Karlsruhe Institute of Technology

Sustainability Efforts at the KIT Scientific Computing Center

Andreas Petzold

KIT Scientific Computing Center

- Provides basic IT services
- Operates large scale IT infrastructures for science
  - GridKa WLCG Tier-1
  - HoreKa HPC (HoreKa-Teal #9 on Green500)
  - Large Scale Data Facility
- Distributed over 3 sites, 4 data centers

Multi Pronged Approach

- Observability
  - All systems must report power consumption
- Optimize usage of compute resources
  - Energy efficient compute
  - Dynamic usage of (non-local) resources
- Data center design
  - Focus on heat reuse
- Generate green power locally

GridKa WLCG Tier-1

- **Optimized** High Throughput Computing Infrastructure
  - 280 worker nodes incl. 15 ARM hosts
  - 91PB disk space, 100 servers
  - Administrative nodes, network infrastructure
- **Power consumption steadily decreasing while pledged resources were increasing**

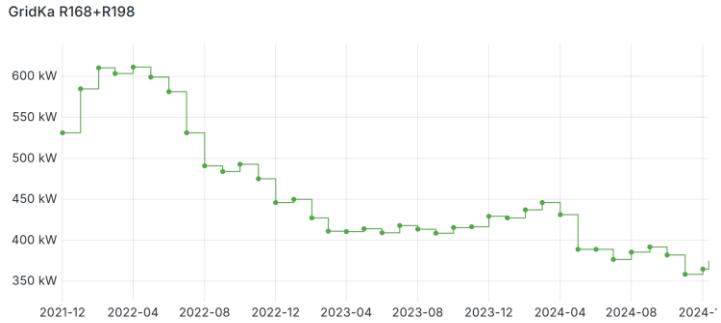

Observability

- There is never enough monitoring
- Power data recorded by SCC
  - Individual Servers: BMCs
  - PDUs: rack level data, UPS + line power
  - Bus bars: rack rows and rooms
  - Building infrastructure (depends on data center)
- Power data recorded by KIT facility management
  - Building infrastructure

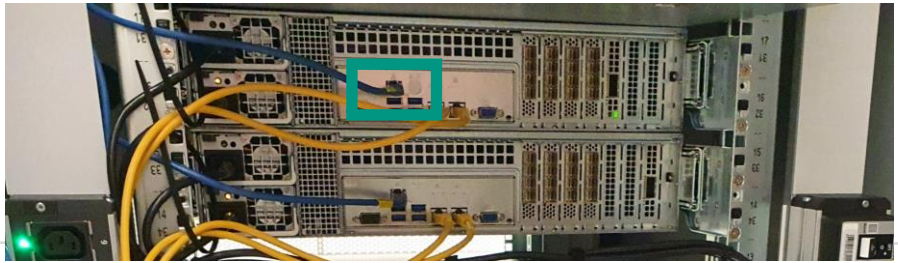

Observability

GridKa Storage Racks

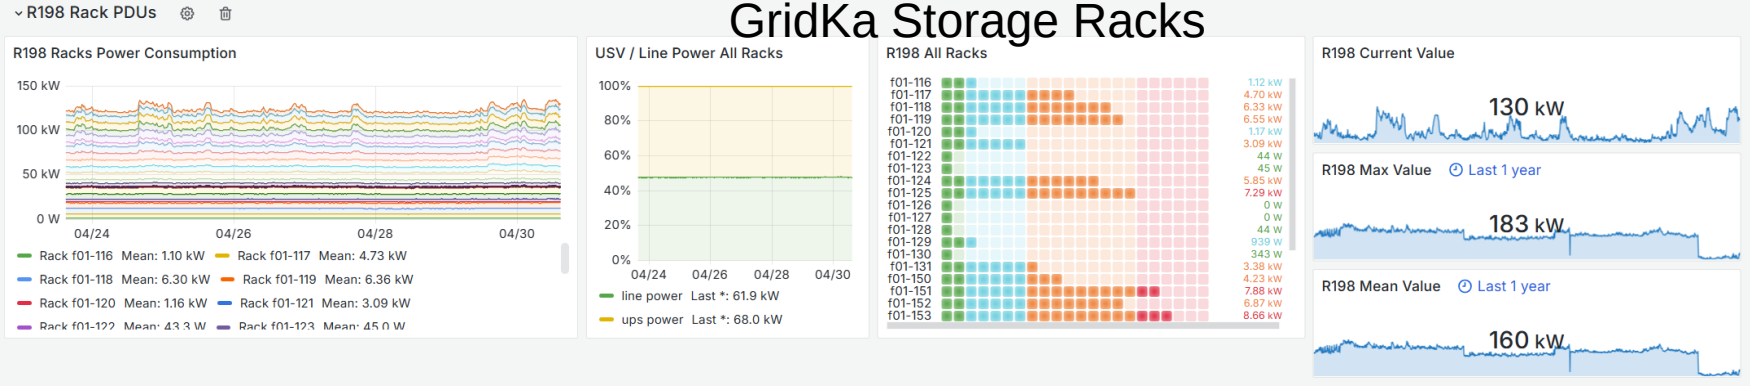

Building Infra (Chillers, Pumps etc.)

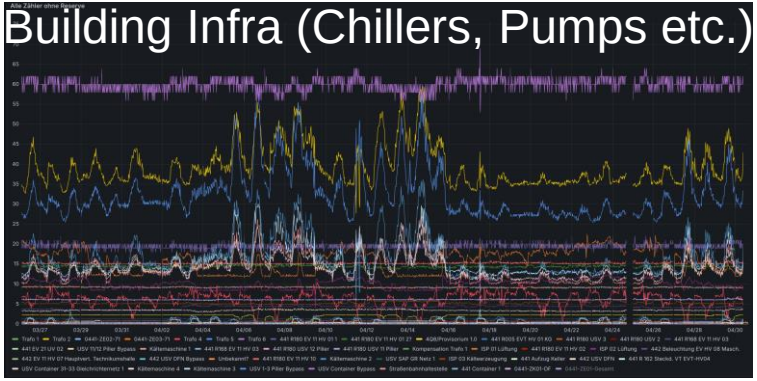

Energy Efficient Compute Resources

- TCO **including power consumption**
  - part of call for tenders since many years
  - Yes, we do measurements to check vendor claims ☺
- ARM compute nodes in production since mid 2024
  - 15 nodes with 3840 cores
  - Intensive benchmarking; **20% better HS/W**
- 56 GPUs available on the grid via GridKa
  - Still mainly used for local development for CMS and Belle II, but also ATLAS pilots

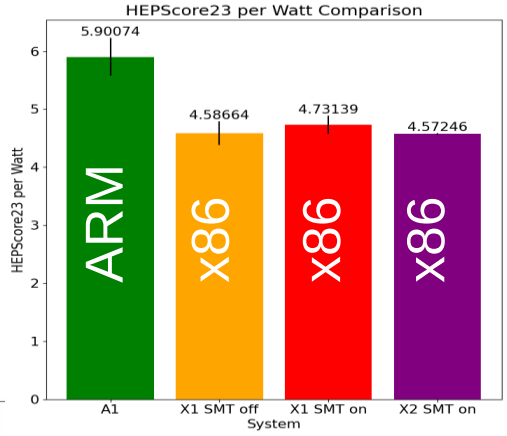

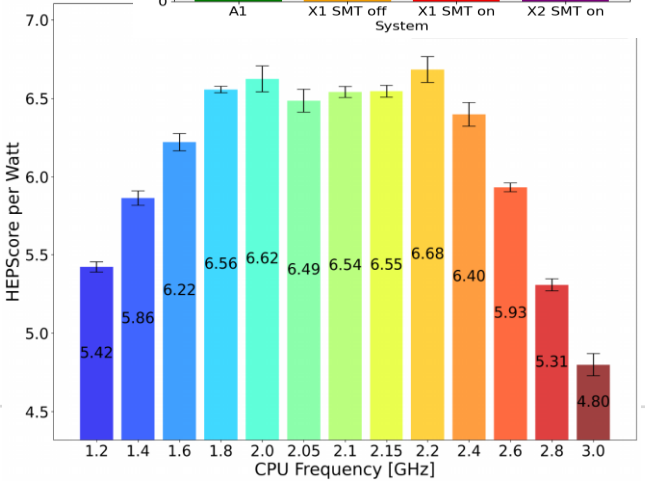

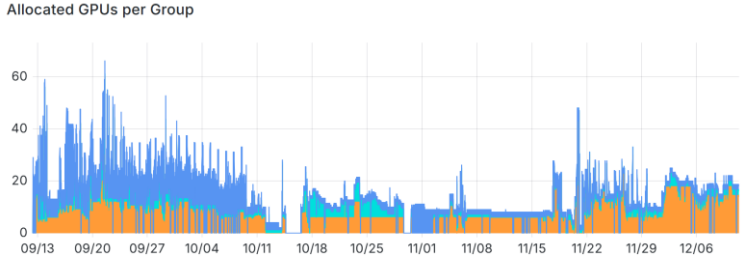

Opportunistic/Dynamic Resources

- Open-source software COBalD/TARDIS for **dynamic, transparent and on-demand integration of remote computing resources**
  - Demonstrated **production scale operation** during scale test together with HoreKa
  - **Central building block** of the Compute infrastructure in PUNCH4NFDI (DFG) and FIDIUM (BMBF)

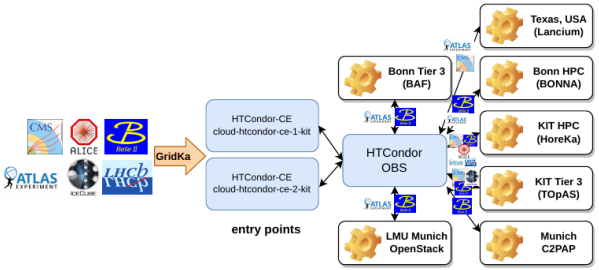

- (Non-local) compute **resources running on renewable energy**
  - Opportunistic resources used when “sun is shining/wind is blowing”
  - Potential use for grid stabilization
  - PoC 2023 with Lancium (US)
    - 100% wind and solar

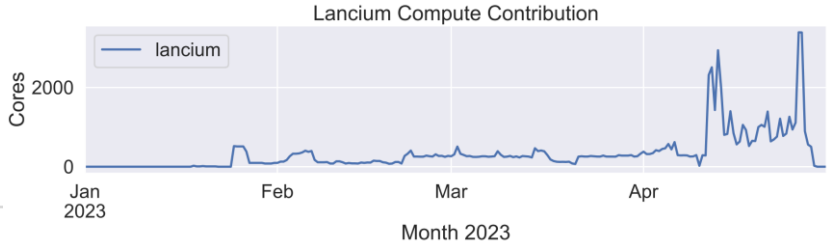

Data Centers

- Most modern building for HPC
  - 2.4MW capacity
  - Warm water cooling for compute; 500kW air cooling
  - PUE < 1.1; free cooling
    <10% heat reuse during winter
- Old GridKa data center
  - 1MW capacity
  - Cold water cooling
  - PUE ~1.3 with compression cooling
  - PUE <<1.3 with cooling provided by local powerplant
- Local combined power/heat/cooling plant (BHKW)
  - Produces cold water from own waste heat

New “Neighborhood” Computing Center

- New 10+MW data center
  - Integrated planning with KIT facility management
  - Focus on heat reuse
  - Located in “neighborhood” on campus close to heat consumers
- Elasticity of usable resources foreseen from the start
  - Availability of green power
  - Stabilization of local power grid
- Engineering firm contracted for concepts very recently

Local Power Generation

- Solar panels installed on GridKa data center and office buildings
- 500kWp ➔ GridKa (today) could be powered 100% by solar on a sunny day