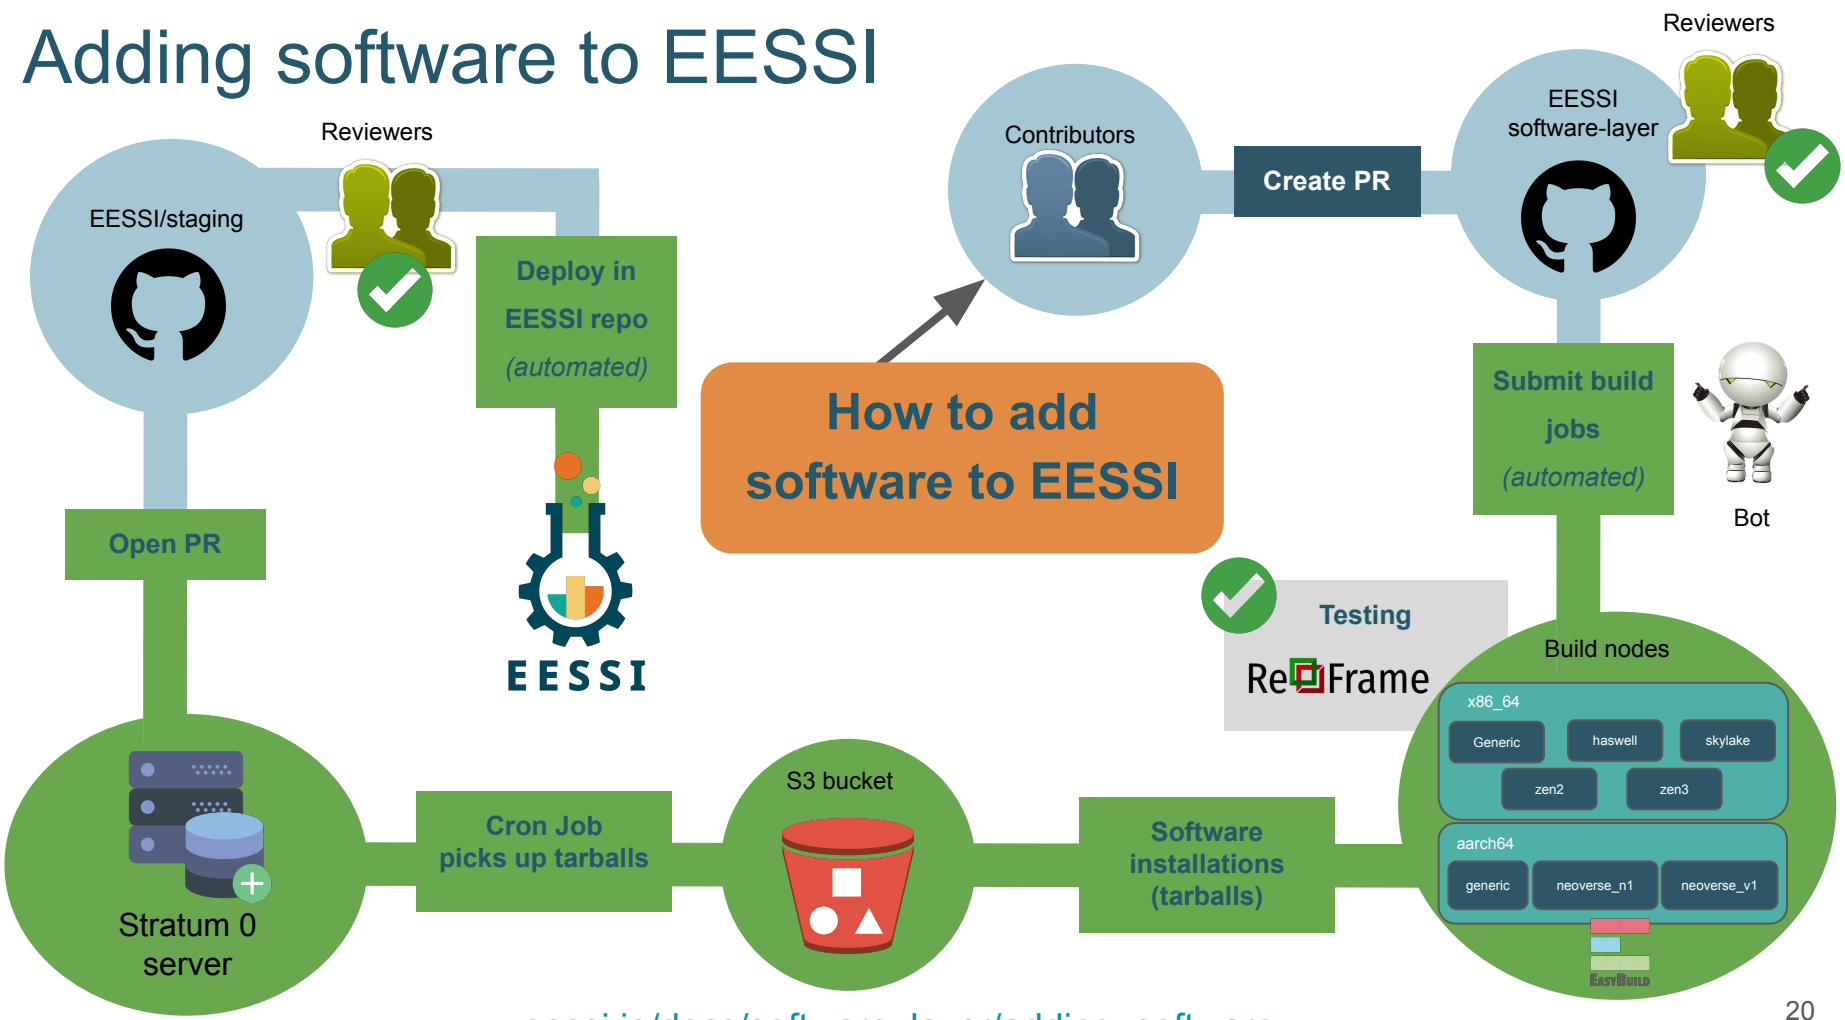

eessi.io/docs/software_layer/adding_software

Software testing in EESSI

- **Software test suite** is run on build host when building software with EasyBuild
- **Sanity check commands** are run to check that installed software is not horribly broken
- **Portable test suite** based on ReFrame to evaluate **functional correctness** + **performance**
  - On host (and OS) that was used to build the software (as part of build procedure)
  - **Also on different host and/or different OS** => software installations should still work!
- Eventually we also want to do **performance monitoring**
  - Periodic runs of EESSI test suite to catch performance regressions

github.com/EESSI/test-suite eessi.io/docs/test-suite

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

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

YouTube channel: youtube.com/@eessi_community

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

EESSI support portal: gitlab.com/eessi/support

Monthly online meetings (first Thursday, 2pm CEST)

Join our mailing list & Slack channel

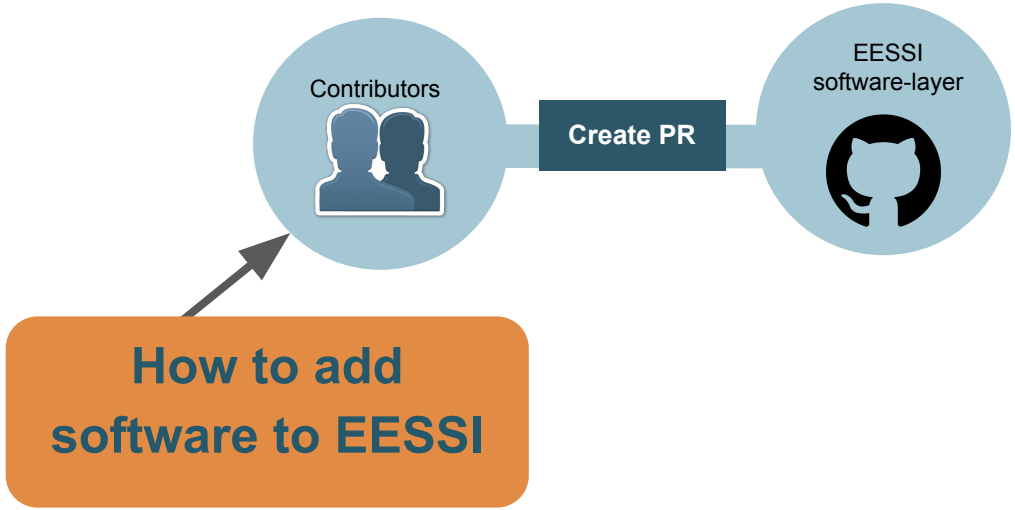
Contributors
Create PR
EESSI
software-layer
How to add
software to EESSI

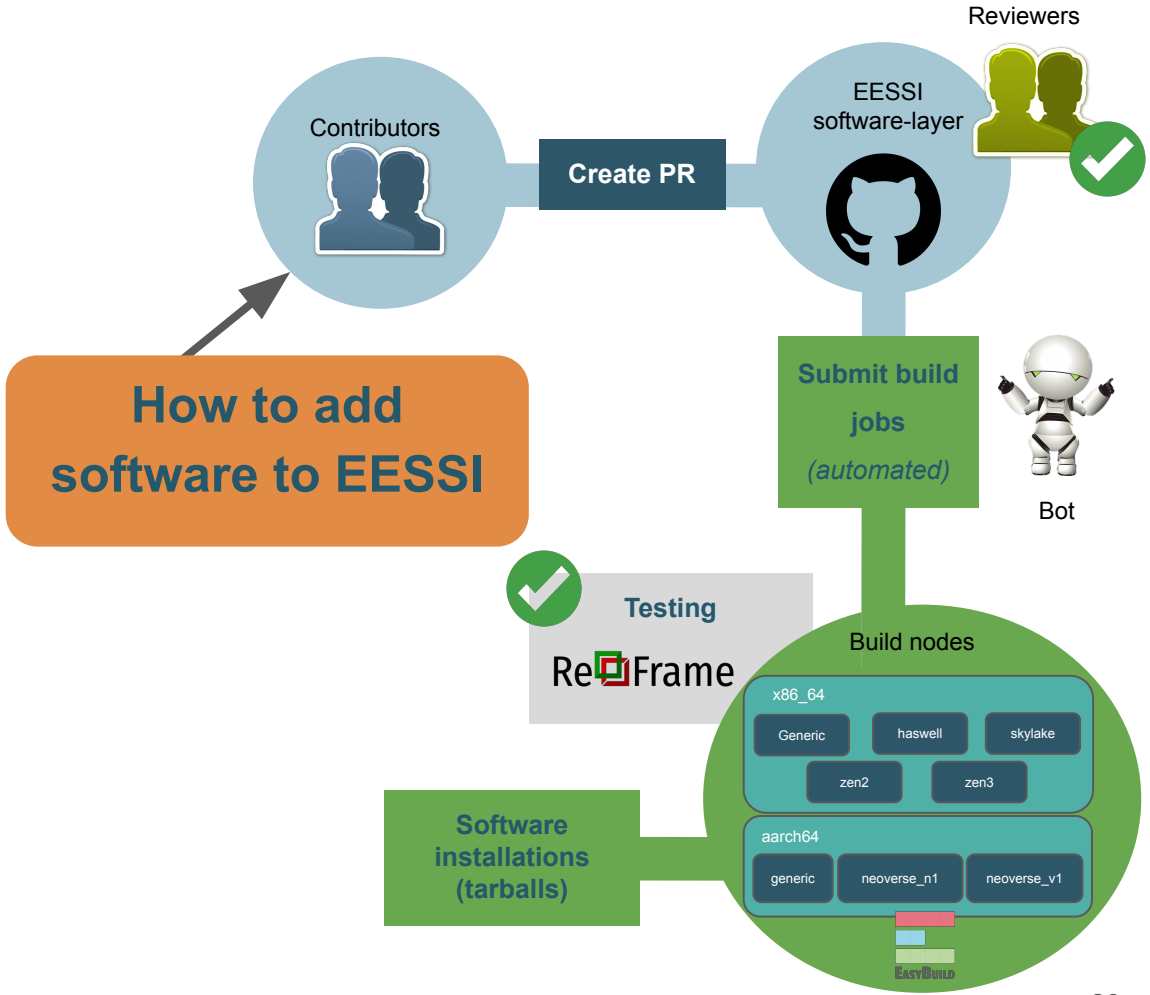
Reviewers
Contributors
Create PR
EESSI
software-layer
How to add
software to EESSI
Submit build
jobs
(automated)
Bot
Testing
ReFrame
Build nodes
x86_64
Generic
haswell
skylake
zen2
zen3
aarch64
generic
neoverse_n1
neoverse_v1
Software
installations
(tarballs)

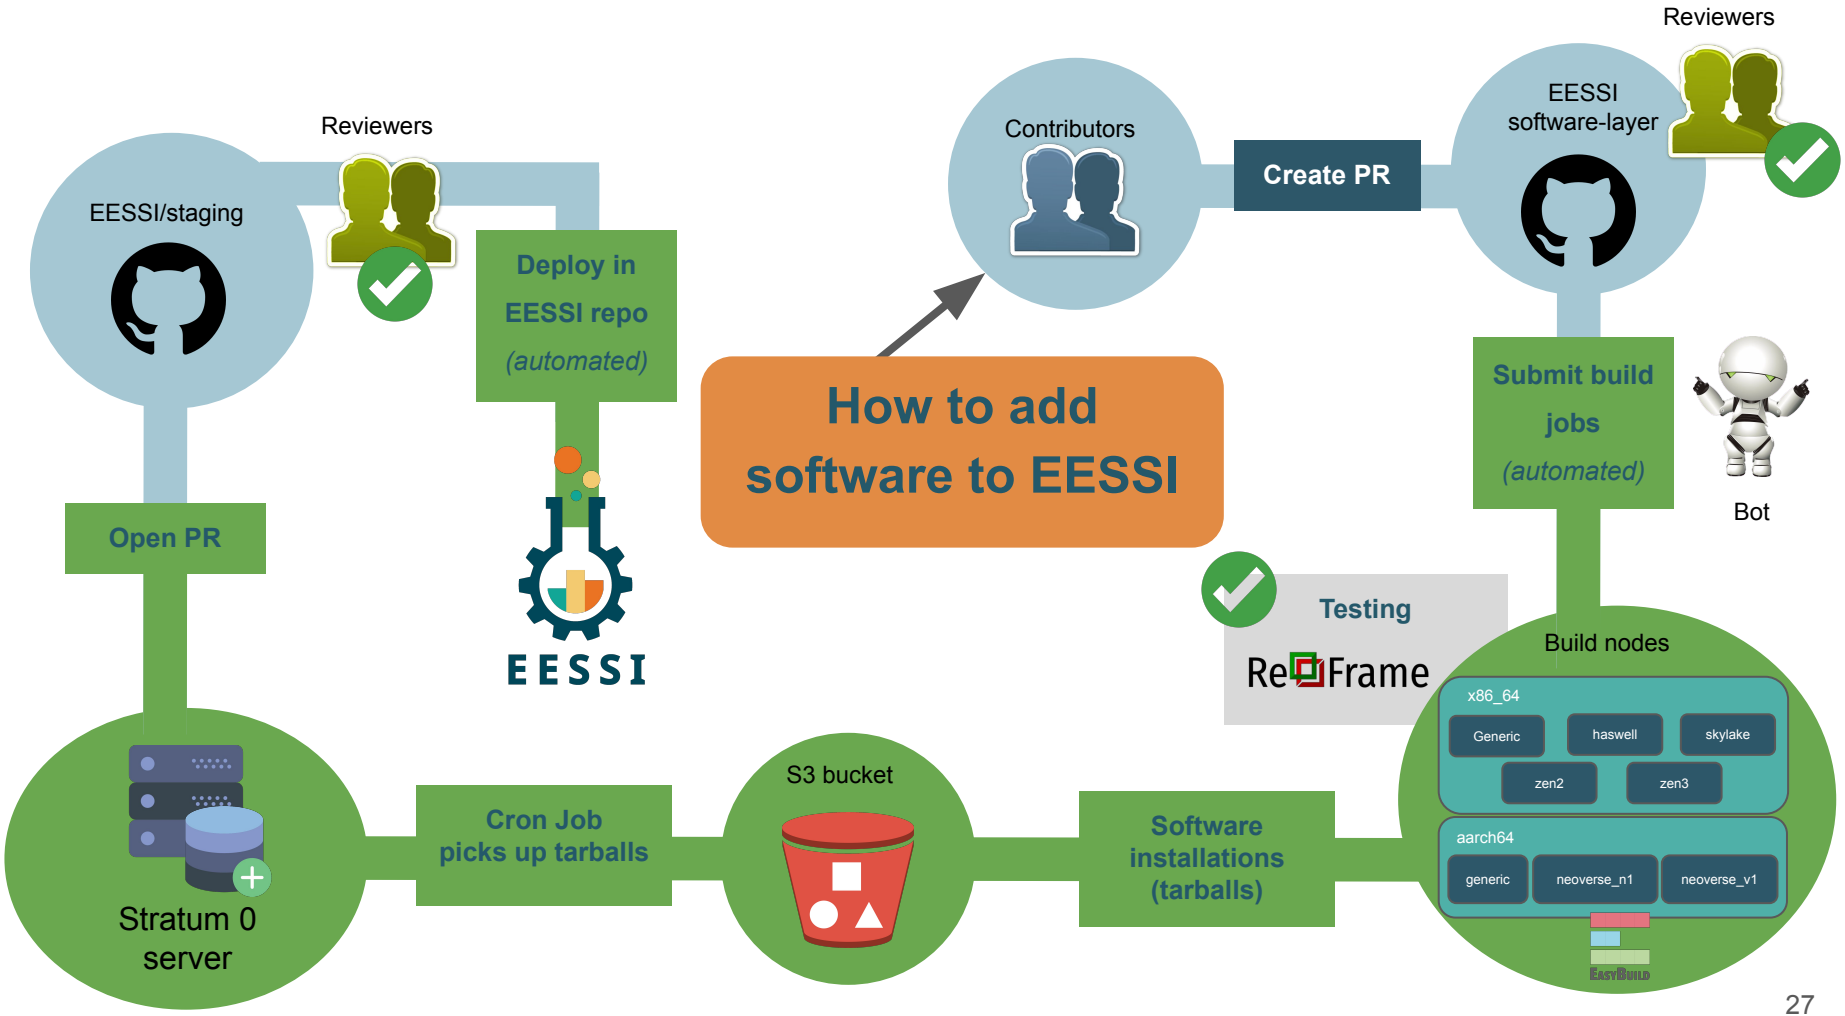
Reviewers
EESSI/staging
Reviewers
Deploy in EESSI repo (automated)
Contributors
Create PR
EESSI software-layer
Reviewers
How to add software to EESSI
Submit build jobs (automated)
Bot
Open PR
EESSI
Testing
ReFrame
Build nodes
x86_64
Generic
haswell
skylake
zen2
zen3
aarch64
generic
neoverse_n1
neoverse_v1
EasyBuild
Stratum 0 server
Cron Job picks up tarballs
S3 bucket
Software installations (tarballs)