


# CernVM-FS makes software EESSI to use

CernVM Workshop 16-18 september 2024  
Thomas Röblitz, University of Bergen

# EESSI in a nutshell

- **On-demand streaming of optimized** scientific software installations
- It's designed to **work on any Linux distribution**
- **Uniform software stack** across various systems: laptop, HPC, cloud, ...
- Community-oriented: **let's tackle the challenges we see together!**



# Using EESSI

[eessi.io/docs/using\\_eessi/eessi\\_demos](https://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
```

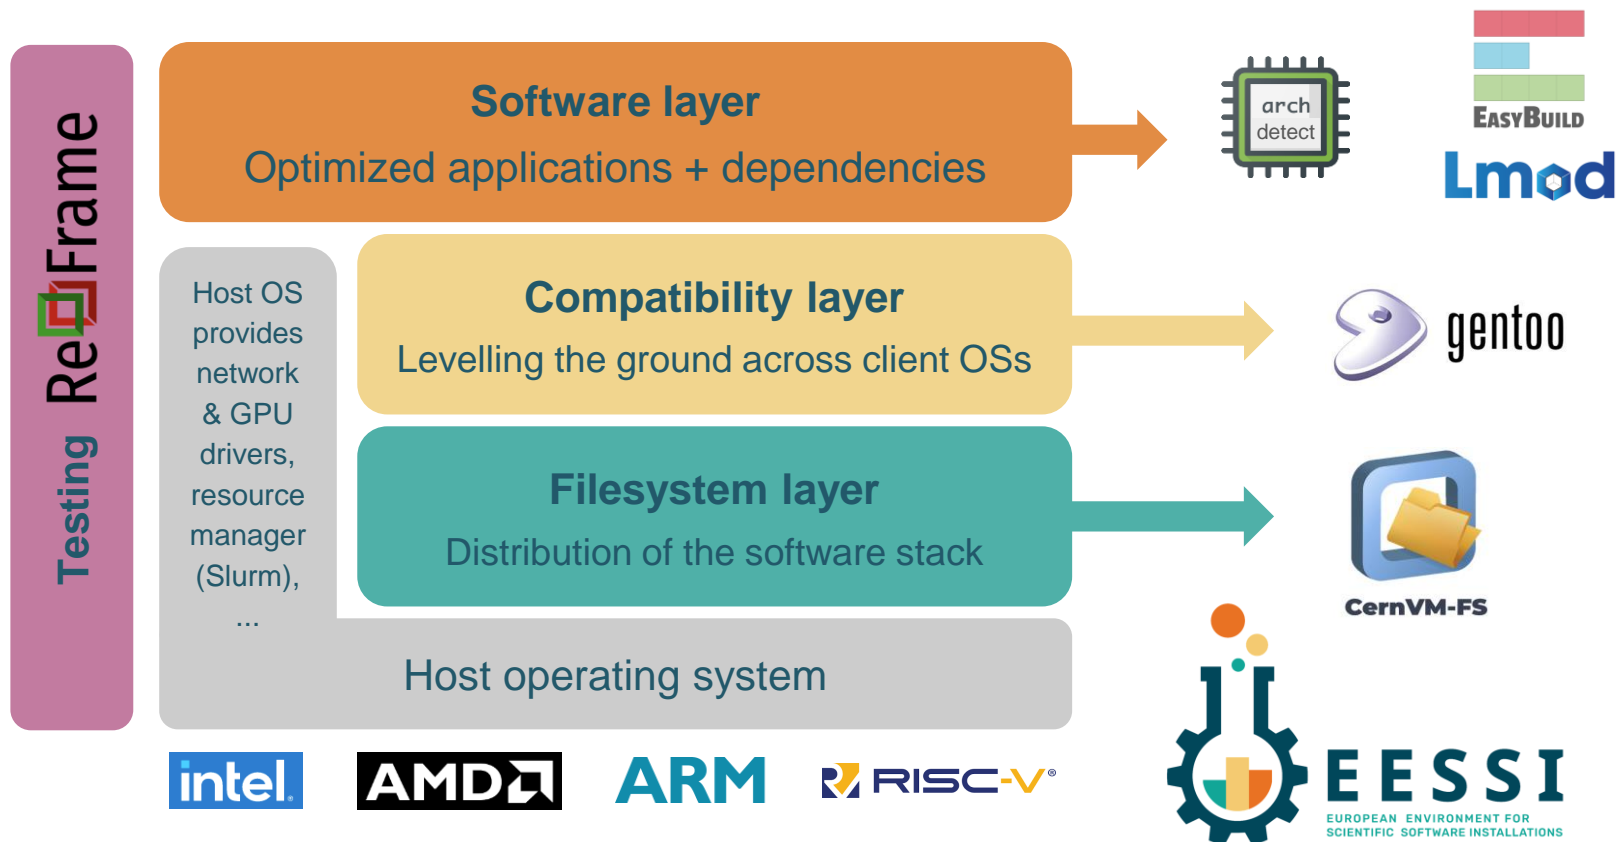
```
...
Environment set up to use EESSI pilot software stack, have fun!
```

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

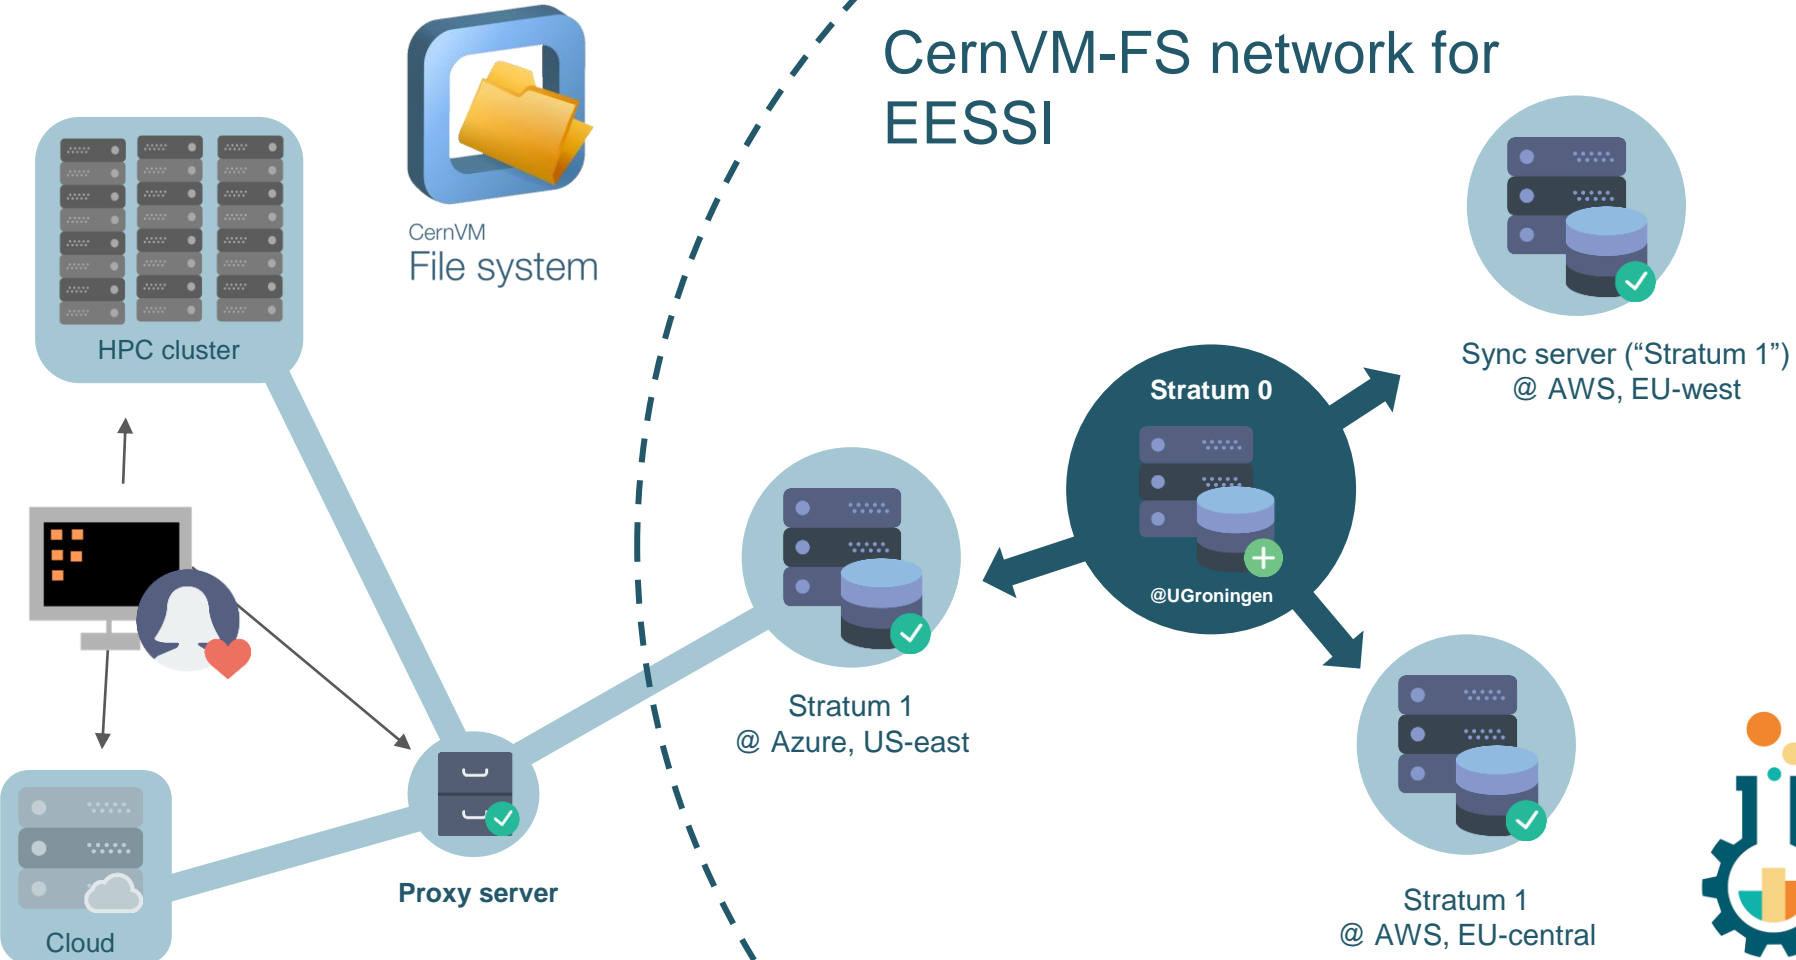
```
{EESSI 2023.06} $ which R
/cvmfs/software.eessi.io/versions/2023.06/software/linux/x86_64/
amd/zen3/software/R/4.3.2-gfbbf-2023a/bin/R
```

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

# EESSI design


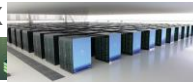


# CernVM-FS network for EESSI



# What happened since the CernVM Workshop '22?

prototype bot | EuroHPC CoE launched | Multi-scale | new compat layer version | gentoo | Best Practices for CernVM-FS in HPC (4 Dec 2023) | start building for zen4 | EESSI-extend module | new Stratum-0 | initial GPU support | regularly run test suite on Vega, Karolina, AWS, Azure | start building for a64fx


2023

improved bot

production repo /cvmfs/software.eessi.io

2024

launch riscv.eessi.io  


EESSI test suite software.eessi.io included in default CernVM-FS config

new AWS build cluster

develop dev.eessi.io


eessi.io/docs

contribution policy

support portal and rotation



Azure k cluster



Automate the building of software

# Adding software to EESSI





# Adding software to EESSI

- adding software is trivial
- [eessi.io/docs/adding\\_software/contribution\\_policy/](https://eessi.io/docs/adding_software/contribution_policy/)

{2023.06}[foss/2023b] Brunсли 0.1, R 4.4.1, R-bundle-CRAN 2024.06 #644

Merged ocaisa merged 14 commits into EESSI:2023.06-software.eessi.io from bedroge:r\_441\_and\_cran\_bundle 34 minutes ago

Conversation 127 Commits 14 Checks 34 Files changed 1

Changes from all commits File filter Conversations Jump to

0 / 1 files v

8 easystack/software.eessi.io/2023.06/eessi-2023.06-eb-4.9.2-2023b.yml

```
@@ -20,3 +20,11 @@ easyconfigs:
20 20 # see https://github.com/easybuilders/easybuild-framework/issues/4400
21 21 + from-commit: 765ba90daf5953e306c
22 22 - HPL-2.3-foss-2023b.eb
23 + - Brunсли-0.1-GCCcore-13.2.0.eb:
24 +   options:
25 +     # https://github.com/easybuilders
26 +     from-commit: 1736a123b16858364525
27 + - R-bundle-CRAN-2024.06-foss-2023b.eb:
28 +   options:
29 +     # see https://github.com/easybuil
30 +     from-commit: 41a2cd83f9fb017b76f0
```

bedroge commented 19 hours ago

bot: build arch:aarch64/generic  
bot: build arch:aarch64/neoverse\_n1  
bot: build arch:aarch64/neoverse\_v1  
bot: build arch:x86\_64/generic  
bot: build arch:x86\_64/amd/zen2  
bot: build arch:x86\_64/amd/zen3  
bot: build arch:x86\_64/amd/zen4  
bot: build arch:x86\_64/intel/haswell  
bot: build arch:x86\_64/intel/skylake\_avx512

eessi-bot bot commented 5 days ago · edited

New job on instance eessi-bot-mc-aws for architecture x86\_64-intel-skylake\_avx512 for repository 2023.06-software in job dir /project/def-users/SHARED/jobs/2024.09/pr\_644/18347

date	job status	comment
Sep 11 14:15:54 UTC 2024	submitted	job id 18347 awaits release by job manager
Sep 11 14:17:02 UTC 2024	released	job awaits launch by Slurm scheduler
Sep 11 14:29:53 UTC 2024	running	job 18347 is running
Sep 11 22:13:32 UTC 2024	finished	▶ 😊 SUCCESS (click triangle for details)
Sep 11 22:13:32 UTC 2024	test result	▶ 😊 SUCCESS (click triangle for details)
Sep 12 05:41:52 UTC 2024	uploaded	transfer of eessi-2023.06-software-linux-x86_64-intel-sky-1726091602.tar.gz to S3 bucket succeeded

Adding support for more CPUs and GPUs

# AMD zen4, Fujitsu a64fx, Intel Sapphire Rapids

- Build infrastructure can easily support more architectures
  - multiple bot instances can run across different clusters with access to specific architectures
- Example: zen4
  - Not all toolchains may be supported, requires GCC 13 or newer
- Need to deal with missing software packages (due to missing toolchain)
  - inform users when they run `module load ...`
  - CI that checks for missing packages across CPU architectures
- Sometimes we hit bugs specific to new architecture: Highway/1.0.4

# Add (better) support for GPUs

- Goal: avoid builds for *all* GPU models / compute capabilities for *each* CPU
  - not all combinations are commonly found, e.g., Haswell + H100
- Approach:
  - only build for a few GPU models for a given CPU architecture
- Implementation:
  - bot: build command requires additional argument (e.g., `accelerator:nvidia/cc80`)
  - installation of GPU builds in subdirectories of CPU builds

`/cvmfs/software.eessi.io/.../zen2/software/PyTorch/2.1.2-foss-2023a-CUDA-12.1.1`

GPU `.../zen2/accel/nvidia/cc80/software/PyTorch/2.1.2-foss-2023a-CUDA-12.1.1`

A diagram consisting of a horizontal curly bracket that spans from the start of the GPU path to the end of the CPU path. An arrow points from the right end of this bracket to the GPU path, indicating that the GPU path is a subdirectory of the CPU path.

CPU-only dependencies `.../zen2/software/Python/3.11.3-GCCcore-12.3.0 (as before)`

# Progress on RISC-V CPUs

# RISC-V ?

- Instruction Set Architecture
- Started in 2010 at UC Berkeley
- BSD license
- Large interest in EU in RISC-V-based processors and accelerators



# RISC-V support in EESSI



- **CernVM-FS client:** built from sources
- **Gentoo Prefix:** GSoC project in 2022 laying the groundwork
- **FOSS toolchain:** 2023b
- **Build environment:** container runtime + overlaysfs
- Separate CernVM-FS repository: **`/cvmfs/riscv.eessi.io`**
- Built on StarFive Vision 2 or SiFive HiFive Unmatched systems
- ~ 150 packages incl R, GROMACS, ESPResSo, ...

# EESSI test suite



# EESSI test suite - Why?

- **Ensure quality of the software installations provided by EESSI**
- Detect system inefficiencies with respect to change in system software at the OS level or system (re)configuration.
- Tests for **CP2K, ESPResSo, GROMACS, LAMMPS, PyTorch, QuantumESPRESSO, TensorFlow, OSU Micro Benchmarks**

# EESSI test suite - Current status

- **Integrated in the automated procedure** to build and deploy software:
- **Periodic tests:** AWS, Azure, EuroHPC, local HPC



# Monitoring

# EESSI status page: <https://status.eessi.io>



- Shows an overall status of the CernVM-FS infrastructure
- Scrapes CVMFS information from all servers
- Exports all statuses as Prometheus metrics

EESSI status!

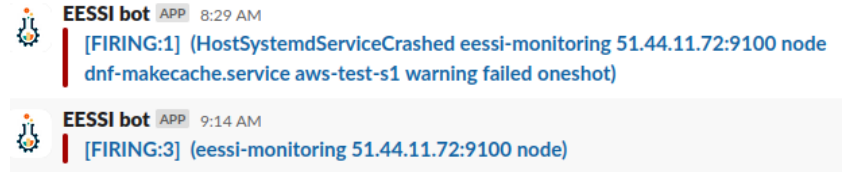
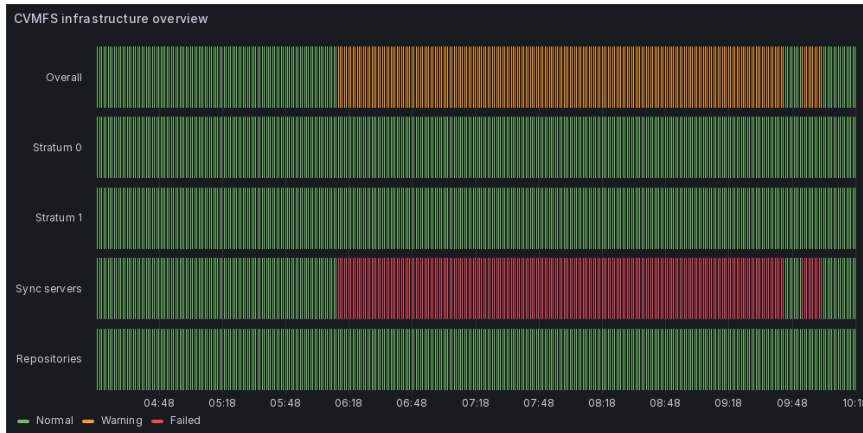


The screenshot shows the EESSI status page. At the top, there is a large green checkmark and the text 'Normal service' with the subtitle 'EESSI services operating without issues.' Below this, there is a navigation bar with five items: 'Normal service' (green checkmark), 'Degraded' (yellow square), 'Warning' (orange triangle), 'Failed' (red circle with cross), and 'Maintenance' (blue arrow). The main content area displays four service status cards: 'Stratum0' (green checkmark), 'Syncservers' (green checkmark), 'Stratum1s' (green checkmark), and 'Repositories' (green checkmark). At the bottom right, there is a footer with the text 'Last updated 2024-09-05T07:10:05Z | support@eessi.io'.

# Monitoring EESSI (WIP)



- Working on a setup of Grafana + Prometheus + Prometheus Alertmanager + several exporters
- Using Ansible for configuring all servers



# Collaboration with software developers

# Problems we have found through software testing (1/2)



## Hang/crash in Open MPI's `smcuda` Byte Transport Layer (BTL) component

- See <https://gitlab.com/eessi/support/-/issues/41>
- Upstream issue: <https://github.com/open-mpi/ompi/issues/12270>
- Causes hanging or failing tests for FFTW, OpenFOAM, ESPResSo, ...
- **Only happens on Arm Neoverse V1 (AWS Graviton 3)**
- Fixed by Luke Robison (AWS), see <https://github.com/open-mpi/ompi/pull/12338>
- Fix will be included in upcoming Open MPI release (v4.1.7+)
- Patch already applied in Open MPI installations included in EESSI 2023.06

# Problems we have found through software testing (2/2)



## Failing tests in GROMACS test suite

- See <https://gitlab.com/eessi/support/-/issues/47>
- Filesystem race when running tests concurrently ([GROMACS PR #4066](#))
- **Bug in 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!**
  - WIP fix in [https://gitlab.com/gromacs/gromacs/-/merge\\_requests/4299](https://gitlab.com/gromacs/gromacs/-/merge_requests/4299)
  - Will be fixed in upcoming GROMACS release (2024.2?)



# /cvmfs/dev.eessi.io ... work in progress



686d26aa318d81a2bc78bb03ba408da0d09eaaabc34537262998a478f2c9bdf2

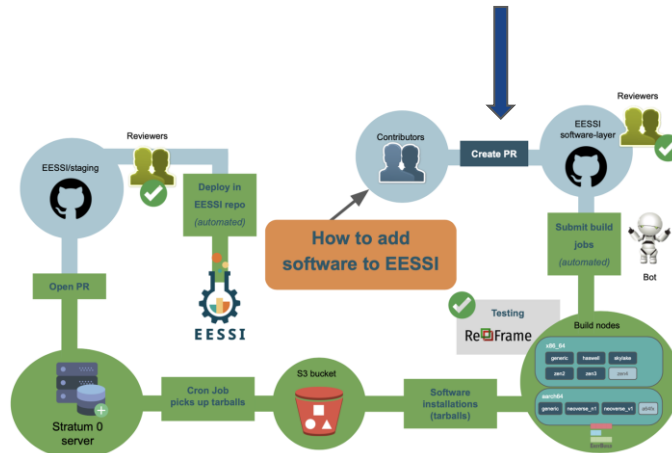
developer repository



project\_ESPResSo/build\_commit:

686d26aa318d81a2bc78bb03ba408da0d09eaaabc34537262998a478f2c9bdf2

EESSI/dev.eessi.io



# 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](https://github.com/EESSI/github-action-eessi)



gitlab-eessi 

 CI/CD Catalog project

# Leveraging EESSI GitHub Action



A screenshot of a GitHub Actions workflow run. The left sidebar shows a navigation menu with 'Summary', 'Jobs', 'Run details', 'Usage', and 'Workflow file'. The 'Jobs' section is expanded to show a single job named 'ubuntu' with a green checkmark. The main area shows the job's status as 'succeeded yesterday in 31m 58s'. Below this, a list of steps is shown: 'Setup EESSI', 'Checkout repository', and 'Install dependencies', all with green checkmarks. The 'Install dependencies' step is expanded to show two log entries: '1 ▶ Run module load ESPResSo/4.2.1-foss-2023a' and '14 Environment set up to use EESSI (2023.06), have fun!'. A search bar for logs is visible in the top right of the main area.

<https://github.com/pyMBE-dev/pyMBE/blob/main/.github/workflows/testsuite.yml>

<https://github.com/pyMBE-dev/pyMBE/actions/runs/8815523092/job/24197651600>

# Training and support activities

# Tutorial “Best Practices for CernVM-FS in HPC”



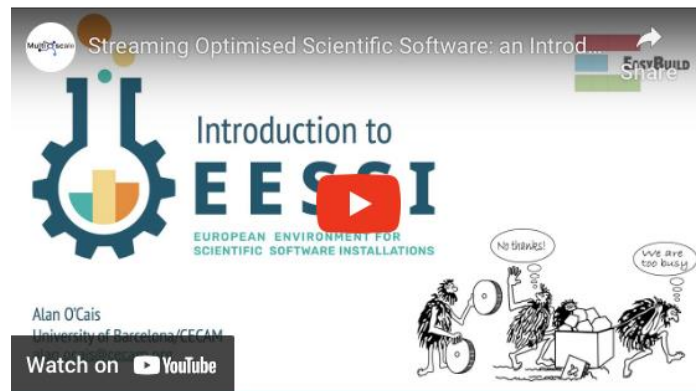
- <https://multixscale.github.io/cvmfs-tutorial-hpc-best-practices>
- Held online on 4 Dec 2023 (~3 hours), **recorded & available on YouTube**
- Over 200 registrations, ~125 attending the meeting
- Lecture + hands-on demos
- Topics:
  - Introduction to CernVM-FS + EESSI
  - Configuring CernVM-FS: client,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



# Tutorials to introduce EESSI to users



- run 3-4 times a year, online or hybrid
- Example: [https://www.eessi.io/docs/talks/2023/20231205\\_castiel2\\_eessi\\_intro/](https://www.eessi.io/docs/talks/2023/20231205_castiel2_eessi_intro/)
- Recordings available on YouTube
- Lecture + hands-on (on temporary cluster)
- Topics:
  - Motivation for EESSI
  - Design of EESSI
  - Starting from an "empty" VM: install CernVM-FS, configure access, do science
  - Adding software, portable workflows, using EESSI in CI, ...
  - Getting support



# Getting support for EESSI

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



[gitlab.com/eessi/support](https://gitlab.com/eessi/support)

Project


SEARCH Search or go to...

EESSI / EESSI support portal

Project

README.md

## EESSI support portal

**MultiXscale**  **EESSI**  
EUROPEAN ENVIRONMENT FOR  
SCIENTIFIC SOFTWARE INSTALLATIONS

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



# Enabling and tuning EESSI on (Euro)HPC sites



# Enabling EESSI on (Euro)HPC systems



- know about 20 systems that make EESSI available
  - HPC clusters, OpenStack environments, k8s platforms, AWS, Azure
  - also some make it available on all Linux clients in an organisation
- standard setups with CernVM-FS clients, Squid proxies, private Stratum-1
- issues:
  - no internet connectivity, changing policies/configuration
  - diskless compute nodes
  - reluctance to install and maintain another service (for various reasons)

# "Tuning" EESSI on (Euro)HPC systems



- EESSI is optimised for CPU architectures
- vendor-specific optimisations (network, math libraries, ...) may be required
- Examples:
  - National HPC in Norway uses parts of HPC-X with OpenMPI
  - EC2 instances with AWS EFA interfaces
- `/cvmfs/software.eessi.io/host_injections` (variant symlink) is searched by EESSI dynamic linker

# Acknowledgements






**EuroHPC**  
Joint Undertaking

- 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.
- Funded by Sigma2 (National e-infrastructure provider in Norway)
- Thanks to Amazon Web Services (AWS) and Microsoft Azure for generously sponsoring the EESSI project with cloud credits, feedback, and guidance.
- Several additional contributors & great collaboration with CernVM devs, The Alliance, Gentoo Prefix maintainers, software developers, ...

**Multi-scale**



- 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 



Website: [eessi.io](https://eessi.io)

GitHub: [github.com/eessi](https://github.com/eessi)

Documentation: [eessi.io/docs](https://eessi.io/docs)

YouTube channel: [youtube.com/@eessi\\_community](https://youtube.com/@eessi_community)

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

EESSI support portal: [gitlab.com/eessi/support](https://gitlab.com/eessi/support)

Bi-monthly online meetings (first Thursday, 2pm CEST)

**Join our mailing list & Slack channel (see website)**