CERN openlab

Evaluating Container Technologies for HPC

Summer Students Lightning Talks

Aleksander Wennersteen

14/08/2018

Contents

- Motivation
- What is a container
- What is High Performance Computing (HPC)
- HPC containers
- Results
- Issues
- Conclusion

ISC 2018 workshops

13th Workshop on Virtualization in High-Performance Cloud Computing

High Performance Container Workshop (Docker)

Getting Started with Containers on HPC through Singularity

SC18 workshops

Container Computing for HPC and Scientific Workflows (Shifter)

Containers, Collaboration, and Community: Hands-On Building a Data Science Environment for Users and Admins

US ECP software stream

Software deployment via containers one of the sub-projects

Why containers

- Reduce the complexity arising from conflicting dependencies in different layers of the OS software stack

- Portability
- Consistency
- Reproducibility

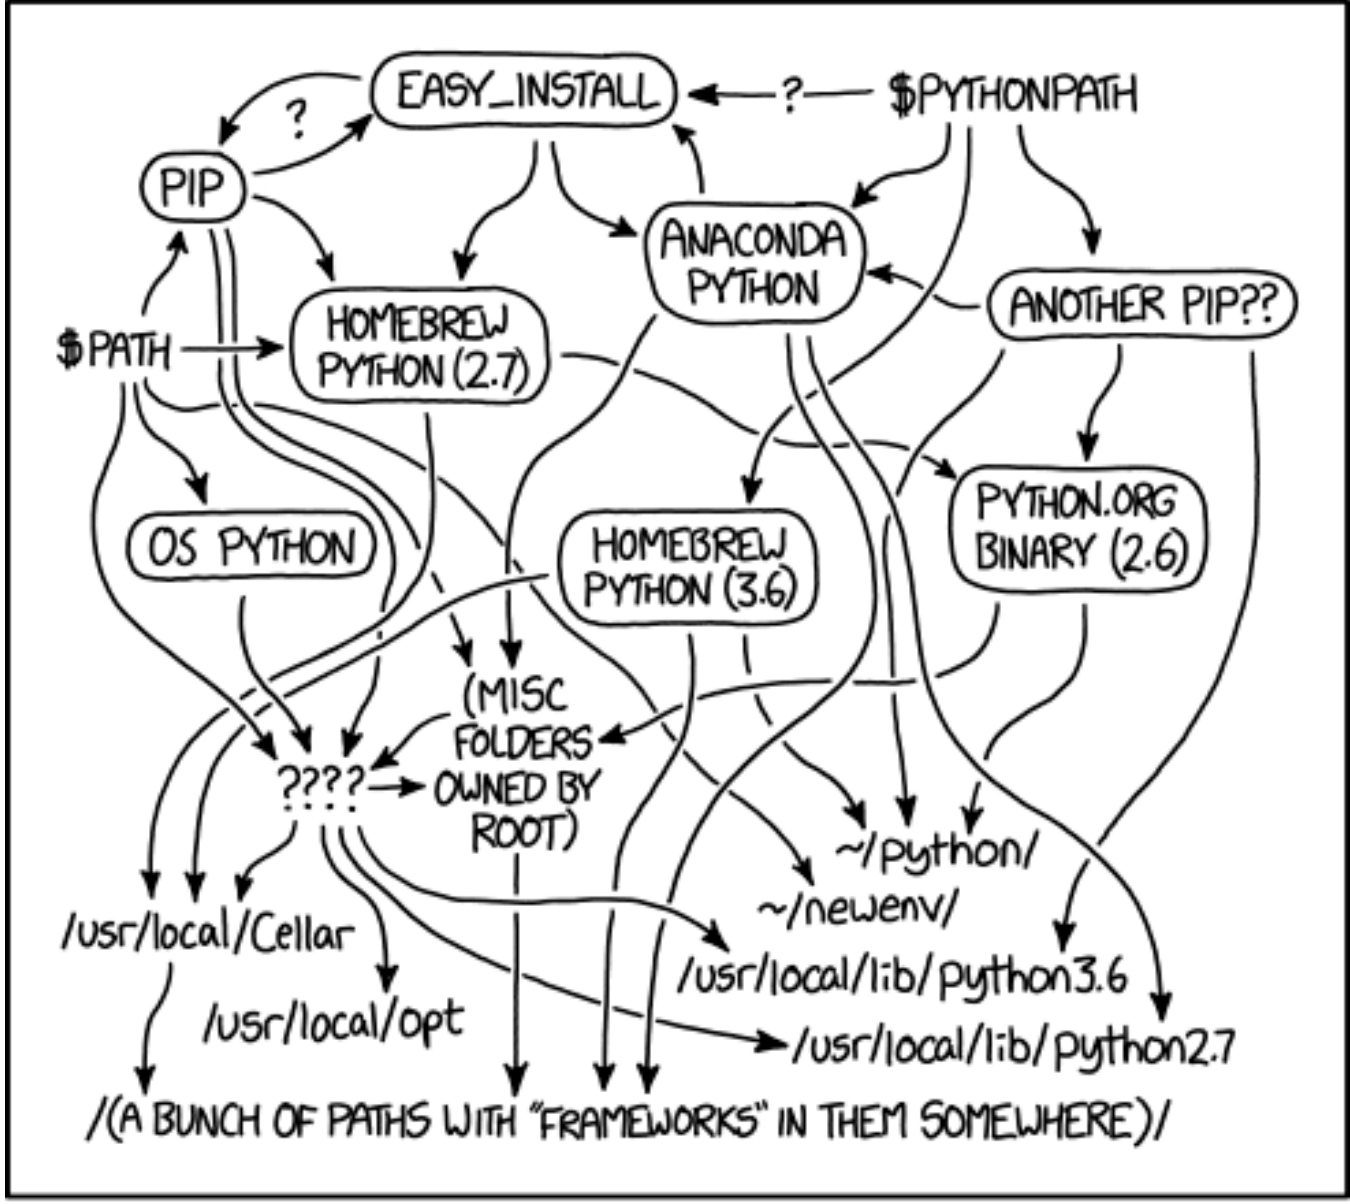

MY PYTHON ENVIRONMENT HAS BECOME SO DEGRADED THAT MY LAPTOP HAS BEEN DECLARED A SUPERFUND SITE.

Comic: xkcd.com/1987


No rubbish here, please
use containers outside.
Pas de déchets ici, mettez
les dans les containers à
l'extérieur svp.

Why containers

Copyright Sergio Pucci

What is a container

- OS level virtualisation
  - More lightweight than VMs
  - Easier to ship around
  - More performant
- You can pack your container with the libraries and software you need

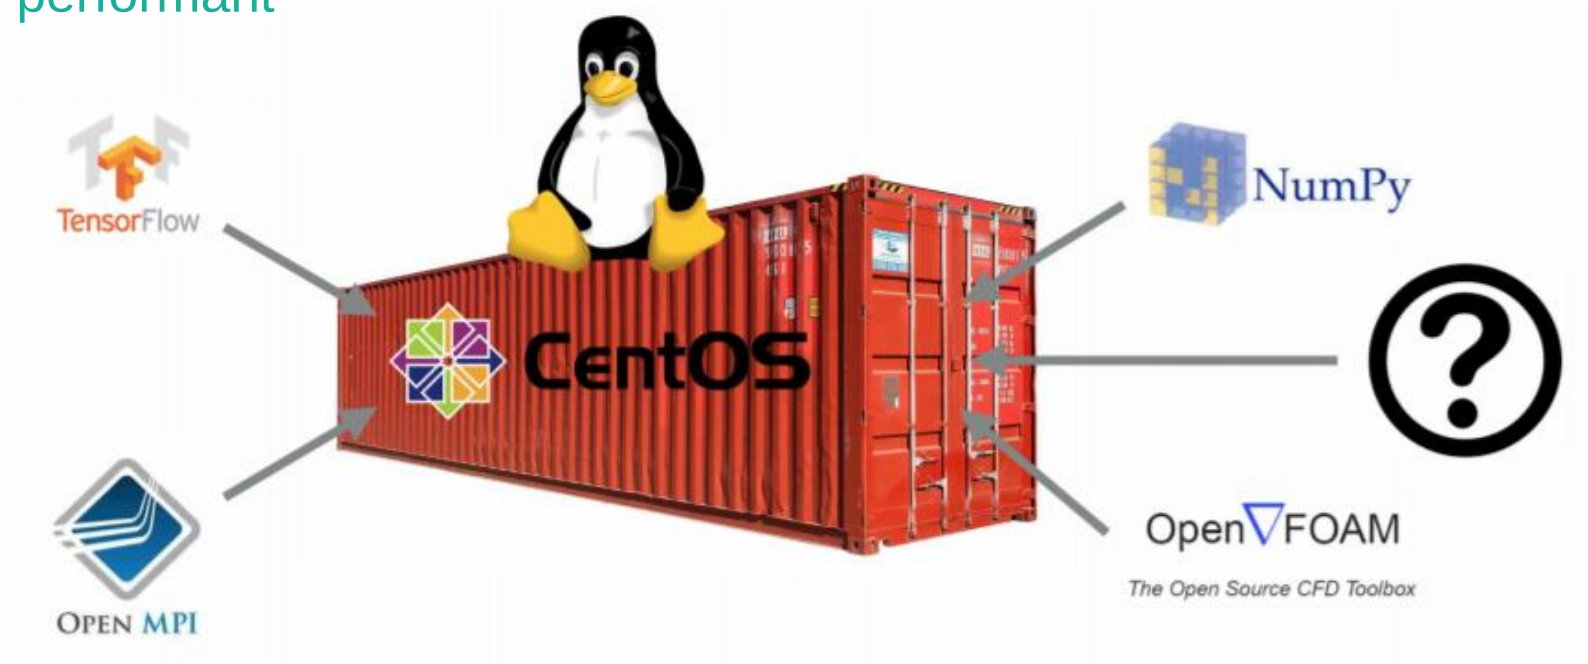

GM Kurtzer

High Performance Computing

Aka Supercomputing

- In the past, mainframes
- These days, a large number of servers connected together by low latency networks
- Distributed memory machines
- Programmed typically using MPI or PGAS languages

Cray X-MP: https://cerncourier.com/computing-at-cern-the-mainframe-era/

HPC containers

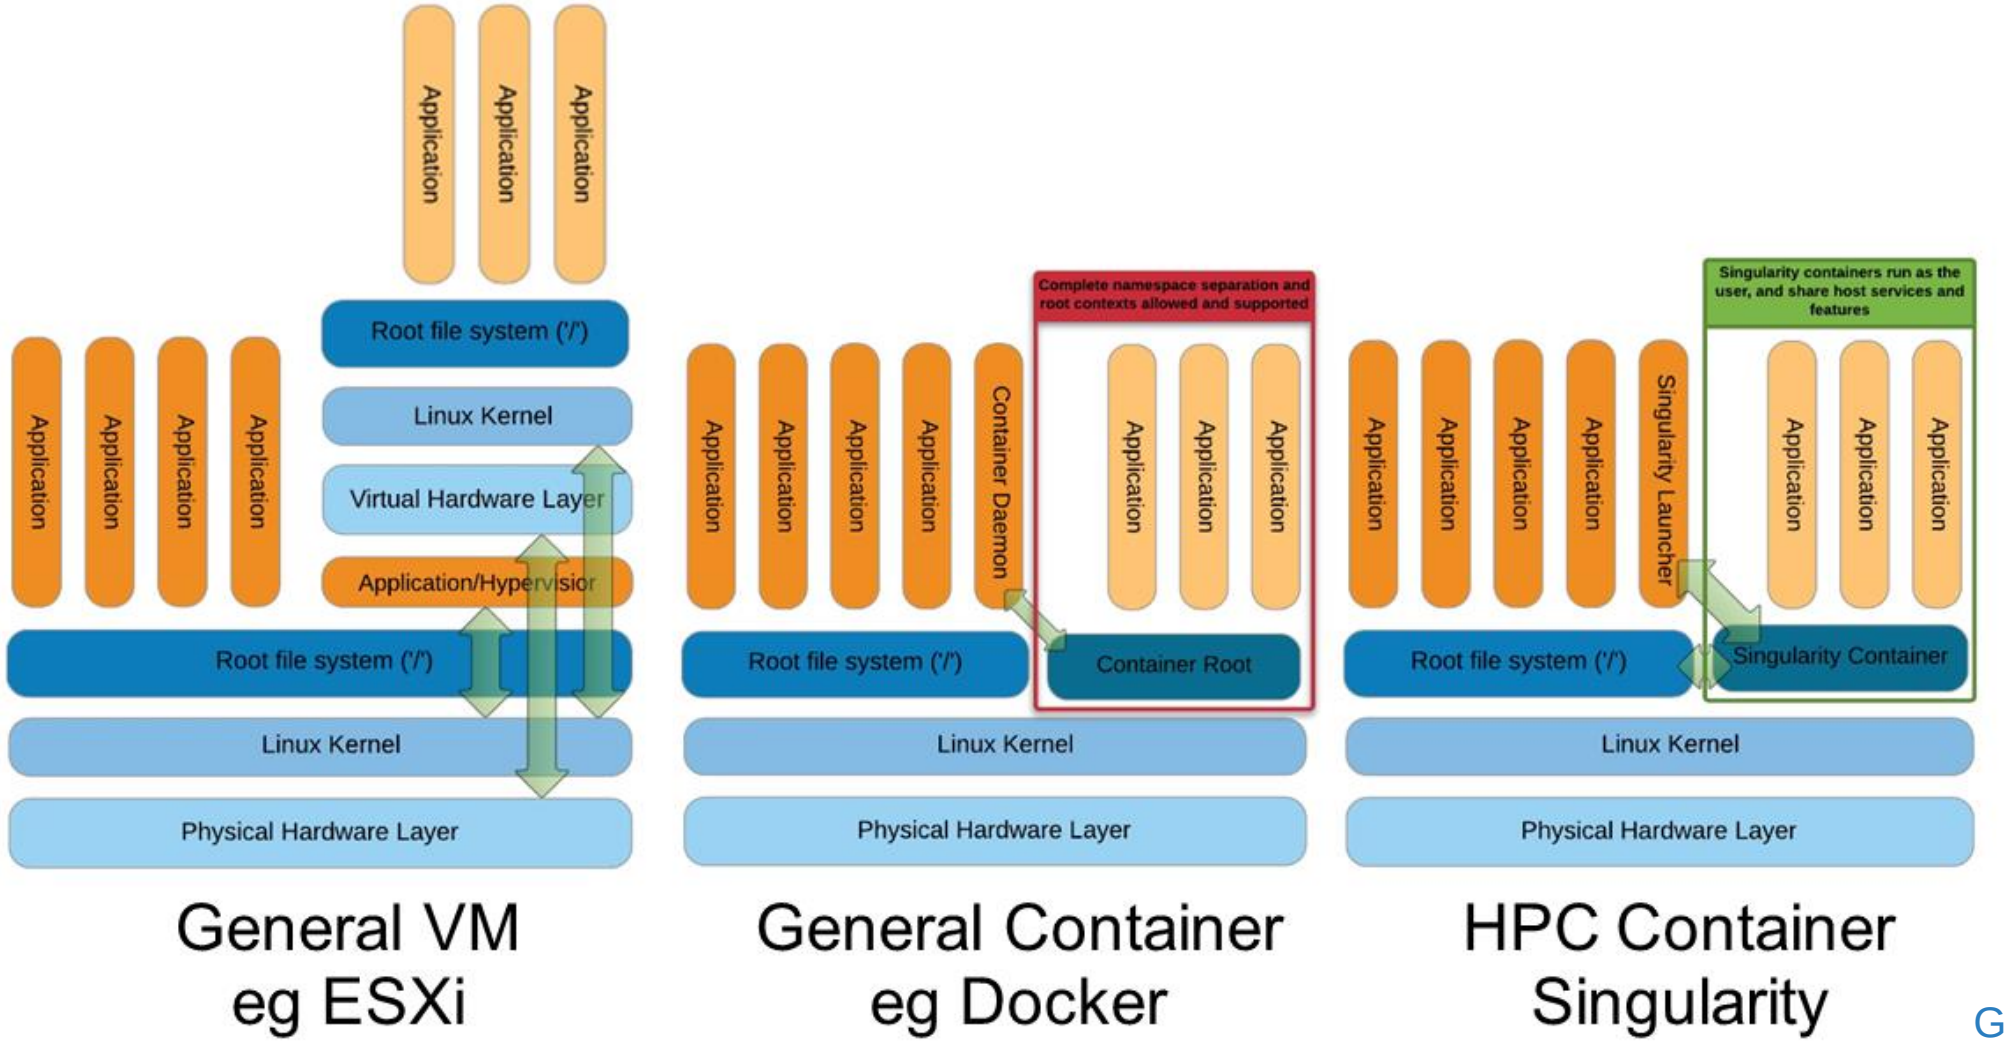

GM Kurtzer

HPC containers

- Unlike “normal” use, HPC does not want microservices and wants as little isolation as possible – bare-metal is “ideal”
- No privilege escalation inside the container
- Needs to be performant and low-latency. No additional latencies induced by additional software network bridging
- Support multi-node and MPI

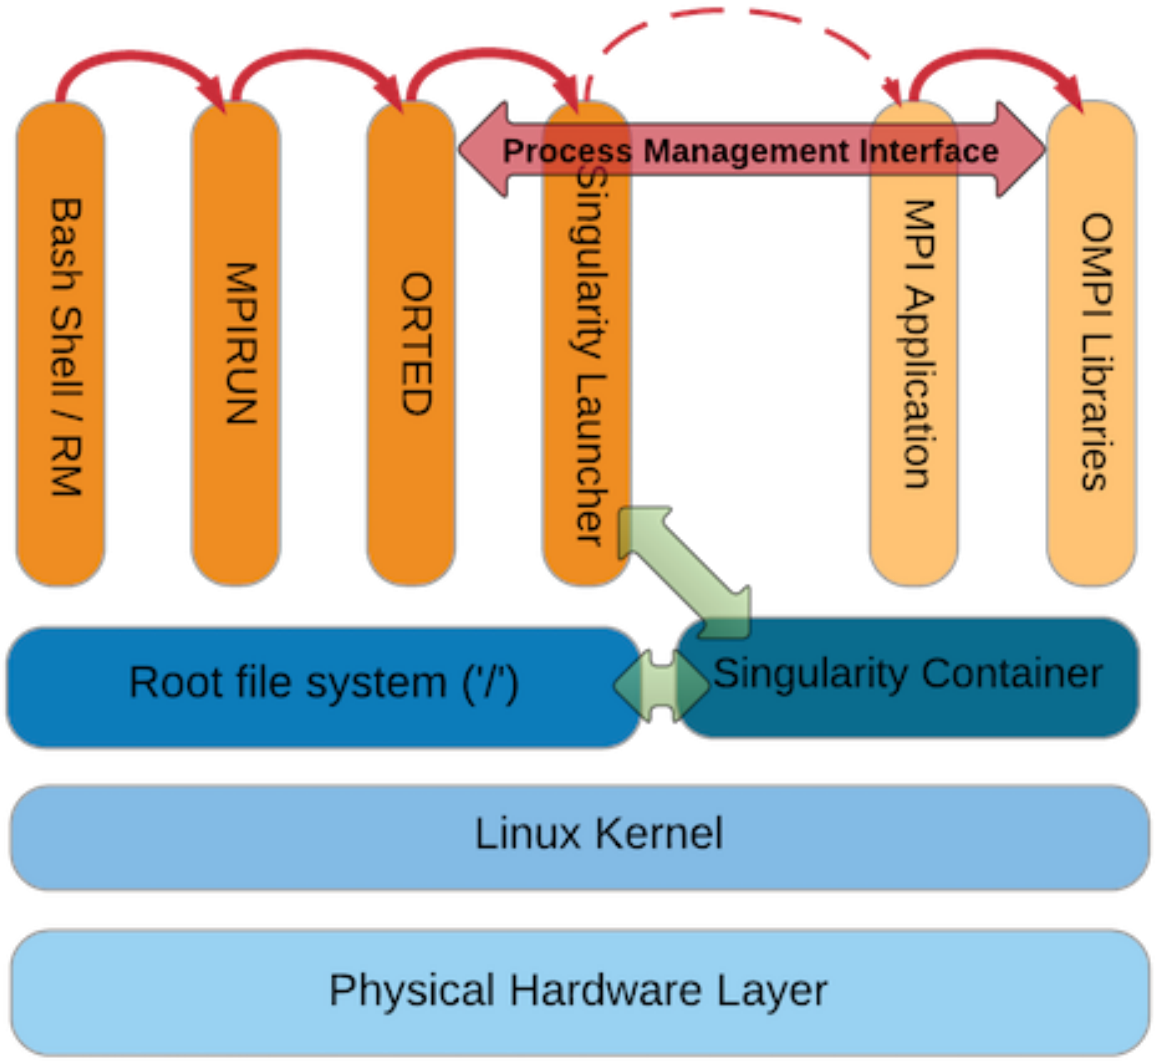

GM Kurtzer

Using Singularity

Serial

- singularity build container.img recipe.def
- singularity exec container.img ./program params
  - E.g singularity exec CentOS.img echo $HOSTNAME
    - hpc1@cern.ch
- singularity run container
  - Executes run script written at build time

MPI

- mpirun singularity exec container.img ./program params
  - E.g mpirun –np 2 singularity exec CentOS.img echo $HOSTNAME
    - hpc001@cern.ch
    - hpc002@cern.ch

Using Singularity

Running a container with Python + MPI using SLURM

```
10:24:03 alwenner slurmgate05 /hpcscratch/user/alwenner/work
→ cat ./launch_mpi4py.sh
srun -t 00:01:00 -p batch-short -N 2 singularity exec -B /hpcscratch/user/alwenner
 CentOS_mpi4py_mvapich.img /hpcscratch/user/alwenner/work/test.py

10:24:07 alwenner slurmgate05 /hpcscratch/user/alwenner/work
→ ./launch_mpi4py.sh
Hello from rank 0!
hostname is hpc002.cern.ch
Hello from rank 1!
hostname is hpc003.cern.ch
```

Issues

- Loop devices
  - Enabling => potential security problem
  - Disabling => limiting use cases
- Need for correspondence of MPI and PMI versions inside/outside
  - Better support with Open MPI >= 2.1
  - Usually host > container works
- Drivers
  - Interconnects
  - GPUs
- Otherwise same issues as installing software normally
  - Except now it’s just once
  - Initially hard to separate issues with container and software

Deployments

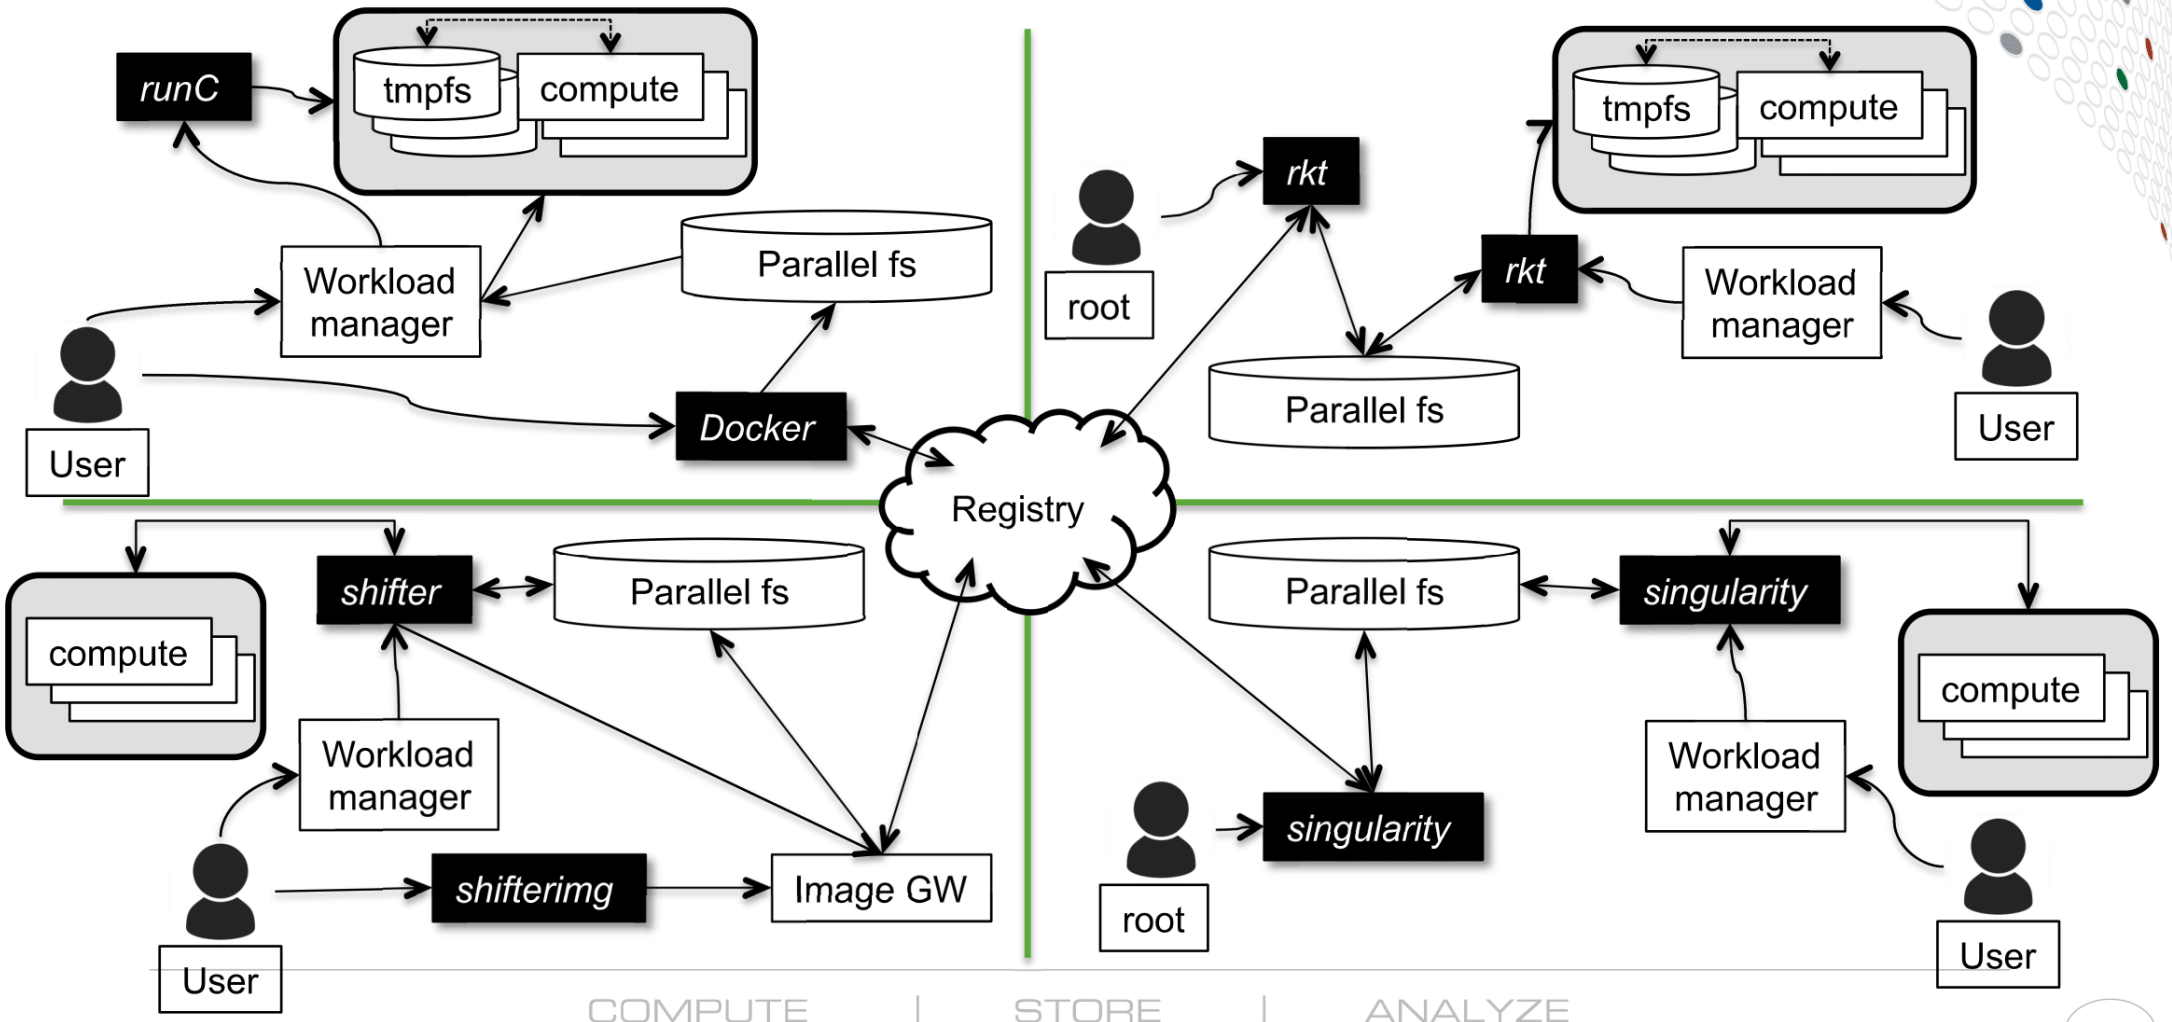

COMPUTE | STORE | ANALYZE

Jonathan Sparks @ Cray Inc: https://cug.org/proceedings/cug2017_proceedings/includes/files/pap164s2-file2.pdf

Preliminary results

- HPC (OSU) benchmarks containerised
  - No measurable performance degradation on our workloads
- Open source engineering applications containerised
- FDS
- Warp (Python 3 + Fortran + C + MPI)

Conclusion, outlook and experiences

- I think containers in HPC are here to stay
  - Just as in all of computing
- Deployment critical for user adaption
- Containerising the first software was time-consuming
  - The second quicker
- Still all the hassle with installing software
  - But now only once!

GM Kurtzer

Suggested reading

[Singularity Presentation by Kurtzer]

https://www.sylabs.io/docs/

Questions?

awennersteen@gmail.com

Supervisors : Pablo Llopis, Carolina Lindqvist