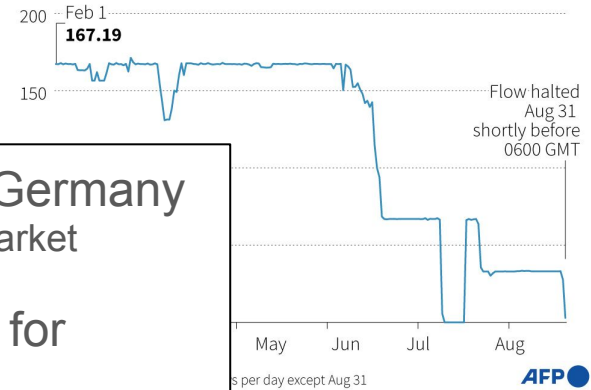


# Scale back compute to save gas

- Rod Walker, LMU Munich,  
Sustainable HEP 5th Sep '22

## Nord Stream 1 gas pipeline

Flow via pipeline from Russia to Germany, in millions of cubic meters per day



# Motivation



- Very particular problem with fossil gas availability in Germany
  - and other countries due to European interconnect/common market
  - not just Europe due to LPG and world market
- Heating & Industry largest consumers, but also used for electricity production
  - proportion gas[1]: Netherlands, UK, US ~40%, Japan, Portugal, Spain ~30%
- Impact of ATLAS computing is small
  - German HEP: 1-2MW from 60GW consumption, but not the point
- demonstrates demand-side balancing based on regional signals
  - trivial to use other signal: variable electricity price or gCO<sub>2</sub>/kWh
  - outreach potential, leading to measures in other areas

# What can HEP Computing do?

- Default has been to keep resources full, even finding low-priority work
  - made sense to get most from hardware & improve statistics, but never very eco-friendly
- Outside urgent periods, can choose where and when to burn cpu
  - could prefer 'cheap' resources or times of day
    - stopping job submission only works for dedicated sites, and only slowly
      - some VOs have some dedicated sites, but also local users
- Quickly clear that energy saving measures should be at the resource level
  - full control of policy and actions
- HEP/WLCG can provide the common tools

# German electricity production

www.energy-charts.info



- Hydro pumped storage consumption
- Biomass
- Fossil gas
- Others
- Solar
- Renewable share of load
- Import Balance
- Fossil brown coal / lignite
- Geothermal
- Waste
- Load
- Nuclear
- Fossil hard coal
- Hydro water reservoir
- Wind offshore
- Residual load
- Hydro Run-of-River
- Fossil oil
- Hydro pumped storage
- Wind onshore
- Renewable share of generation

# Ideal system

- Variable electricity price per kWh paid by compute centre
  - based on CO2 price, or other motivation, e.g. save gas.
- HEP VO sees variable compute price, whether \$ or other consumable
  - maybe a weight on the fairshare
- VO can choose when and where to run to conserve \$ or fairshare
- If all users restrict usage, site can drain nodes to sleep/suspend/off
  - not to forget factor ~2 from cooling
- Many parts of this are missing, starting with a variable kWh price
  - So what can we do quickly?

# Approximation to ideal system



- Assume VOs ok with up to 6hr total pause per day
  - max N pauses of max M minutes
  - removes the pricing/priority component to give sites free choice when to reduce power
- Sites opt-in to modulate power consumption according to external signal
- When to pause?
  - Merit-Order scheme means expensive gas only used when necessary
    - less electricity consumption means less gas burned
  - Not all gas production equal in cost or efficiency
    - some gas generated electricity feeds 'waste' heat into district heating
  - Reduce consumption during gas usage peaks

# Information sources

- Must be a forecast, based on
  - projected load, weather(sun,wind), maintenance schedules
  - exists and public for production planning, network stability, and spot market
- Projected gas usage and gCO2/kWh
  - <https://github.com/EnergieID/entsoe-py> covers European countries
    - API `fetch_generation_forecast('DE')`
      - gas component in GW
      - weight production type by gCO2/kWh to get average, e.g. solar=0.
  - Or conveniently for test: `curl "https://api.awattar.de/v1/marketdata"`

# Market spot price vs Production/Load


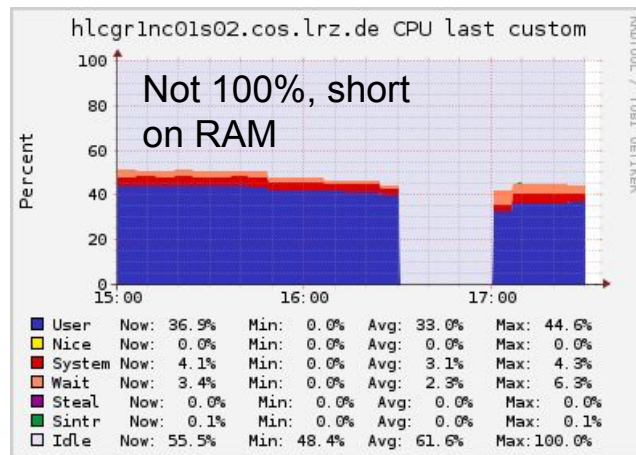
< 03.09.2022 >





# Proof of principle

- LRZ-LMU ATLAS T2, with Slurm
- Create reservation to drain single node 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
    - tried it but no power measurement yet.



# Proposal

- At least ATLAS jobs would survive SIGSTOP up to 2hrs
  - Longer if heartbeat check loosened
  - If scaling frequency down has same power saving, then that is preferable
    - but needs root action on each node, rather than only slurm privilege
- Proposing new simple WLCG component
  - collection of electricity prod data: various sources, depending on country/region
    - convert to standard timeseries. Maybe run centrally, to publish and plot.
  - Range of standard actions (contributed)
    - Slurm, condor: reservations and suspend
    - Cpu frequency scaling(cluster via ssh/puppet/ansible)
      - or stand-alone package WN, desktop, laptop(run from battery when plugged-in)
  - WLCG community could develop, test and roll this out very quickly
- I started a python package
  - can share but wouldn't insult real developers, who would design from scratch.

# Summary & Outreach

- Can modulate power usage down by ~64%
  - reduced cpu cycles for VOs, but can choose by how much
- HEP computing can have only marginal effect
  - certainly worth doing because easy, also for potential outreach effect
- Ask general public to save electricity during particular hours
  - automatic systems are good, but manual actions also save power
    - delay washing machine or dishwasher. When to charge EV
  - people **would** act - if this is our only war sacrifice
- Info in weather forecast after evening news broadcast
  - similar to pollen/air-quality/UV/Ozone extras, shown when relevant
  - very correlated with sun and wind prognosis
- Interest from resources/developers? Take to WLCG?