



GAIA : Software deployment through CI

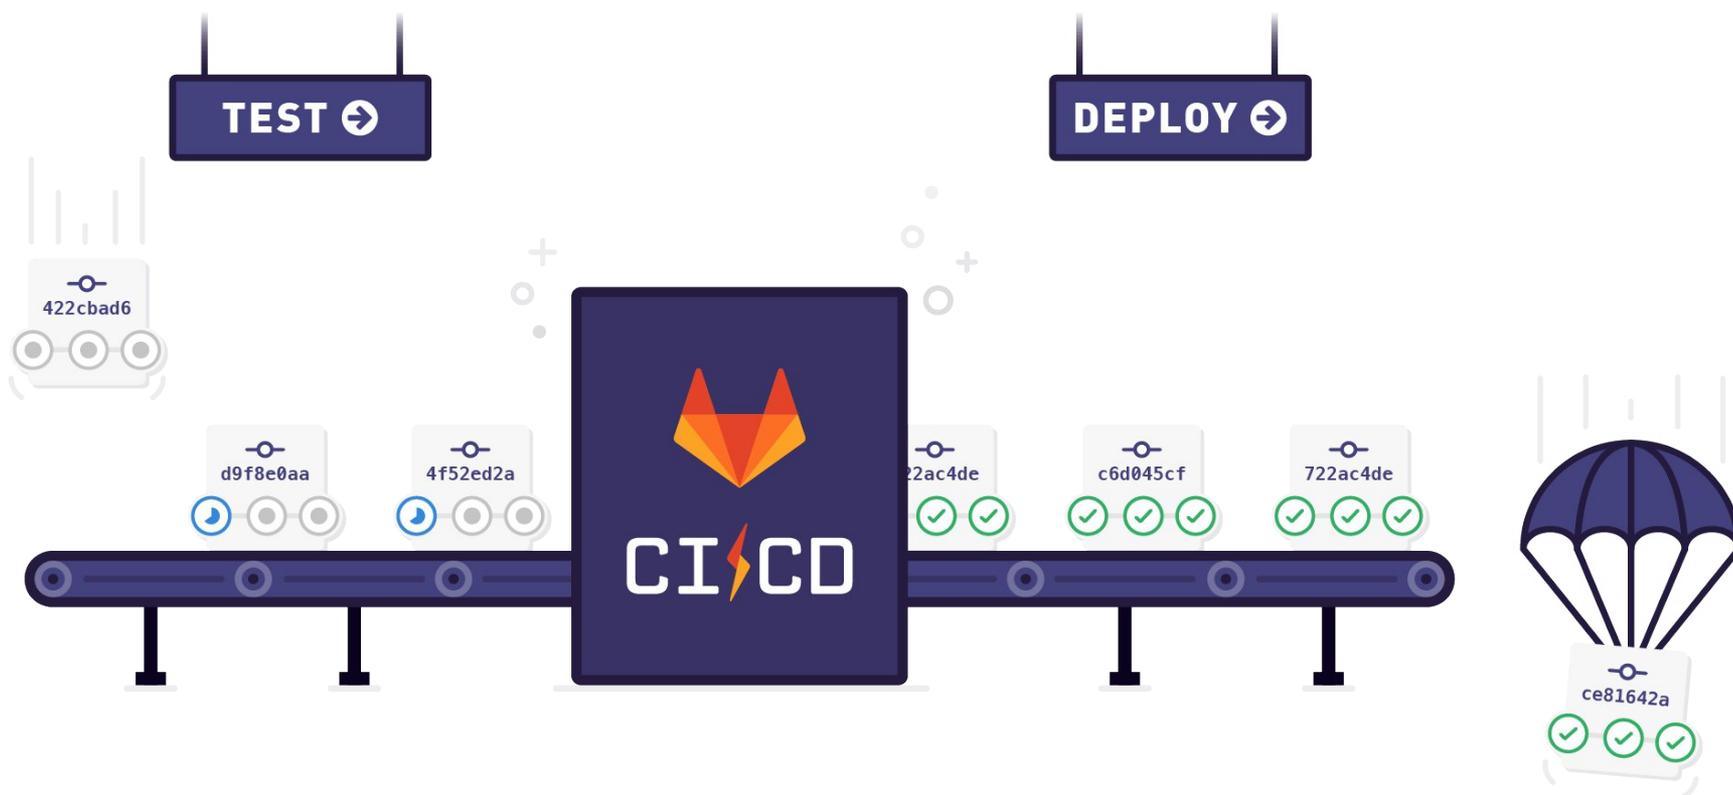
Cédric Jourdain : HPC expert @CINES

Presented by Gabriel Hauteux, head of HPC department

jourdain@cines.fr ; hauteux@cines.fr

• Objectives

- Automate **software deployment** and module generation.
- Automate **validation and testing** of functionality and performance.
- Deployment **reproducibility**



• Project Overview In Figures

The Gaia project, initiated in 2019 on OCCIGEN, now manages all software stacks for ADASTRA.

Commit statistics

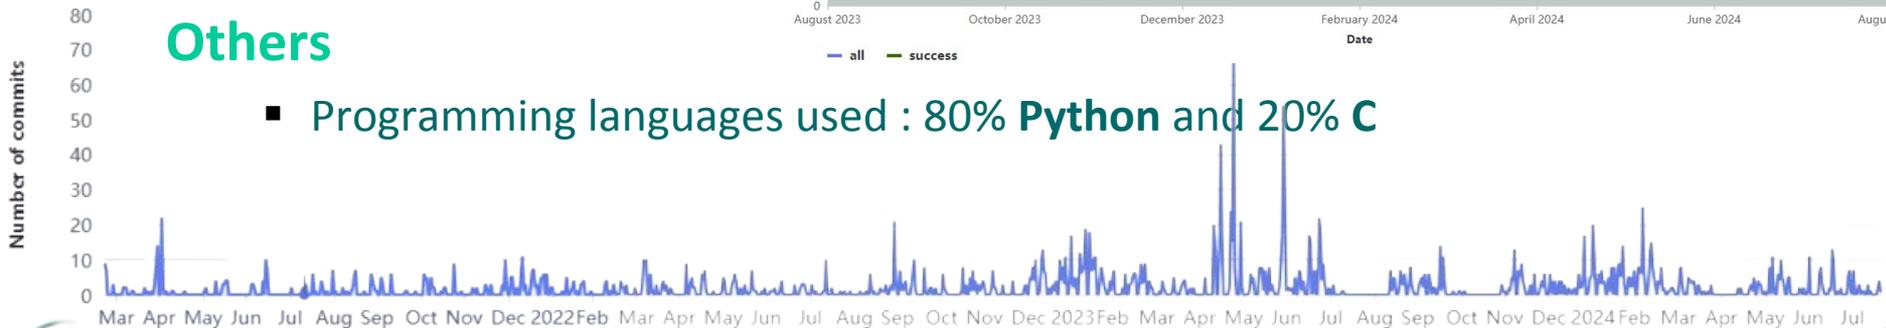
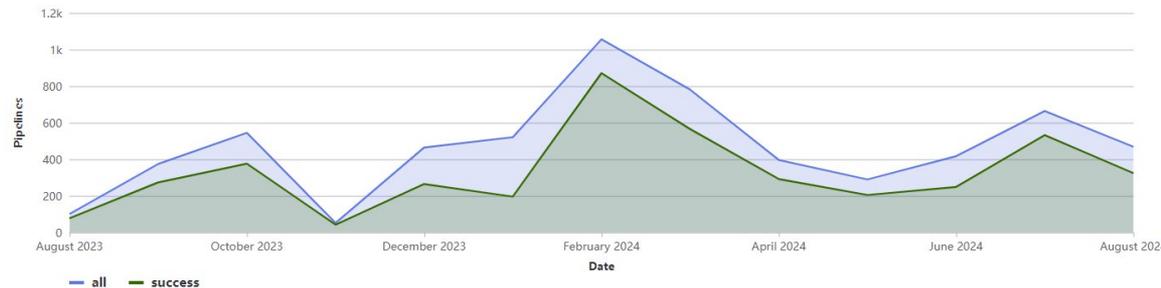
- Total: 2000 commits
- Average per day: 2.4 commits
- Authors: ~ 15

CI/CD Analytics

- Total Pipelines: 4041
- Success ratio: 52.84%

Others

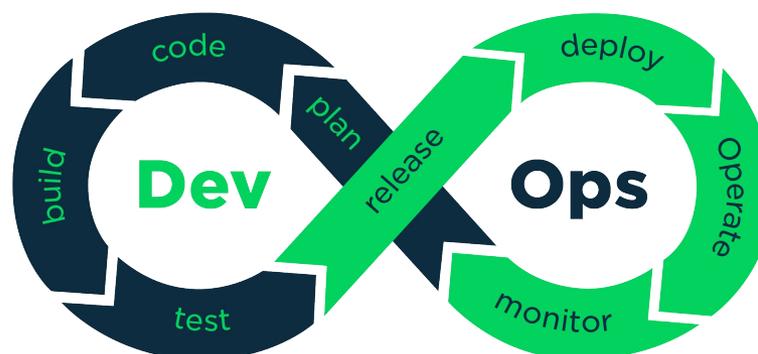
- Programming languages used : 80% **Python** and 20% **C**



- **Generality :**

How we work ?

- Issues, branch, merge request (Milestone and Sprint)
- DevOps:



Which tools ?

- **Gitlab** : Collaborative tools for **Code Repository** and **CI/CD**
- **Spack** : Software deployment tool : **Packages** and **Modules**
- **Reframe** : Framework for writing regression tests

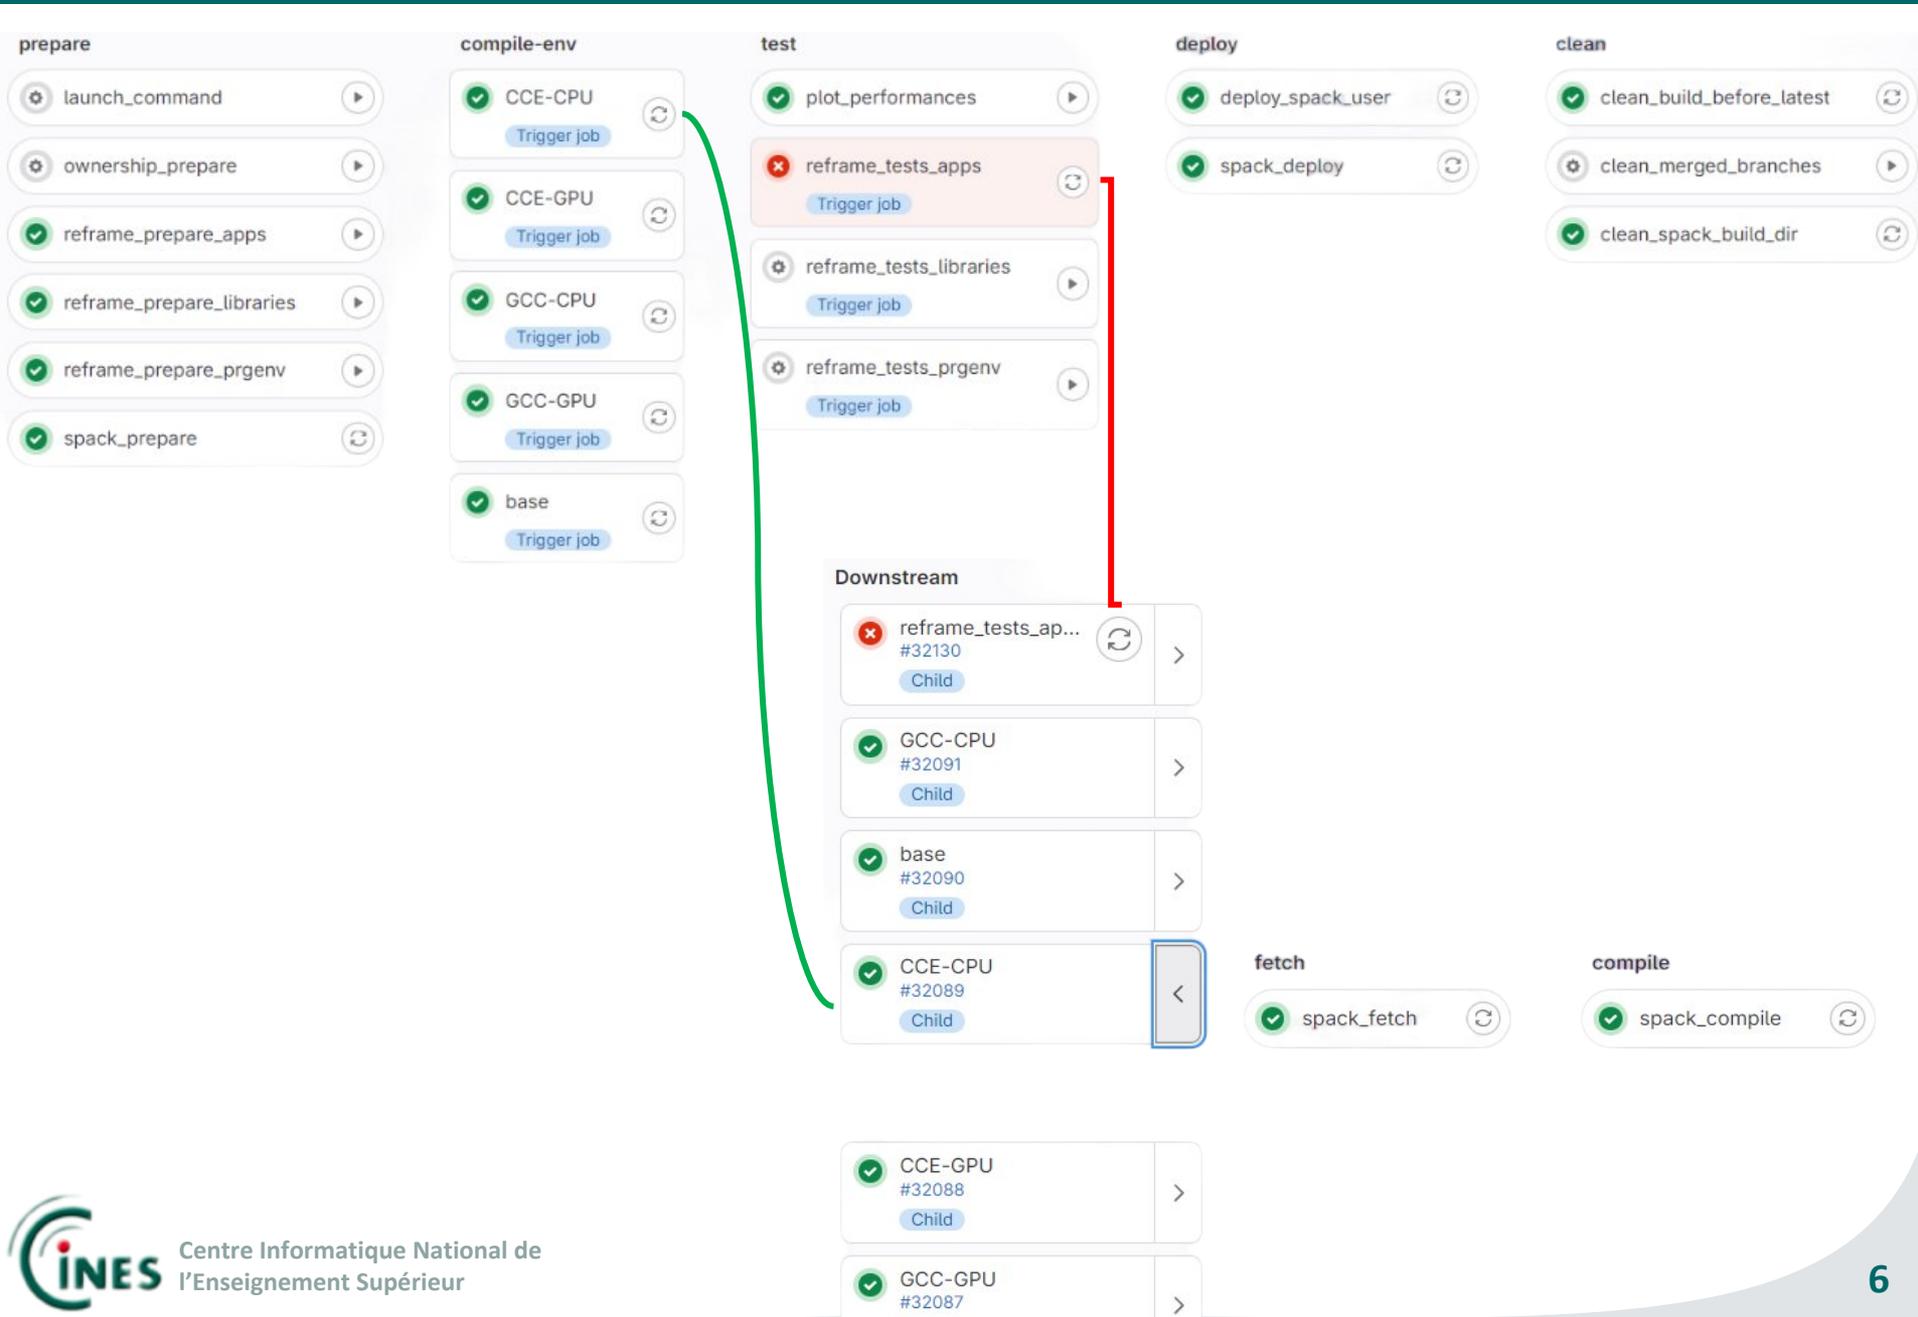




• Gitlab

- CINES Internal repository: [Gaia](#)
 - CI/CD: DevOpS methodology
 - [Pipelines](#): Define and automate workflows
 - Child pipelines:
 - 1 per Spack environment
 - 1 per type of tests (App, Lib, prgEnv)
 - Multiple jobs in one stages
- } ~ 5 parallels jobs per pipelines
- [Runners](#): agents that execute CI jobs
 - Utilizing **20 concurrents runners**
 - Running on the ``/tmp`` filesystem of the **login node** (cross-compilation for AMD GPUs)

TOOLS: Gitlab : pipelines





- # Spack

- **What is Spack ?**

“Spack is a package management tool designed to support multiple versions and configurations of software on a wide variety of platforms and environments”

- **What is a Spack package ?**

*“Spack packages are installation scripts, which are essentially **recipes** for building (and testing) software.”*

- ✓ *Where to find and how to retrieve the software*
- ✓ *Its dependencies*
- ✓ *Options for building from source*
- ✓ *Build commands.*

- **Spack package list ?**

- More than 7000 packages and recipes available : [\[list\]](#)

• Spack in GAIA

■ *How ?*

- **Git submodule** (*git clone --recursive*)
 - [.gitmodules](#) based on a version [commit](#) (updated every stable release)
- **Internal recipe** ([local-repo](#))
 - [repos.yaml](#) : Specifies repositories for preferred ``packages.py``
- **Modules**: Generated by Spack with an internal overlay to simplify the view of the modules
- **Spack-user**: We provide users with a module ``spack-user`` with our configuration (no need for small installations).

■ *Why?*

- Benefit from contributions of other HPC centers (LLNL, Pawsey, KTH, etc.)
- Promote CINES Through Spack Contributions:
 - Enhance the visibility of CINES by sharing recipes (``package.py``) with the Spack community. (~ 15 PR (thks to Etienne, Mohand, Mathieu and more to comes)

- Spack in GAIA : Module view

```

----- /opt/software/gaia/prod/latest/Core -----
CCE-CPU-3.1.0      GCC-CPU-3.1.0      develop          (L)      spack-MI250-3.1.0
CCE-GPU-3.1.0      GCC-GPU-3.1.0      spack-GEN0A-3.1.0

----- /opt/software/gaia/deprecated -----
CCE-CPU-3.0.0      CCE-GPU-3.0.0      GCC-CPU-3.0.0   GCC-GPU-3.0.0   archive (L)      user-CPU-3.0.0      user-GPU-3.0.0

----- /opt/software/gaia/prod/latest/Base -----
anaconda3/2023.09-0  ffmpeg/6.0          hpcviewer/2023.07  metis/5.1.0-int64  papi/6.0.0.1
benchmark/1.8.3     gdb/13.1            julia/1.10.0       metis/5.1.0        (D)      poppler/21.09.0
blitz/1.0.2         gnuplot/5.4.3      libxml2/2.10.3     mmg/5.7.1          qt/5.15.11
bzip2/1.0.8        gsl/2.7.1          libxsmm/1.17       mpifileutils/0.11.1-mpi  zlib/1.3
cmake/3.27.7       hpctoolkit/2023.08.1-mpi  likwid/5.2.2      nsimd/3.0.1

----- /opt/software/gaia/dev/latest/Core -----
CCE-CPU-4.0.0      CCE-GPU-4.0.0      GCC-CPU-4.0.0 (L)  GCC-GPU-4.0.0   spack-user-4.0.0

----- /opt/software/gaia/dev/latest/Base -----
gdb/14.1          (D)      hpcviewer/2024.02 (D)  paraview/5.13.0    r/4.3.0          sox/14.4.2
gnuplot/6.0.0 (D)      papi/7.1.0          poppler/23.04.0 (D)  rust/1.78.0

----- /opt/software/gaia/dev/3.1.0-05ab50b1/GCC-CPU-4.0.0 -----
adios2/2.10.0-mpi  mumps/5.6.2-int64-omp-mpi  petsc/3.21.1-omp-mpi  siesta/5.0.1-elpa-omp-mpi  (D)
boost/1.85.0-mpi  octopus/14.1-elpa-mpi      plumed/2.9.0-mpi
cdo/2.4.2-omp-mpi  openblas/0.3.26-omp       quantum-espresso/6.8-omp-mpi  sundials/7.0.0-mpi
hypre/2.31.0-mpi  openfoam/2312-int64-mpi    siesta/4.1.5-elpa-omp-mpi  superlu-dist/8.2.1-int64-omp-mpi

----- /opt/software/gaia/deprecated/.archive -----
CCE-CPU-2.0.0      CPE-22.11-rocmmc-5.2.3-softs  CPE-23.02-rocmmc-5.3.0-GPU-softs  user-CPU-2.1.0
CCE-CPU-2.1.0      CPE-23.02-aocc-3.2.0-CPU-softs  GCC-CPU-2.0.0                      user-GPU-2.0.0
CCE-GPU-2.0.0      CPE-23.02-cce-15.0.1-CPU-softs  GCC-CPU-2.1.0                      user-GPU-2.1.0
CCE-GPU-2.1.0      CPE-23.02-cce-15.0.1-GPU-softs  GCC-GPU-2.0.0
CPE-22.11-cce-15.0.0-softs  CPE-23.02-gcc-12.2.0-CPU-softs  GCC-GPU-2.1.0
CPE-22.11-gcc-11.2.0-softs  CPE-23.02-gcc-12.2.0-GPU-softs  user-CPU-2.0.0

```

• Configuration

▪ Use of spack environments :

- Allows users to work within isolated groups of independent packages.
- **Objective**: To offer virtual environments (similar to [Python venv](#)) that enable users to share common installations.
- Designed for **easy sharing and reuse** by others and across different systems.

▪ Configurations files :

- [compilers.yaml](#)
- [packages.yaml](#)
- [concretizer.yam](#)
- [config.yaml](#)
- [modules.yaml](#)
- [repos.yaml](#)
- [mirrors.yaml](#)

Most of the work has to be done here. It contains the default preferences for compilers and for providers of virtual packages, highly dependent on the system being used."

Global configuration, with minimal dependence on the system.

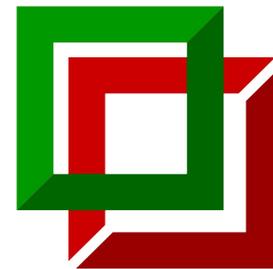
- **Interesting Spack features**

- **Currently in use:**

- **Upstreams**: allows you to point your Spack installation to another installation in order to use all the packages installed there
- **Parallel builds**:
 - « *child pipeline* » ([*spack ci generate*](#))
 - *Package level (make -j)*
- **Spack binaries** (pre-built binary caches ~binary packaging: .rpm, .deb)
 - Installation times are reduced by around **20 times!!**

- **To test again:**

- **Spack containerize**: «*creates recipes to build images for different container runtimes*»
- **Spack test**: «*run spack's tests for an install*»



• Reframe

■ *What is Reframe?*

“ReFrame is a powerful framework for writing system regression tests and benchmarks, specifically for HPC systems. Written in python”

■ *How and where?*

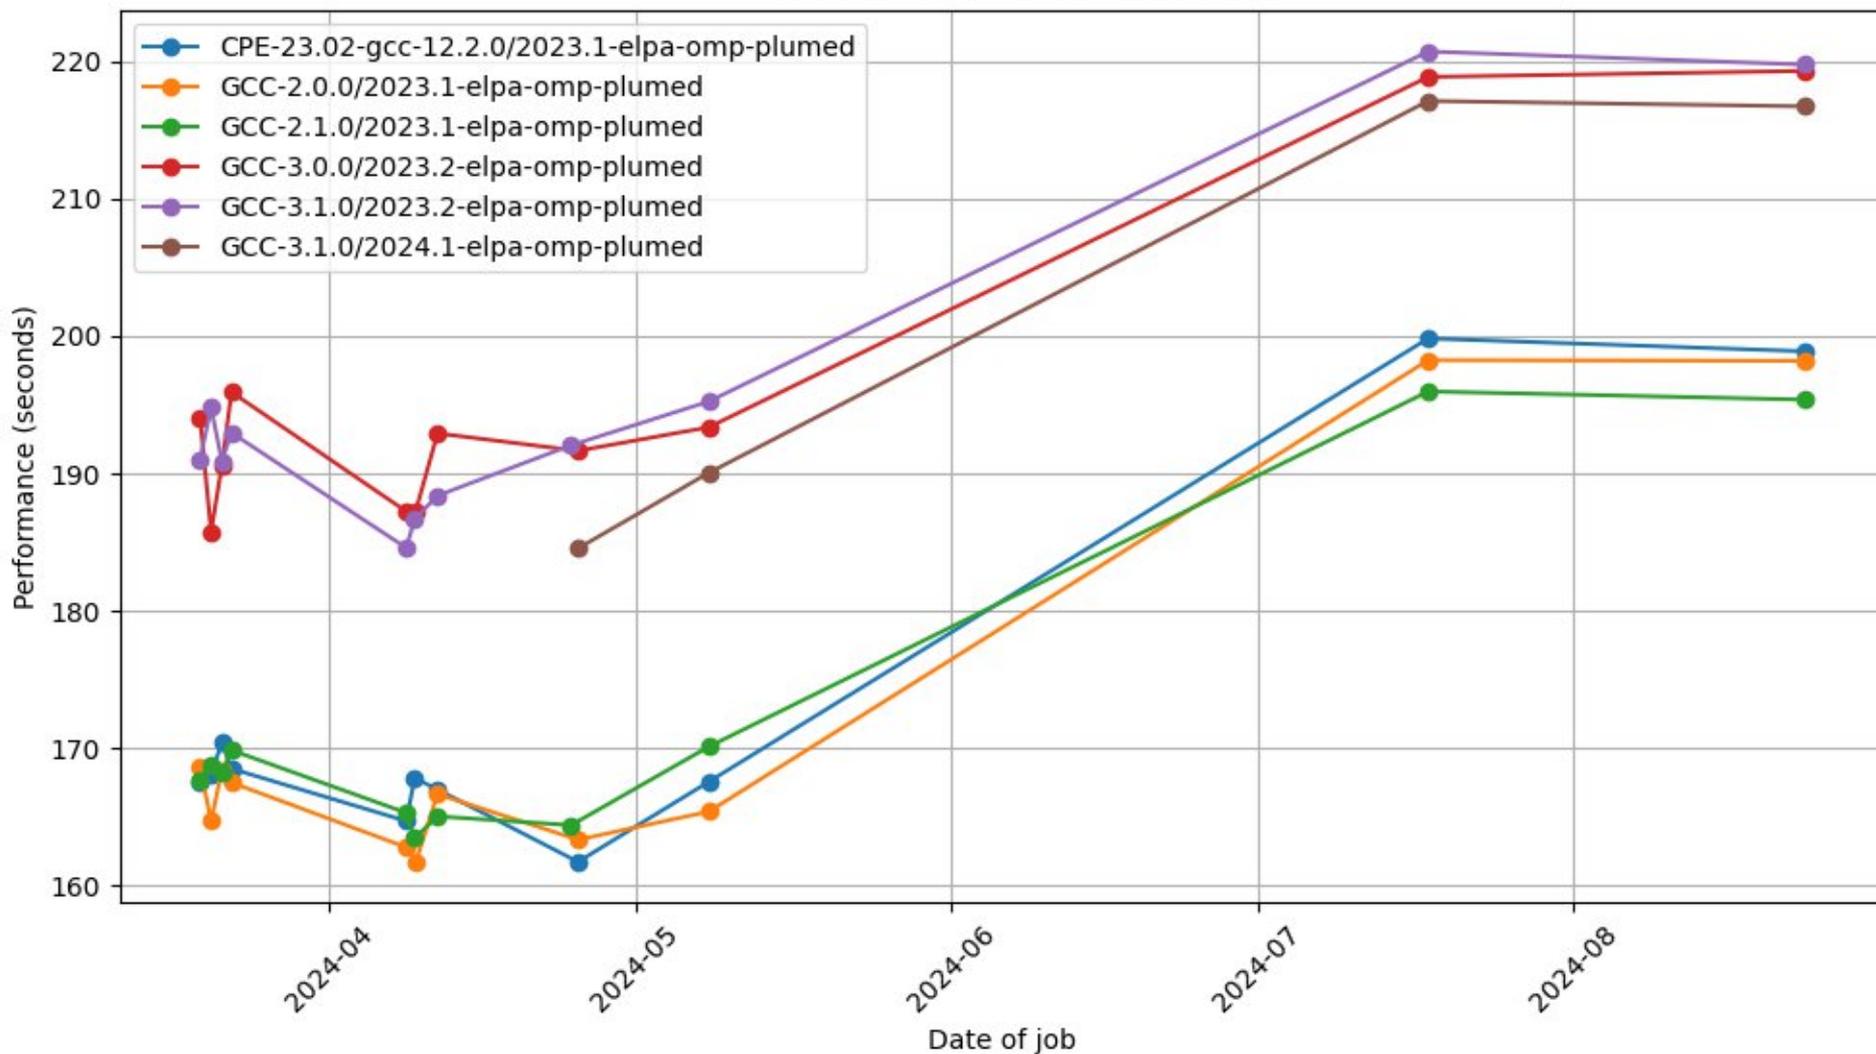
- **Physically installed**, but intended to be used as a sub-module [coming soon]
- **Configuration**: [setting.py](#) files: describes the systems used (partition, the scheduler, etc.) and the "environments" (application environment to be tested)
- Internal **git** repository: [cines-checks](#)

■ *For what purpose?*

- **Apps/Libs**: Monitoring and testing software installations deployed by spack
- **Prgenv**: verification of the programming environment deployed by HPE
- **System** (IO, Slurm)
- **Microbenchmarks** (HPL, OSU, HPCG, gpu_burn, ...)

Reframe : performance monitoring graphs

CP2K GPU Performance Over Time



Thank you for your attention!

Questions?

Cédric Jourdain (Jourdain@cines.fr)

