

State of the union

Kenneth Hoste


Ghent University, Flemish Supercomputer Centre (VSC)

kenneth.hoste@ugent.be

***10th* EasyBuild User Meeting**

25 March 2025 - Jülich (Germany)

<https://easybuild.io/eum25>



State of the union

Kenneth Hoste

Ghent University, Flemish Supercomputer Centre (VSC)

kenneth.hoste@ugent.be


***10th* EasyBuild User Meeting**

25 March 2025 - Jülich (Germany)

<https://easybuild.io/eum25>

State of the Union - Outline

- Recent developments in EasyBuild (technical, community, uptake, ...)
- Results of 8th EasyBuild User Survey
- Brief look at EESSI & EuroHPC Federation Platform
- Short + long term goals & challenges
- Crazy ideas...



EasyBuild in a nutshell

easybuild.io

docs.easybuild.io

youtube.com/c/easybuilders

bsky.app/profile/easybuild.bsky.social

[tutorial .easybuild.io](https://tutorial.easybuild.io)

easybuild.io/eum

- **Framework to get scientific software installed** (preferably from source)
- Focus on Linux & HPC systems, specific attention to testing & performance of installed software
- Implemented in Python, requires Python 3.6+ (*Python 3.9+ recommended*)
- Integrates with environment modules tool (Lmod or Environment Modules)
- **Open source since 2012** (GPLv2), available via GitHub and Python Package Index (PyPI)
- Supports different compiler toolchains + over 2,500 software packages (excl. extensions)
- Originally created by HPC-UGent, now supported & developed by a **worldwide community...**

EESSI in a nutshell




- **Community project** that evolved from the EasyBuild community, started in 2020
- Main goal: build a **publicly accessible, shared, optimized software stack** for scientific computing
- Focus on easy-of-use, performance, testing, automating, collaboration, ...
- Powered by open source software, incl. EasyBuild, Lmod, CernVM-FS, Gentoo Prefix, ReFrame, ...
- Effort is currently funded through MultiXscale EuroHPC Centre-of-Excellence (2023-2026)
- **EESSI is being integrated in EuroHPC Federation Platform as Federated Software Stack!**
- We will focus extensively on EESSI during 3rd day of EUM'25


eessi.io | github.com/EESSI | eessi.io/docs | eessi.io/docs/blog




EasyBuild User Survey



- Anonymous survey, **8th edition** (was held in Jan-Feb 2025)
- Feedback from + insight into EasyBuild community
- 52 questions, covers wide variety of aspects
- 117 people participated (1 short of record of 118 in 2022)
- Full survey results available soon at <https://docs.easybuild.io/user-survey>
- Highlights are covered in this talk - questions with “boring” results are skipped
- 8 years of survey results allows for observing evolution and trends...




EasyBuild User Survey

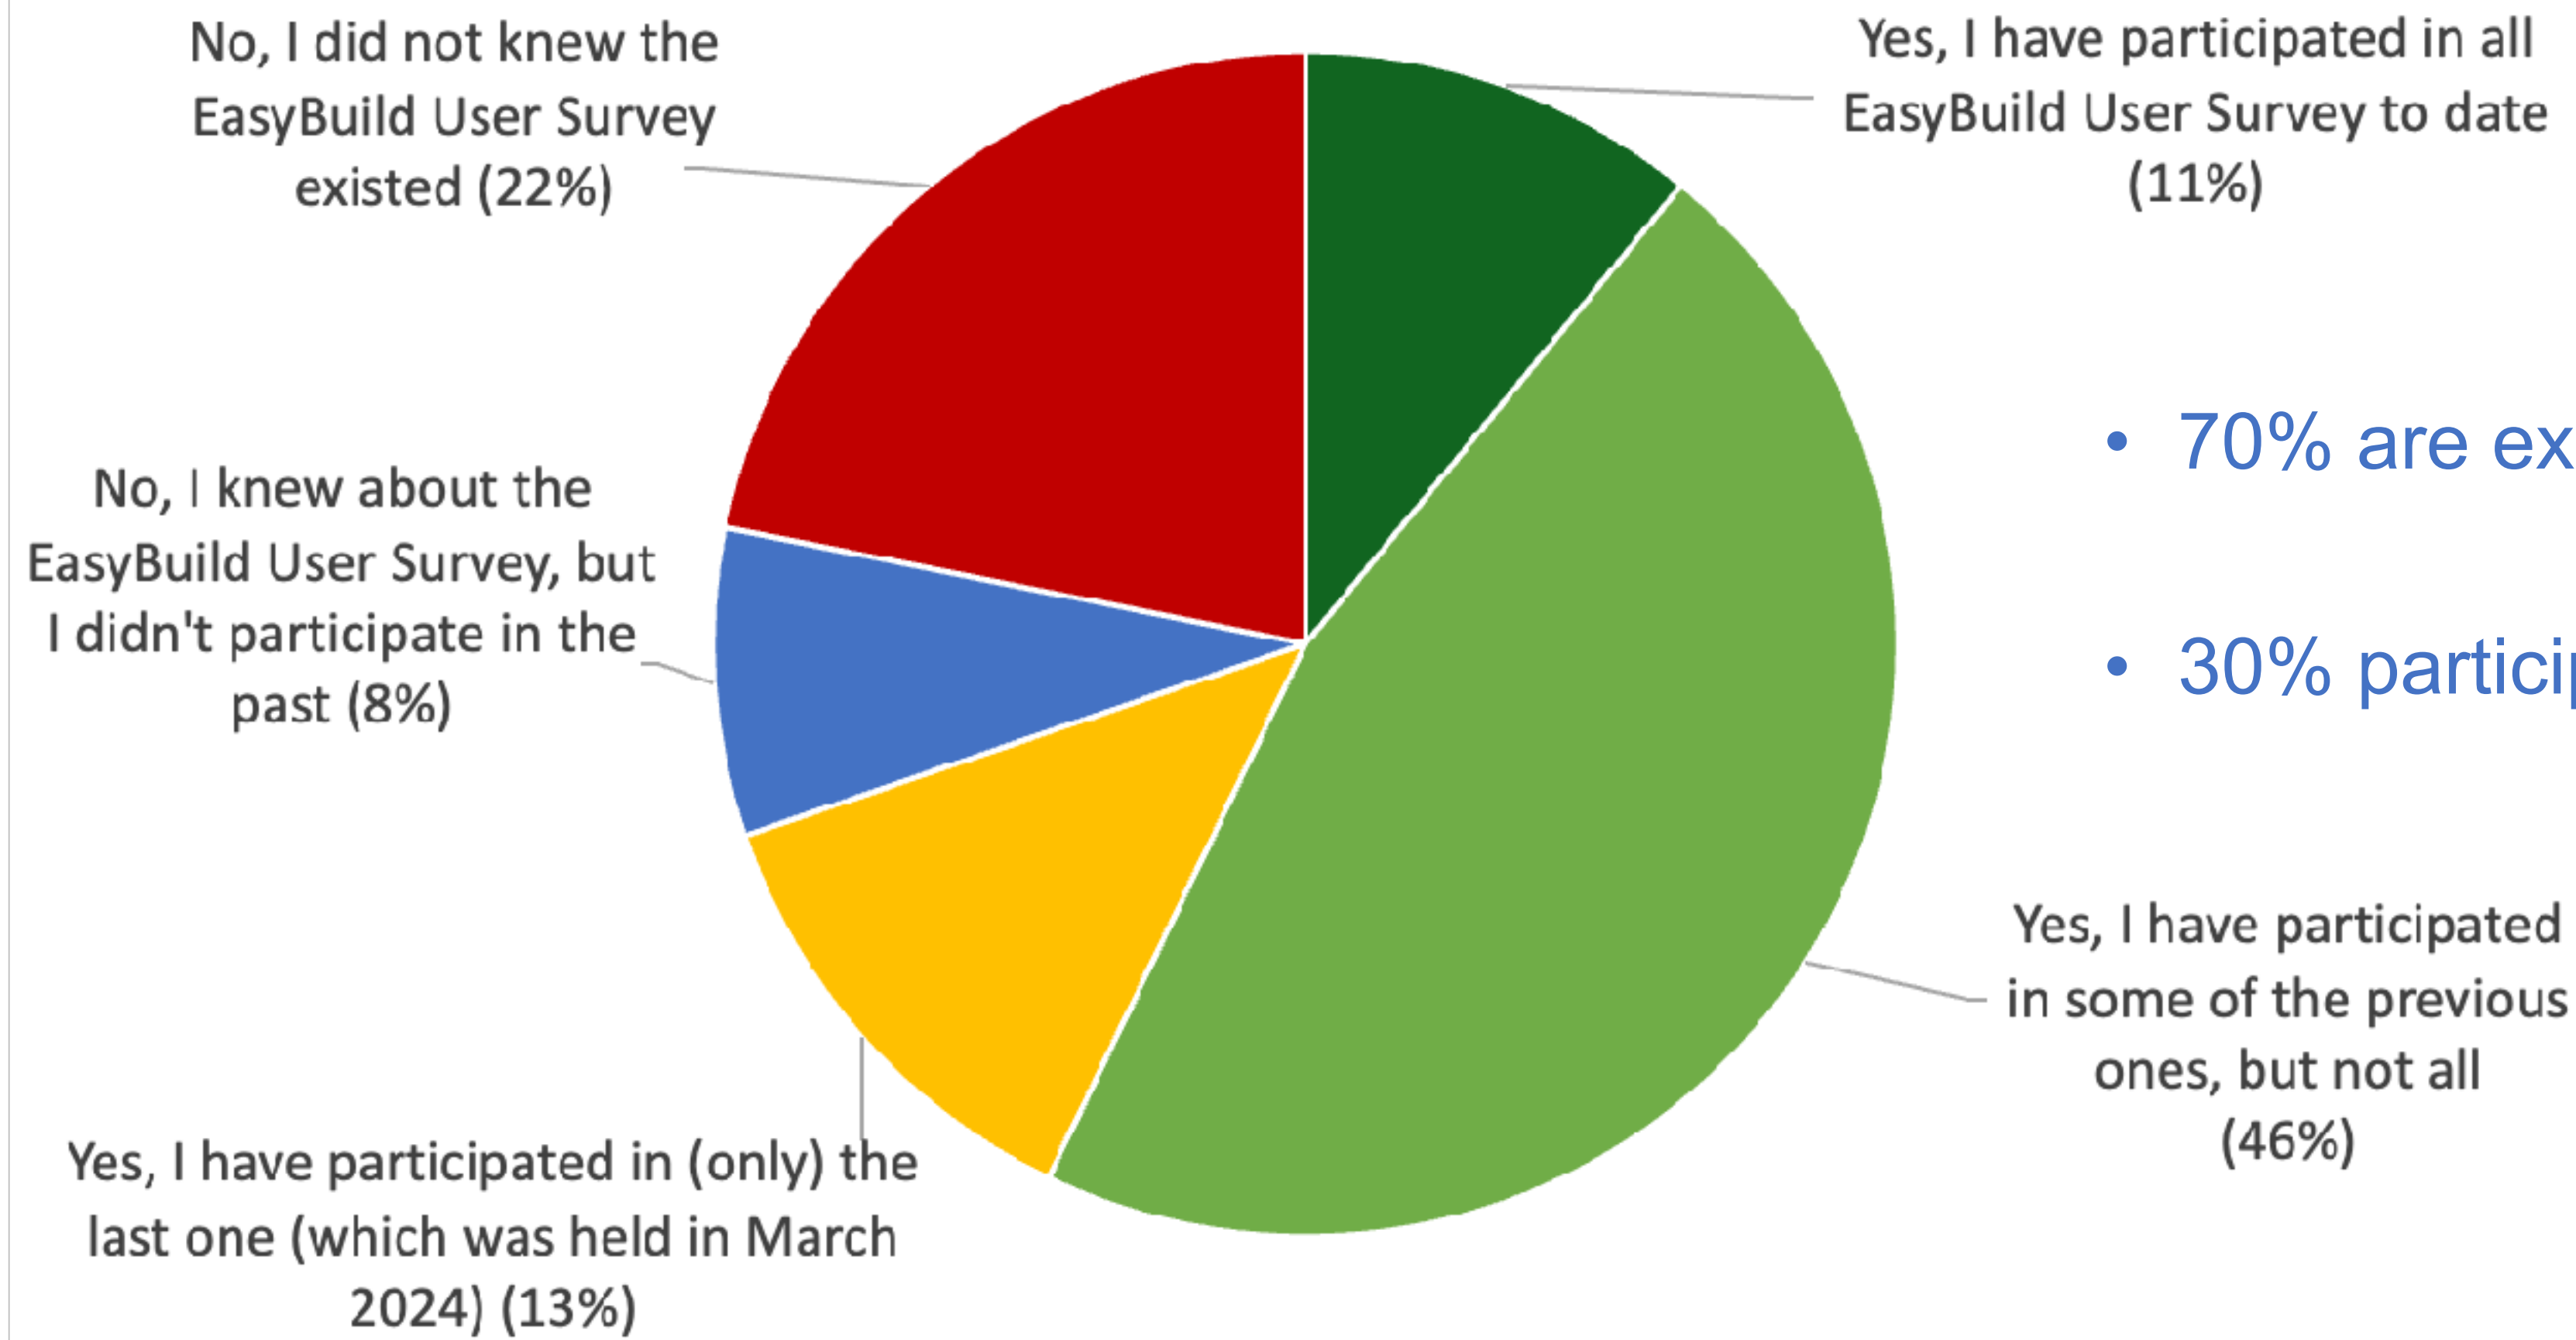


Survey results are *assumed* to give a representative view on EasyBuild community,
but results should be interpreted with a grain of salt...

EasyBuild User Survey




Have you participated in the EasyBuild User Survey before?




- 70% are experienced survey participants
- 30% participate for the first time

Community insights through user survey



- Vast majority of community still has sysadmin (~43%) or user support (~28%) role
- Mostly European (~83%), small part from North America (~13%)
- Large mix in terms of experience with EasyBuild, over 50% has 5+ years of experience
 - 10+ years: 10% | 5-10y: 42% | 2-5y: 27% | 13% 1-2y: 13% | 6m-1y: 2% | < 6m: 6%
- Word-of-mouth remains main way of how people learn about EasyBuild (>50%)
 - ~27% learn about EasyBuild when starting a new job

Mix of new & experienced EasyBuilders



- **We should be aware of people new to EasyBuild** (or even installing software on HPC systems)
- They can learn a lot from experienced EasyBuild users, **but also the other way around!**
- We should try and look beyond the questions they ask: **what can we do to improve EasyBuild?**
- **Feedback by newcomers is very valuable**, experienced EasyBuilders “go blind” to certain things
- Influx of new people is partially fueled by EESSI & MultiXscale
 - They are funded to work with/on EasyBuild (but in a specific context)



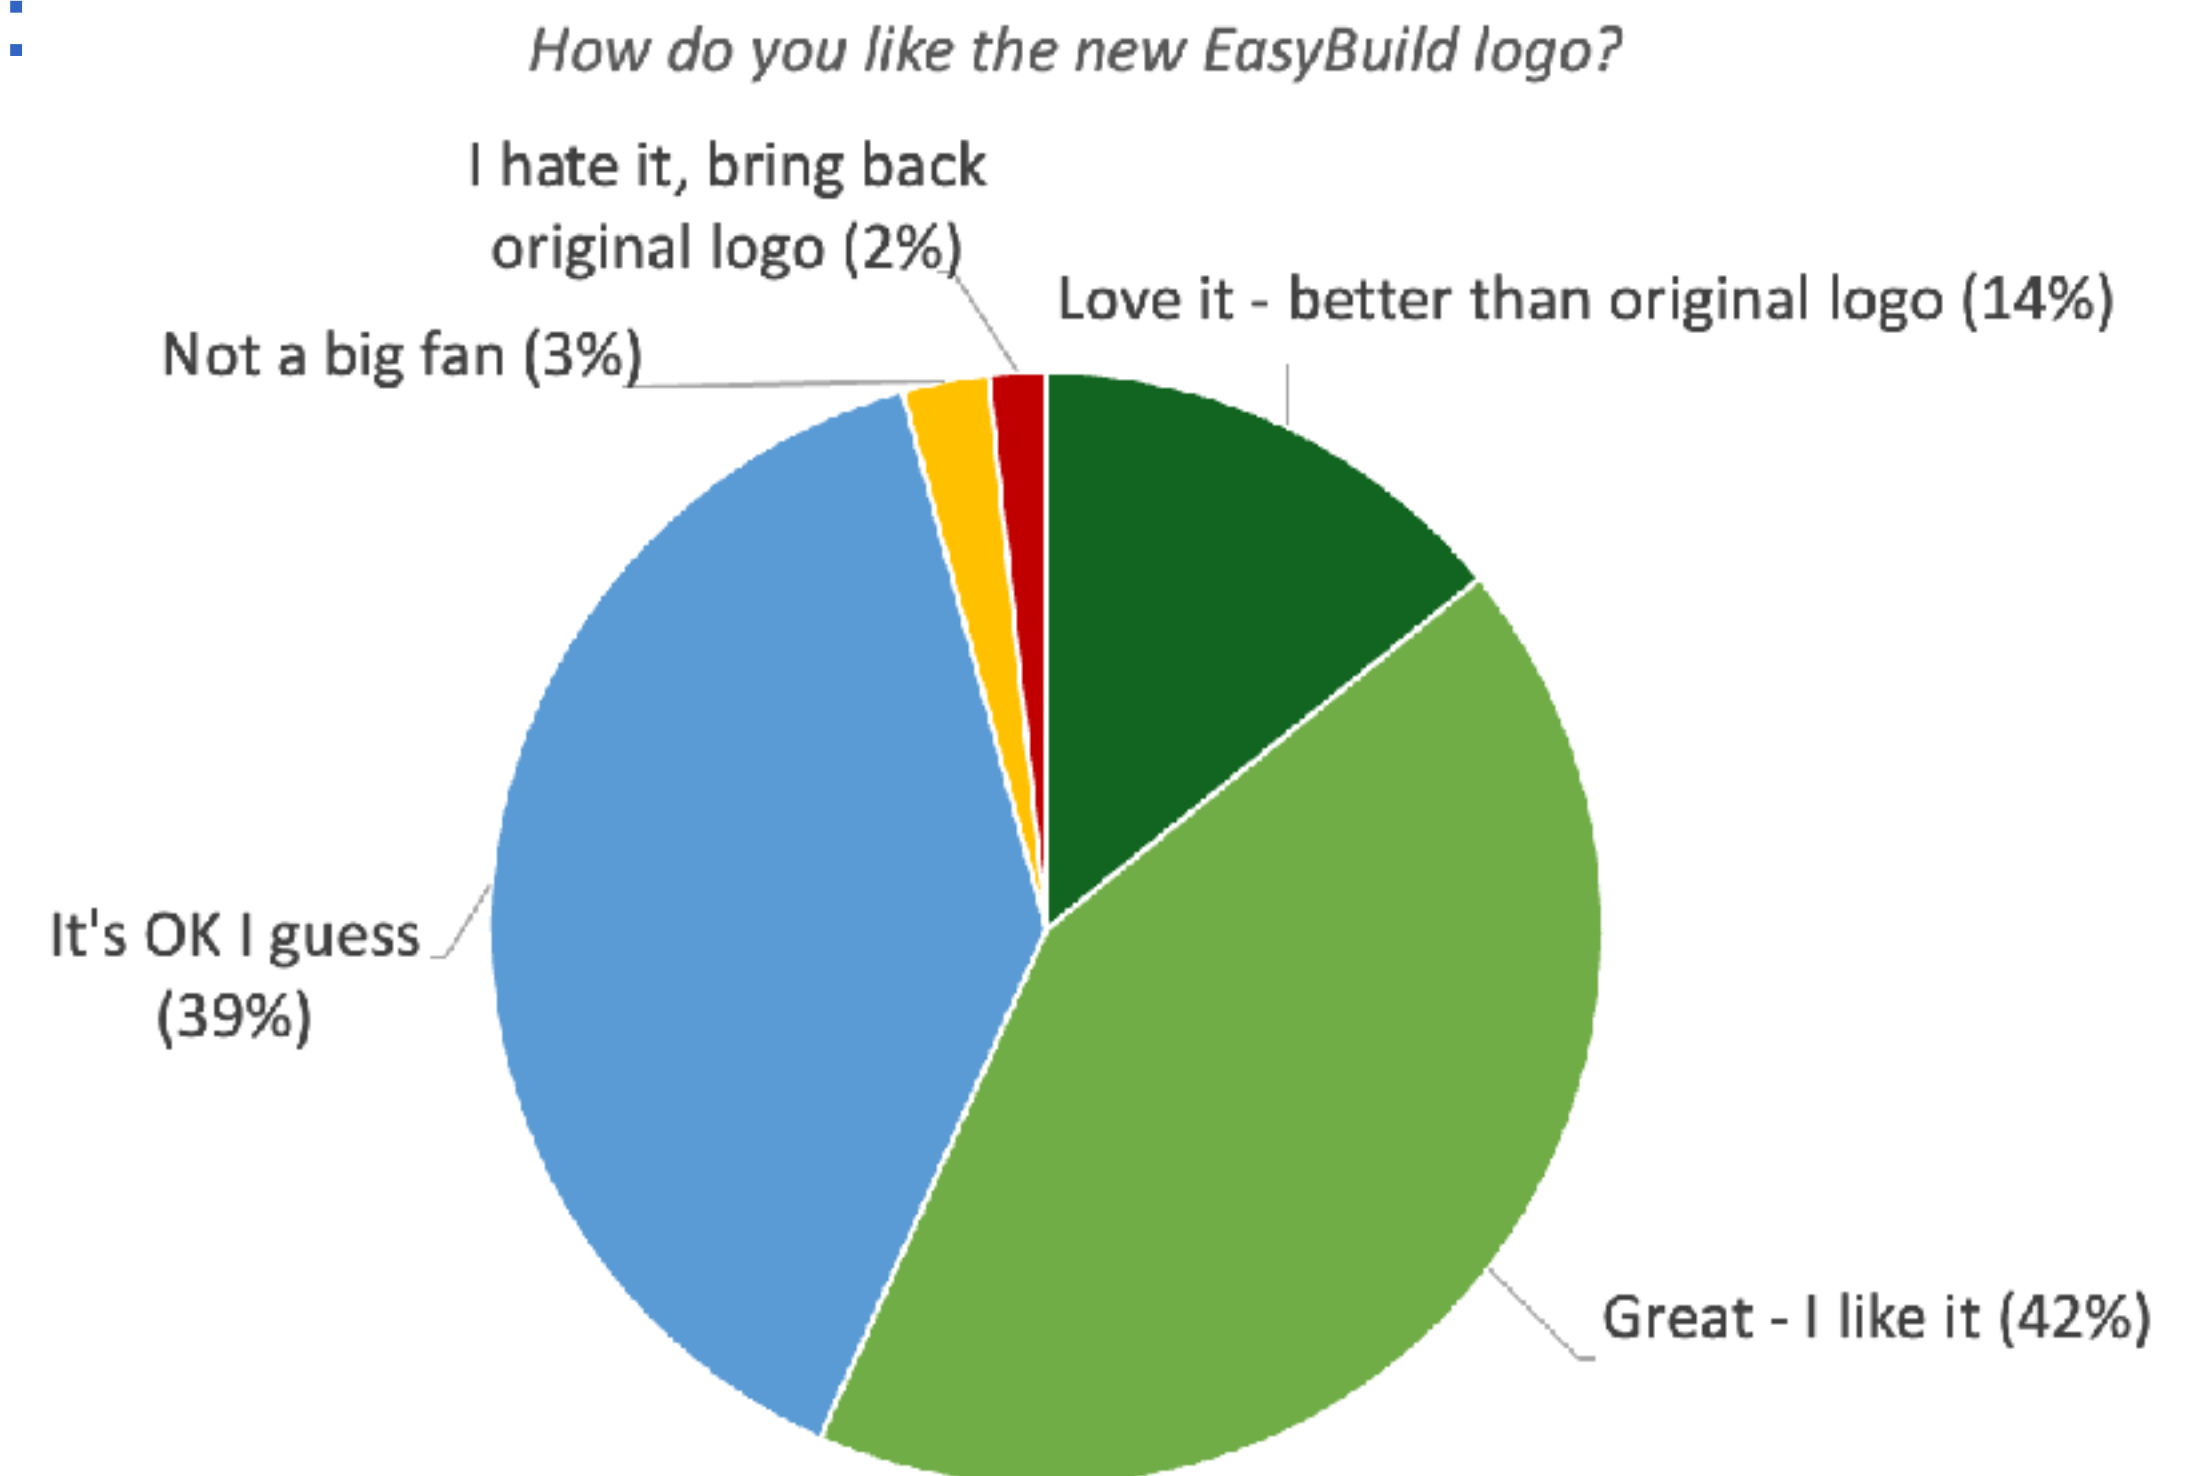

New(ish) EasyBuild logo

A new logo for **EasyBuild** was introduced on 14 Nov 2022, exactly 10 years after the release of EasyBuild v1.0 .

More information on (the motivation for) new logo:

<https://easybuild.io/new-logo-2022.html>

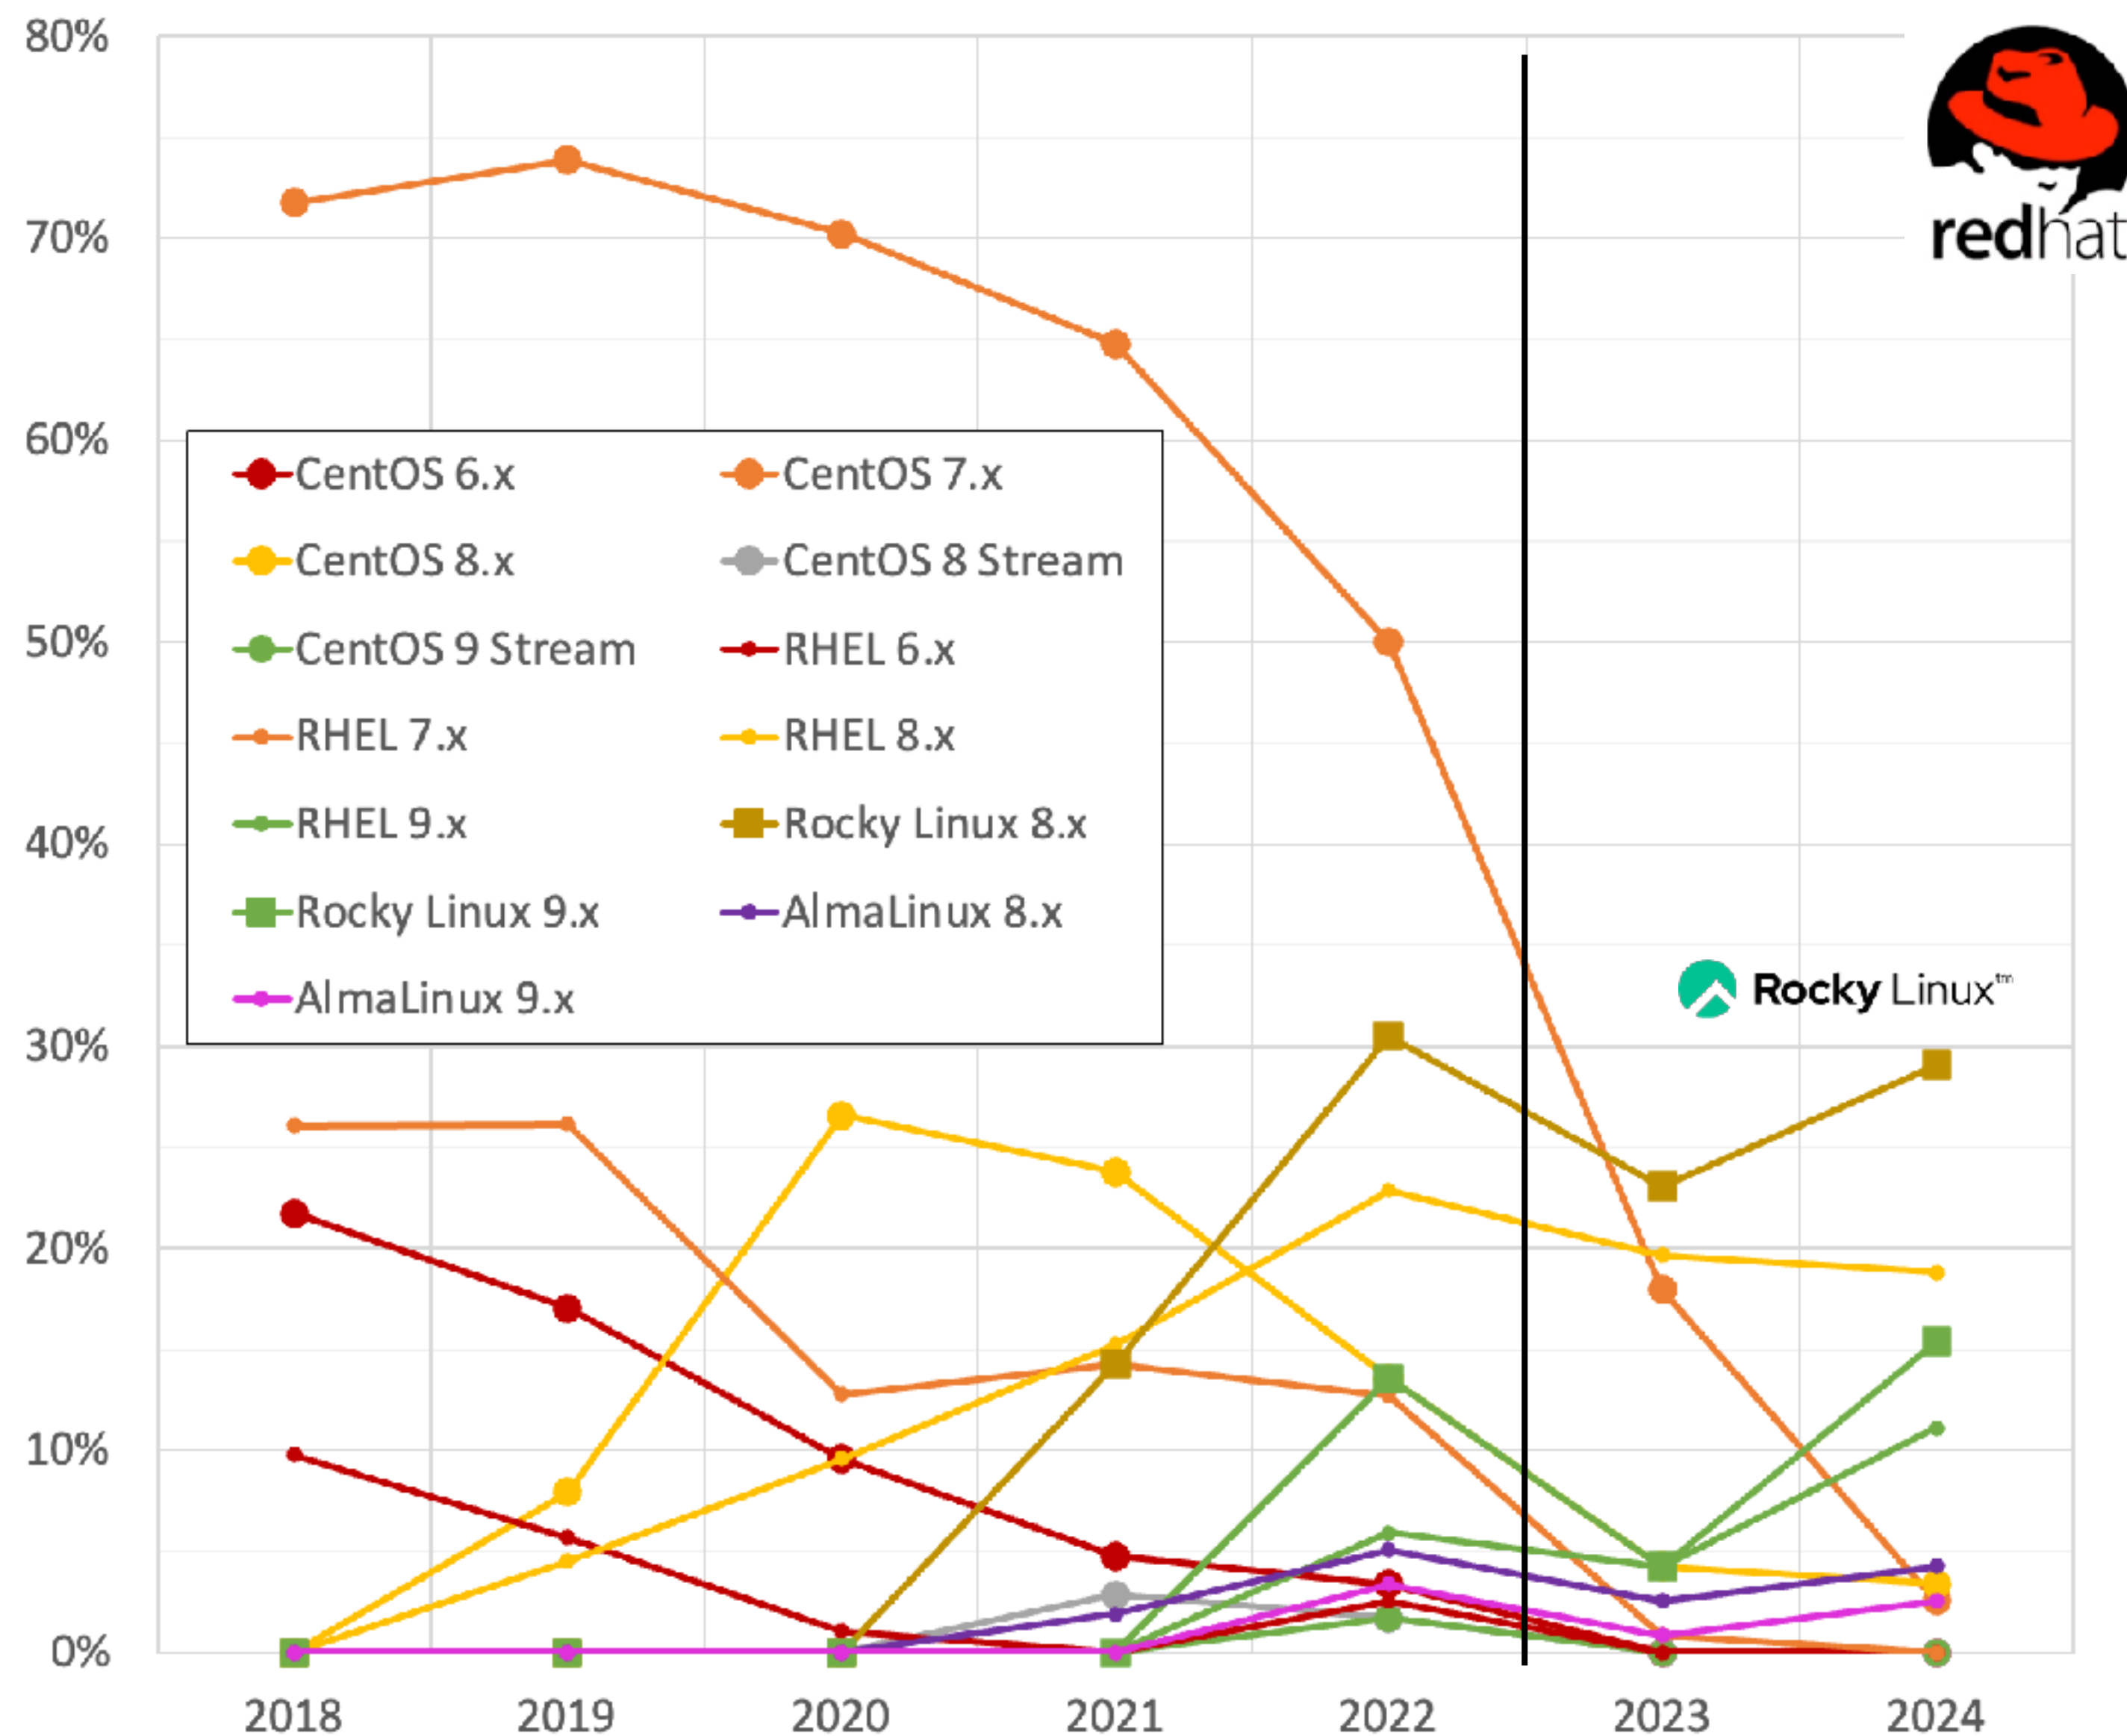
New logo “approved” by majority,
it’s growing on people who didn’t like it initially,
small stubborn group (now ~5%, was ~12%)
doesn’t really like it.



Operating systems in EasyBuild community





On which operating system(s) do you use EasyBuild?



- Vast majority of EasyBuild community is using a RHEL-based OS (unchanged)
- **CentOS 7 is basically dead** (EOL June 2024), ~18% last year, now ~2.5%
- Rocky Linux 8 (~29%) + RHEL8 (~19%) replaced CentOS 7
- Adoption of RHEL9-based OS growing: ~15% RHEL 9, ~11% Rocky Linux 9

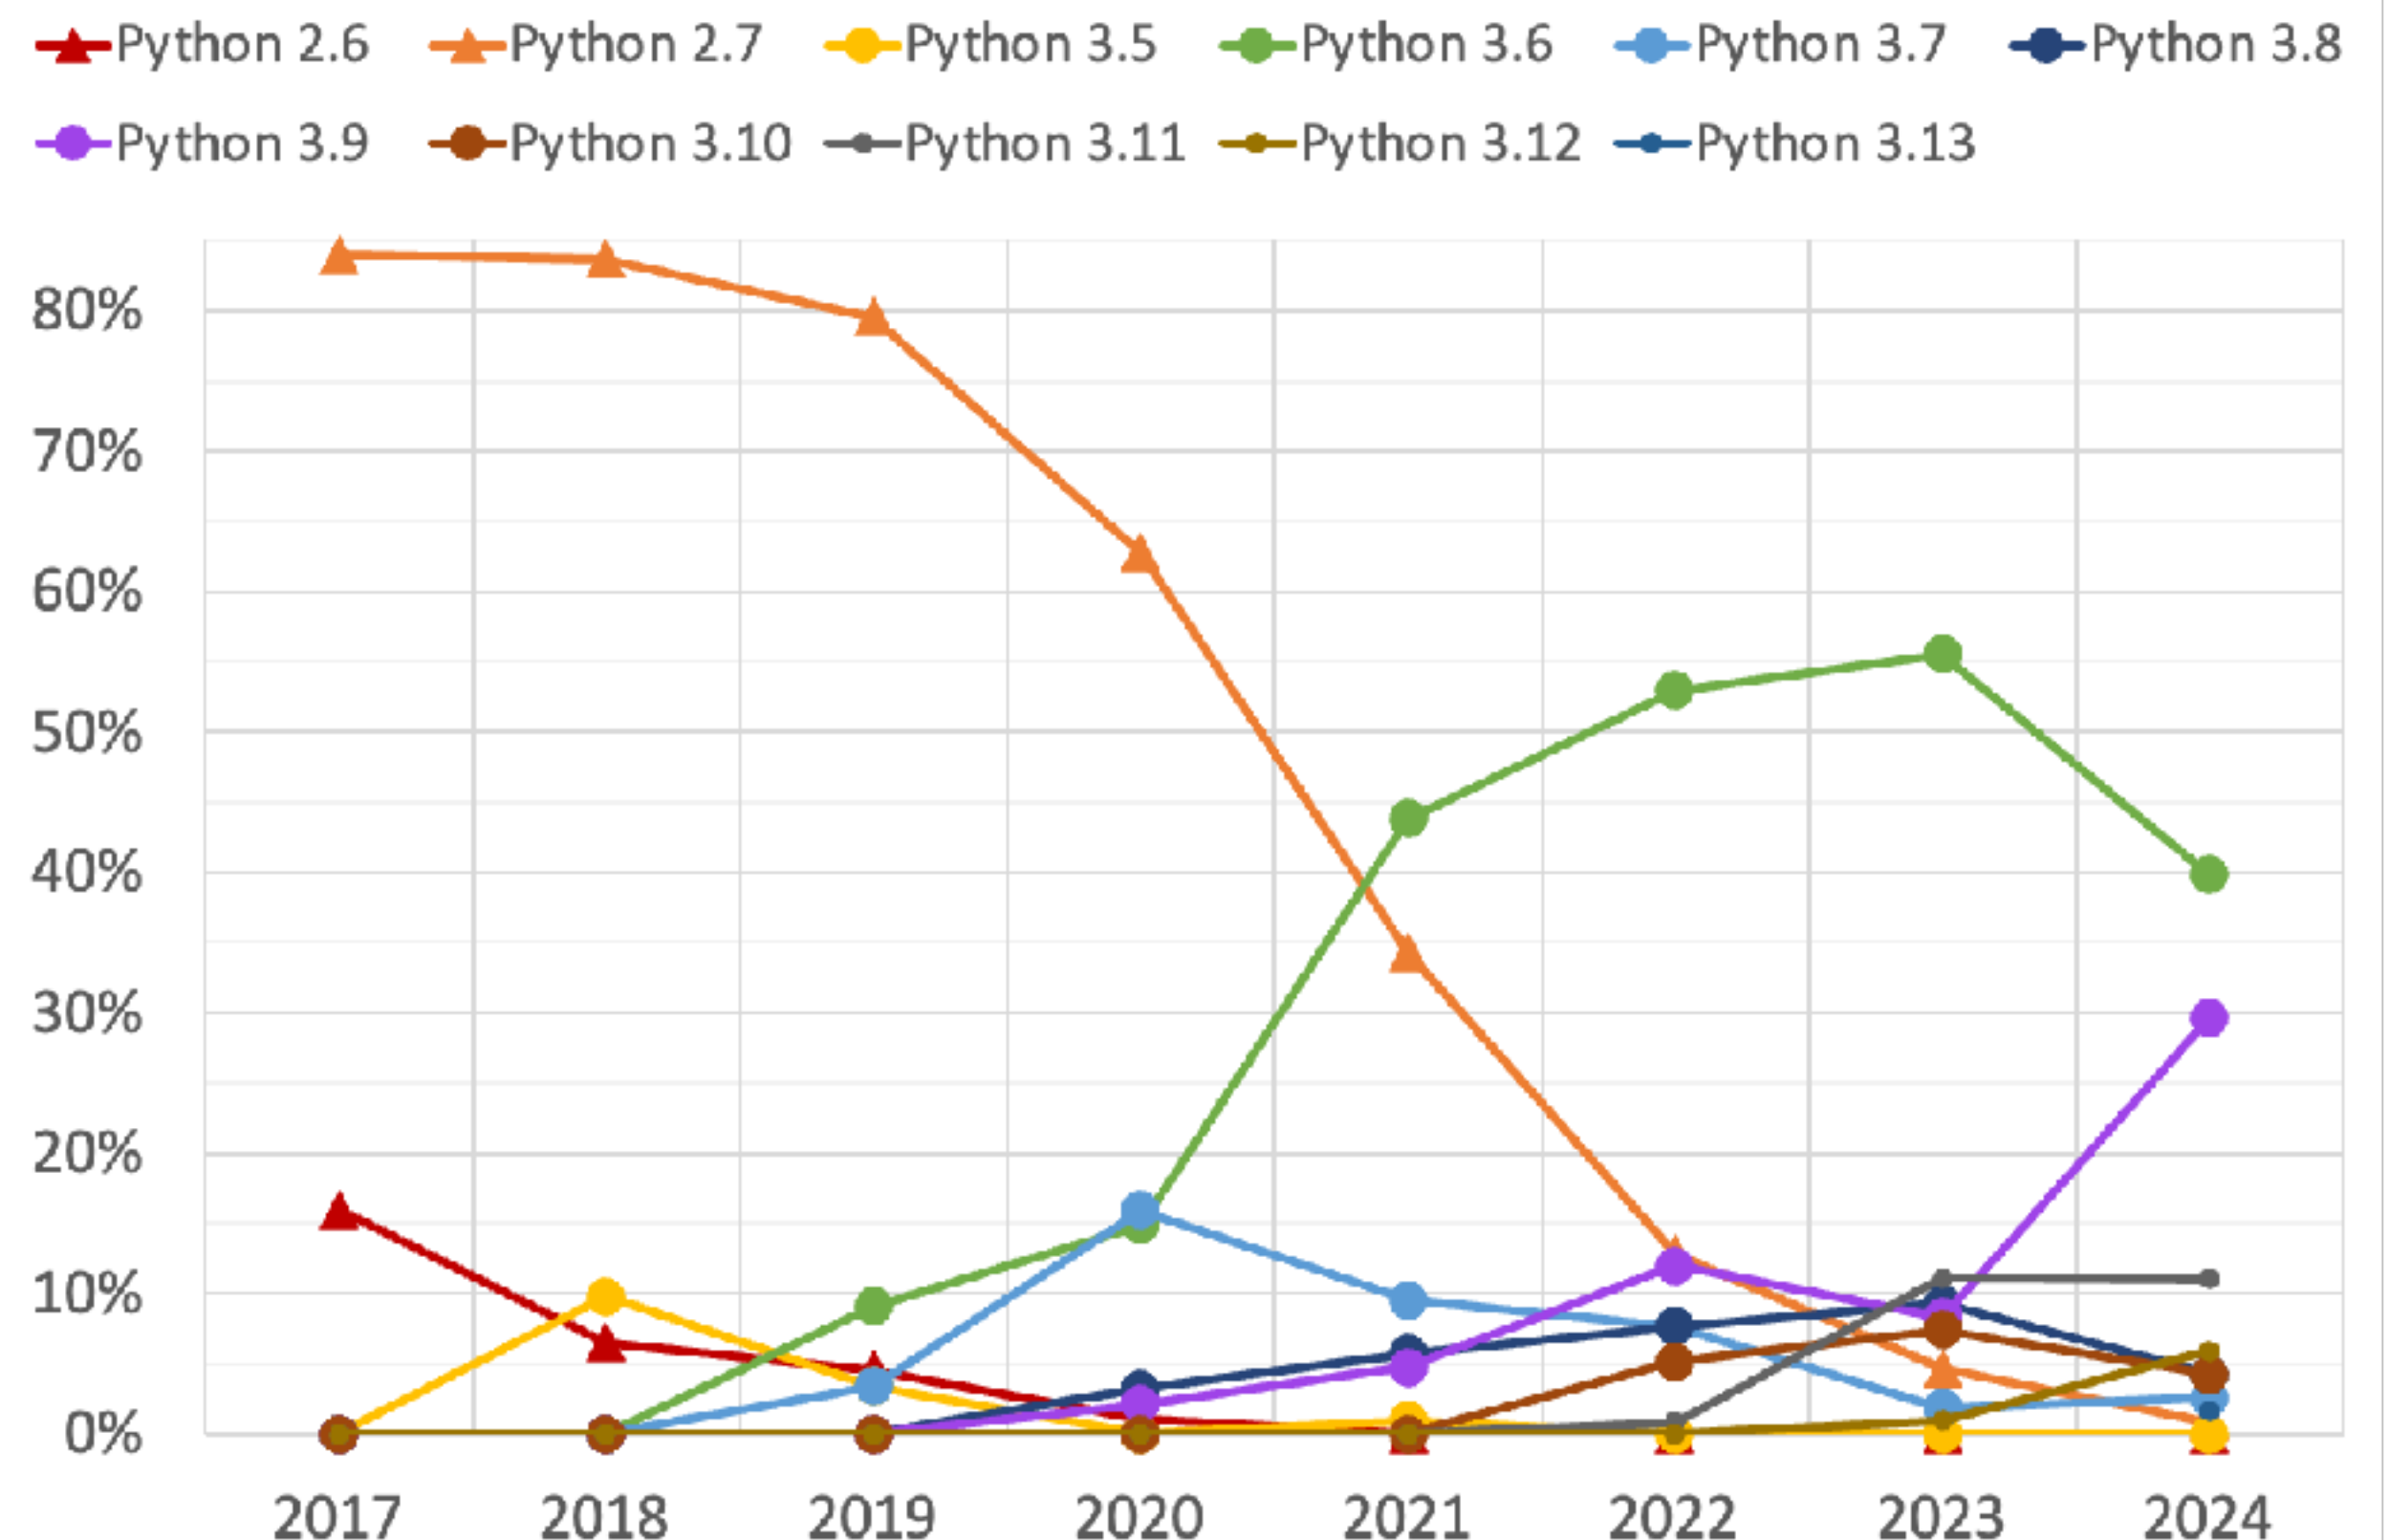


*Note: since previous survey (2023 results), you could only pick *one* option*

Operating systems in EasyBuild community



Non-RHEL based OSs are a minority in the EasyBuild community, only Ubuntu (barely) over > 5%

Python versions used to run EasyBuild



Python 2.7 (EOL since 2020!) is finally dead... No longer supported in EasyBuild v5.0.0 !

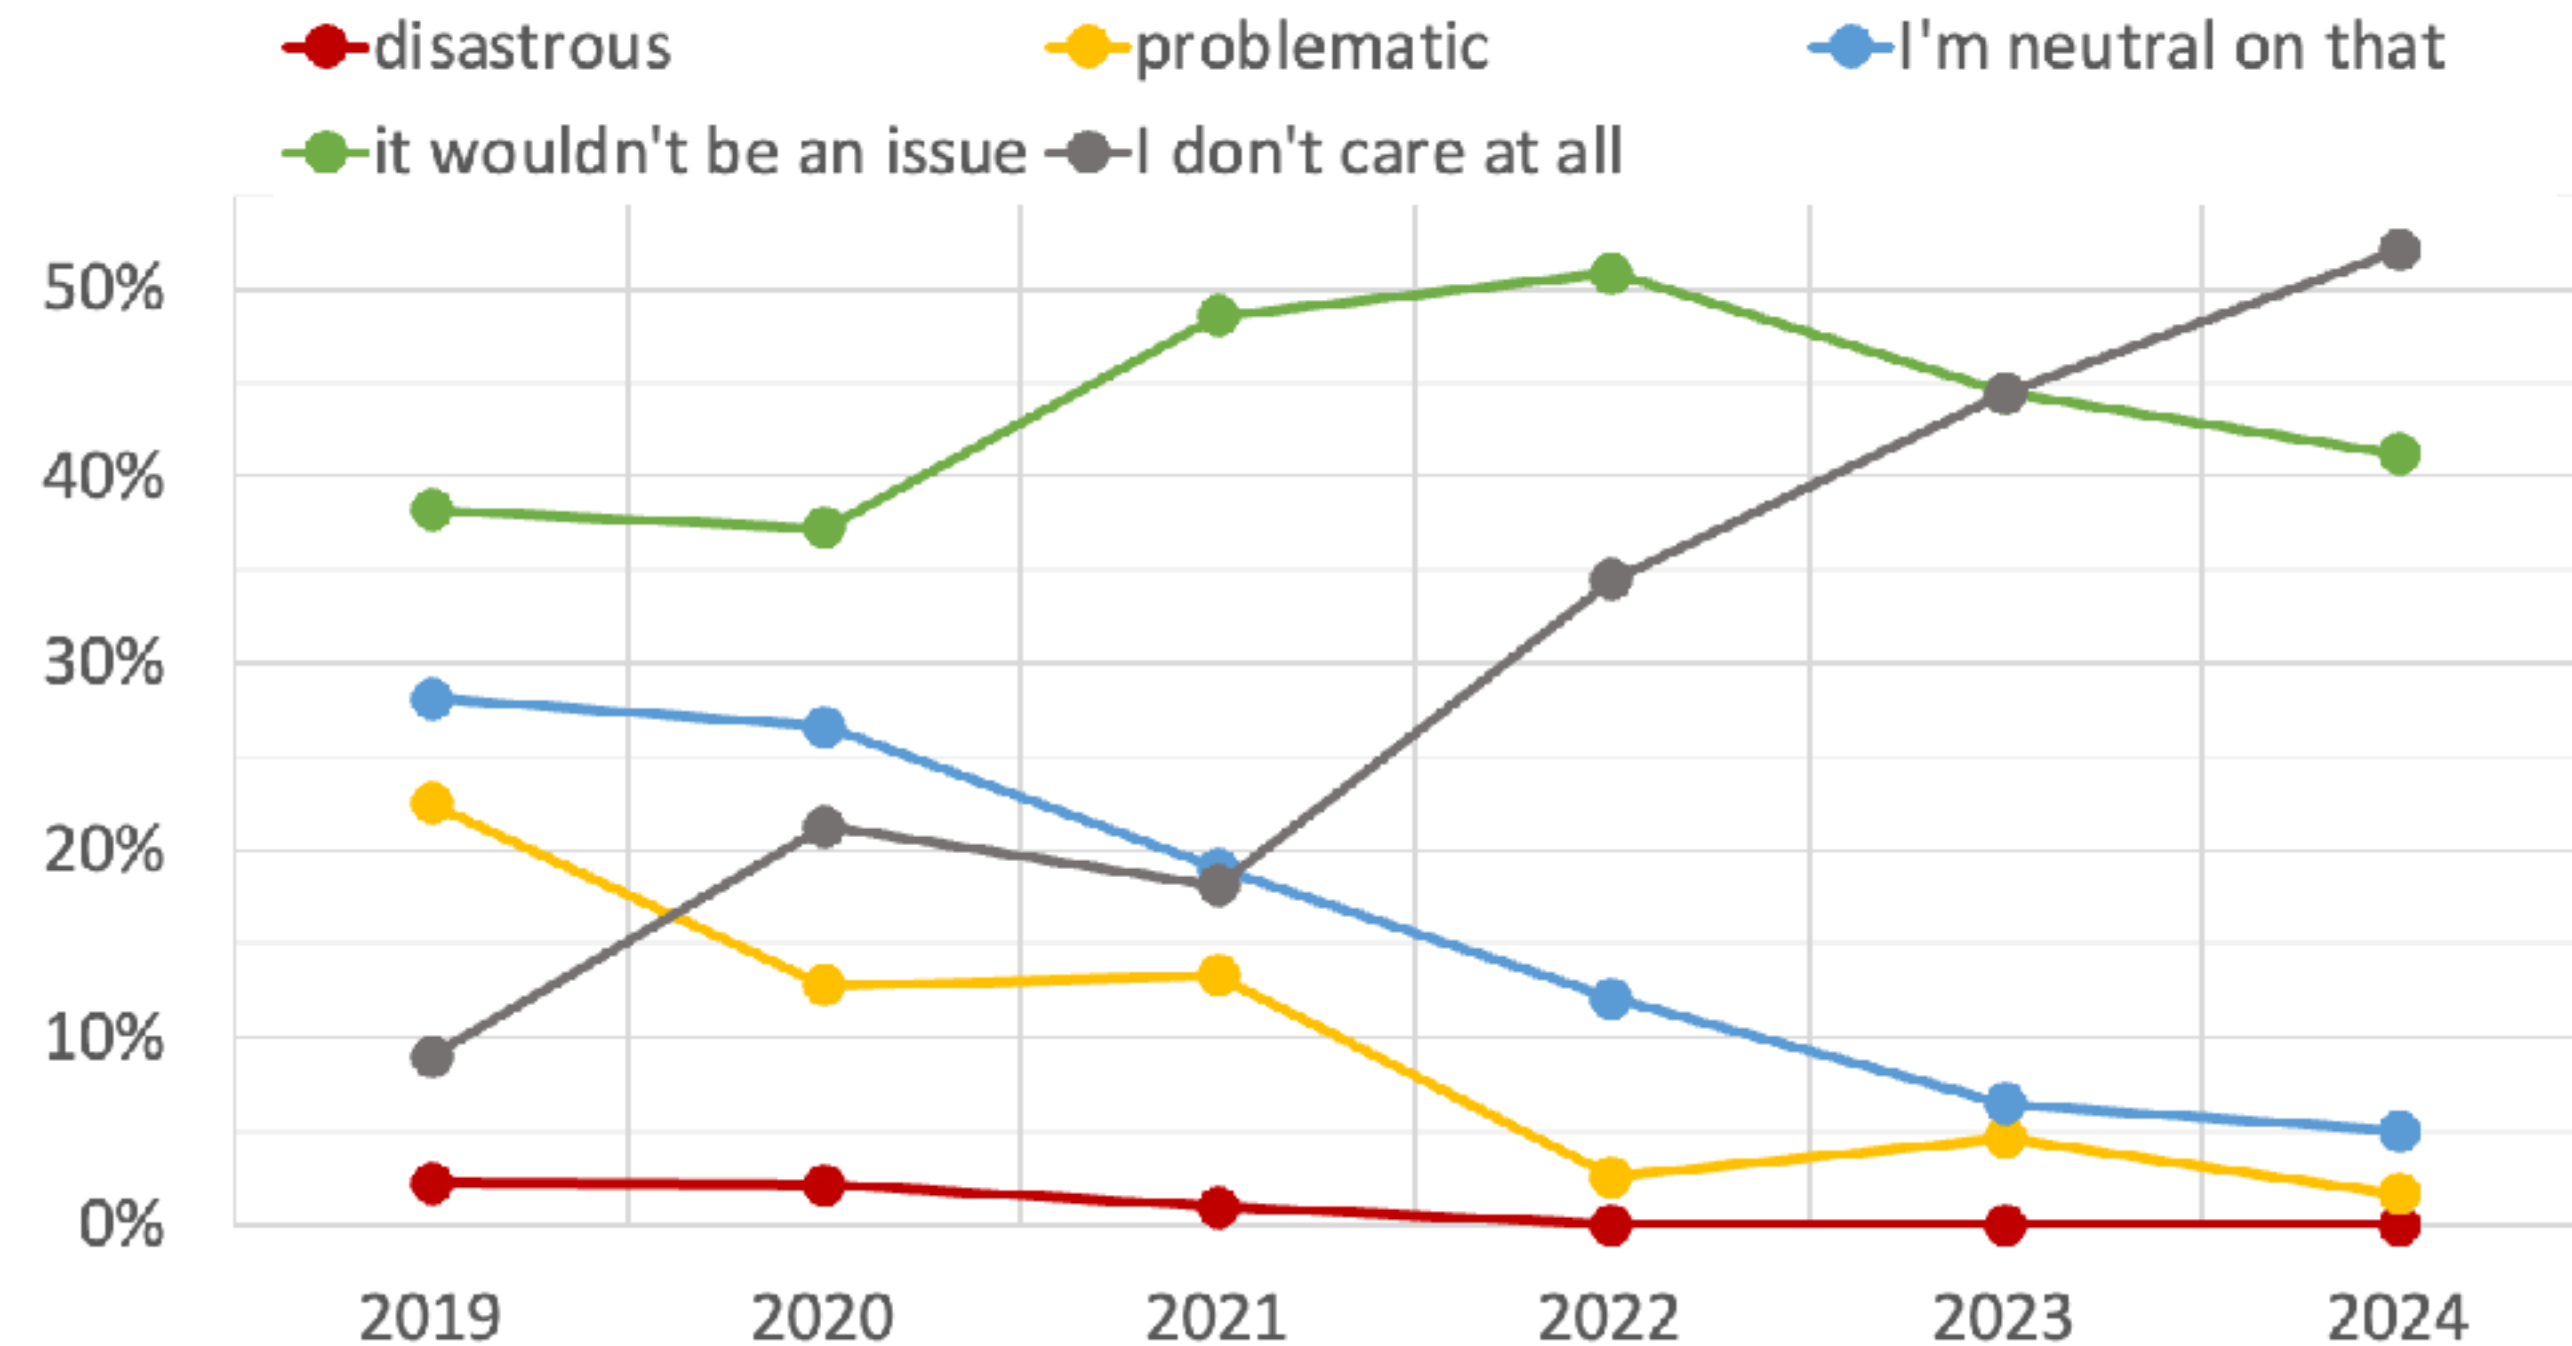
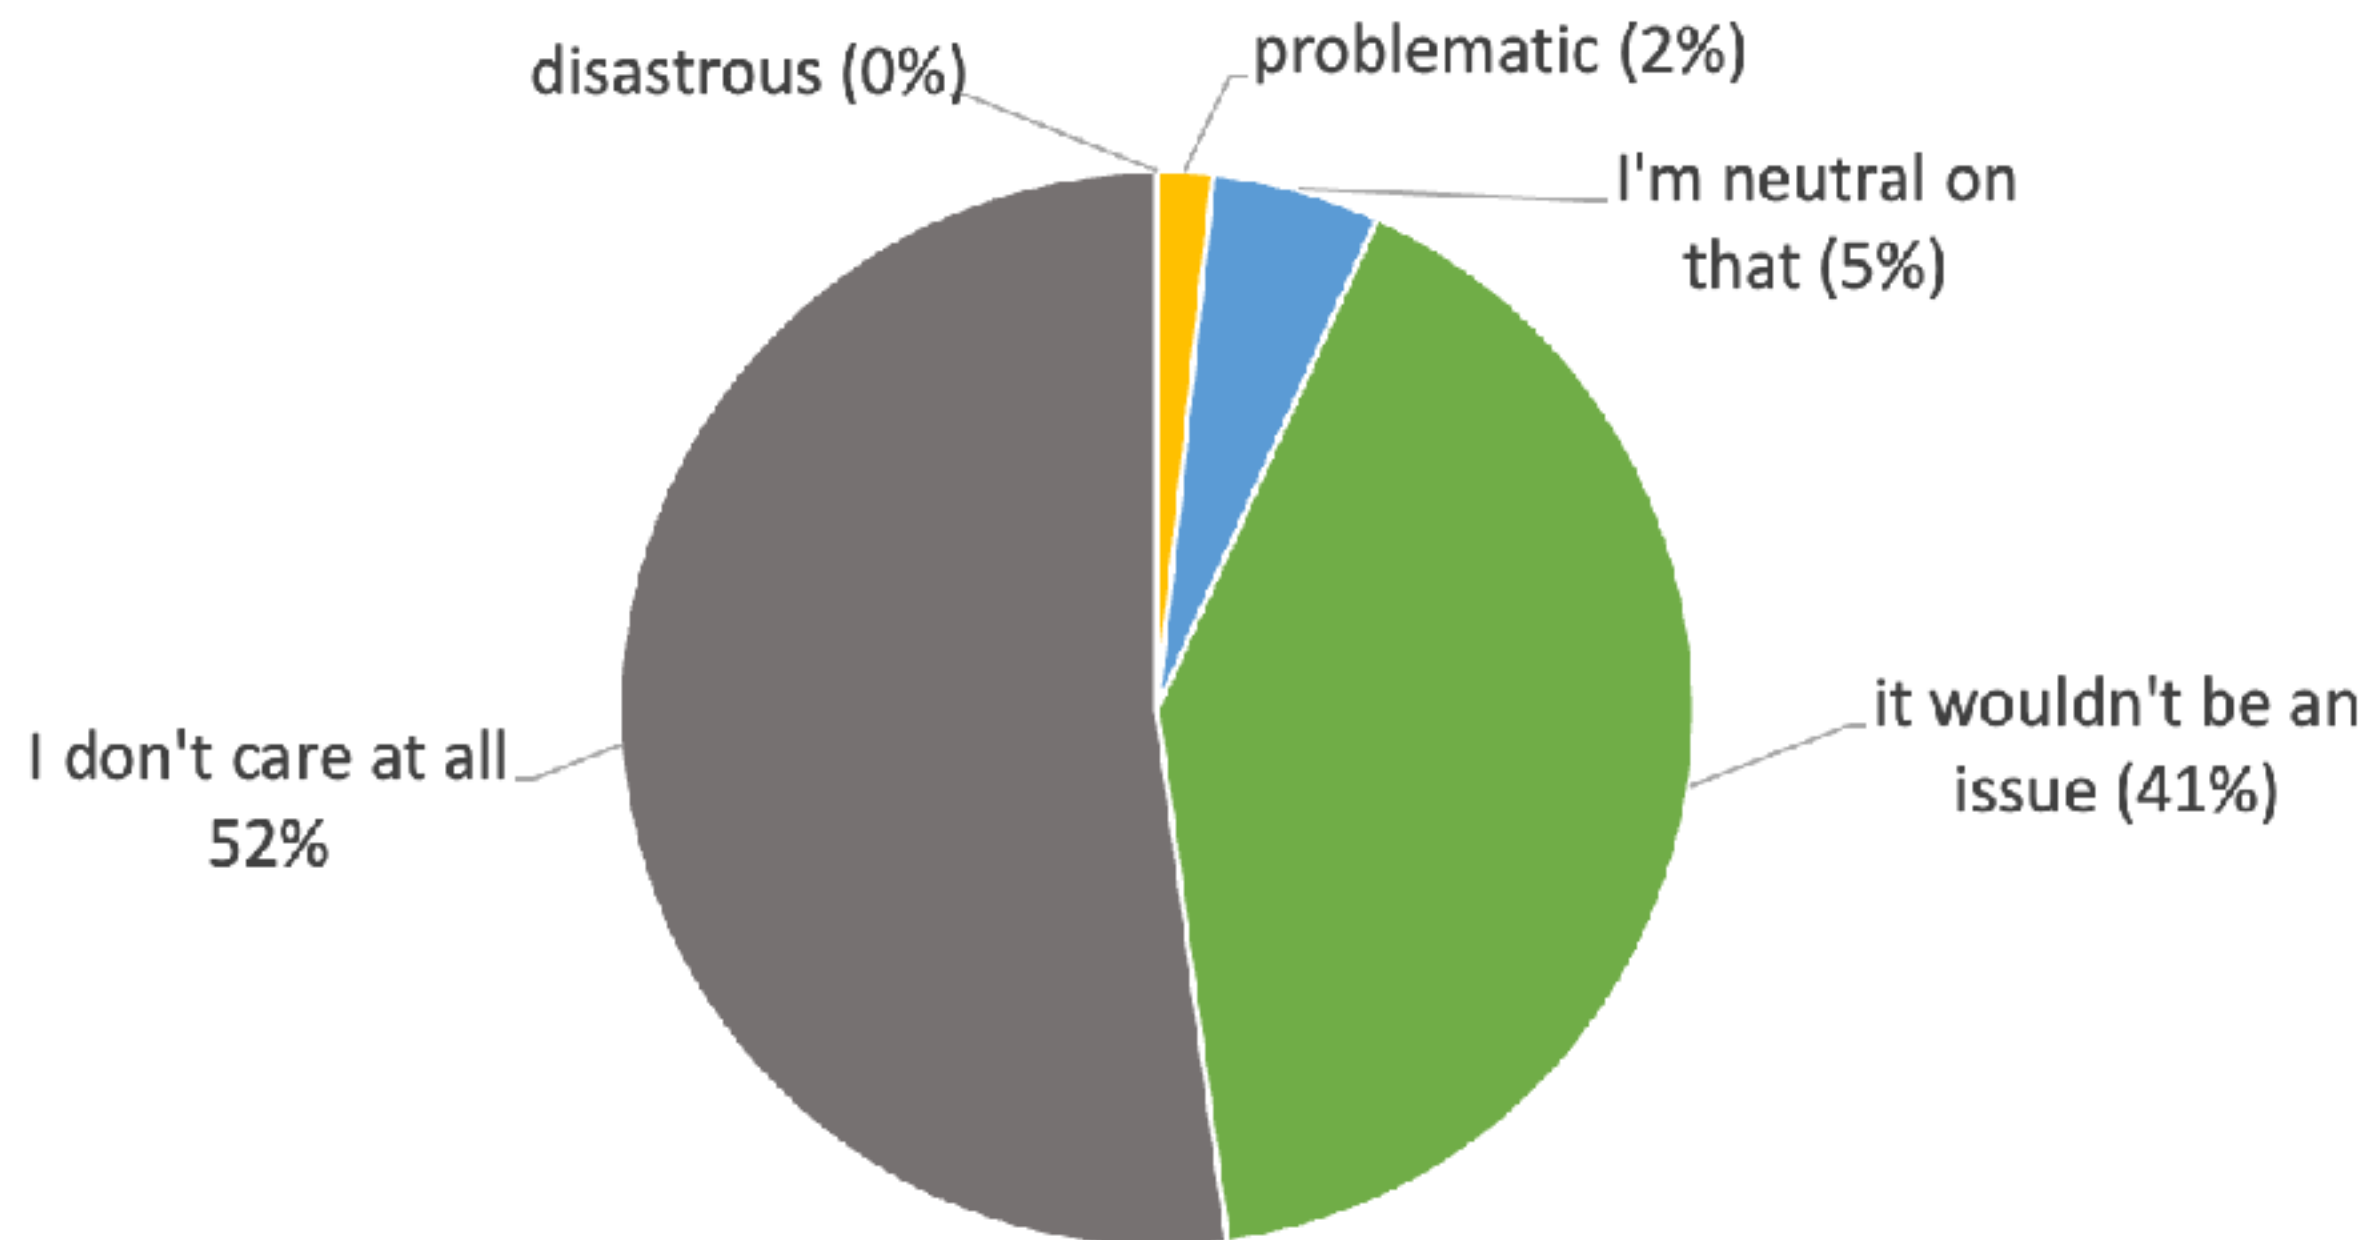
Keeping support for using Python 3.6 (default Python in RHEL8-based OSs) is important, used by ~40%

Python 3.9 on the rise (~30%), splintering across various Python 3 versions is a challenge going forward...



Community is fine with getting rid of Python 2.7

How troublesome would it be for you if EasyBuild becomes incompatible with Python 2.7?




Support for running on top of Python 2.7 was deprecated since EasyBuild v4.7.0 (Jan'23),
and has been **removed** in EasyBuild v5.0.0

98% of EasyBuild community is totally fine with that (rest will need to adapt)

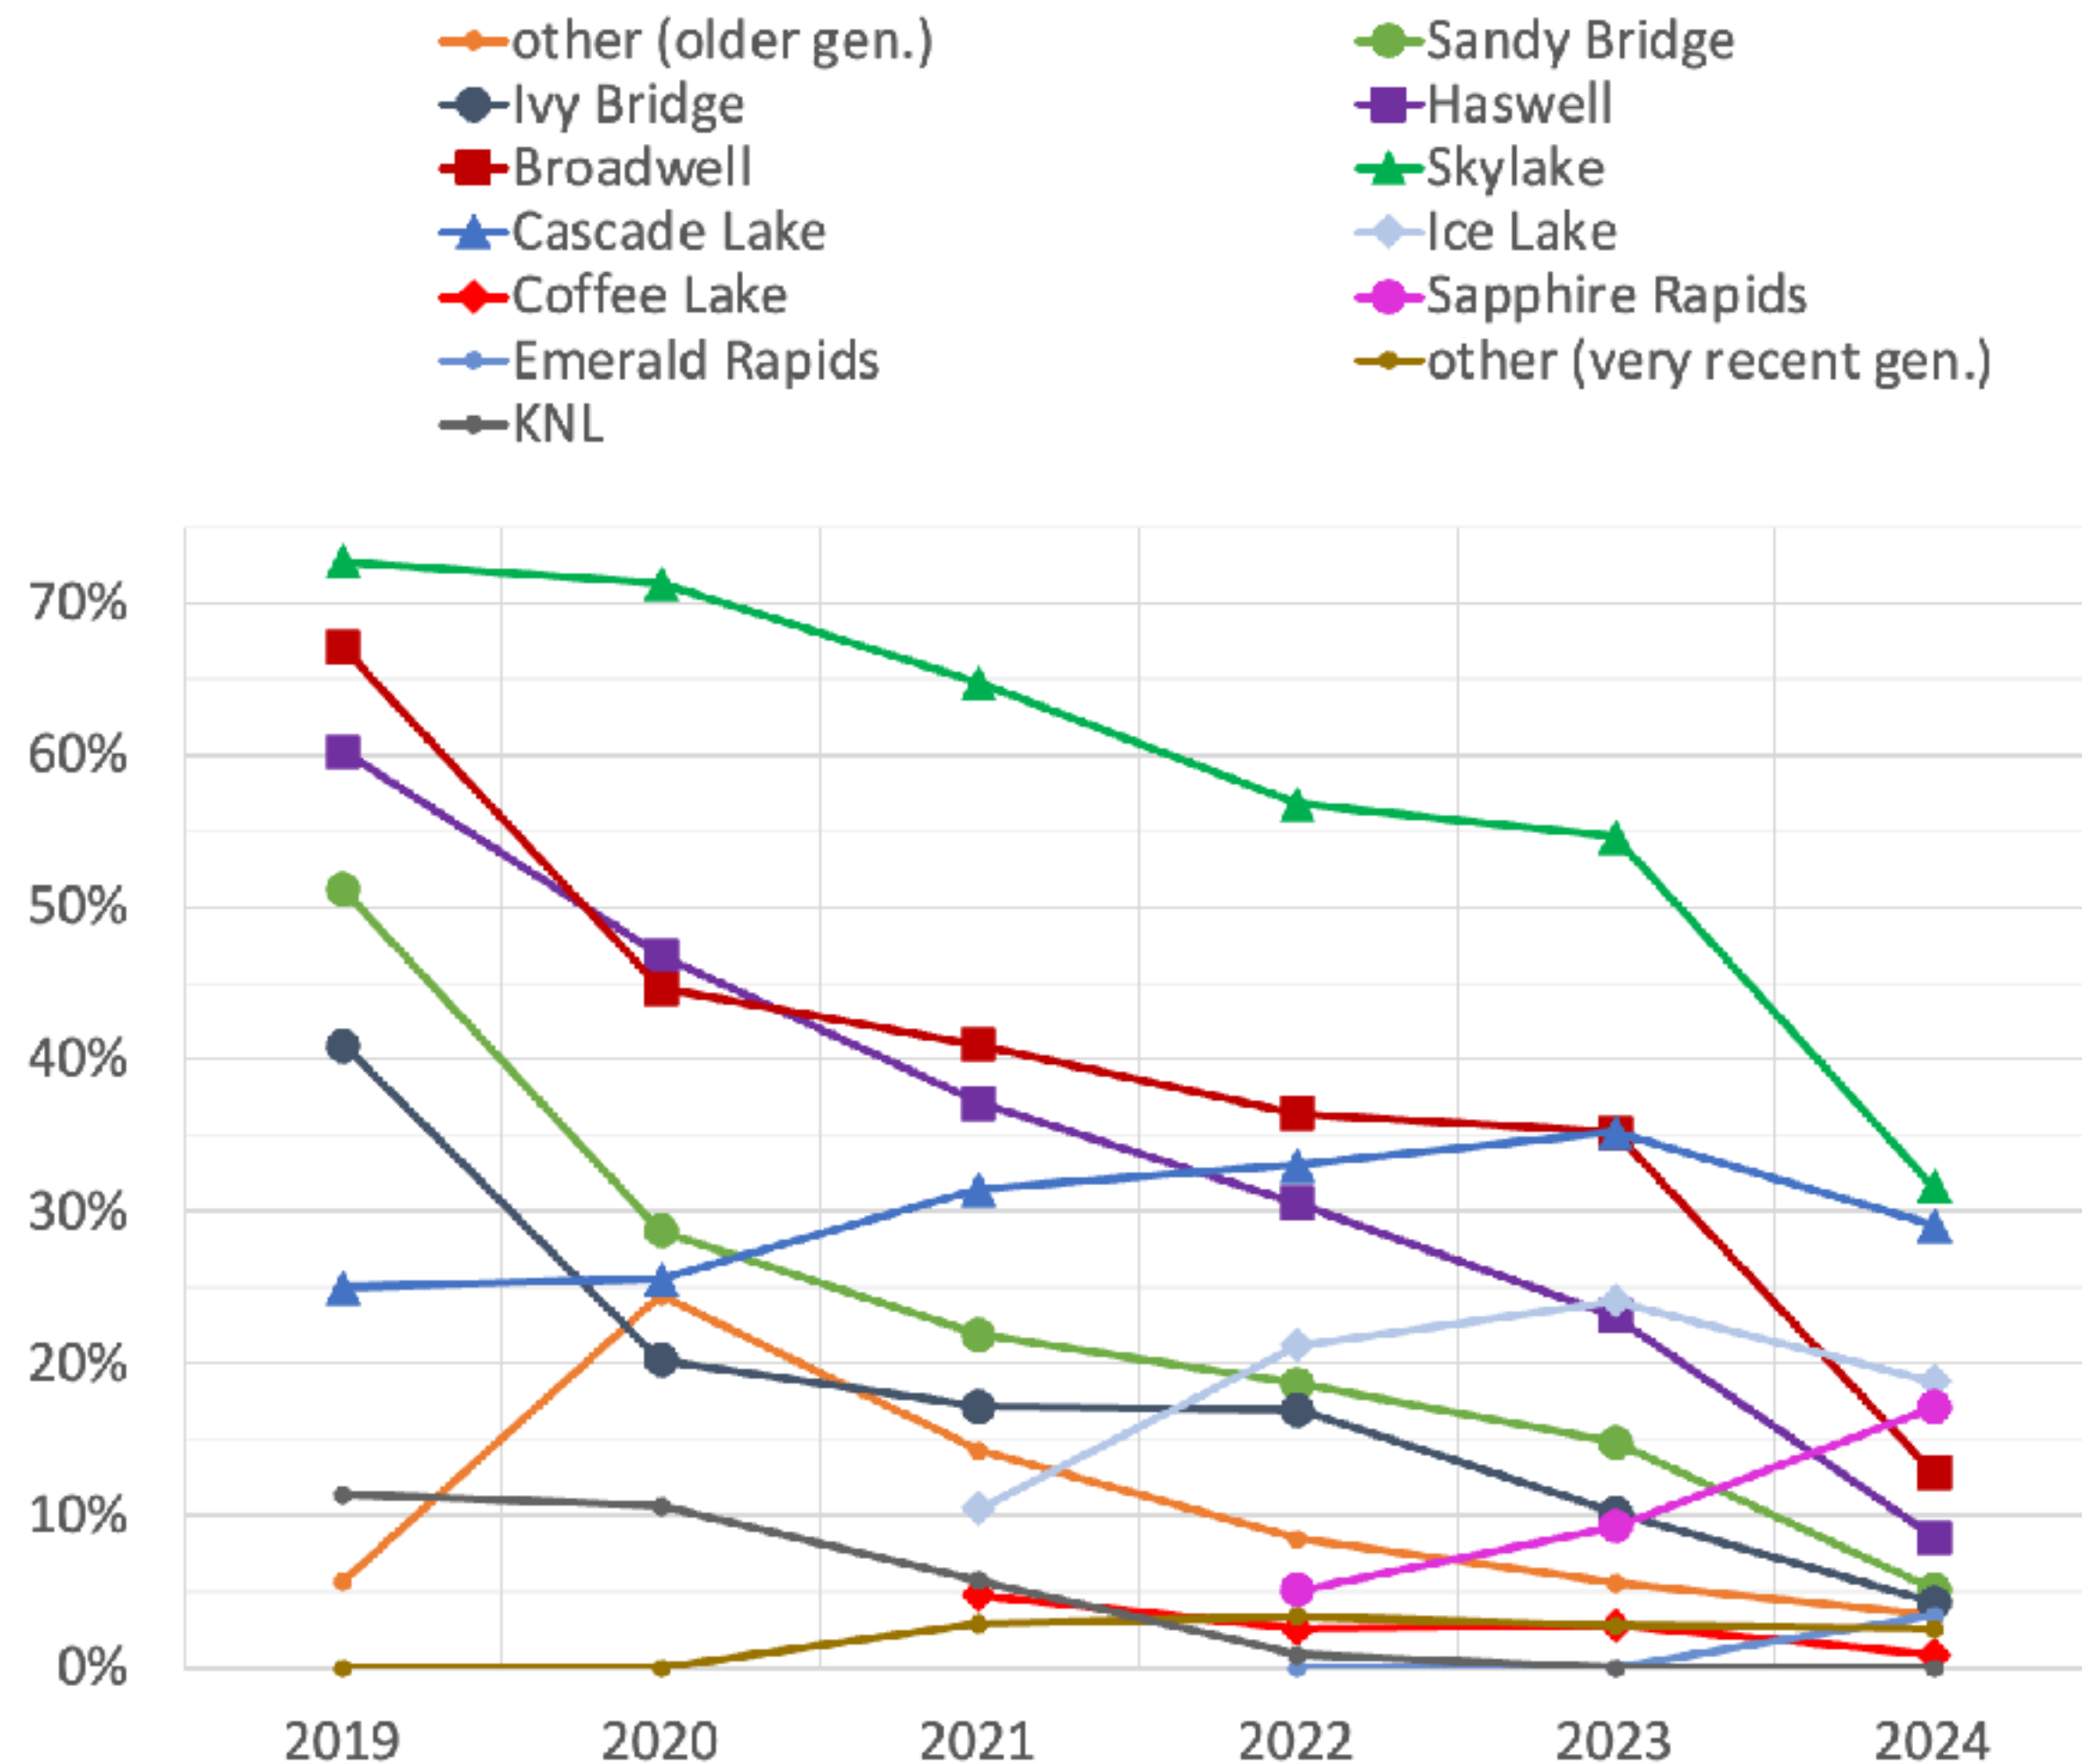


Processor families (Intel)




For which processor families are you installing software using EasyBuild?

*For which CPU types are you installing software using EasyBuild?
(only Intel CPUs)*



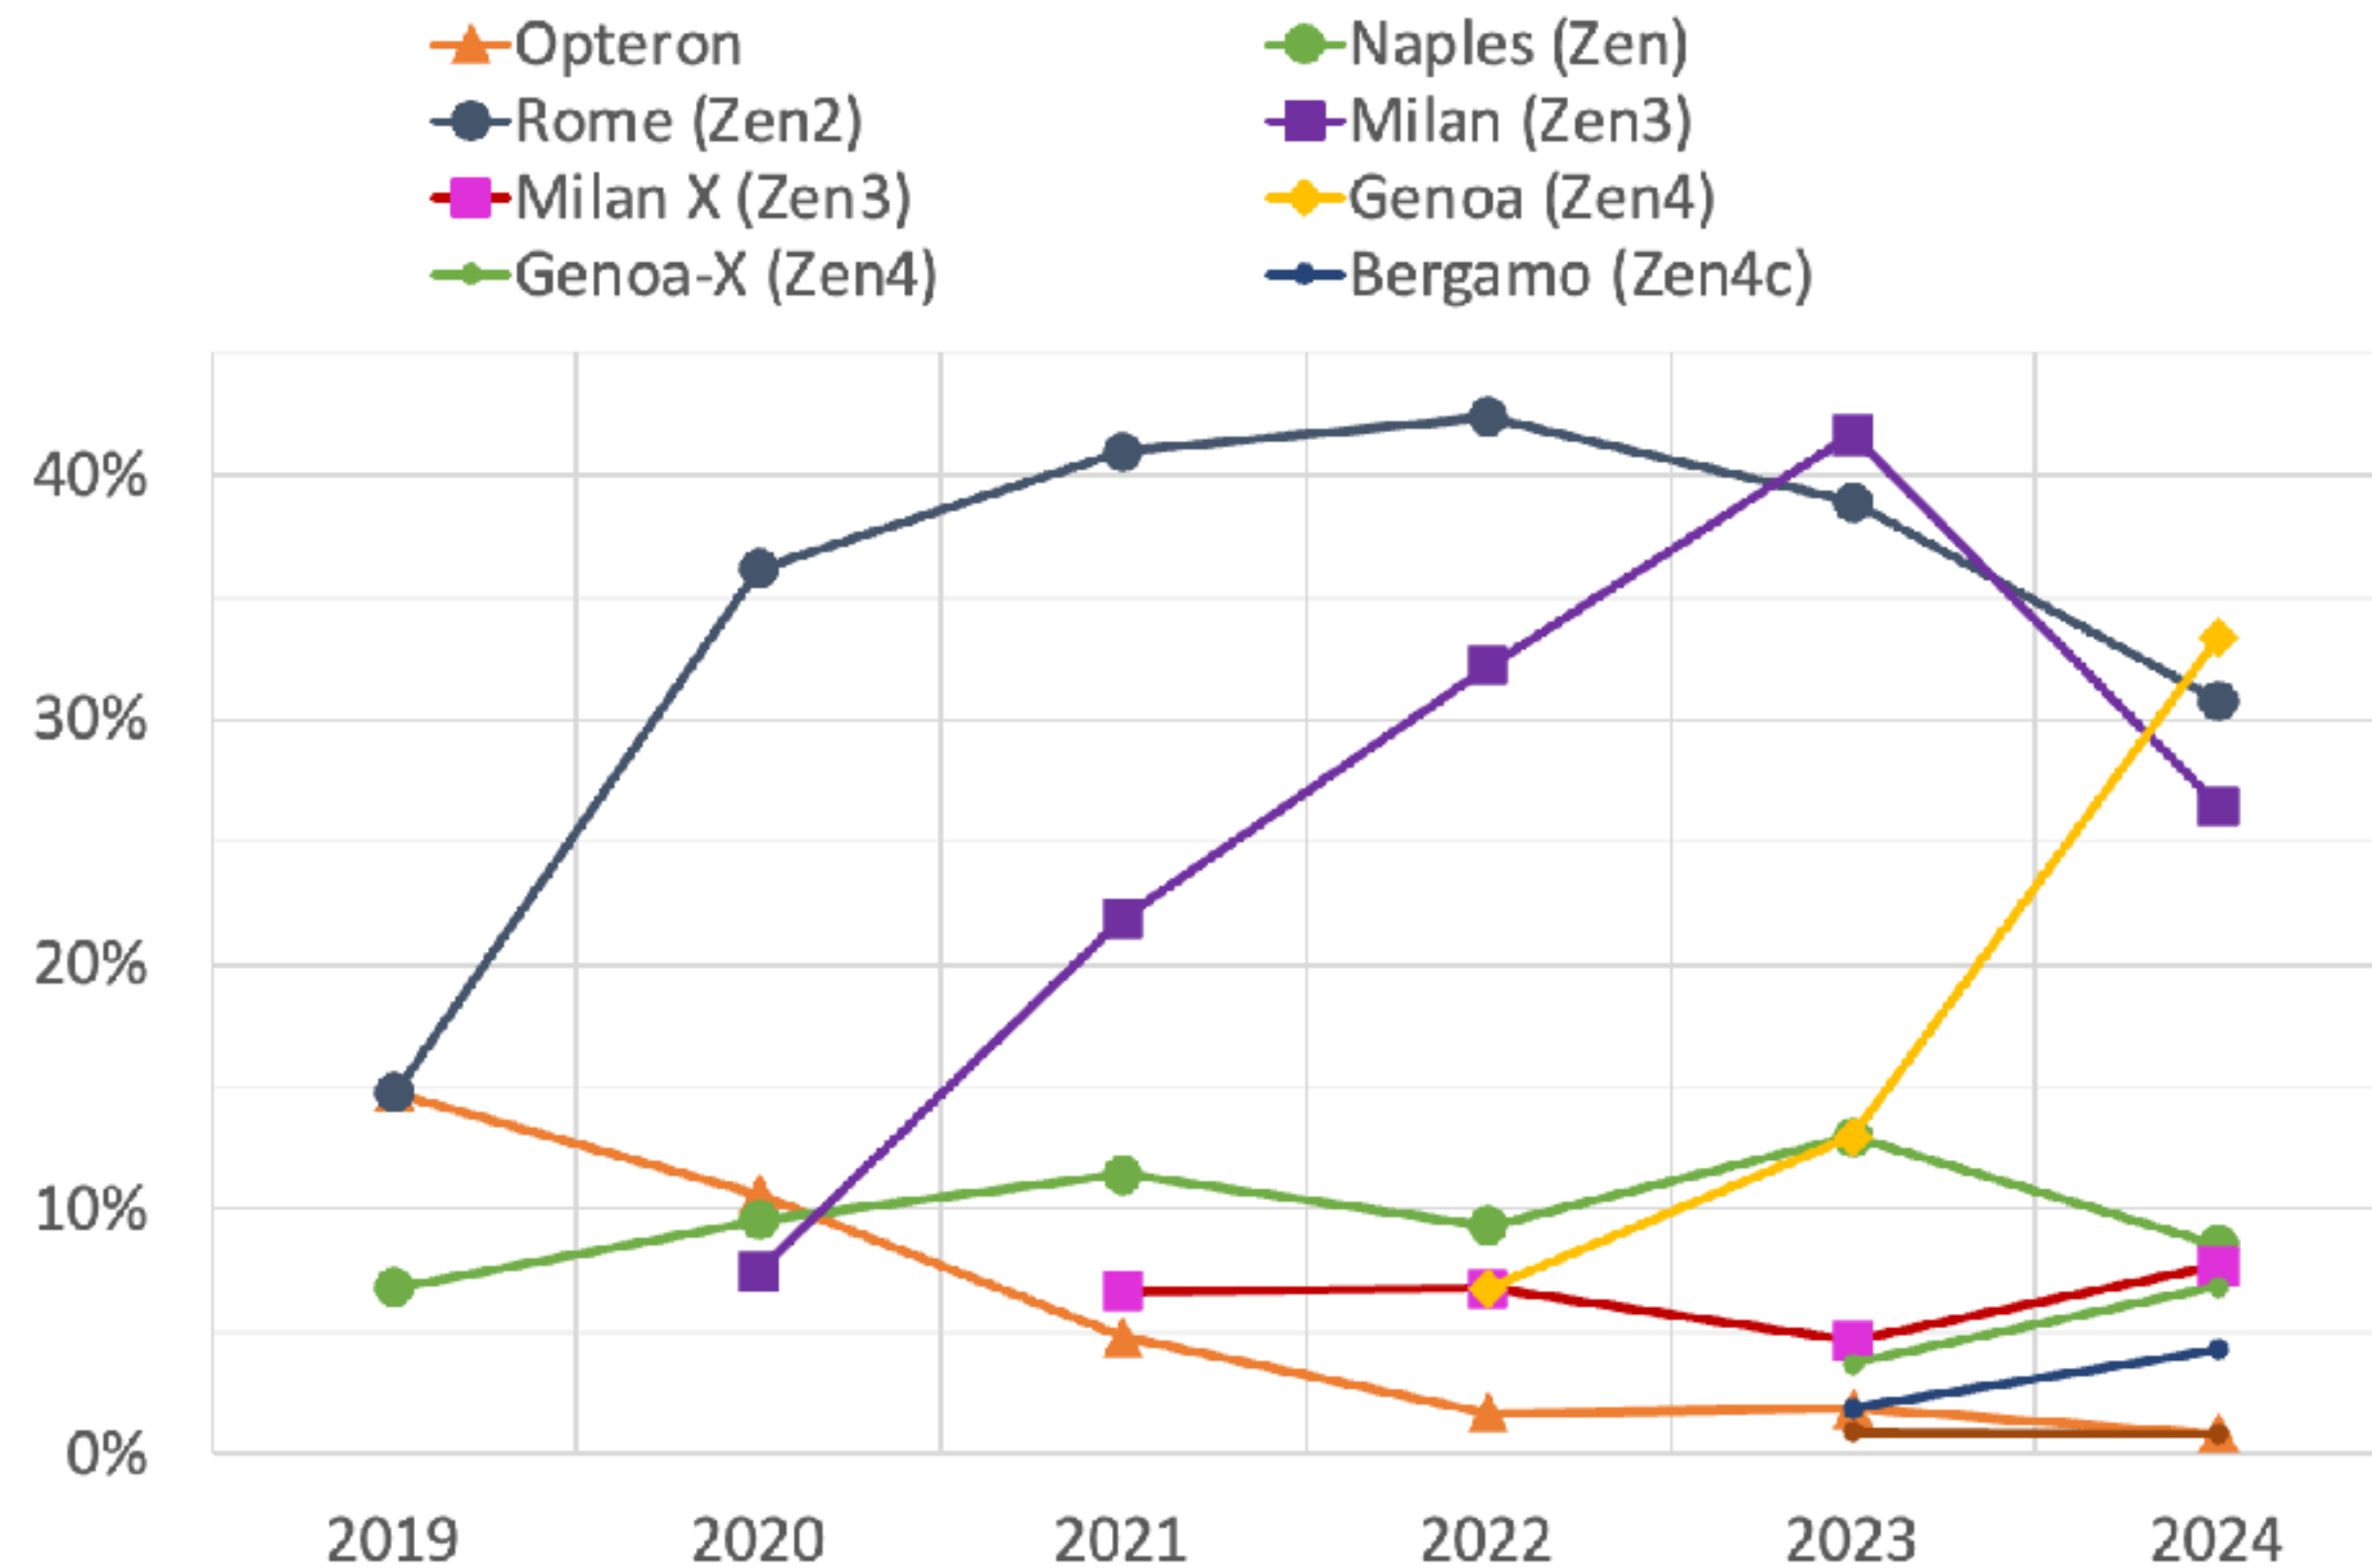
- Skylake is *still* most popular (Intel) CPU at 32%, but only barely...
- Cascade Lake (29%), Ice Lake (19%)
- Sapphire Rapids is new kid on the block (now ~17%)
- Haswell dropped to 8.5%, Broadwell tanked to ~13%
- Ivy Bridge, Sandy Bridge: **still** at 5%...

Processor families (AMD)




For which processor families are you installing software using EasyBuild?

*For which CPU types are you installing software using EasyBuild?
(only AMD CPUs)*



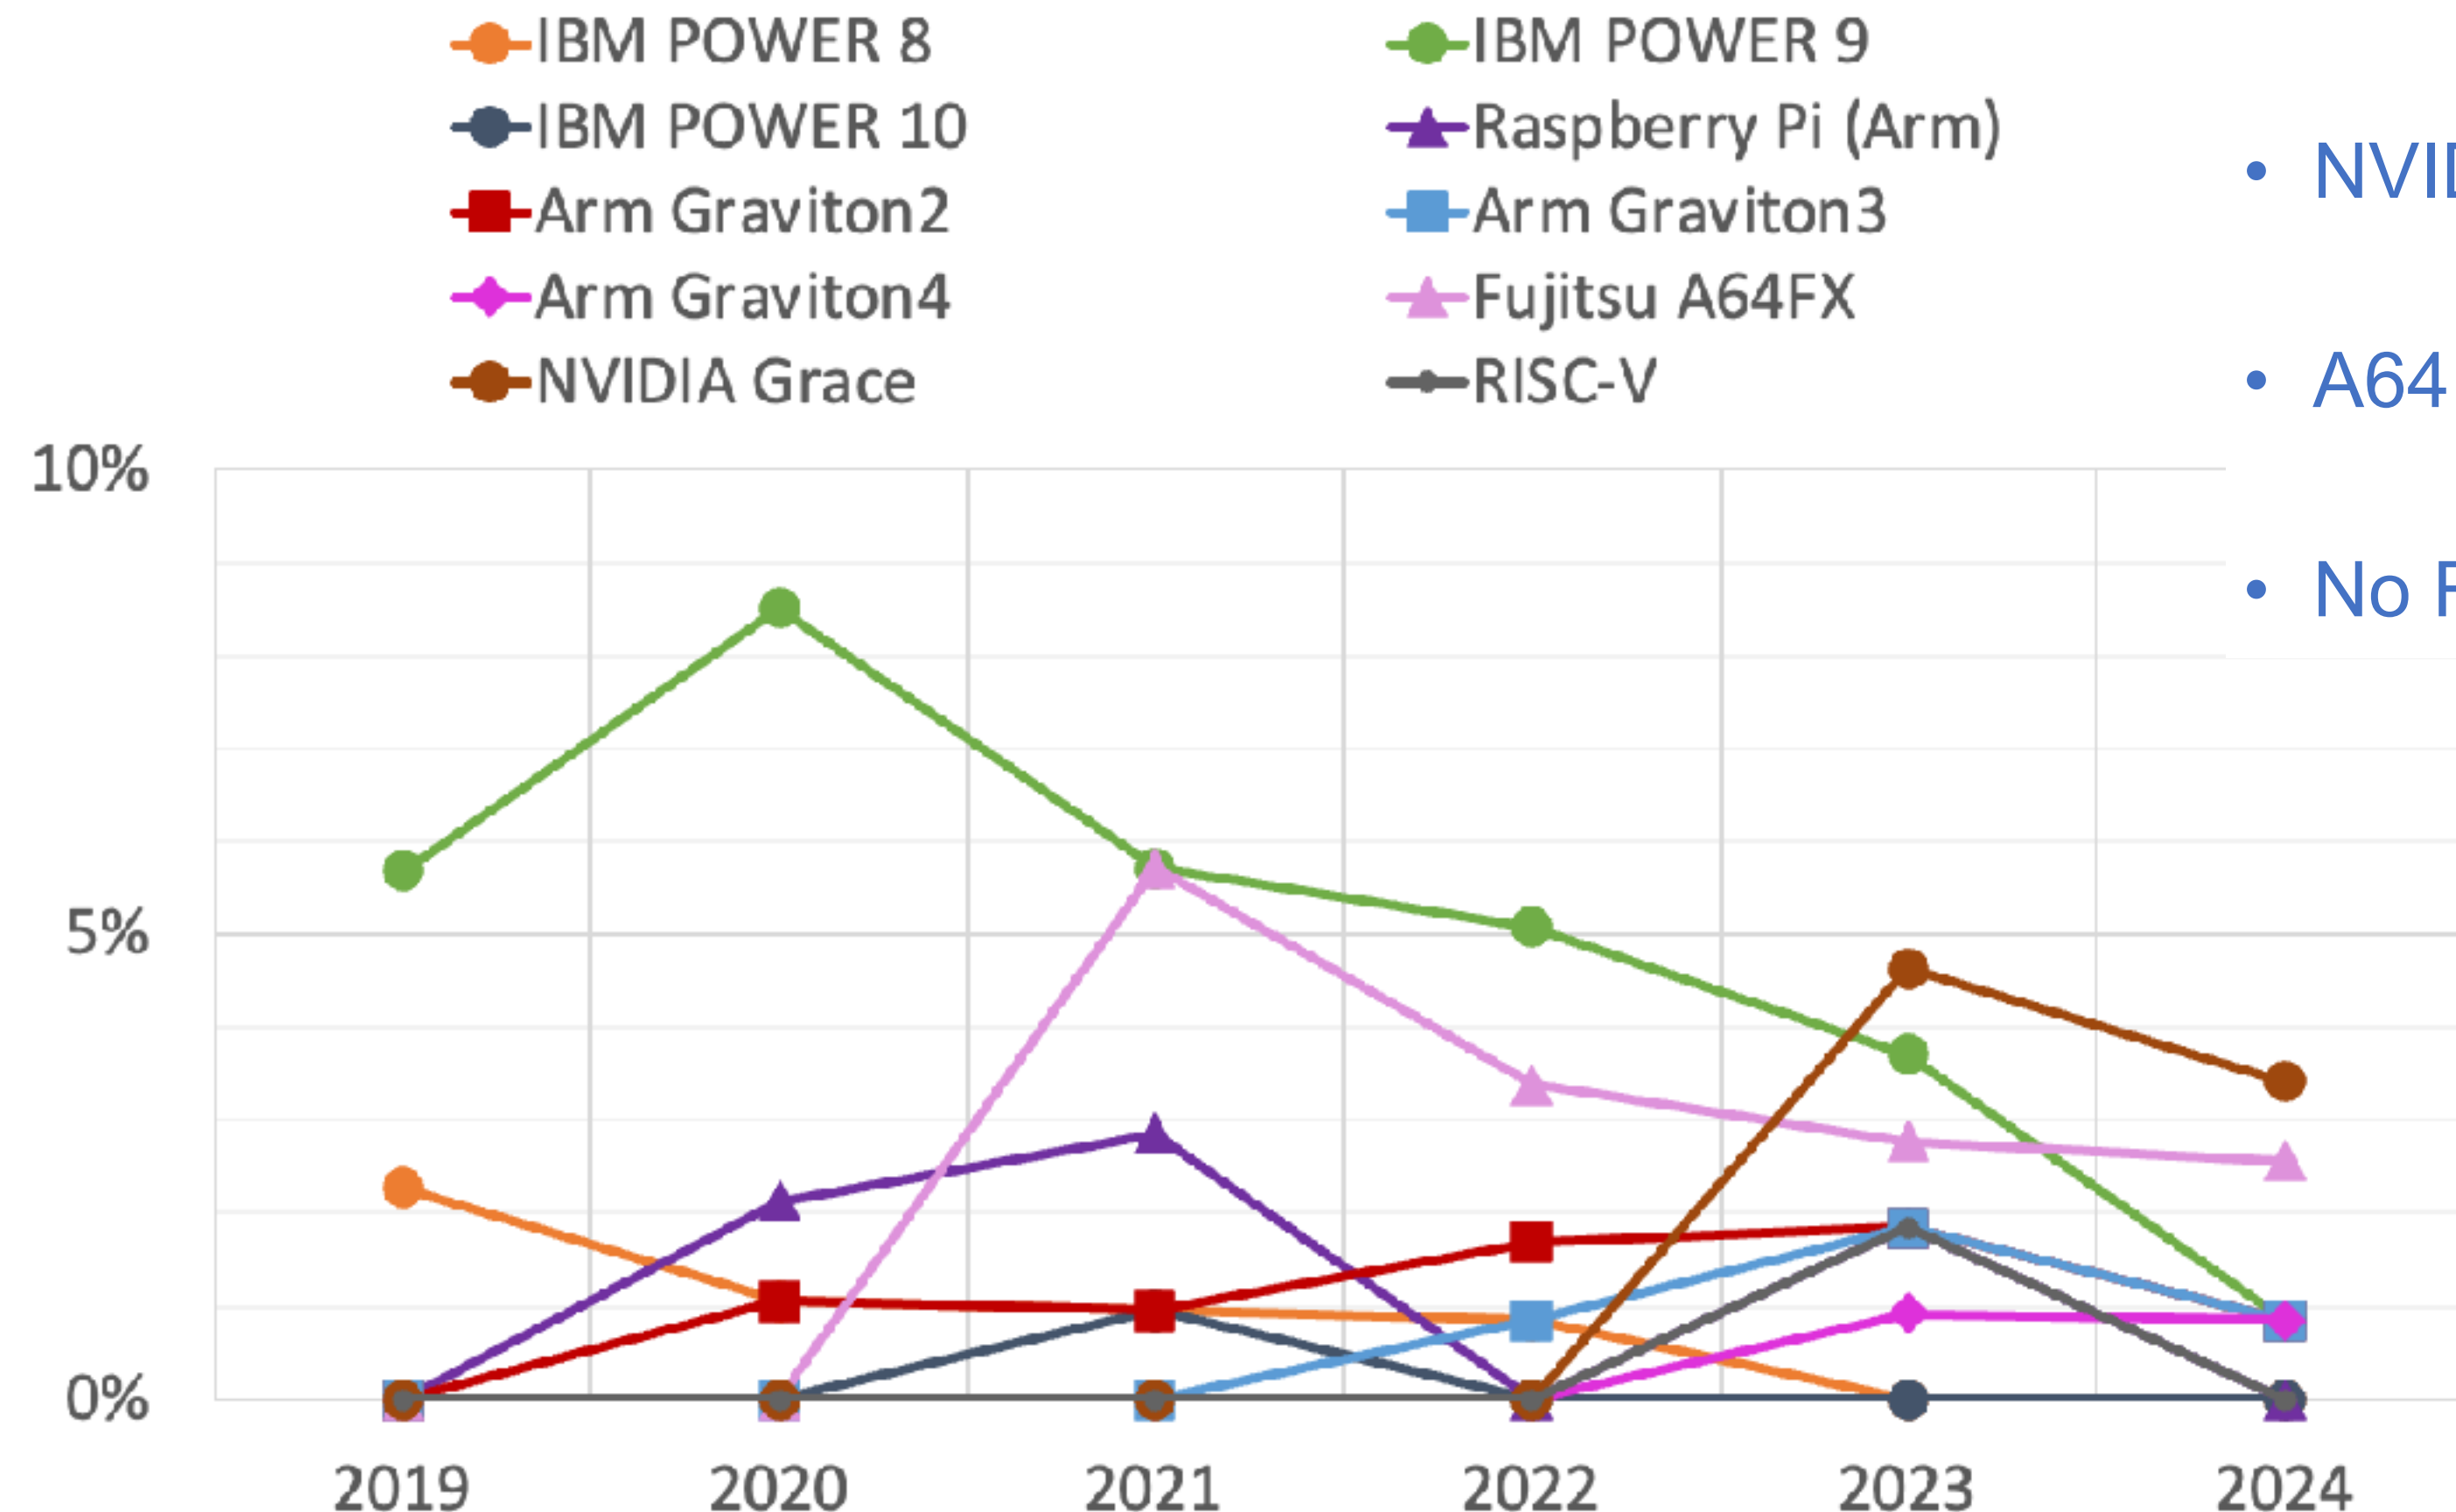
- 3 generations of AMD CPUs doing well at ~30%
- **Genoa (Zen4) is most popular at ~33%** (> Skylake), especially when also counting Genoa X (~7%)
- Decrease for Rome (Zen2): 39% -> 30%, similar for Milan (~27%), small increase for Milan-X (~8%)
- Naples (Zen) melting away to ~8.5%

Processor families (non-x86_64)




For which processor families are you installing software using EasyBuild?

*For which CPU types are you installing software using EasyBuild?
(excl. Intel + AMD CPUs)*



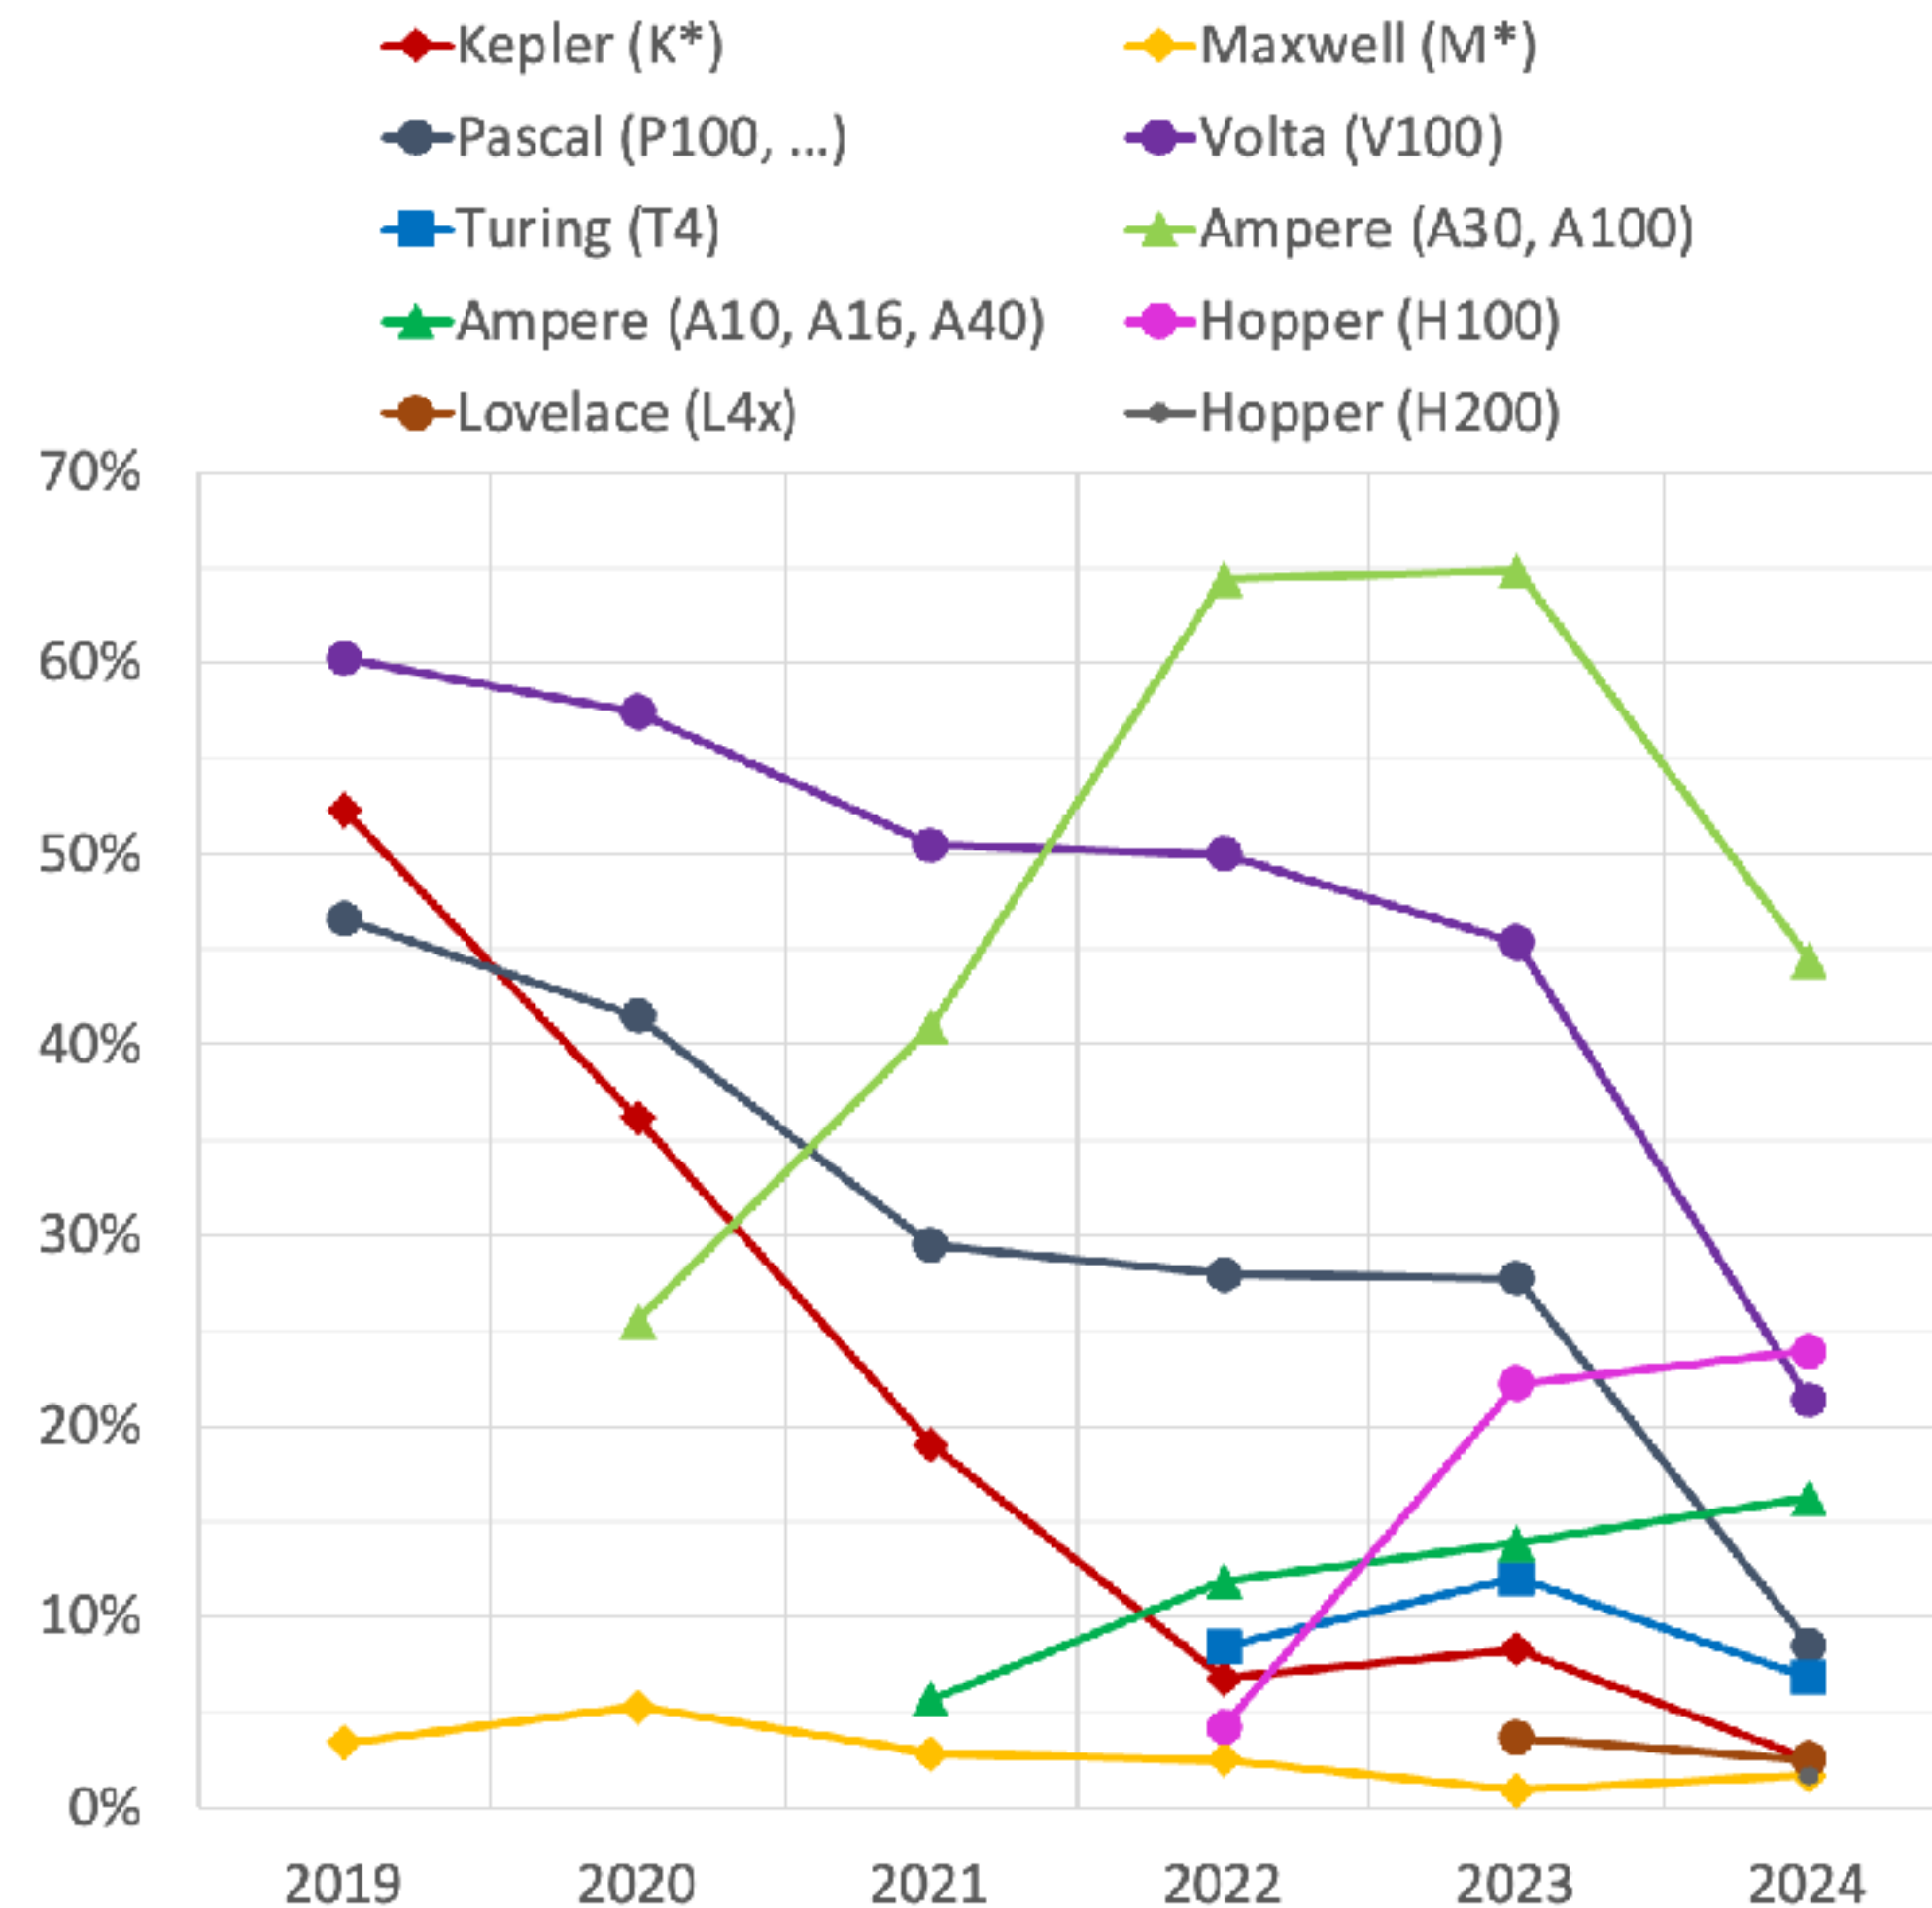
- Limited adoption of other CPU families beyond x86_64
- NVIDIA Grace most popular, but still small minority < 5%
- A64FX as runner-up (mostly due to Deucalion?)
- No RISC-V (yet), but coming soon?

Accelerators (GPUs)




For which accelerators are you installing software using EasyBuild?

*For which accelerators are you installing software using EasyBuild?
(NVIDIA Tesla only)*



- Ampere (A100, A30) still most popular GPU family at ~44%
- Small increase for H100 at 24%
- V100 in freefall, down to 21%
- P100 barely hanging on at 8.5%
- **AMD MI210 GPUs at ~7.5% (not shown in plot)**


EasyBuild tutorials




<https://tutorial.easybuild.io>

- Last EasyBuild tutorial was at ISC'22 (half-day), submission for ISC'23 was not accepted :(
- For ISC'24/SC'24/ISC'25 we focused on EESSI (+ CernVM-FS), but no success either...
- Most recent (extensive) tutorial: EasyBuild + EESSI UK workshop: 27+28 April 2023 (right after EUM'23)
 - **Slides + recording available via <https://tutorial.easybuild.io/2023-eb-eessi-uk-workshop>**
- Introductory tutorial on EasyBuild through EPICURE on 13 Dec'24 was recorded:
<https://epicure-hpc.eu/2024/11/29/webinar-building-software-with-ease-an-introduction-to-easybuild>
- A series of online webinars in May 2025 is being planned in context of EESSI
 - Will cover CernVM-FS + various aspects of EESSI, and also EasyBuild
- Time for another (extensive) online tutorial on EasyBuild now that EasyBuild v5.0.0 is out?

Supported software




Evolution of number of supported software packages in EasyBuild (without extensions)





Close to 2,835 different software projects supported by EasyBuild v5.0.0 (plus 2,000+ extensions)

Significant reduction in EasyBuild v5.0.0 due to archiving of easyconfigs using unsupported toolchain...

EasyBuild mailing list




Are you subscribed to the EasyBuild mailing list?




- Active participation remains quite limited, just ~14%
- Interest in mailing list is slowly declining
- Mailing list remains an important communication channel!

Subscribe to EasyBuild mailing list via
<https://lists.ugent.be/wws/subscribe/easybuild>

EasyBuild mailing list (traffic)




Traffic on EasyBuild mailing list peaked in 2017,
due to start of EasyBuild Slack channel, but remains active.





340 subscribers to EasyBuild mailing list currently, small uptick in activity in 2023 was temporary?

EasyBuild Slack



Do you use the EasyBuild Slack channel?




Only ~36% of community *actively* participates in Slack.

Active & passive participation is slowly increasing.

Small part of community (~14%) is not on Slack.


Join EasyBuild Slack via
<https://easybuild.io/join-slack>

EasyBuild Slack




- Steady growth in activity since start of EasyBuild Slack in 2017
- 993 members on Slack channel, steady increase of 100-150 per year since 2020
- 100-200 people active per week, ~50 people posting per week
- Has quickly become the central hub for discussing EasyBuild & getting help

Join EasyBuild Slack via <https://easybuild.io/join-slack> !




Lmod + Environment Modules




Do you use any other tools/projects in combination with EasyBuild?

*Do you use any other tools/projects in combination with EasyBuild?
(only modules tools)*




- Lmod remains most popular modules tool (by far)
- ~82% is using Lmod 8.x
- ~13% is still using Lmod 7.x
- ~10% is still using Lmod 6.x
 - Partially due to outdated Lmod package in Debian/Ubuntu
- **Lmod 6 + 7 are no longer supported by EasyBuild 5.0 !**
- Tcl-based Environment Modules is small (growing) minority, version 5.x used by ~9%

Other tools & projects





Do you use any other tools/projects in combination with EasyBuild?




- Slurm (now default job backend) used by ~52%
- Apptainer is most popular 3rd party tool used, now at 55%
- Conda is used by over 36% (growing...)
- CernVM-FS is slowly increasing, now at 18%
- Small stable minority (~11%) uses both EasyBuild and Spack?
- Steady but somewhat limited: EESSI (17%), ReFrame (20%), Rich (22%), Open OnDemand (28%)

European Environment for Scientific Software Installations





- Awareness of EESSI (slowly) growing in EasyBuild community
- Slight increase in people who have hands-on experience with EESSI (31% ➡ 35%)
- Significant involvement of EasyBuild community, but does not replace EasyBuild!
- **More on EESSI on 3rd day of EUM'25**

EasyBuild releases



How do you like the frequency of EasyBuild releases?

How do you like the frequency of EasyBuild releases?




Fewer EasyBuild releases in 2024 due to effort on EasyBuild v5.0
(not counting beta releases)

Expected to increase again in 2025 and beyond


No negative response to fewer releases in 2024 👍

EasyBuild 5.0: Highlights




- Drop support for running EasyBuild on top of Python 2.7
- Python 3.6+ and Lmod 8+ or Environment Modules 4.3.0+ required
- Deprecate `run_cmd` + `run_cmd_qa`, replaced by new `run_shell_cmd` function
- New mechanism to control contents of generated module files: `module_load_environment`
- **Easy way to start interactive shell to debug failing shell commands**
- Reproducible source tarballs for easyconfigs using `git_config` in sources
- Changed defaults: use RPATH linking, enable `--trace`, use `depends_on`, use `pip`, ...
- Default maximum parallelism set to 16 (can be changed in EasyBuild configuration)
- **Full overview of changes in EasyBuild documentation: docs.easybuild.io/easybuild-v5**

EasyBuild 5.0: Stats



- 103 weeks after branching 5.0.x from develop (in March 2023)
- 25.5 weeks after release of EasyBuild 4.9.4 (previous stable)
- 59 sync meetings (via Zoom) to discuss EasyBuild 5.0 development
- **1,364 merged pull requests**
 - framework: 245 (24 in develop branch, 220 in 5.0.x branch)
 - easyblocks: 345 (41 in develop branch, 303 in 5.0.x branch)
 - easyconfigs: 804 (628 in develop branch, 175 in 5.0.x branch)

EasyBuild 5.0: Lessons learned





- **Why did it take us so long to get EasyBuild 5.0 ready to release?**
 - Branching off 5.0.x and dropping support for Python 2.7 early on resulted in overhead
 - Scope creep throughout the development process
 - Lack of coordination (got significantly better in last couple of months)
 - Limited number of people actively involved (but too many is a problem too)
 - Tension with ongoing activities in EESSI
- **Major EasyBuild versions are not happening often enough, should be every ~2 years?**
 - “Once in a lifetime” kind of feeling for breaking changes, leads to delays
 - Ensuring that deprecated functionality actually keeps working is not realistic
 - We shouldn’t hold back so much on making breaking changes?

Contributors, contributors, contributors, ...





Overview of contributions (framework)





- Steady increase in number of framework PRs since 2022 (largely due to effort towards EasyBuild v5)
- ~71% of framework PRs in 2023+2024 by contributors outside HPC-UGent (was ~55% in 2020-2022)
- Growing share of framework PRs by non-maintainers (~34% in 2023, ~38% in 2024)

Overview of contributions (easyblocks)





- Reduced activity in easyblocks in 2022-2023 (shift to using generic easyblocks?)
- Huge surge in activity in 2024 due to migration to `run_shell_cmd` and `module_load_environment`
- In 2024: ~85% of PRs by contributors outside HPC-UGent, incl. ~30% by people by non-maintainers

Overview of contributions (easyconfigs)



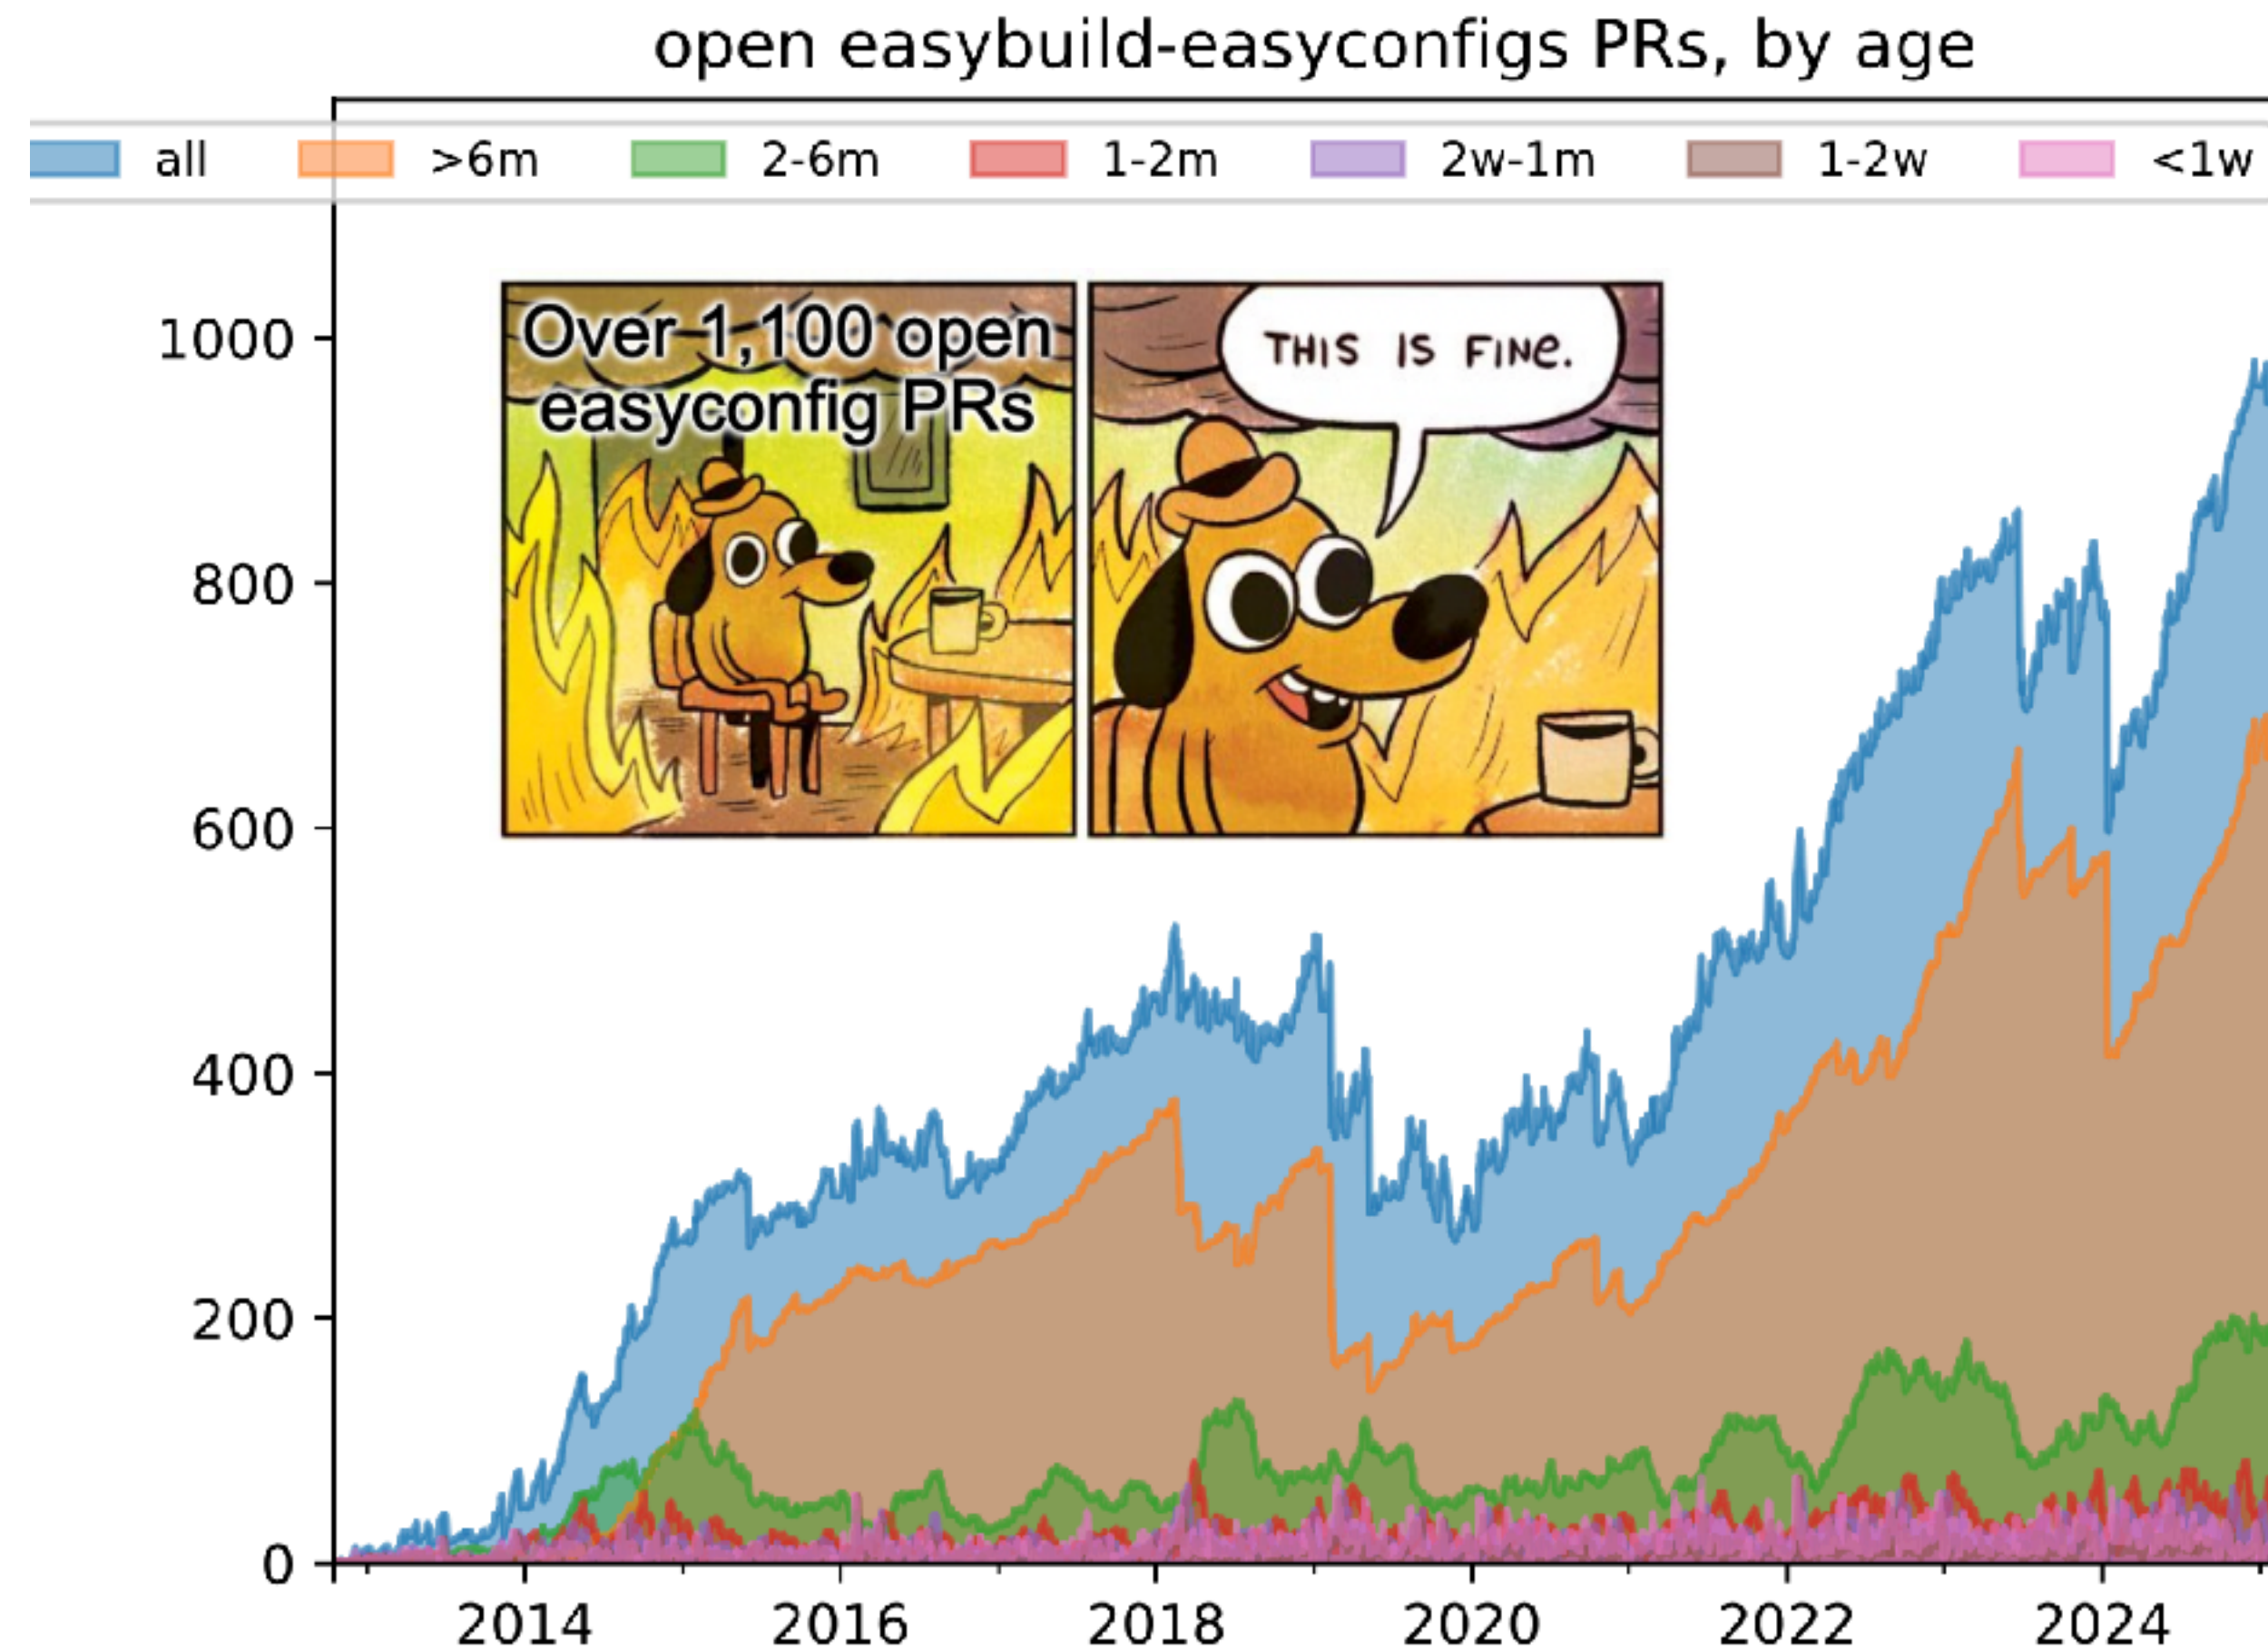

- Over 2,500 easyconfig PRs in 2024 (for the first time!)
- Shrinking ratio by boegel (now ~10%)
- Increasing ratio by HPC-UGent (~15%), of which ~13% by INUITs (up from 4-5% in 2022-2023)
- ~75% of PRs from contributors outside of HPC-UGent

Overview of contributions (easyconfigs)




- Vast majority of easyconfig PRs merged by maintainers other than boegel (>75% since 2021)
- Significant number of easyconfigs PRs from 2022-2024 are still open (20% of 2024 PRs)
- Ratio of manually created PRs (not using `--new-pr`) increased from ~12% in 2023 to ~33% in 2024 ?!

Overview of contributions (open easyconfigs PRs)




- Backlog of open easyconfigs PRs is bigger than ever: currently ~1,100 PRs (~70% opened > 6 months ago)
- Clear impact from focused effort towards EasyBuild v5.0.0; thorough cleanup is long overdue
- High rate of incoming PRs: > 10 new easyconfig PRs per working day (>2,500 per year)

Overview of contributors






- 38 new contributors to easyconfigs repo in 2024, now over 470 in total
- Continued growth in unique contributors for all 3 major EasyBuild GitHub repos

Overview of contributors (per year)



- In 2024: 28 unique contributors to EasyBuild framework, 49 to easyblocks
- 115 unique contributors for easyconfigs in 2024 (down from 130 in 2023)
- Growth in backlog for easyconfig PRs may be discouraging contributors?

Overall quality of EasyBuild



Very positive evaluation of overall quality is stable

(nobody ever answered "pretty bad")

Integration of EESSI in EuroHPC Federation Platform




- Ghent University was invited to join consortium led by CSC.fi that will implement the EuroHPC Federation Platform (EFP, see Henrik's EUM'25 talk)
- EESSI will be integrated in EFP as Federated Software Stack
- Quoting from invitation to tender:

Federated scientific software catalogue across systems

Federation of scientific software catalogues across HPC systems as well as simplification of their use on HPC systems is not a new concept. The result of early attempts is the implementation of the environment setting modules to load the appropriate libraries and dependences through “module load” commands and then use the adequate software. This method is currently widely used on most of the HPC systems across Union.


Another relatively new development is EasyBuild and EESSI, a software build and installation framework that hides the complexity of compiler specific adjustment and allows management of the software installation of HPC systems in an easy way. It is even capable of installations of workflows on HPC systems.

Goals & challenges (short term)




- **Get back on top of backlog of contributions** (especially for easyconfigs)
 - Cleanup of stale/old PRs (cfr. toolchain policy)
 - Reinstate Maintainer-of-the Week role
 - Regular merge sprints (monthly?): focused effort on processing recent contributions
- **Revamp testing infrastructure** (more on that in Sebastian's EUM'25 talk)
 - Test in containers (diff. OSs), EESSI build environment, consistent testing on Arm 64-bit (aarch64)
- Further **improvements to error reporting** in EasyBuild
 - Sanity check failures
 - Context-aware error reporting (config.log, Cmake logs, output produced by pip, ...)
 - And some more crazy ideas...
- AMD GPU support (ROCm, ...)

Goals & challenges (longer term)




- **Keep up with incoming contributions + software installation requests**
- Synergy / tension between EasyBuild & EESSI
- Increasing diversity of hardware: Arm is here to stay, RISC-V is on horizon
- LLVM-based toolchain (see Davide's talk on day 2 of EUM'25)
 - Mature Fortran frontend (flang) now also available:
<https://blog.llvm.org/posts/2025-03-11-flang-new>
 - Integrate LLVM as additional compiler in foss toolchain?
- *What else? Let us know what's important to you!*

Towards proper governance



- **Initial Steering Committee (ISC)** for EasyBuild was created a couple of weeks ago
 - As discussed during last (virtual) EasyBuild Maintainers Summit (10+13 Feb'25)
 - Members: Kenneth Hoste, Simon Branford, Sebastian Achilles, Adam Huffman, Sam Moors, Caspar van Leeuwen
 - See also docs.easybuild.io/governance
 - First meeting 3 Mar'25, next meeting 18 June'25 to discuss draft proposal of governance structure
- **Main goal: set up proper governance for EasyBuild** (beyond current BDFL + do-ocracy model)
 - Inspired by governance done for other open source software projects, but tailored to EasyBuild
 - Scope, governance structure (SC, advisory board?), frequency of meetings, voting procedure, etc.
- Envisioned responsibilities for SC: roadmap, vision statement (long-term goals & ambitions), funding, ...



Plans for EasyBuild 6.0



- Initial work on EasyBuild 6.0 has already started!
- Main goal would be to **improve consistency in naming**:
EasyBuild configuration options, easyconfig parameters, toolchain options, etc.
 - In EasyBuild v5.0, you can already use `configure_opts` instead of `configopts`
- Current plan is to release EasyBuild 6.0 in ~2 years (before EUM'27 ?)
 - Less time between major releases of EasyBuild
 - Quicker removal of deprecated functionality



Crazy ideas...



Crazy ideas (part 1)

Set up a blog for EasyBuild (<https://blog.easybuild.io>)

- Similar to EESSI blog: <https://eessi.io/docs/blog>
- Aim for at least one post per month
 - New releases
 - Feature highlights
 - User stories
 - Reports on relevant events
- Different authors, anyone who's up for writing on a topic relevant to the EasyBuild community
- **First (planned) post: 10th EasyBuild User Meeting (Helena + Kenneth + ??)**

Crazy ideas (part 2)

Create non-profit / foundation for EasyBuild & EESSI

- Provide support services, training, etc.
- Potential “customers”: EuroHPC Hosting Entities, commercial companies, ...
- Avenue to let organisations sponsor EasyBuild and/or EESSI
- Help with sustainability of both EasyBuild & EESSI, organise events, etc.
- Maybe through Linux Foundation & High Performance Software Foundation (HPSF) ?


Crazy ideas (part 3)

Integration of Large Language Models (LLMs) into EasyBuild

- Generate easyconfig files
- Explain error messages in failing builds to humans
- Provide suggestions to fix the problem
- Agent mode: try out fixes *in sandbox environment* to eventually come up with a working easyconfig

May be controversial for some people, so **must be done opt-in**:

- Only when EasyBuild is *explicitly configured* to use LLMs (--llm option)
- Only when a separate `easybuild-llm` package was installed (`import easybuild.llm`)
- By default only using local LLMs, separate configuration option for enabling use of remote LLMs



Crazy ideas (part 4)


Official shop for EasyBuild + EESSI swag

- Fair price, incl. small donation for future free swag (~\$1 per item)
- You also pay for shipping + tax
- Currently only for EasyBuild + EESSI coffee mugs
- What else should be available?
- See link on <https://easybuild.io>




**Free coffee mugs for all
EUM'25 in-person attendees!**

Sponsored by NVIDIA/Microsoft/INUITs




More crazy ideas?





EASYBUILD



E E S S I

EUROPEAN ENVIRONMENT FOR
SCIENTIFIC SOFTWARE INSTALLATIONS

State of the union

Kenneth Hoste

Ghent University, Flemish Supercomputer Centre (VSC)

kenneth.hoste@ugent.be

**10th* EasyBuild User Meeting*

25 March 2025 - Jülich (Germany)

<https://easybuild.io/eum25>