

LHConCray project @ CSCS

■ Consolidation project to run LHC jobs on Piz Daint


- Partners: CSCS, CHIPP (ATLAS, CMS, LHCb)
- Started ~2 year ago
- Started production in April 2017: 25 nodes/1600 cores (ATLAS:CMS:LHCb - 40:40:20)
- **The goal is to run ALL VO jobs without changes to the experiments' workflows**

■ Normal workflow:

- Jobs submitted via ARC, SLURM LRMS
- Running in containers (Shifter)
- CVMFS Native on Cray nodes
- **Plugs transparently in to the experiments' WMSs**

■ Roadmap

- Measure performance in production environment, produce a cost study (until Dec. 2017)
- Decision due: continue or revert to invest on Phoenix (x86_64 T2 cluster @CSCS)



LHConCray project - Concept

- **The growth of resources will be provided using shared resources at CSCS**
 - (e.g. yearly increase: compute 30% and storage 20%)
- **In general move from “HW investment mode” to “paying for services”**
 - **Compute**
 - Dedicated number of nodes
 - Resources will be charged whether they are being used or not, including downtimes
 - Compute unit: HS06
 - No opportunistic usage (*some underlying negotiation in place*)
 - **Storage**
 - Allocation of disk space via the SAN (as today)
- **CHIPP funding will be used to pay operational costs including HW investments, maintenance and personnel. Costs for electricity & cooling are covered by ETHz**

LHConCray project - Allocation and expectations

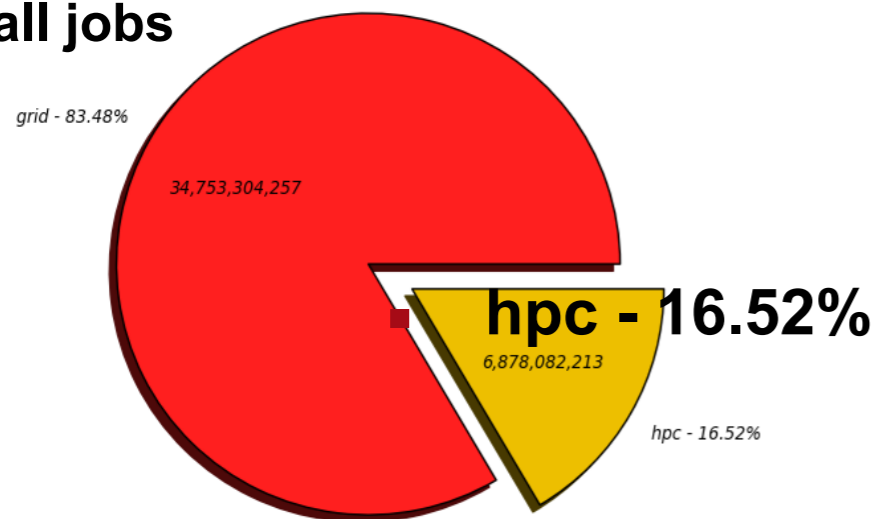
- **APRIL 2017: 25 Piz Daint nodes in production (CSCS-LCG2-HPC* PanDA queues)**
 - Use 64 out of 72 cores / node, 128 GB RAM / node
 - Total: 1600 cores, approx 20500 HS06 - *24% of the total installed capacity*
 - Recently switched to use 68 cores / node (1700 total), total HS06 is ~the same
- **DECEMBER 2017: expect decision**
 - Move to Piz Daint or revert back to the x86_64 cluster (Phoenix)
 - Based on experience and performance during 2017
 - Use as baseline Phoenix's performance (with care)
- **if [DAIN] ; then**
 - Phase out Phoenix at the end of 2018
 - Predicted compute HS06 growth rate 30% year-on-year, assuming:
 - flat funding
 - same performance of Piz Daint and Phoenix normalised to the HS06
 - no opportunistic usage (*some negotiations taking place at higher levels*)
- **else**
 - Revert to investing on Phoenix in 2018
 - Predicted compute HS06 growth rate 20% year-on-year, assuming:
 - flat funding

LHConCray project - Usage overview



Wall Clock consumption All Jobs in seconds (Sum: 41,631,386,470)

WC all jobs



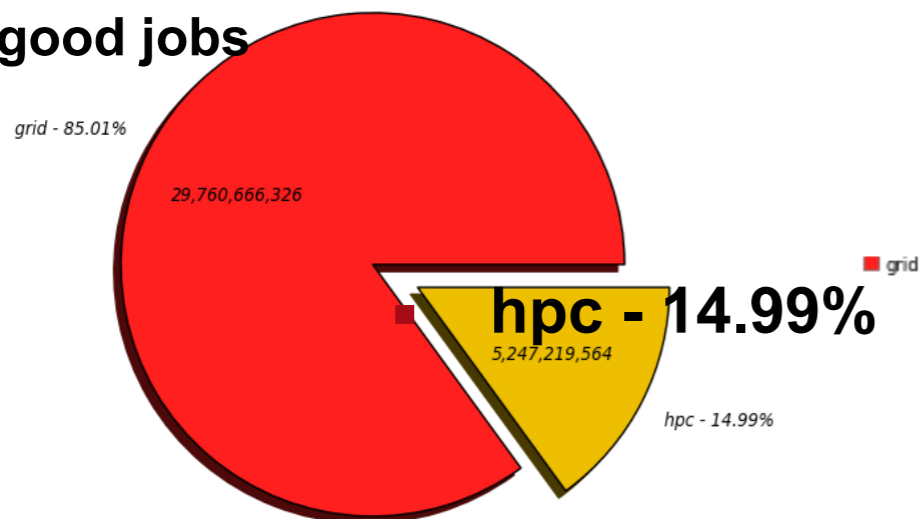
<http://cern.ch/go/sC8v>

■ grid - 83.48% (34,753,304,257) ■ hpc - 16.52% (6,878,082,213)



Wall Clock consumption Good Jobs in seconds (Sum: 35,007,885,890)

WC good jobs



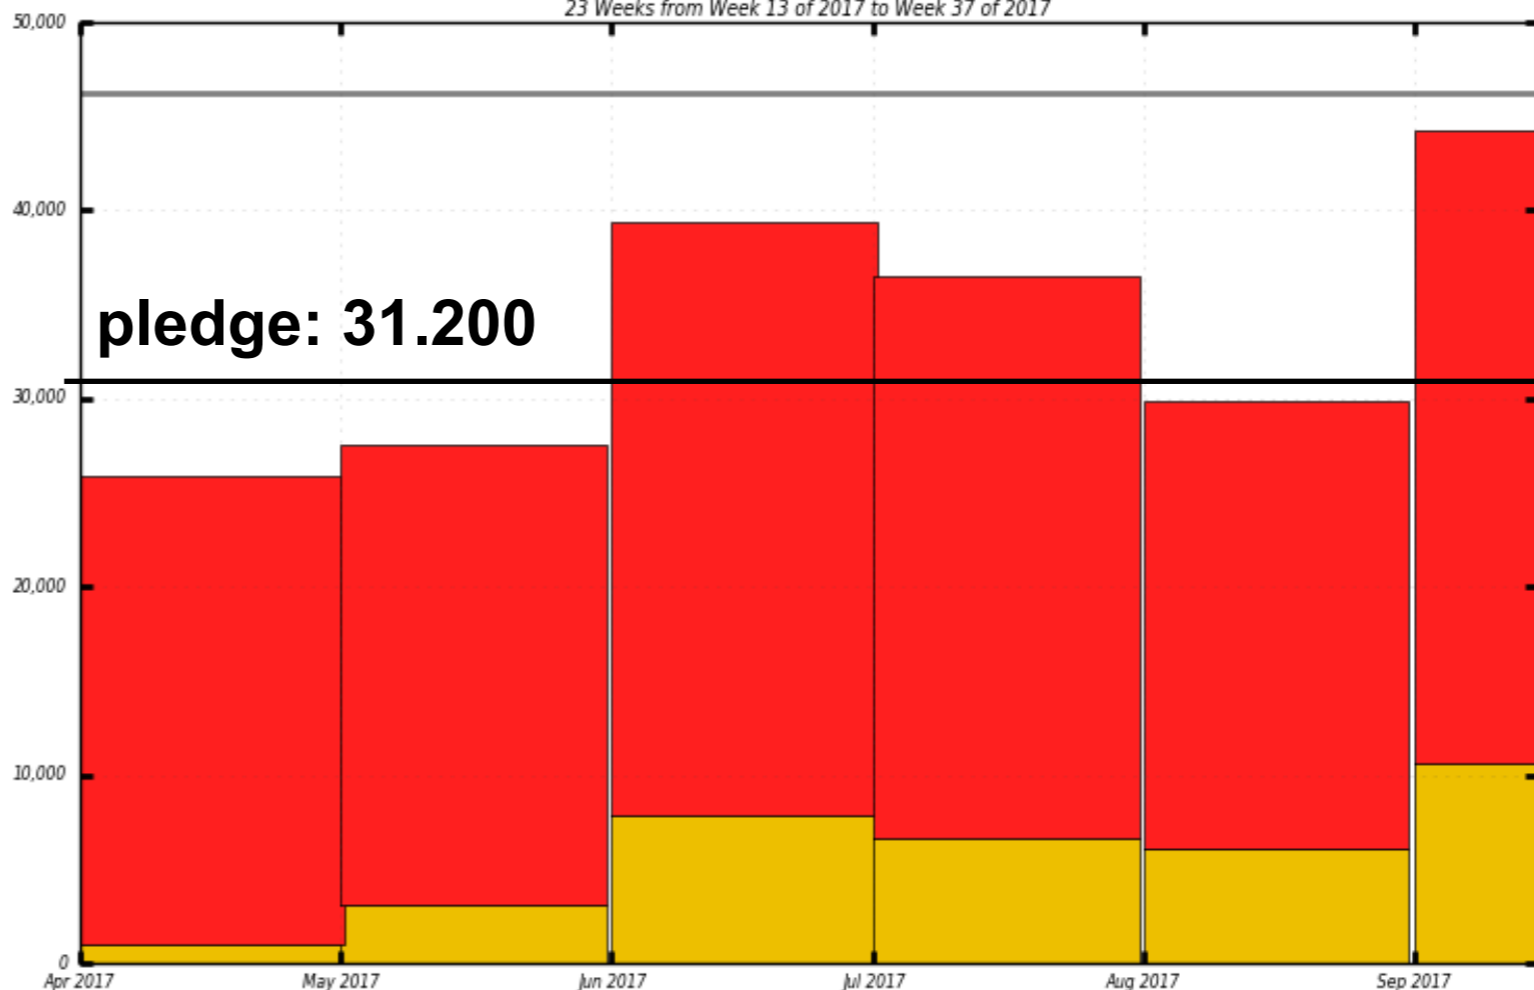
<http://cern.ch/go/sC8v>

■ grid - 85.01% (29,760,666,326) ■ hpc - 14.99% (5,247,219,564)



WallClock HEPSPROC6

23 Weeks from Week 13 of 2017 to Week 37 of 2017



<http://cern.ch/go/FkV7>

Maximum: 44,232 , Minimum: 0.00 , Average: 29,057 , Current: 44,232

- CPU/WC eff good jobs:
 - grid: 0.78
 - hpc: 0.85