


**EASYBUILD**



---

---

# EasyBuild Maintainers Summit 2023

**Mon 27 Feb 2023 + Wed 1 Mar 2023**

---

---

# Practical info

- Central topics:
  - Where are we now, where are we going
  - What works well, what should change
- Two sessions
  - Mon 27 Feb 2023 – 12:00-14:30 UTC
  - Wed 1 Mar 2023 – 12:00-14:30 UTC
  - Hopefully (significantly) shorter than 2.5h per session...
- Not a one-man show – please join the discussion!
- Live notes @ [https://hackmd.io/\\_5hRC8IHSoWQE8mt-kLfgw](https://hackmd.io/_5hRC8IHSoWQE8mt-kLfgw) (help is welcome!)
  - Will be copied to EasyBuild wiki after each session
  - If you don't want something mentioned in the notes, tell us!

**Slides to guide discussion, part I (Mon 27 Feb'23)**

# Things to discuss (not necessarily in order)


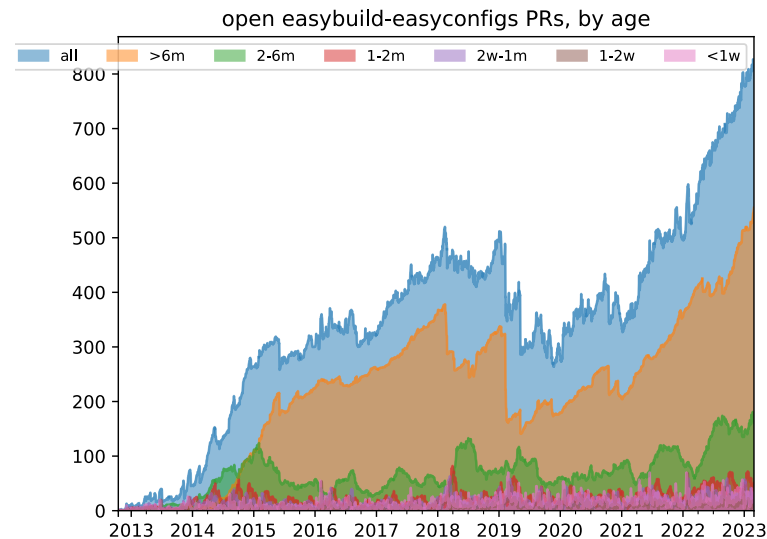
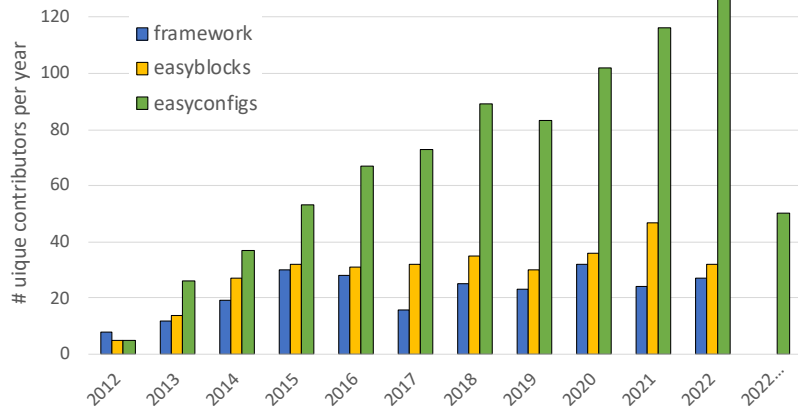
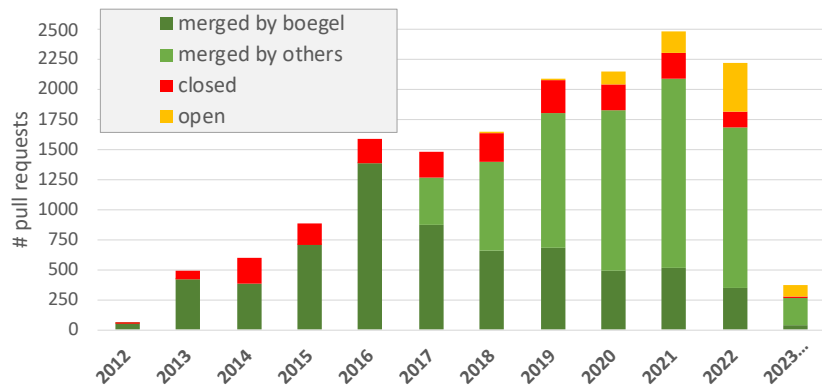
- Easyconfig PRs
- EasyBuild v5.0
- EasyBuild docs
- 6th EasyBuild User Survey
- EasyBuild Tech Talks in 2023
- EasyBuild User Meeting 2023
- **Additional topics still welcome!**

# Review of previous Maintainers Summit

<https://github.com/easybuilders/easybuild/wiki/Notes-from-EasyBuild-maintainer-summit-2021>

- “Missing” contributions from large sites => very little change?
- Using FlexiBLAS in common toolchain => done!
- Keeping up with incoming contributions => very little change?
- Revamp of EasyBuild docs => done!
- EasyBuild v5.0 => limited progress (deprecating Python 2)
- generoso test cluster => actively used for testing easyconfig PRs!
- Proper EasyBuild code-of-conduct => still WIP

# Easyconfig PRs – backlog (plots)




# Easyconfig PRs – backlog (stats)

- PRs per year:
  - 2019: 2,094 – 2020: 2,156 – 2021: 2,483 – 2022: 2,230
  - => ~10 PRs per work day, ~50 PRs per week
- Open PRs to date:
  - 2019: 11 - 2020: 107 – 2021: 181 – 2022: 416 – 2023: 93
  - ~18% of PRs opened in 2022 is still open ☹️
- Over 125 unique easyconfig contributors per year (and growing?)
  - New contributors are still popping up (good)
  - May not be aware of contribution policy

# Easyconfig PRs – backlog

- What can we do to push back on backlog + keep up with incoming PRs?
- Does Maintainer-of-the-Week (MotW) still work (well)?
  - How often do you manage to reach the “1h/day, on average” MotW target?
  - Do others feel like they shouldn’t be looking into PRs when they’re not MotW?
  - **Would another MotW role to look into oldest open easyconfig PRs make sense?**
- Maintainer “buddy” for easyconfig PR whales?
  - Close collaboration (temporary, voluntary) with contributors that have lots of open PRs
  - Could include 1-on-1 call(s) to discuss blocked PRs
- Potential candidates for additional maintainers?
  - @sassy-crick (269 merged PRs) - @SimonPinches (121) - @DirectorSloan (17)
- Separate session to exchange “best practices” for reviewing/testing/merging PRs?

# Easyconfig PRs – whales



Currently ~19 contributors with  $\geq 10$  open easyconfig PRs

# Easyconfig PRs – testing

- Trend towards more attention for tests & sanity checks (good!)
- Tension between testing and keeping PRs open longer
  - More testing reveals more problems, which we want to fix
  - Results in PRs being kept open for longer (see scipy, TensorFlow, PyTorch, ...)
- We should do more structured testing...
  - Clearly define minimal requirements w.r.t. test builds for PRs + Tier-1 vs Tier-2 test envs
  - Using containers to test for different OSs + also test on aarch64
  - More automation – bot being developed for EESSI could be a big help!
  - Also test in special environments like EESSI (RPATH + alternate sysroot + filter-deps + ...)
  - We should also leverage sponsored AWS/Azure credits from EESSI project!

# EasyBuild 5.0

- Drop support for Python 2 (code cleanup)
- Archive easyconfigs using deprecated toolchain
- Get rid of `distutils` dependency (required for Python 3.12 compatibility)
- Implement + adopt new `run` function with cleaner API
- Improved error reporting
- + various smaller changes...
- Not ETA yet, but hoping to have EasyBuild 5.0 released in 2023...
- **Working group to actively work towards EasyBuild 5.0 in coming months?**

# EasyBuild docs

- Port to Markdown + MkDocs finally done! \o/ *(thanks Simon + James)*
- Let's use this as a starting point for making EasyBuild docs better...
- Start reviewing and updating/improving/extended existing docs
  - Review one page/week => Full documentation reviewed in ~1 year
  - Update it (open PR) or identify shortcomings (open issue)
  - Fill in the gaps (module naming schemes, recently implemented features, ...)
- Integrate EasyBuild tutorial into EasyBuild docs?
- Separate group of maintainers for docs?

# 6th EasyBuild User Survey

- Currently preparing next EasyBuild User Survey
- Hope to have it open March 2023
- Present results during *State of the Union* talk at EUM'23
- Need to update existing questions + add a couple on ported docs
- **Need a volunteer to help review survey once it's ready...**

# EasyBuild Tech Talks in 2023

- Suggestions for topics?
- Jacamar CI (<https://gitlab.com/ecp-ci/jacamar-ci>) ?
- OpenMPI 5.0?
- Intel oneAPI?

# EasyBuild User Meeting 2023

- Who will be joining us in London, Mon-Wed 24-26 April 2023?
- Ideas for EasyBuild talks/tutorials – any volunteers?
  - *EasyBuild 5.0: coming soon!*
  - *Best Practices for EasyBuild Contributors*
  - *So you want to implement a bug fix of feature in EasyBuild framework...*
- Ideas for additional topics/speakers/tutorials?

**Slides to guide discussion, part II (Wed 1 Mar'23)**

# Things to discuss (not necessarily in order)

- EasyBuild v5.0
- Easyconfig PRs + MotW (revisited)
- Action items
- Follow-up meeting in ~1 month?

# EasyBuild 5.0

Different categories of things to implement:

- Backwards-incompatible changes
- Code cleanup
- Improvements
- Features
- Policies

# EasyBuild 5.0 – drop support for Python 2

- Only support Python  $\geq 3.6$
- Remove dead code only required to support Python 2
  - Like deprecation warning when using Python 2, `easybuild.tools.py2vs3` module
- Clean up `requirements.txt`
- Ensure there's a clean error message when trying to use Python  $< 3.6$ 
  - Tweak `eb` wrapper script
- Actively clean up code to leverage Python 3.6+ features
  - Both in framework and `easyblocks`

# EasyBuild 5.0 – changes (1/2)

- Changing defaults
  - Enable `download_dep_fail`, `use_pip`, `sanity_pip_check` by default in `PythonPackage`
  - Run `prepare` step before `source` step (load deps before unpacking sources) ([issue #1376](#))
  - Use Slurm as default backend for `--job` (+ deprecate support for GC3Pie)
- Don't run sanity check commands in installation directory
  - Mainly to avoid side effects when running sanity check
  - This may break some (incorrect) checks that assume they're being run from installation
- Change name of run method used to install extensions
  - Total misnomer, too confusing... => add `install_ext` method, deprecate `run` method

# EasyBuild 5.0 – changes (2/2)

- Stop setting \$CPATH, set \$C\_INCLUDE\_PATH & co instead ([issue #3331](#))
- Drop support for Lmod 6.x (last release Nov'16) – also Lmod 7.x? last release Mar'19)
- Use different Python package to support --dep-graph (networkx?)
- Cleaner way to specify toolchain in easyconfig
  - `toolchain = {'name': 'foss', 'version': '2022b'}`  
=> `toolchain = ('foss', '2022a')`
- Better way to specify alternate name for extension load checks
  - `options = {'modulename': 'foo'}` => `load_ext_as = 'foo' (?)`
- Make specifying software license required in easyconfigs (?)

# EasyBuild 5.0 – code cleanup

- Archive easyconfigs using a deprecated toolchain
  - + ancient software versions when system toolchain is used
  - Avoid (significant) decline in supported software, as much as possible
- Remove experimental support for `.yeb` easyconfigs
  - Dead feature, no improvements for years, it's just not going to happen...
- Stop relying on distutils ([issue #3963](#))
  - Required to support Python 3.12 (distutils will be removed from Python stdlib)
  - First step already done: `ingest(distutils.version.)LooseVersion`
- Remove dead code in easyblocks (ancient software versions)
- Use `LooseVersion` from `easybuild.tools` in easyblocks

# EasyBuild 5.0 – improvements

- **Cleaner `run` function**

- API of `run_cmd` and `run_cmd_qa` functions is a mess...
- We should implement a new `run` function to replace *both*
- Should be significantly easier if only Python 3.6+ needs to be supported!
- Keep `run_cmd` and `run_cmd_qa` for now, but *deprecate* them (and stop using them)
- Possible inspiration: <http://amoffat.github.io/sh>

- **Improved error reporting**

- More meaningful error messages that provide useful actionable info (path to log file, cmd, etc.)
- Context-aware error reporting (`config.log` when `configure` fails, error patterns, ...)
- Catch EasyBuild crashes, suggest to report issue?
- Support for easily starting an interactive shell in build environment could help a lot too...
- Using colors (if optional `rich` dependency is available)

# EasyBuild 5.0 – features

- Stable support for installing (R) extensions in parallel
- Native support in framework for bundle of components (replacement for Bundle)
- Proper support for runtime patching (no more abusing sed)
- Support for querying libraries.io for available software versions?
- Support for starting interactive shell to make debugging failing installations easier
- Support for resuming builds (checkpoint build dir + env per installation step)

# EasyBuild 5.0 – policies

- Yearly review cycle of EasyBuild docs
- Start to actively archive/deprecate old toolchain versions
  - Archive easyconfigs using a 20XY-5 toolchain
  - Deprecate easyconfigs using a 20XY-4 toolchain
  - Stop accepting PRs for new easyconfigs that use a 20XY-3 toolchain
- Change tactics in reviewing/testing of easyconfig PRs
  - Labels to help maintainers prioritize PRs to review (waiting-for-contributor, new/update, etc.)
  - Automate as much as possible (GitHub App / bot) + test in virtual Slurm clusters (AWS/Azure)
  - Using containers + different test environments (like EESSI)
  - Test environments matrix (OS + CPU arch + EasyBuild configuration)
  - Different CPU architectures (Intel, AMD, aarch64)
  - Public dashboard (?)

# EasyBuild 5.0 – battle plan

- Fork off 5.0.x branches after EasyBuild v4.7.1 release
  - Should be regularly synced with deve1op branch (also between releases)
- Start implementing planned changes in 5.0.x branches (via PRs)
- Keep making regular 4.x.y releases until 5.0.0 release is out
- Add documentation page with overview of (breaking) changes
- Bi-weekly (or monthly?) sync calls to discuss progress, next steps, ...
- Present progress report + overview of planned changes at EUM'23
- **Group of volunteers to actively work towards EasyBuild v5.0**
  - Prioritize tasks (create issues + project board?), and gradually tackle them...

# Easyconfig PRs – revisited

- Additional MotW role to reduce backlog
  - Tackle oldest open PRs (aim for closing/merging ~10 PRs per week?)
- Labels to help prioritizing PRs
  - Ideally added/removed automatically by a bot...
- Revamp testing of easyconfigs PRs: containers, better bot(s), ...
- More “aggressive” closing of old PRs
- Automatically submit review-pr-like overview of changes in PRs

# Action items – volunteers wanted!

1. Introduce additional MotW role to tackle oldest open easyconfig PRs
2. Set up working group for EasyBuild v5.0
3. Checklist/guidelines for reviewers
4. Templates for issues/PRs
5. Create issues for yearly review cycle of docs + find docs maintainers
6. Labels for PRs, semi-automatically managed by a GitHub App bot
7. Revamp testing of easyconfig PRs
8. Automatically submit review-pr-like overview in easyconfig PRs
9. Proper EasyBuild code-of-conduct
10. Plan follow-up meeting mid April'23 (somewhere week of 17-21 April)