



EESSI behind the scenes: software layer

EESSI Community Meeting @ Amsterdam

15 Sept 2022

Thomas Röblitz (Univ. of Bergen) + Kenneth Hoste (HPC-UGent)

What is the software layer?



- Contains all the (scientific) software installations
 - Built with Easybuild, incl. module files
 - Every CPU type has its own subdirectory with optimized binaries
- Lmod (from the compat layer) is used to make the apps available via the module files
- archspec is used to determine the most suitable subdirectory
 - If a direct match is not found, it falls back to the best compatible CPU type
 - Generically built binaries are available as a last resort
 - This functionality is part of the init scripts:
<https://github.com/EESSI/software-layer/tree/main/init>

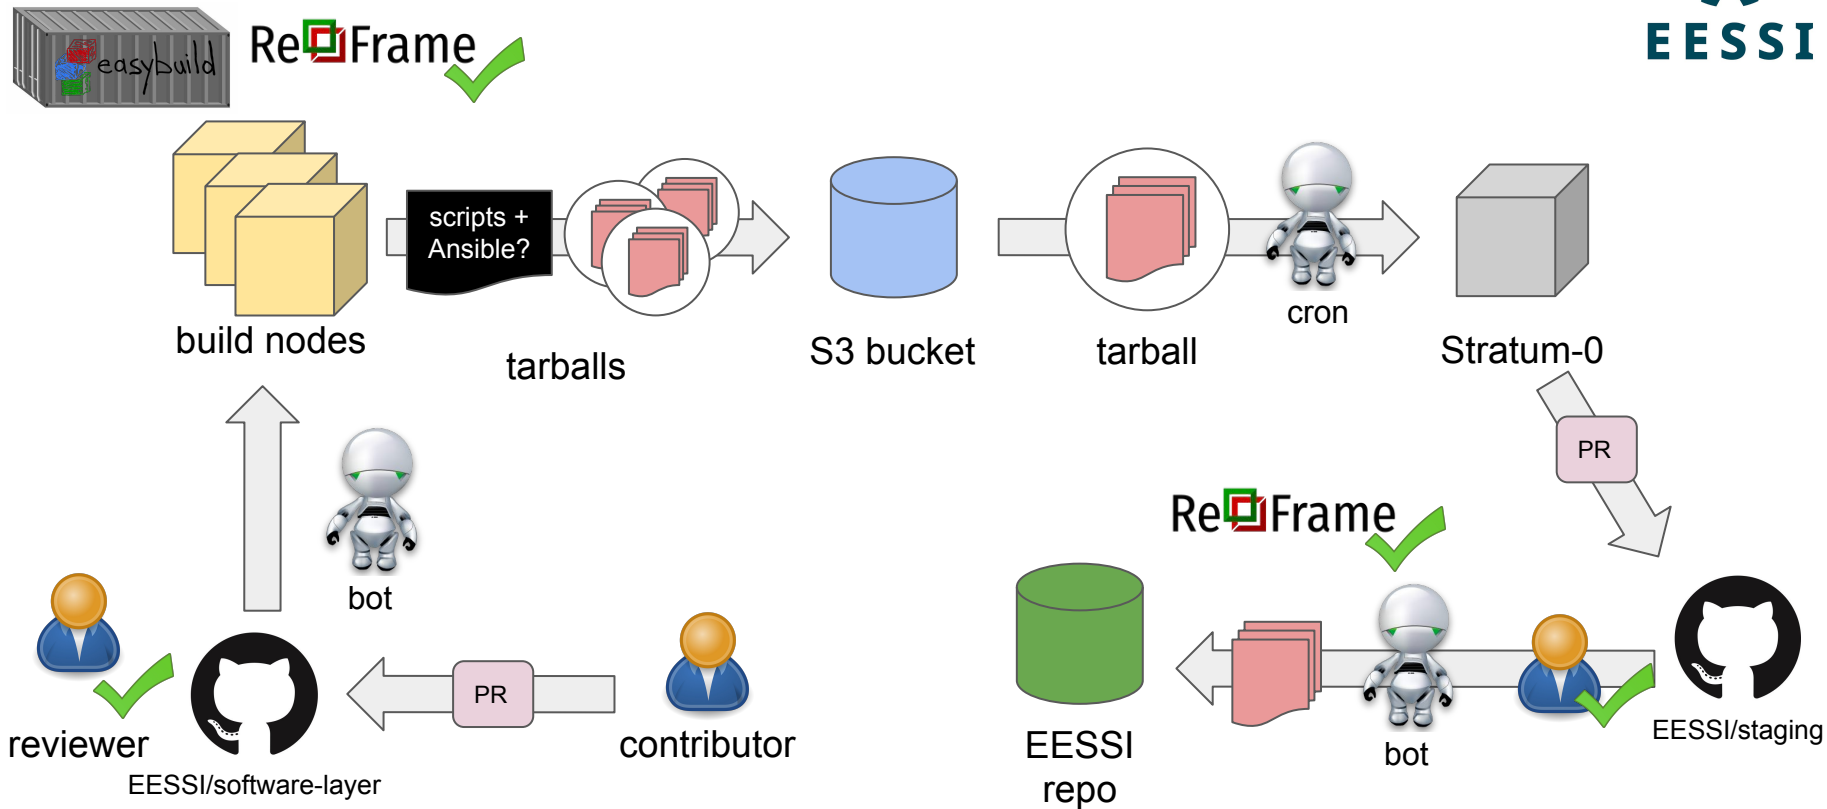
Current workflow for building the software



- Long and ugly bash script that calls EasyBuild for every app to be installed
 - <https://github.com/EESSI/software-layer/blob/main/EESSI-pilot-install-software.sh>
- For each target CPU, this script is run on a machine with that type of CPU
- Singularity build container is used to always have the same, controlled environment
 - Minimal number of packages, to prevent host libraries from accidentally being picked up
 - Also provides a writable overlay on top of `/cvmfs` using `fuse-overlays`
 - Available on GHCR: <https://github.com/EESSI/filesystem-layer/pkgs/container/build-node>
- Scripts are available to facilitate running stuff in build container with EESSI compat layer
 - See `build-container.sh` + `install_software_layer.sh` scripts in [software-layer](#) repo
- Tarballs are created for every build/stack and ingested to the repository

Target workflow for building the software

Goal: automated, with human oversight




Target workflow for building the software



- Replace the bash script by easystack files
 - Defines the entire stack, incl. criteria for pkgs being excluded/included for specific CPU types
 - Additional EasyBuild configuration options where needed (`--from-pr`, etc.)
 - WIP PRs are open to enhance support for easystack files in EasyBuild
- Use GitHub pull requests to propose additions to the stack
 - PRs propose changes to easystack files that define the EESSI software layer
- Have bot(s) running on build nodes/clusters that react to these PRs
 - Automatically start building + testing software, deploy after approval by a EESSI maintainer
 - Report back to the PR via comments
 - See Thomas' presentation

Discussion

- Criteria/policy for adding new software: easyconfig + tests?
- Which set of CPU types do we want to support for the next (beta?) version?
- Module naming scheme: flat vs hierarchical vs ...
- Move more packages to the compat layer (see Bart's presentation)?



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

Website: <https://www.eessi-hpc.org>

Join our mailing list & Slack channel

<https://www.eessi-hpc.org/join>

Documentation: <https://eessi.github.io/docs>

GitHub: <https://github.com/eessi>

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

youtube.com/channel/UCKLS5X7_oMWhUrAZuzSwBxQ

Monthly online meetings (first Thursday, 2pm CEST)