

Continuous Integration for the FairRoot Software Stack

Dennis Klein, Christian Tacke, Florian Uhlig, Mohammad Al-Turany
GSI Helmholtz Centre for Heavy Ion Research GmbH, Darmstadt, Germany



ACAT 2022 Conference, Bari, Italy

Abstract

In this work we present recent improvements of the Continuous Integration (CI, see beside) for the FairRoot software stack (see beside). We discuss relevant development workflows and how they were improved through automation. Furthermore, we present our infrastructure detailing its hardware and software design choices. The entire toolchain is composed of free and open source software. Finally, this work concludes with lessons learned from an operational as well as a user perspective and outlines ideas for future improvements.

Saving Energy (Ongoing work)

- Typically, the CI build farm has a low overall utilization with long idle periods and short peak loads
- Deployed an energy monitor
- Implemented the Slurm Power Save Suspend and Resume hooks to shut off idle worker nodes and restart them on-demand (via BMC signals)
- No long term data available yet, evaluation in the proceedings paper



Continuous Integration

Continuous Integration (CI) is a modern software engineering method to efficiently assure software quality. The underlying idea is to continuously rebase (pending) feature branches onto the latest version of the software repository and re-validate them by a set of checks in an automated fashion. This approach allows to naturally and efficiently resolve the merging of concurrent development activities (by one or multiple authors).

The (partial) automation of the software release process (e.g. generation of release assets, documentation, extended testing) including the automated deployment is sometimes described by the term Continuous Deployment (CD). However, the term CI is often understood to subsume CD as well. In this work we adopt this more general definition of CI.

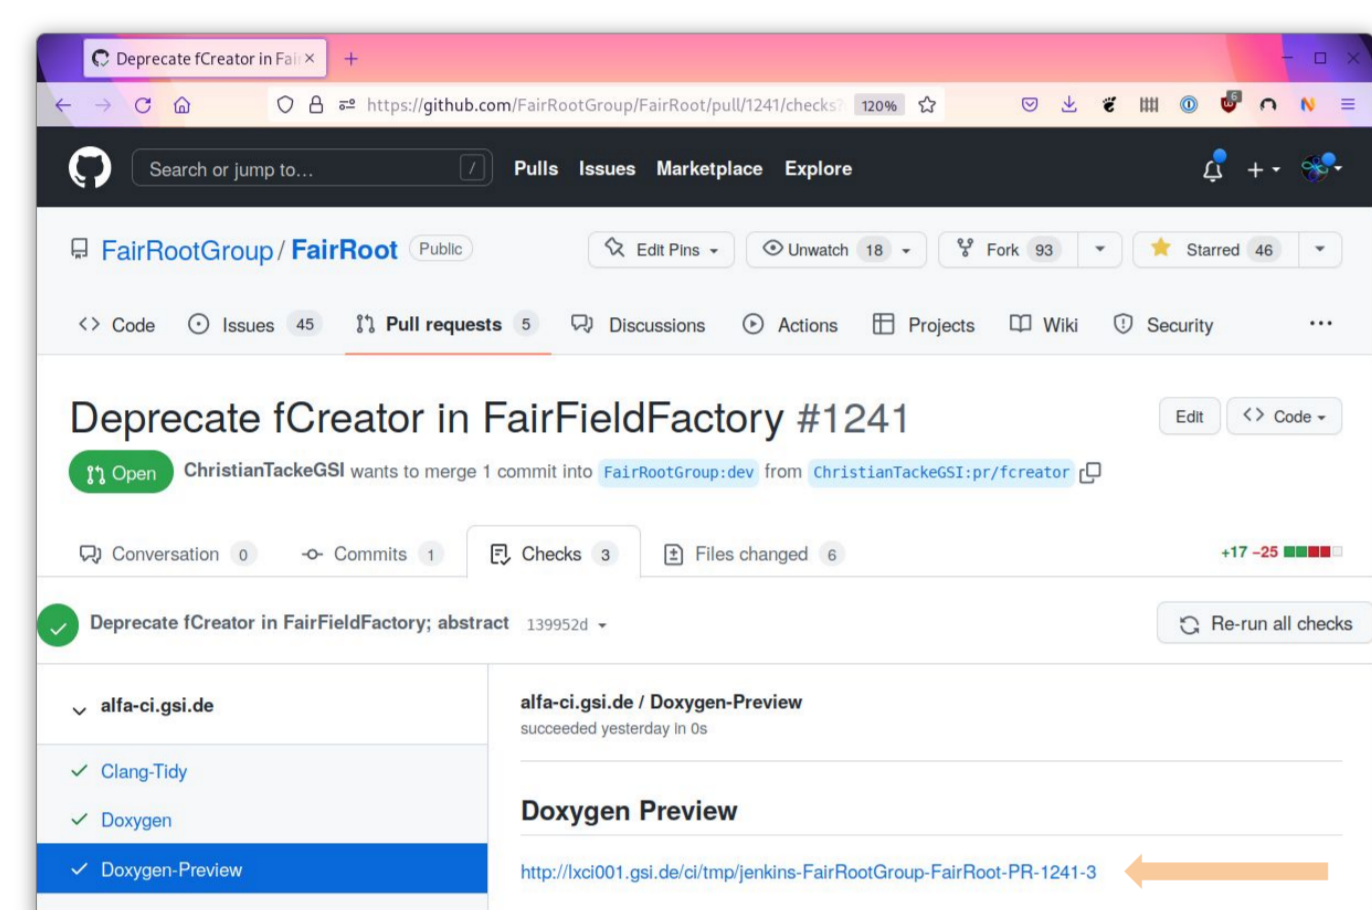
The FairRoot Software Stack

The FairRoot¹ software stack is a toolset for the simulation, reconstruction, and analysis of high energy particle physics experiments, currently used i.e. at FAIR/GSI, and CERN.



In the FairSoft² source distribution the most important software dependencies needed to run a FairRoot-based experiment software framework are released for Linux and macOS while further "system dependencies" are taken from the Linux distribution and the Homebrew project³ on macOS. Notable packages in a FairSoft release are

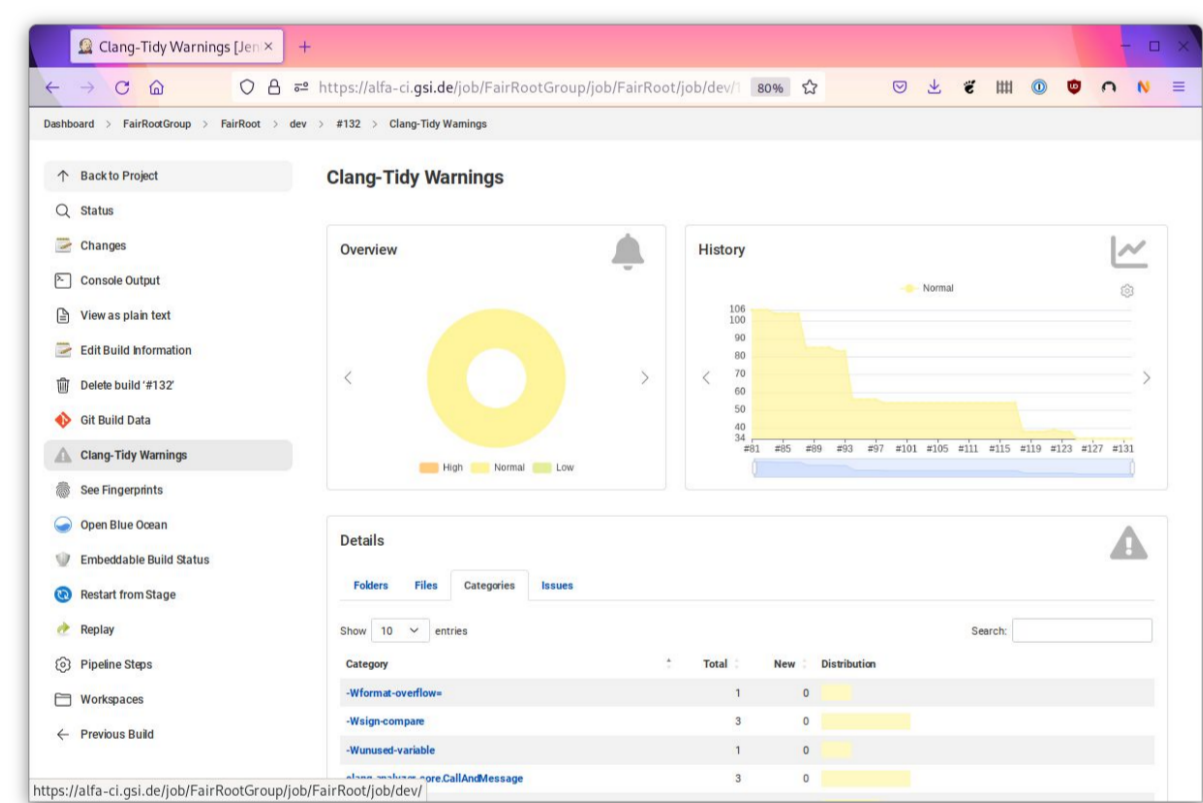
- ROOT (<https://root.cern/>),
- Boost (<https://www.boost.org/>),
- Geant3 (<https://github.com/FairRootGroup/geant3>),
- Geant4 (<https://geant4.web.cern.ch/>),
- VMC (<https://github.com/vmc-project/vmc>),
- HEPMC (<http://hepweb.cern.ch/>),
- Pythia6 (<https://github.com/alisw/pythia6>),
- Pythia8 (<https://pythia.org/>),
- Vc (<https://github.com/VcDevel/Vc>),
- VGM (<https://github.com/vmc-project/vgm>),
- CLHEP (<http://proj-clhep.web.cern.ch/>),
- DDS (<http://dds.gsi.de/>), and
- FairMQ (<https://github.com/FairRootGroup/FairMQ>).



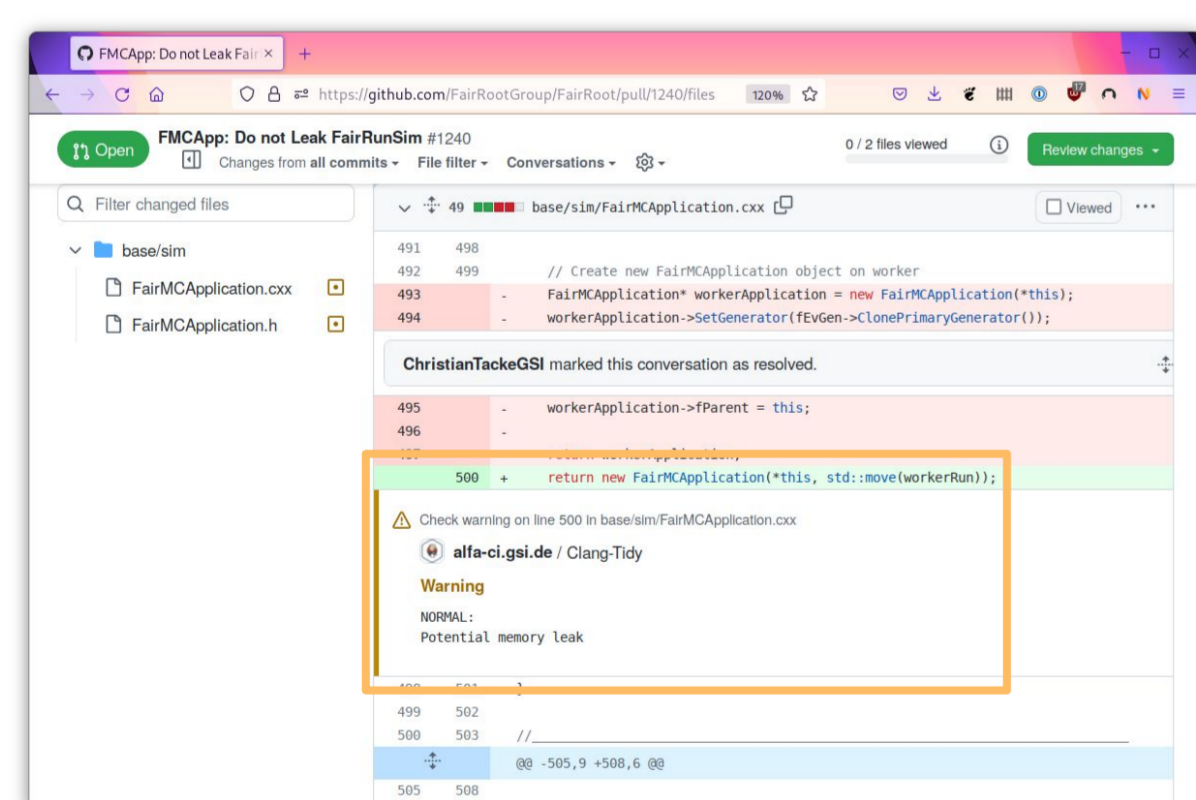
This example shows how a link to the doxygen based documentation preview in a FairRoot pull request is provided to the reviewers/developers.

- Automated generation of documentation

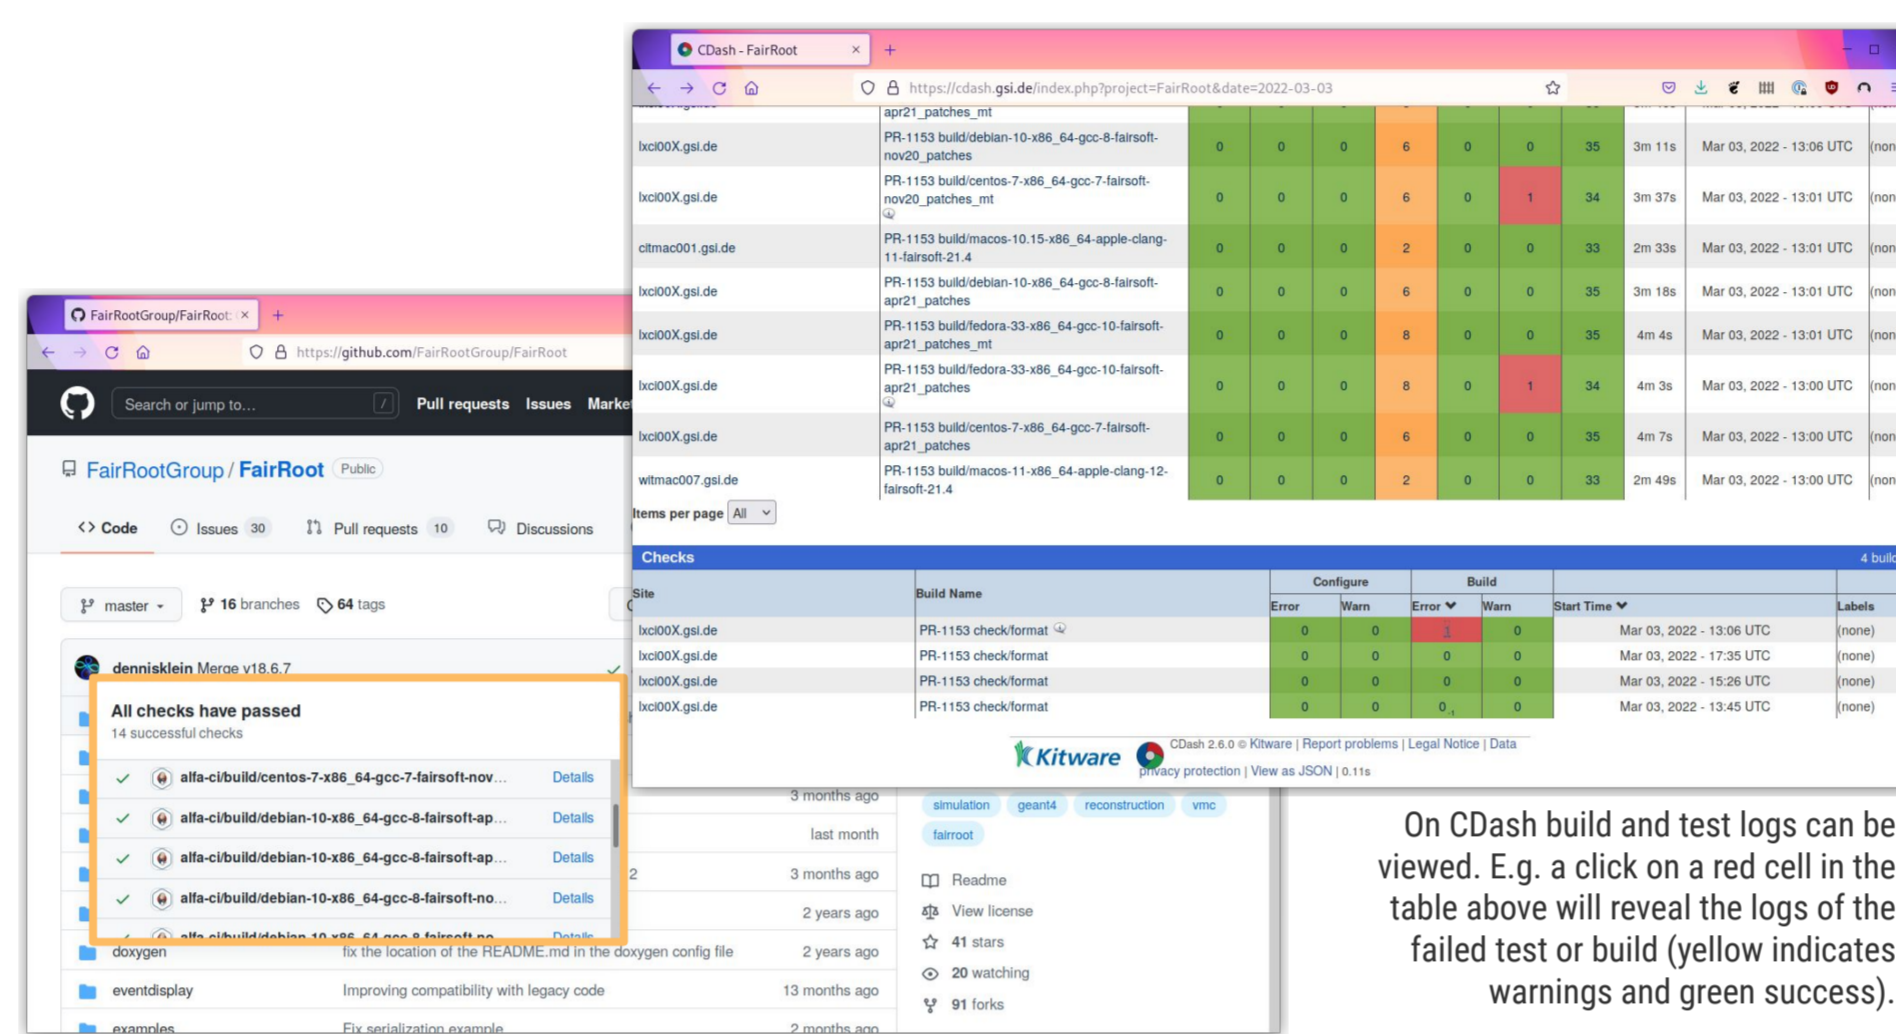
- Advanced build log analysis with CDash⁵ and the Warnings Next Generation Jenkins⁶ Plugin



This screenshot depicts a breakdown dashboard for the captured clang-tidy warnings for the FairRoot dev branch realized with the Warnings Next Generation Jenkins Plugin



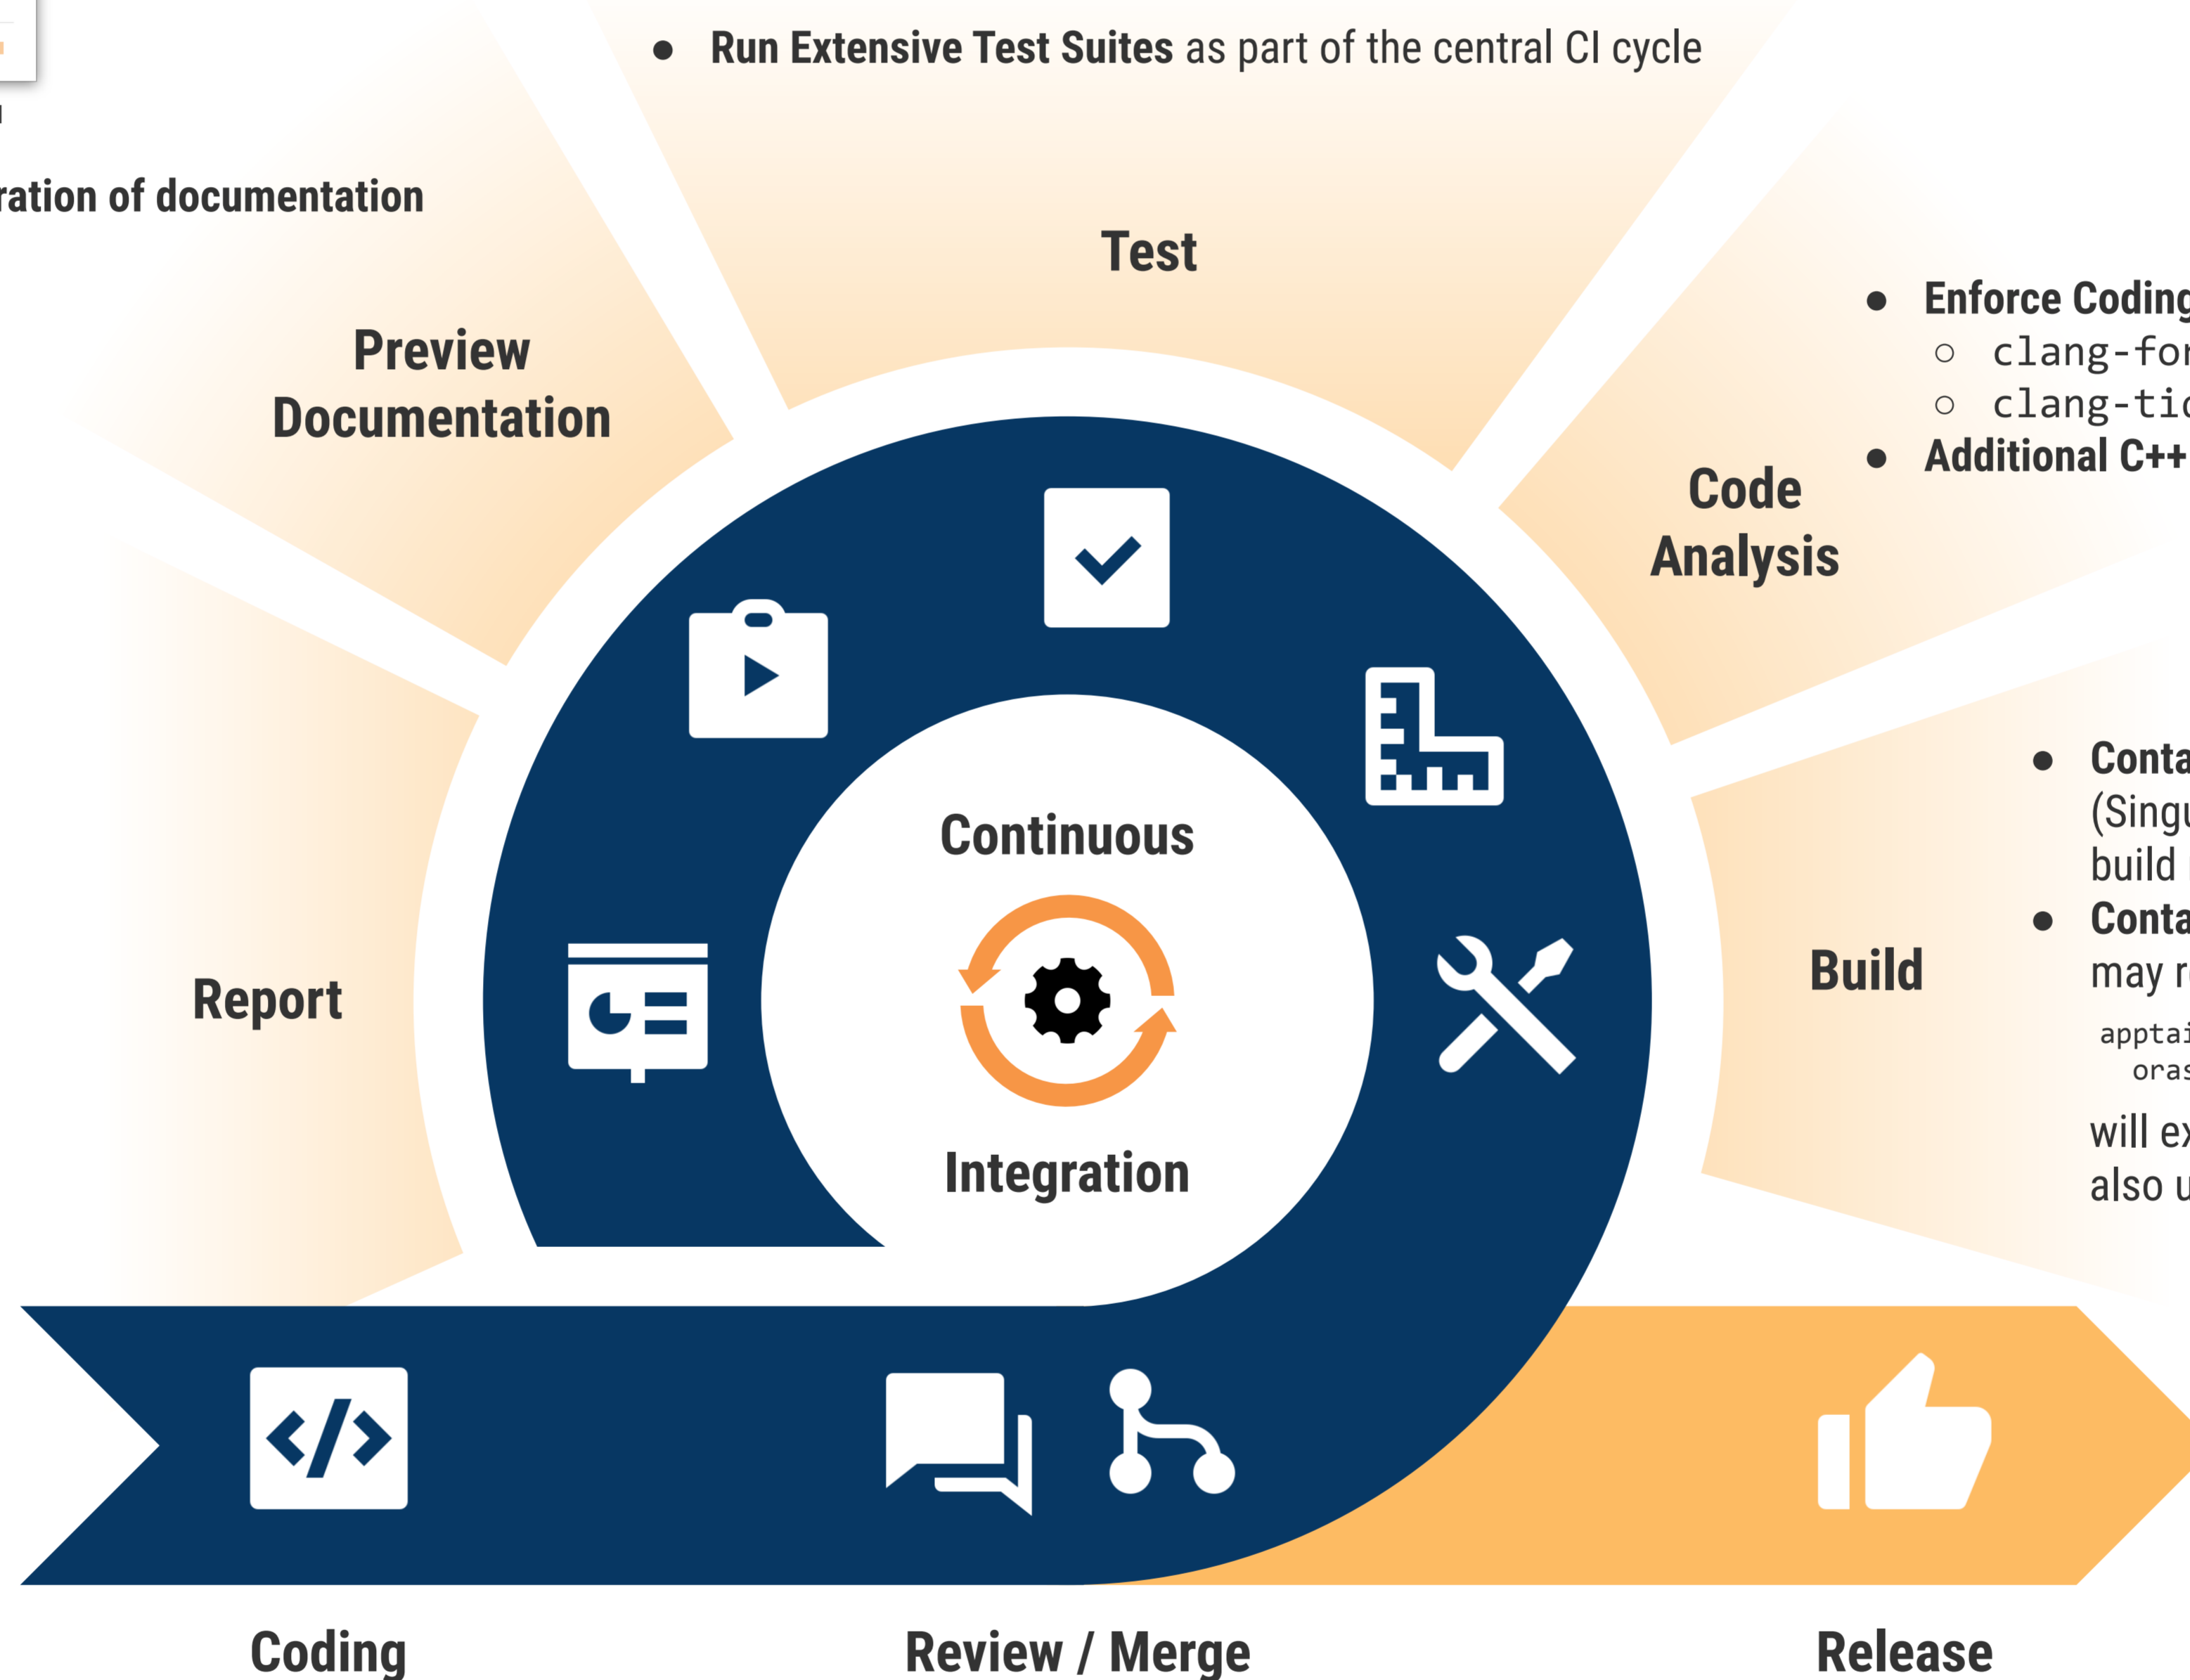
This example shows the "changes" tab on a FairRoot pull request with a source code annotation about a clang-tidy warning directly in the github.com user interface (realized with the Warnings Next Generation Jenkins Plugin).



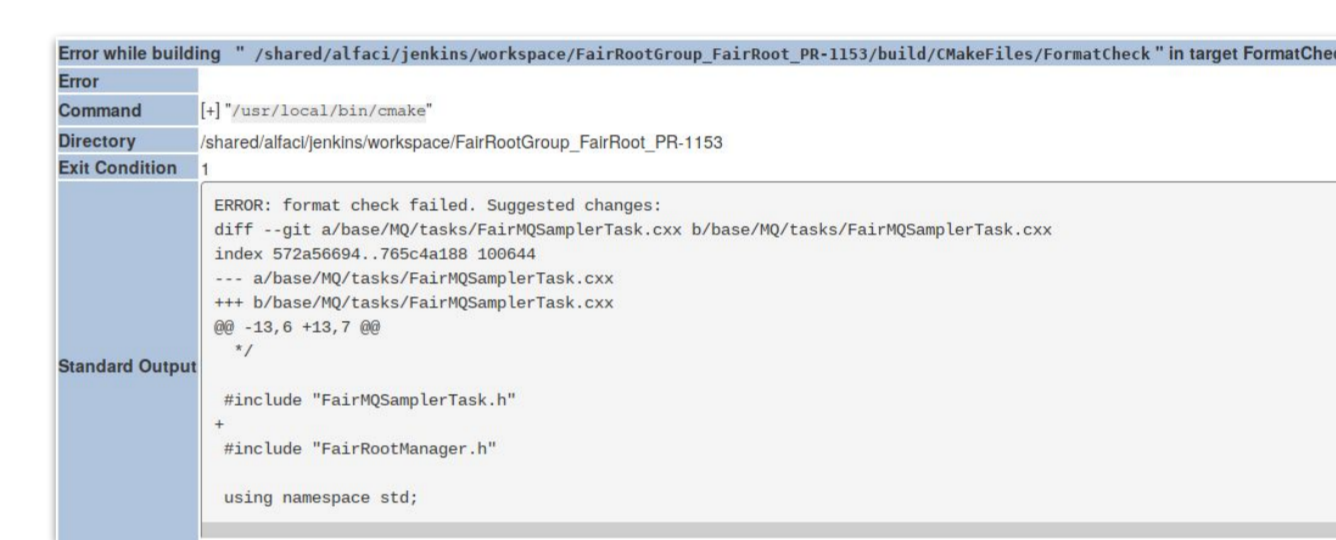
On CDash build and test logs can be viewed. E.g. a click on a red cell in the table above will reveal the logs of the failed test or build (yellow indicates warnings and green success).

Developers can discover performed checks directly integrated in the Repository User Interface (here github.com)

- Run Extensive Test Suites as part of the central CI cycle



- Enforce Coding Conventions via
 - clang-format
 - clang-tidy
- Additional C++ Analysis via Address, Thread, UB, Leak Sanitizers



The screenshot shows a failed clang-format check viewed via CDash.

- Containerized Builds (on Linux) with Apptainer⁹ (Singularity) to enable heterogeneous and large build matrices
- Container Definitions public so the developer may re-use the build environment locally, e.g.



```
apptainer exec \oras://ghcr.io/fairrootgroup/fairmq-dev/fedora-35-sif:latest <cmd>
```

will execute the given <cmd> in the same environment that is also used by FairMQ's CI jobs based on the Fedora 35 image.

Build Farm

- 5 Linux hosts (20-28 CPU cores, 128-256GB RAM each) + 3 Mac minis (2x x86_64, 1x aarch64) for **public, untrusted** CI jobs
- 3 Linux hosts (20-28 CPU cores, 128-256GB RAM each) for **private, trusted** CI (mainly CD) jobs
- All hosts have direct attached SSD storage suitable for **cross job build caches**
- Job run times range from a few minutes to multiple hours < 6h depending on repo and change.

Infrastructure

The relevant source code repositories are hosted in a heterogeneous environment, some are on github.com, some on various GitLab⁶ instances. GitLab natively comes with a CI service. As of ~2021 only, github.com started to offer a CI service. Because we started to use the CI method before, we operate a Jenkins CI service as well. In anticipation of the still quickly changing CI service landscape, we aim to keep a **weak** dependency on those CI services only. In addition, we want to support multiple instances in parallel and still share our build farm. To achieve this, CI jobs are scheduled onto the build farm via a shared queue - realized with the Slurm⁴ resource manager.

Lessons Learned

- Long Running CI jobs are acceptable
- Design CI checks cleanly from the bottom up and avoid unnecessary dependencies on the CI system

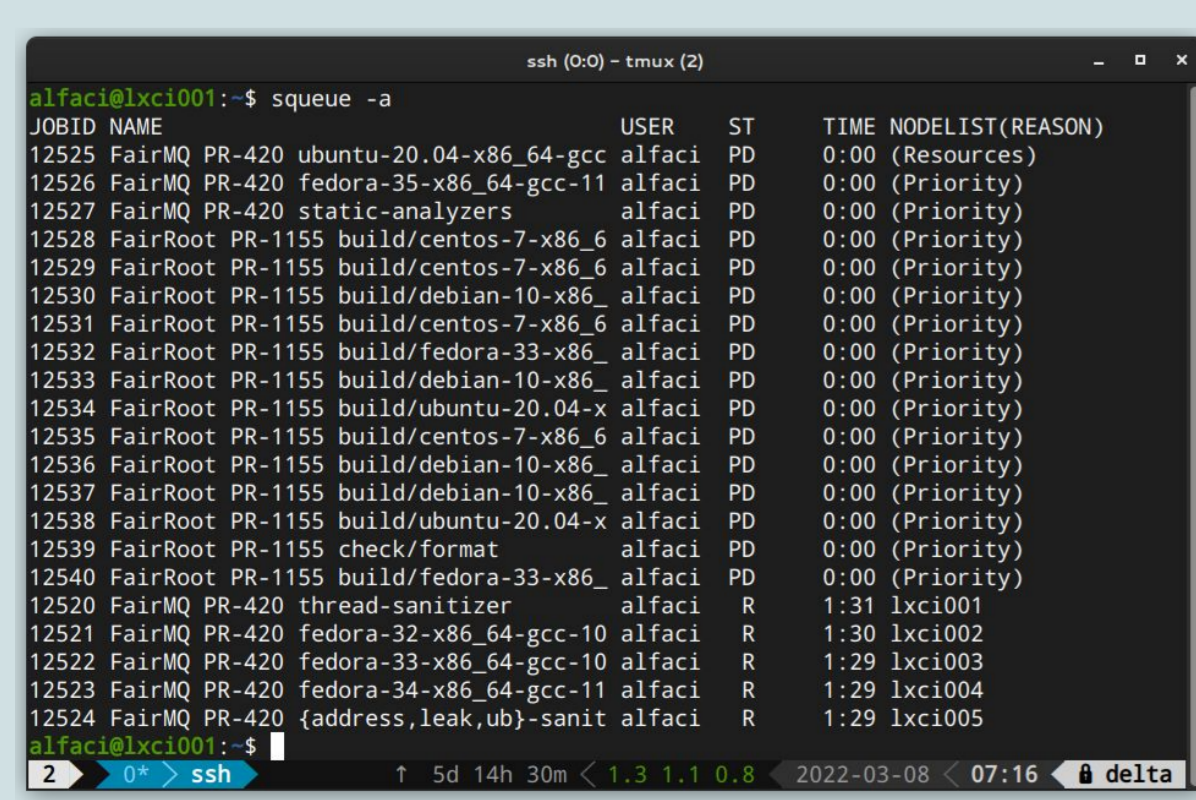
Ongoing and Future Work

- Publish pre-built CI container images on a registry and/or CVMFS⁸, so the developer can re-create a CI environment for a local interactive debugging session.
- Deploy and Operate a Github Actions Runner and explore if Jenkins can be phased out

References

- [1] FairRoot - <https://fairroot.gsi.de/>
- [2] FairSoft - <https://github.com/FairRootGroup/FairSoft>
- [3] Homebrew - <https://brew.sh/>
- [4] Slurm - <https://slurm.schedmd.com/>
- [5] CDash - <https://www.cdash.org/>
- [6] GitLab - <https://about.gitlab.com/>
- [7] Jenkins - <https://www.jenkins.io/>
- [8] CVMFS - <https://cernvm.cern.ch/fs/>
- [9] Apptainer - <https://apptainer.org/>

All URLs have been last visited on 19th October 2022.



The screenshot shows the status of the SLURM job queue filled with various CI jobs (R for Running, PD for Pending)

