UNIVERSITÄT BONN

OPPORTUNISTIC RESOURCE MANAGEMENT WITH COBALD/TARDIS AT U BONN

M. FISCHER[1], O. FREYERMUTH[2], M. GIFFELS[1], M. SCHNEPF[1], P. WIENEMANN[2]

[1]KIT
[2]UNIVERSITY OF BONN

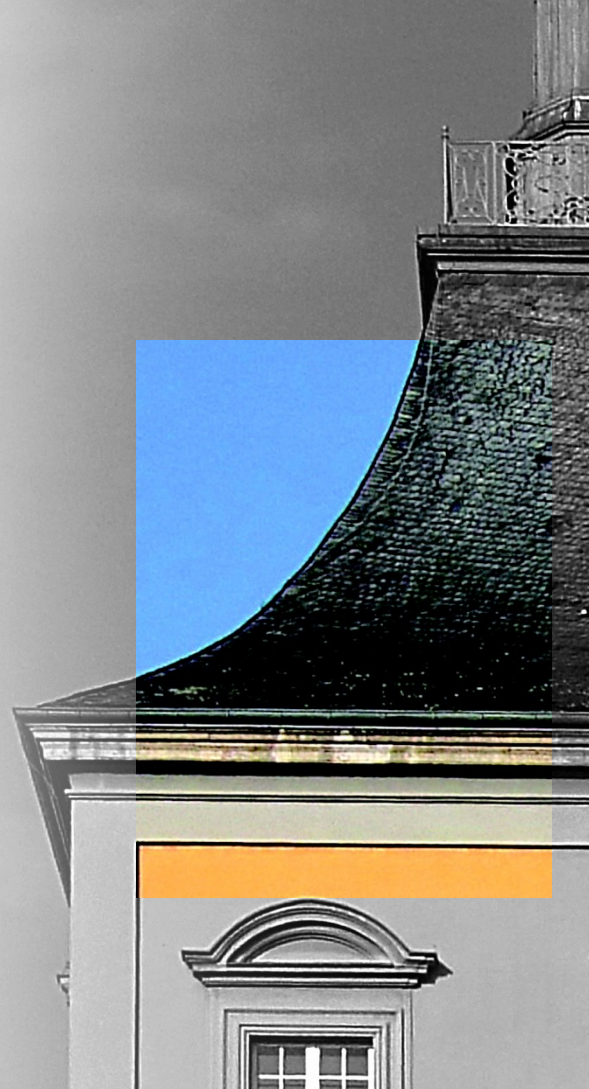

UNIVERSITÄT BONN

OPPORTUNISTIC RESOURCES

- Computing demands by HL-LHC are way beyond what can be provided by dedicated resources (even if performance/price evolution is taken into account)
- Multiple approaches to address challenge
  - Temporarily use non-dedicated resources (HPC/HTC clusters, cloud resources, container orchestration suites, …)
  - Improve software/algorithms
  - Investigate new technologies (quantum computing, …)
  - etc.

UNIVERSITÄT BONN

THE CHALLENGE

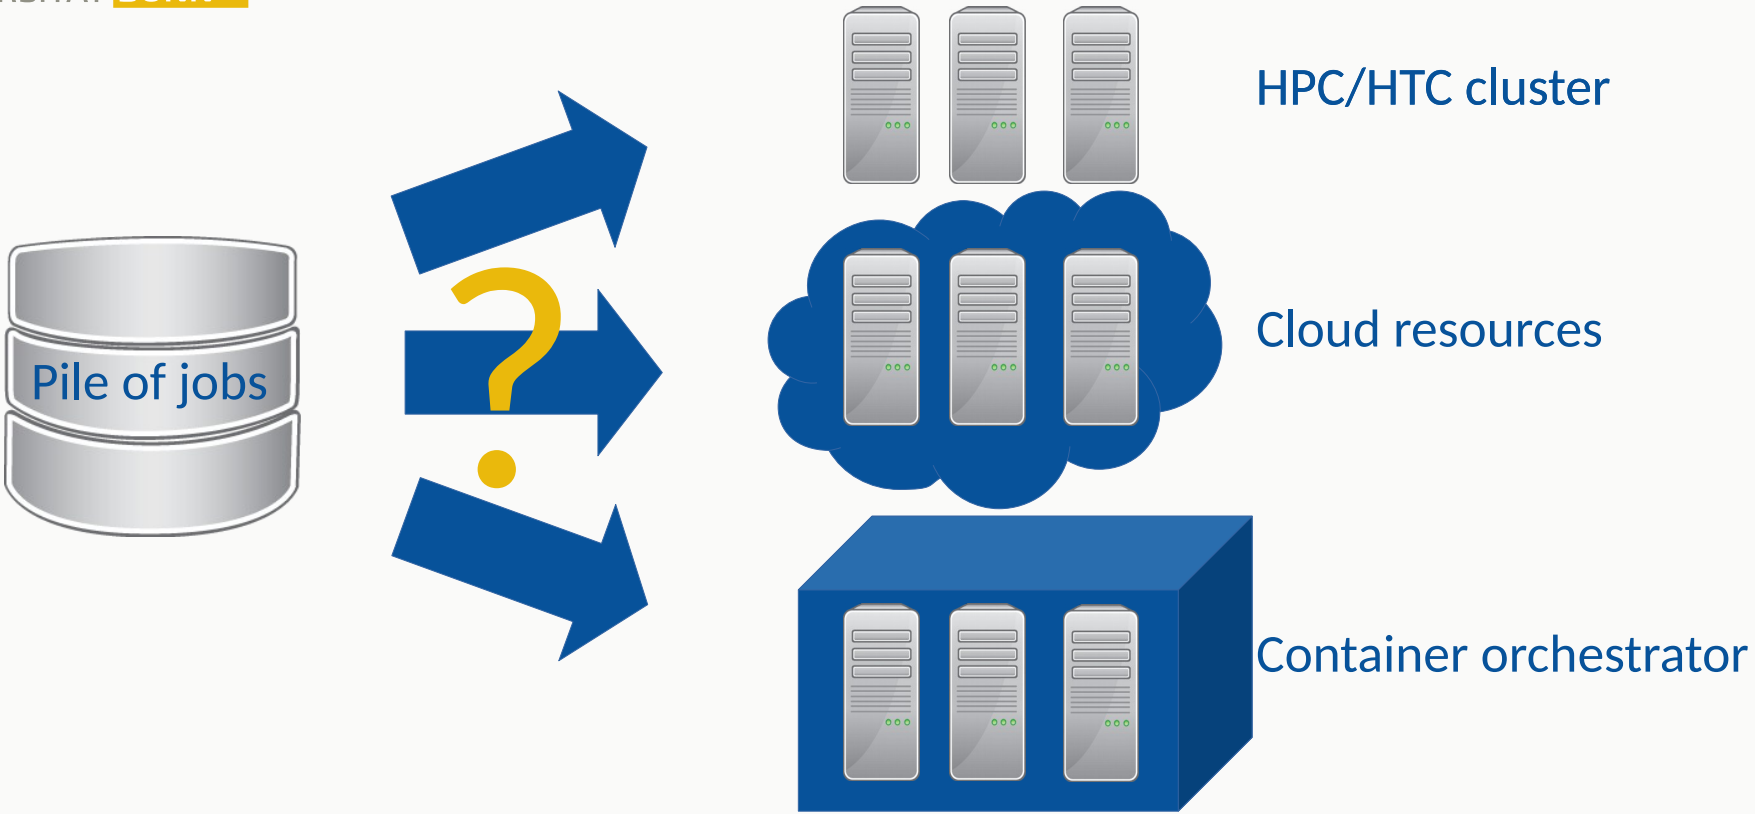

UNIVERSITÄT BONN

THE CHALLENGE (CONT'D)

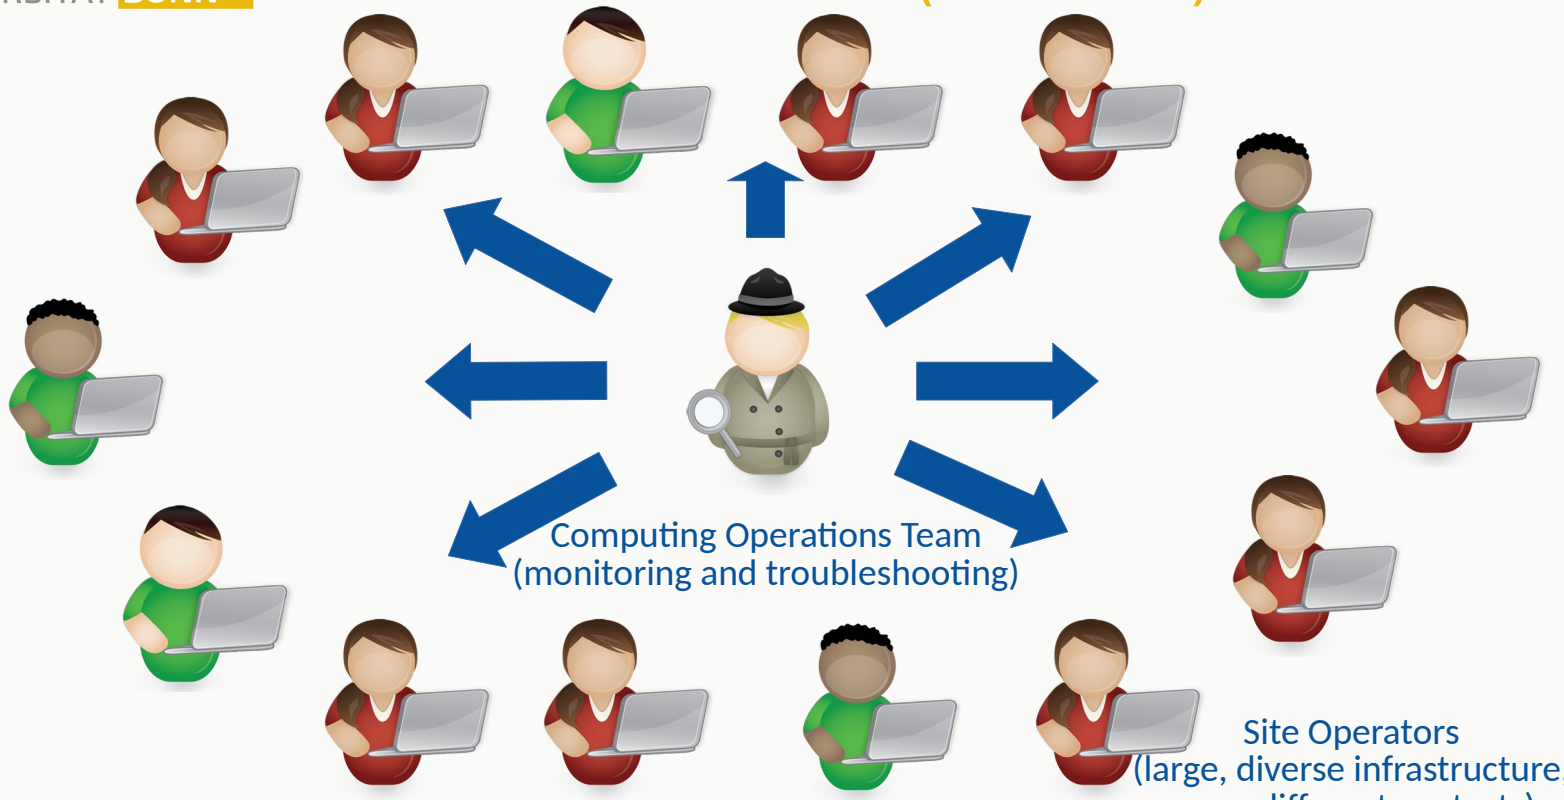

UNIVERSITÄT BONN

COBALD/TARDIS APPROACH

- Add abstraction layer between resource users and resource providers
- Hide heterogenous resources behind a single point of entry
- COBalD (COBalD 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 (OPS)
- Developed by KIT (M. Fischer, M. Giffels, E. Kühn, M. Schnepf, et al.)

UNIVERSITÄT BONN

BONN ANALYSIS FACILITY (BAF)

Desktops (~170 PCs)

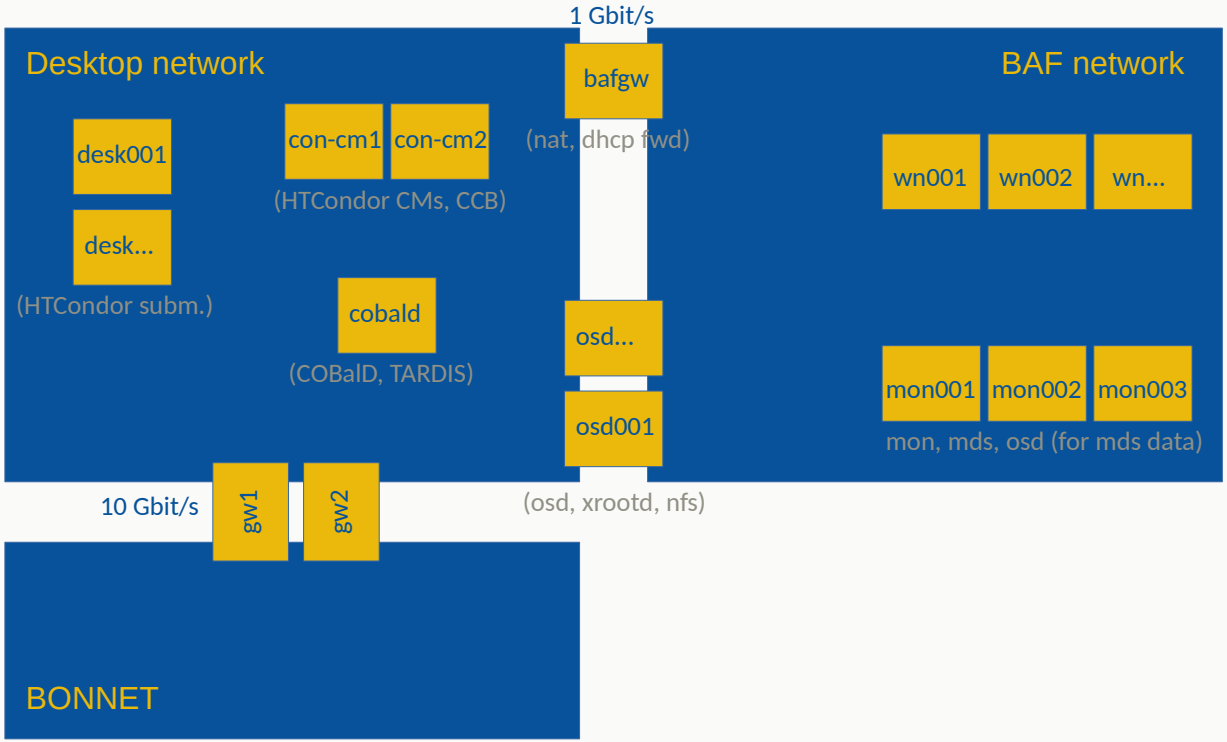

HTCondor WNs (1120 cores)

CephFS storage (500 TB)

UNIVERSITÄT BONN

PILOTS AND DRONES

- Pilot = „Empty“ job skeleton which allocates resources (# cores, RAM)
  - ATLAS pilot: Also transfers input/output, pulls in payload when CPU becomes available (reduces latency)
- Drone = Generalised pilot started by TARDIS
  - Starts HTCondor master, provides software environment, adds resources to overlay batch system (OBS)

COBALD/TARDIS SETUP @ U BONN

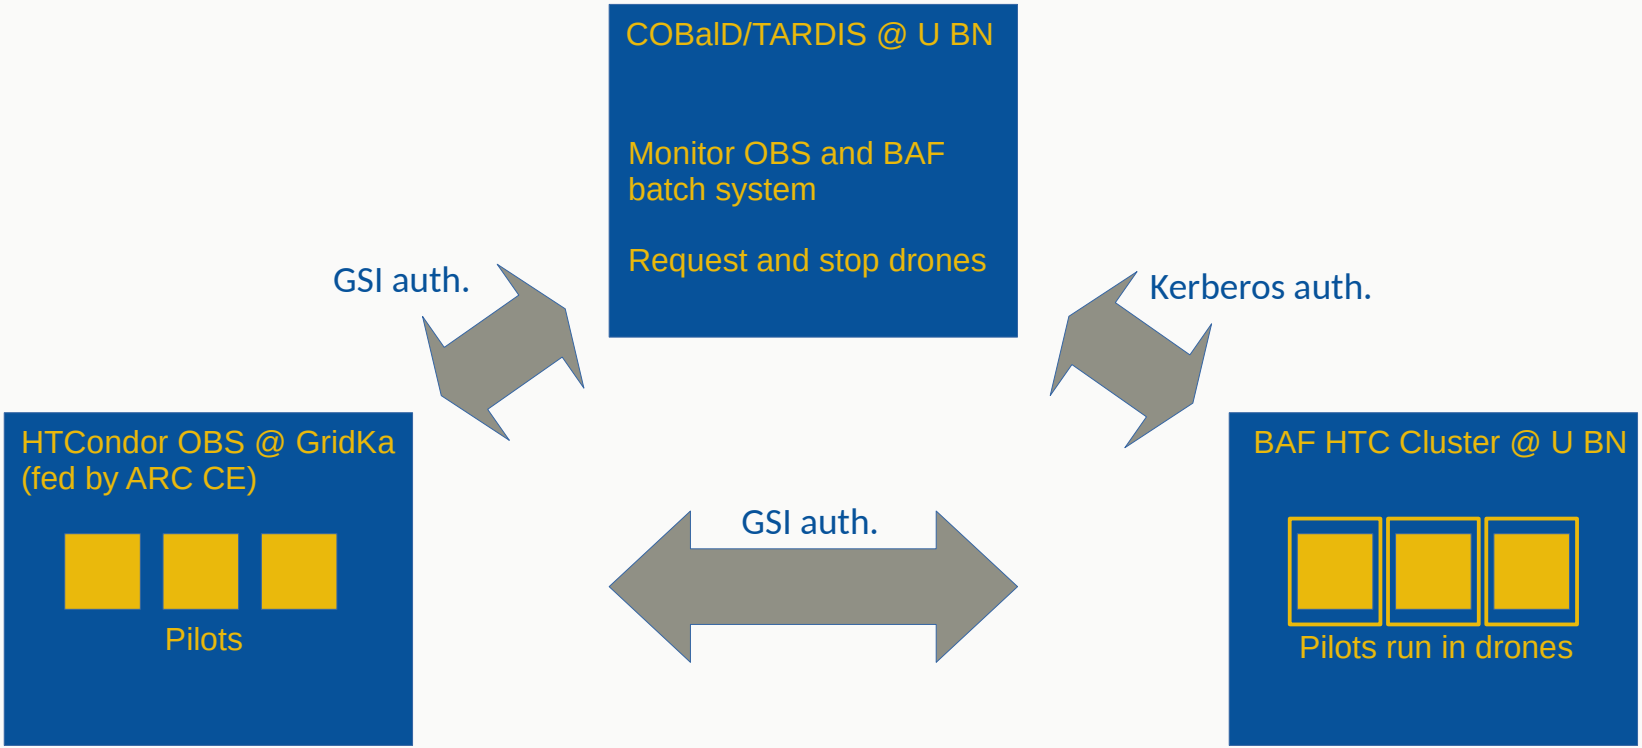

UNIVERSITÄT BONN

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

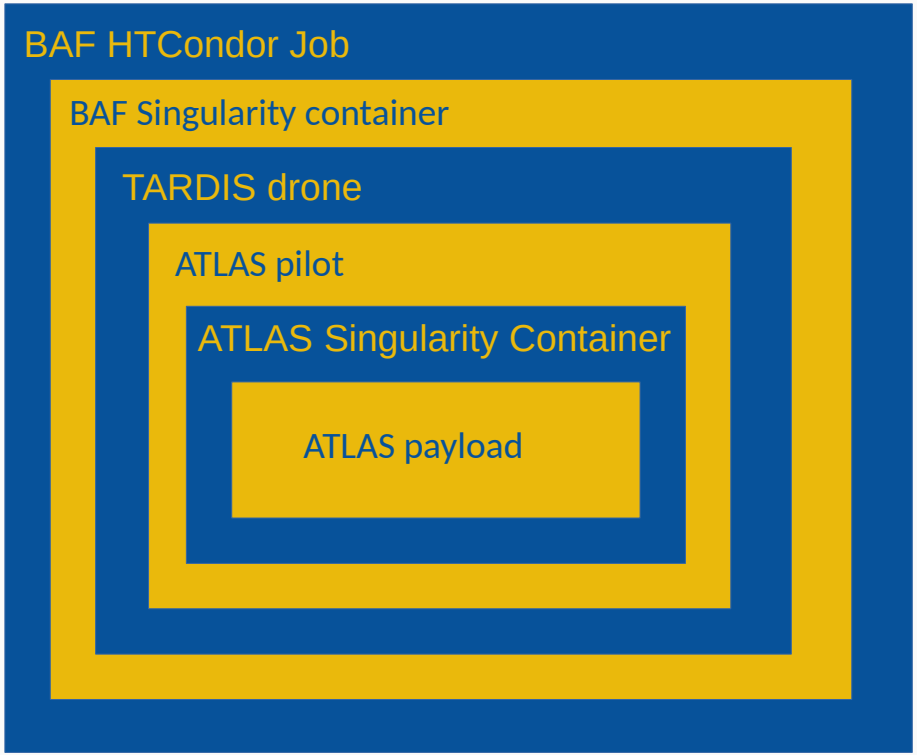

UNIVERSITÄT BONN

OPPORTUNISTIC USAGE OF BAF CLUSTER

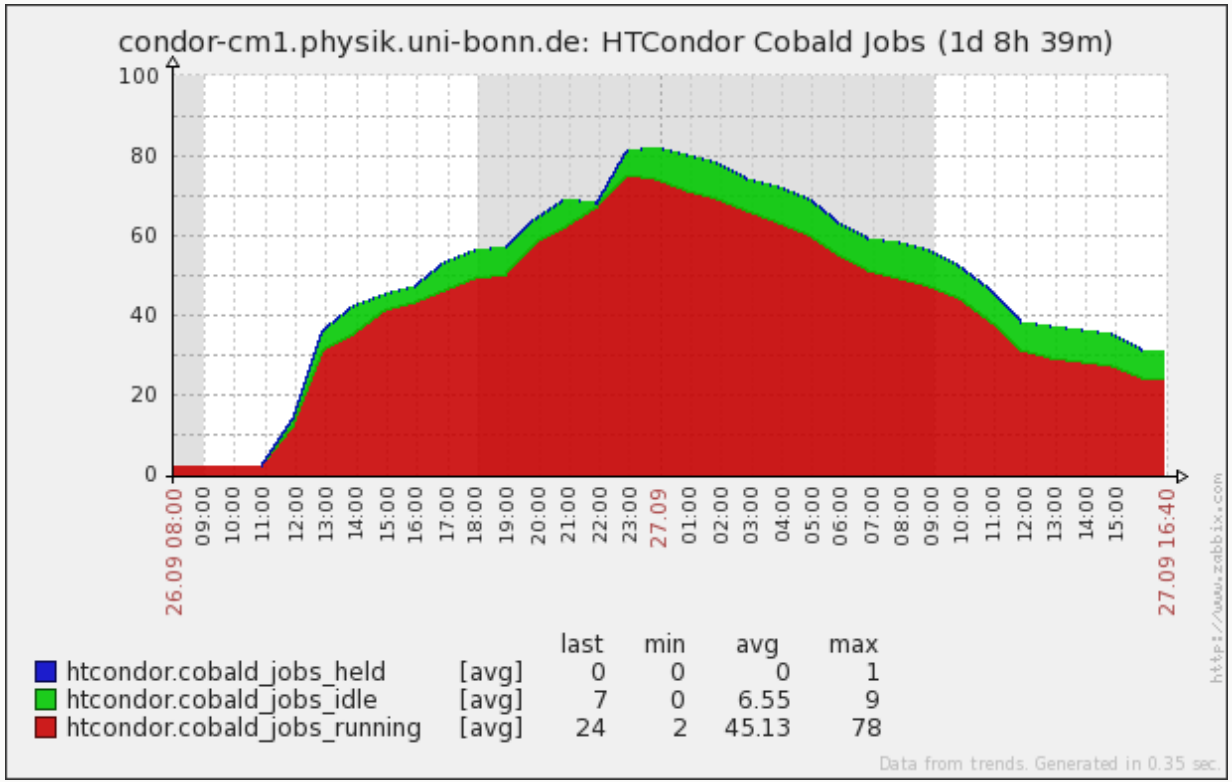

OPPORTUNISTIC USAGE OF BAF CLUSTER

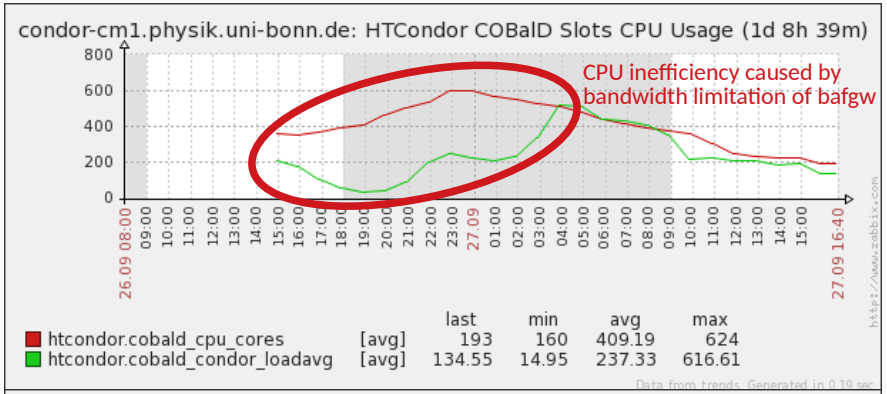

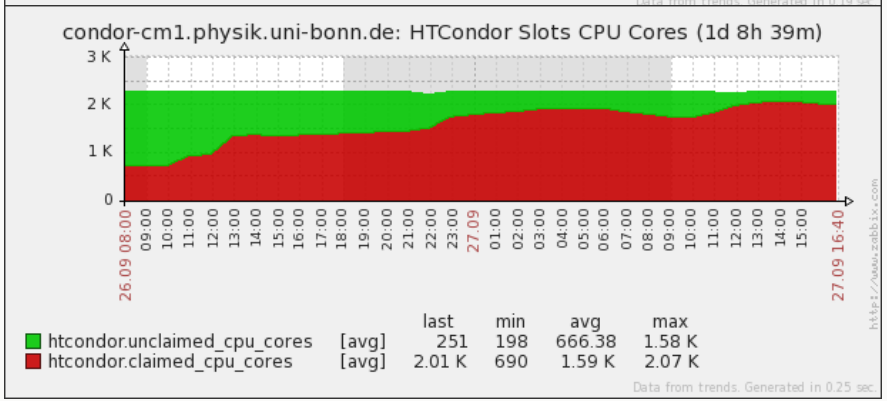

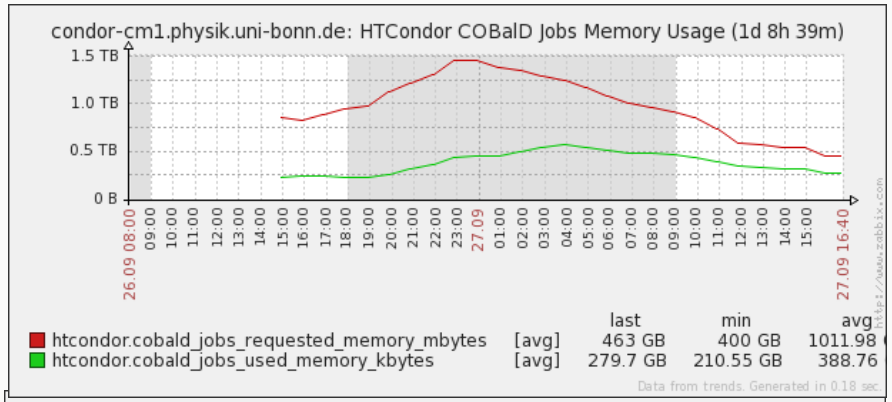

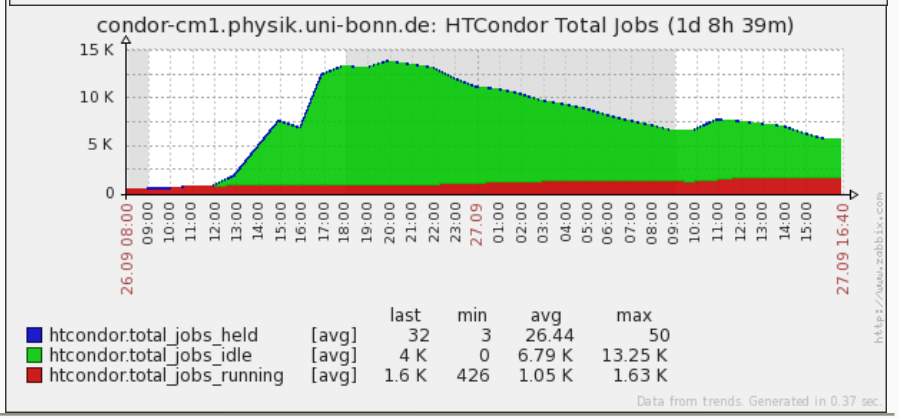

UNIVERSITÄT BONN

SUMMARY

- COBalD/TARDIS manages diverse opportunistic resources in a unifom way
- Service successfully deployed a U Bonn, management fully puppetised
- First ATLAS production jobs are running in Bonn
- To-Dos:
  - Need to understand/optimise the various knobs
  - Study impact on local users
  - Refactor puppet code to make it maximally re-usable for other sites
  - Exploit more resources

UNIVERSITÄT BONN

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