Energy Crisis and WLCG

- Rod Walker, LMU,GDB 14th Sep

Crisis? What crisis?

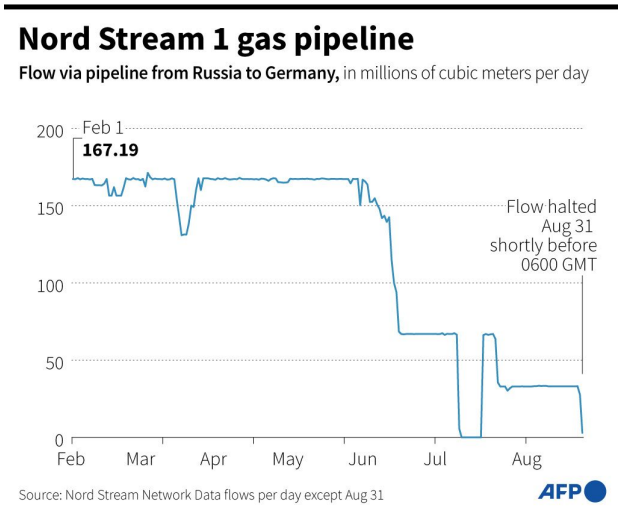

- Gas deliveries from Russia slowed then stopped
- Electricity price will triple (DE, Jan 2023)
  - power budgets do not
- Many WLCG sites worried about the power cost
  - turning off older servers is not enough
- CERN considering voluntary pause of accelerators, inc. LHC schedule
- EU policy ".. mandatory target for reducing electricity use at ***peak*** hours."
  - peak demand covered by generation from gas

A Case Study: DESY-ZN

- Factor 3 price increase for 50% electricity not yet bought for 2023
- No increase in funding for energy so far
- Computer centre uses 50% of site electricity
  - rest is offices, workshops, canteen … no savings
  - planning with 30% reduction for computer centre
    - worst-case 50%.
- Shutting down older, high W/HS06 servers
  - full impact unclear until electricity contract is made

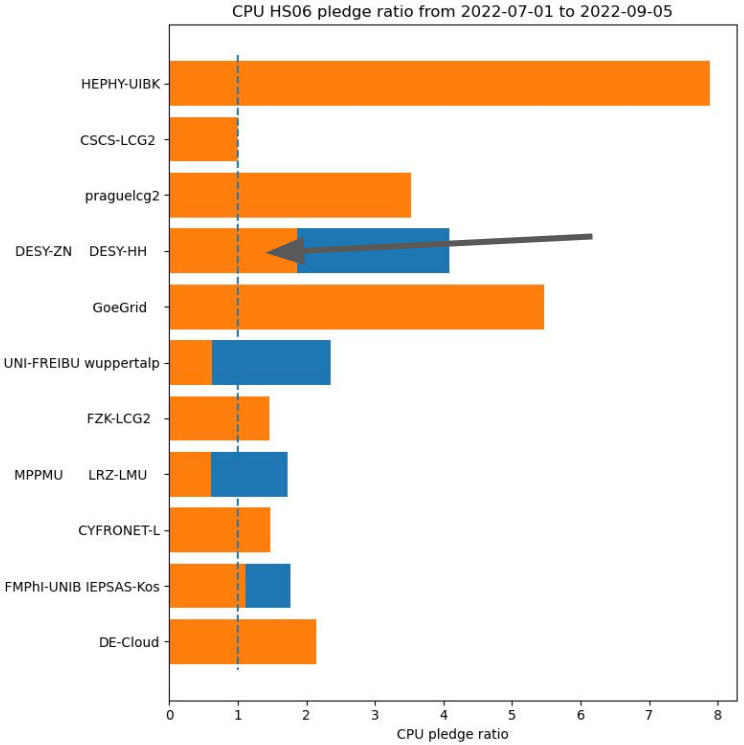

"We have a high risk in the energy budget and it is probable that we have to reduce our computing power in the grid (and the local computing clusters)" - Kai Leffhalm, DESY-ZN.

Mitigation for energy budget limitation

- Less compute means less power usage
  - power down servers
  - VOs unlikely to volunteer, but could ask for temporary reduction of pledge
    - many sites are over pledge and could cut to just meet pledge
- Outsource compute to region not in crisis, i.e. not using world market gas
  - such WLCG sites already fully used, so must be new and quick to realize
  - Lancium offers CPU at 1ct/core/hr - comparable to just the power in DE
  - Extend lifetime of well-placed resource, paying just the power
    - needs solidarity since any hint of FTE cost would sink the economics
- Positive mitigation: solar panels on computer centre roof!
  - invest the power budget

LRZ-LMU 28 core 300W,
cooling factor=1.3
@25ct/kWh 0.4ct/core/hr
@75ct/kWh **1.2ct/**core/hr

Electricity crisis is in peak load, not base

- Only need to reduce power usage for a few hrs per day: Load shaping
  - from a blackout avoidance perspective, and saving gas, if not budget (lack of variable tarif)
    - Germans react to gas crisis by buying electric heaters!
    - voluntary grid-balancing, that is worth real money
    - should be fast and reversible without job loss
      - suspending processes allows cpu to idle, at low frequency

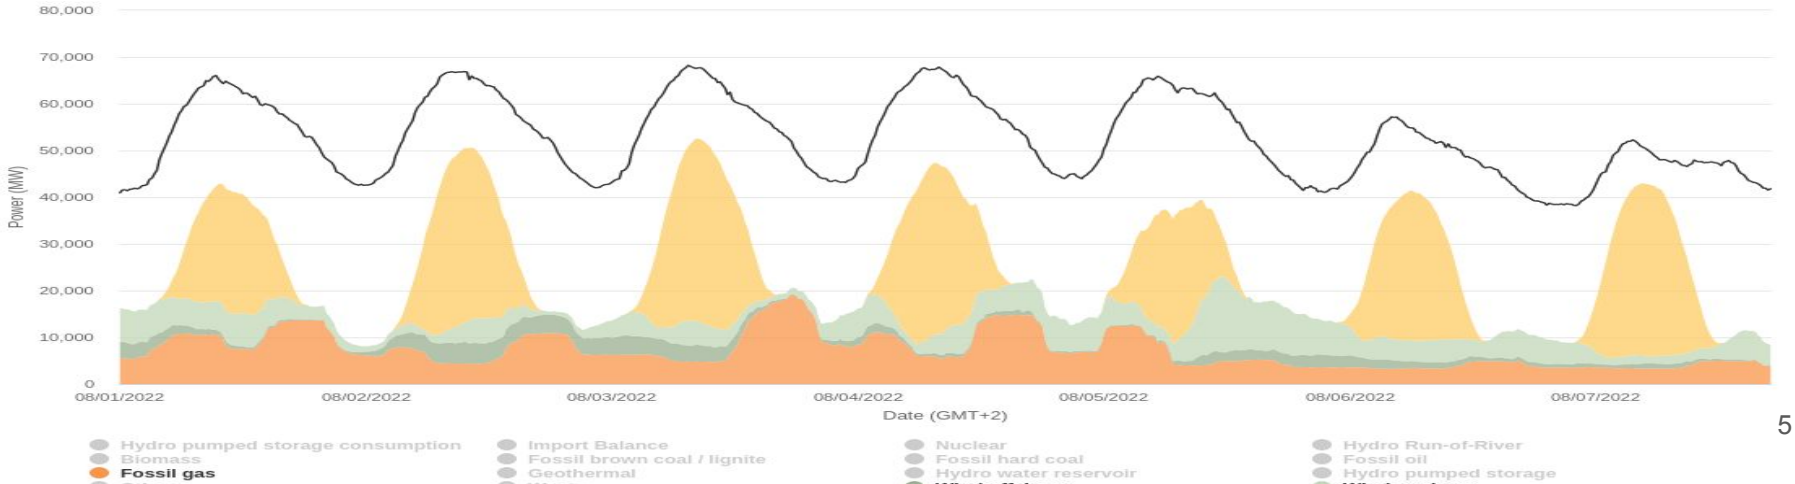

Load shaping: Proof of principle

- LRZ-LMU ATLAS T2, with Slurm

- Create reservation to drain 1hr before suspending slurm jobs, resume after 1hr
  - suspend sends SIGSTOP
  - draining is slow and leaves cpu idle
  - contention on resume due to new job starts
- Reservation only at suspend time(to stop new jobs)
  - no wasted cpu or contention after resume
- Blade enclosure power measurement
  - Loaded: 265W
  - Suspended: 96W (36%)
- Observe CPU frequency governed to minimum
  - Just do that? cpupower frequency-set -g powersave
    - 40% power usage, but needs more tests

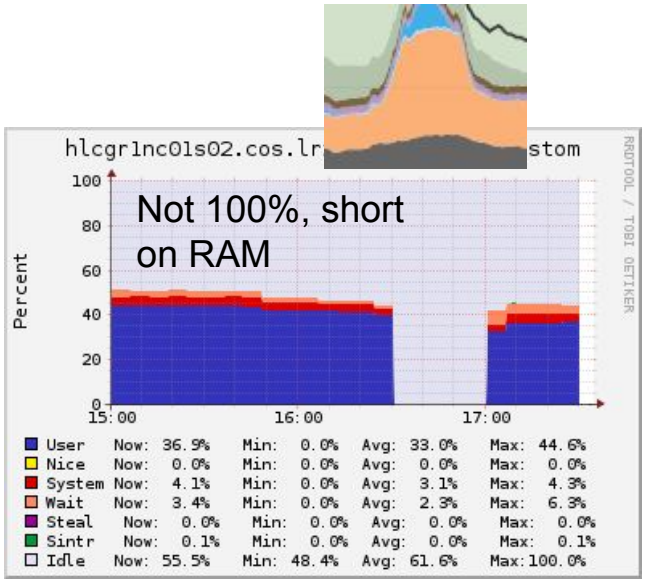

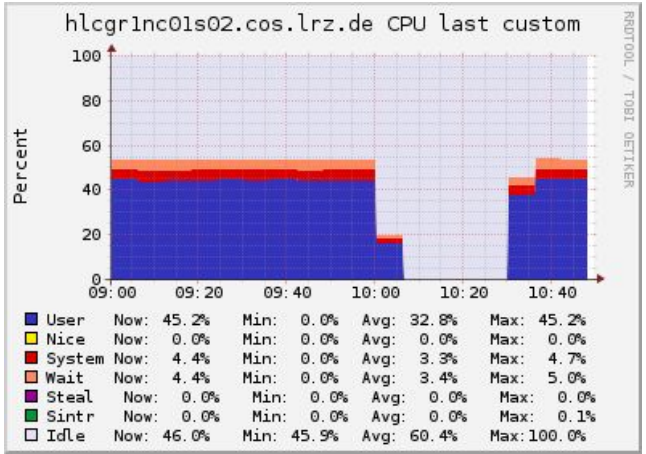

Load shaping: proposal

- Proposing new WLCG component
  - collection of electricity prod forecast: various sources, depending on country/region
    - convert to standard timeseries.
    - library of standard actions
      - e.g. Slurm, condor: reservations and suspend jobs
      - cpu frequency scaling(cluster via ssh/puppet/ansible)
  - WLCG community could develop, test and roll this out very quickly
- I started a python package
  - https://gitlab.cern.ch/walkerr/computethrottler
  - shared but wouldn’t insult real developers, who can start from scratch.

Conclusions

- Energy budget constraints could lead to temporary reduction in resources
  - collate to have some overview of mothballed HS06
- EU call to arms! Any interest in developing/testing load shaping?
  - estimate power usage vs time, to show modulations and saving
  - tag participating sites, but less worried about strict accounting
- Possible session in Lancaster workshop: Facilities block
  - review progress, hear more ideas and discuss steps
    - target should be January 2023
- Clear overlap with longer term sustainability topics - new area for WLCG?
  - store dynamic gCO2 per kWh/HS06s/hr, regional or per site
  - enable VOs to set job eco-profile, prefer low carbon sites/times