

Getting started with containers

Singularity

Why singularity?

On a HPC:

- Users have different levels of privileges
- Users submit jobs with time/cpu/memory restrictions

These are two things Docker doesn't facilitate:

- Requires superuser privileges
- Docker commands are an API of a daemon (parentless)

Specific singularity features

- UID + permissions are inside the container always the same as outside
- No daemon – container is a child process

Singularity without docker

- The singularity `Dockerfiles` : Singularity recipes (i.e. definition file)
- Singularity hub: <https://singularity-hub.org/>
- **Note:** installing as root inside container requires root privileges outside container!
- Solutions: `--fakeroot` or build with external runner

Pipeline development

- Containers are particularly useful for pipelines
- Both snakemake and nextflow support conda, docker and singularity
- Containers are built based on conda environments
- Once your cluster environment is set: ultimate scalability/reproducibility