


OPPORTUNISTIC RESOURCE MANAGEMENT WITH COBALD/TARDIS AT U BONN

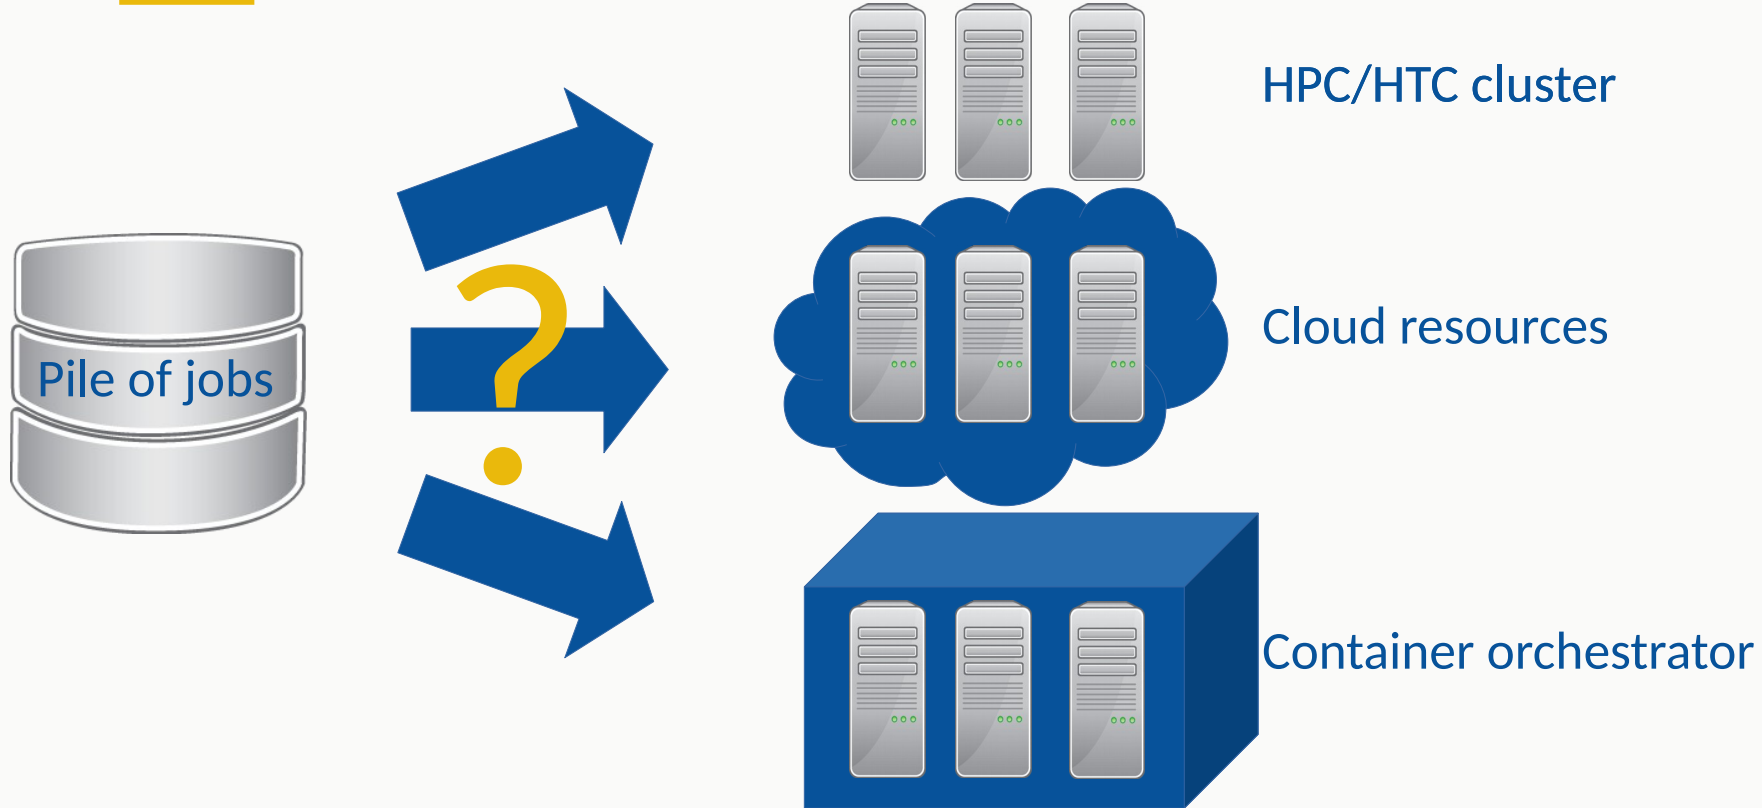
M. FISCHER¹, O. FREYERMUTH²,
M. GIFFELS¹, M. SCHNEPF¹,
P. WIENEMANN²

¹KIT

²UNIVERSITY OF BONN



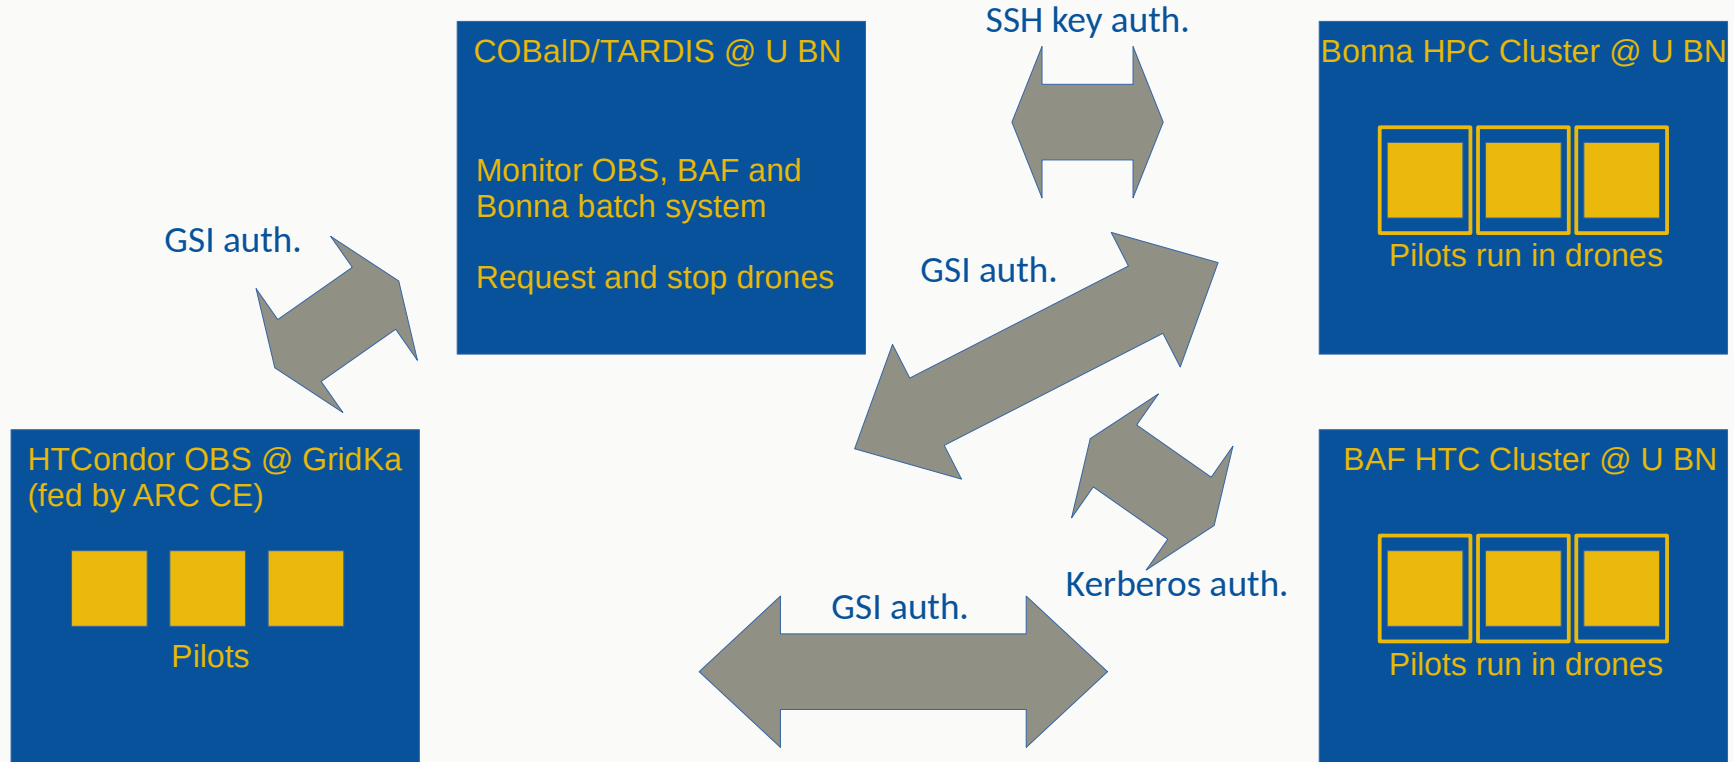
THE CHALLENGE



COBALD/TARDIS APPROACH


- Add abstraction layer between resource users and resource providers
- Hide heterogenous resources behind a single point of entry
- COBaID (COBaID Opportunistic Balancing Daemon)
 - Monitors usage of booked resources
 - Ramps up booked resources if they are well utilised
 - Reduces booked resources if they remain unused
- TARDIS (Transparent Adaptive Resource Dynamic Integration System)
 - Implements integration and management of resources provided by different systems (currently supported: OpenStack, CloudStack, Moab, Slurm, HTCondor) into overlay batch system (OBS)
- Developed by KIT (M. Fischer, M. Giffels, E. Kühn, M. Schnepf, et al.)

SETUP FOR CLUSTERS @ U BONN



JOB STRUCTURE @ U BONN


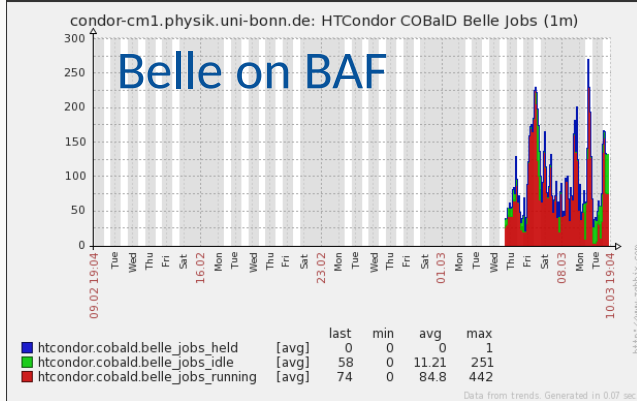
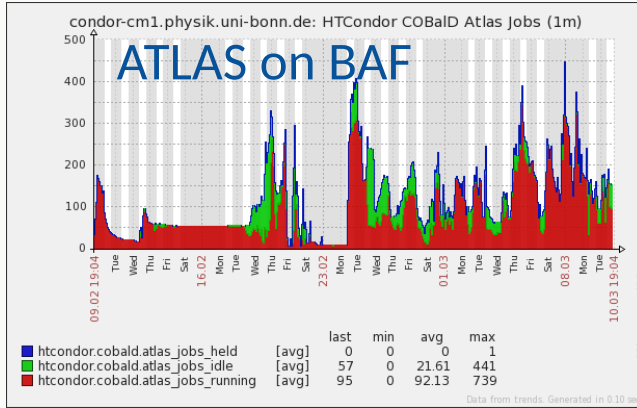
- Nested structure
- BAF containers to decouple cluster operation from user requirements (convenient for operators)
- ATLAS containers to reduce site requirements (convenient for ATLAS)
- ATLAS pilots to improve throughput of ATLAS production system
- Recently started to accept Belle jobs



BONNA SETUP

- Extend COBalD/TARDIS puppet module to handle remote submissions via ssh
→ <https://github.com/unibonn/puppet-cobald>
- Fully unprivileged setup (no root privileges on Bonna cluster)
 - use cvmfsexec to mount CVMFS repositories (cached by BAF squids)
 - local repository for unpacked drone container image
 - CERN repositories for experiment software
 - <https://github.com/cvmfs-contrib/cvmfsexec>
 - use lightweight Charliecloud container runtime (uses unprivileged user namespaces) → <https://github.com/hpc/charliecloud>

BAF & BONNA USAGE



SUMMARY

- COBaLD/TARDIS successfully deployed a U Bonn, management fully puppetized
- Service feeds *two* clusters
 - local submission to physics cluster (BAF)
 - remote submission to university-wide HPC cluster
- Routinely run ATLAS production jobs
- Recently started to accept Belle jobs

Thank you for your attention!

Max Fischer: max.fischer@kit.edu

Oliver Freyermuth: freyermuth@physik.uni-bonn.de

Manuel Giffels: manuel.giffels@kit.edu

Matthias Schnepf: matthias.schnepf@kit.edu

Peter Wienemann: peter.wienemann@uni-bonn.de