

- ▶ Performance overview July 2018 - June 2019



ATLAS T2 VO REPORT


Gianfranco Sciacca

AEC - Laboratory for High Energy Physics, University of Bern, Switzerland

CHIPP-CSCS face 2 face - 21 January 2019

WallClock HS06 - CH-CHIPP-CSCS Federation

- **CSCS-LCG2 HS06 ATLAS *installed (*)*: 43970** (66%) - ***pledged: 36000*** (67%) (2018)
of total
- **UNIBE-LHEP HS06 ATLAS *installed (**)*: 22000** (34%) - ***pledged 18000*** (33%) (2018)
of total



total installed

total pledged

Pledge Increase over 2017:

*CSCS +15%
(supposed to be +27%)*

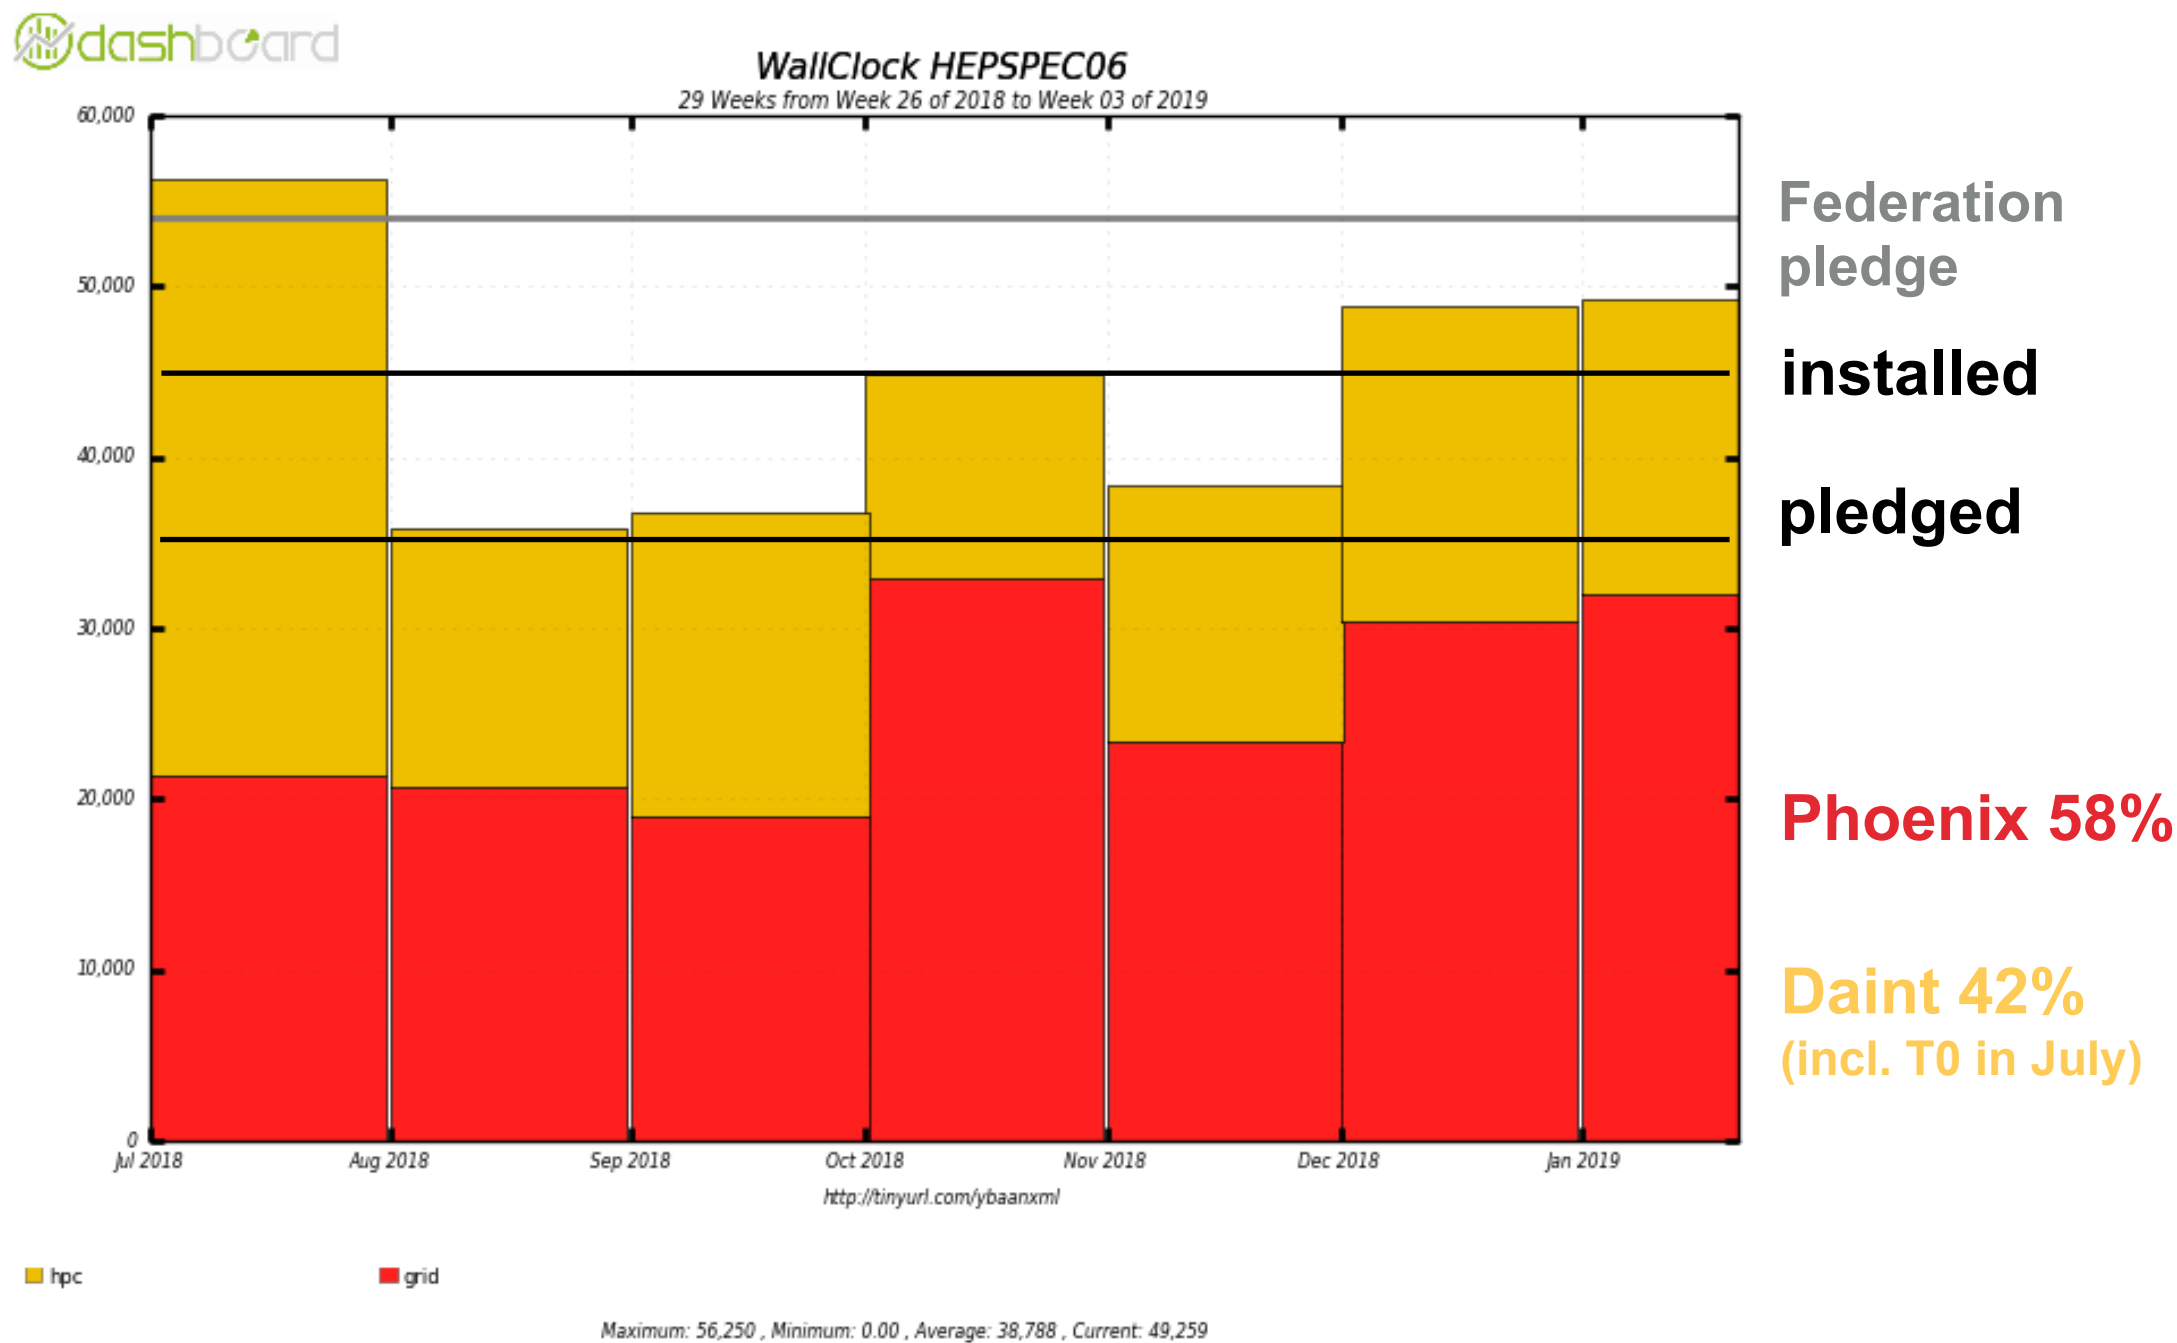
Bern +20%

(*) - ATLAS share

(**) - also serving t2k.org and fermilab/uboone
- some opportunistic usage on Ubelix

WallClock HS06 - CSCS-LCG2

- **CSCS-LCG2 HS06 ATLAS installed: 43970 - pledged: 36000** (2018)



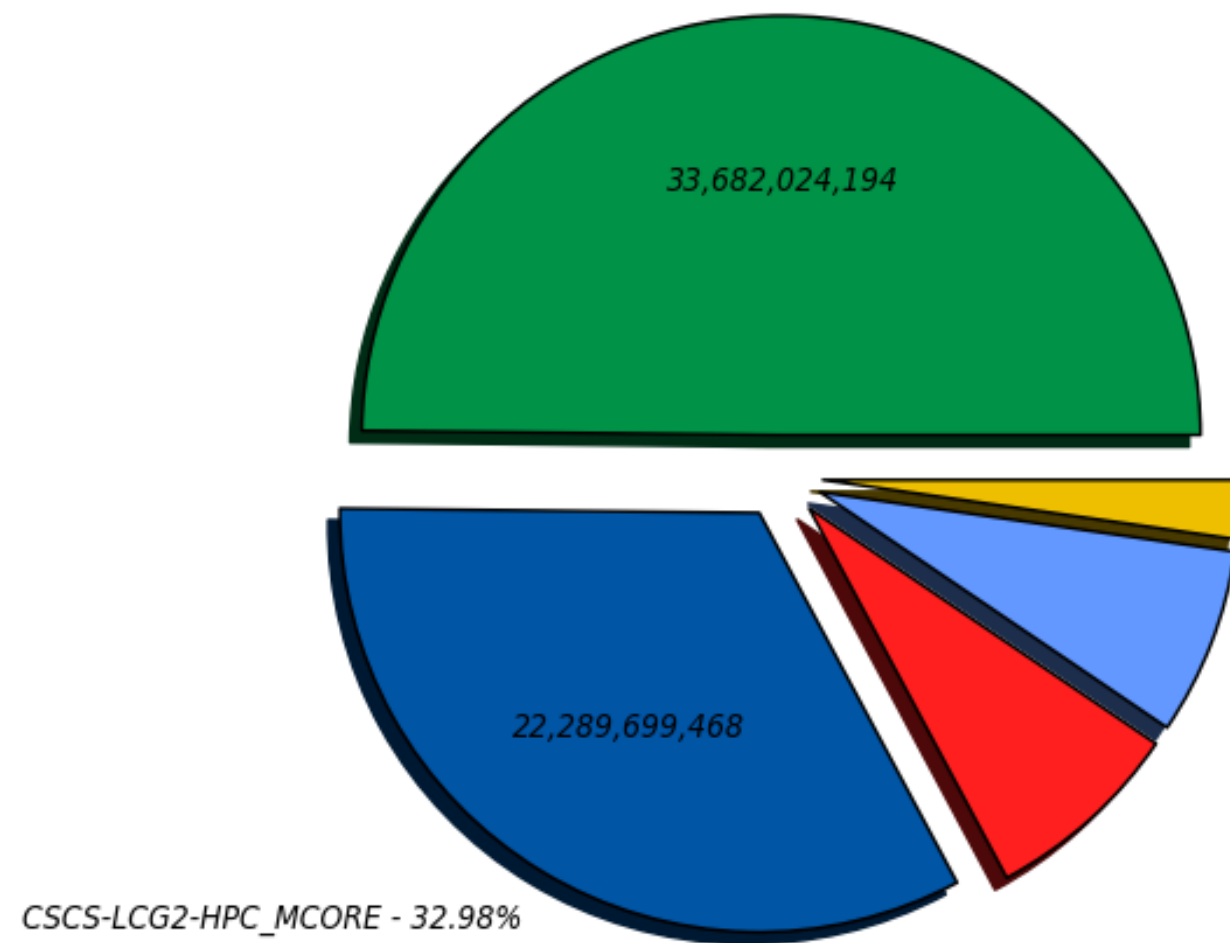
(*) - ATLAS share

WallClock HS06 - CSCS-LCG2

- **CSCS-LCG2 HS06 ATLAS installed: 43970 - pledged: 36000** (2018). (*)



Wall Clock consumption All Jobs in seconds (Sum: 67,585,113,137)
CSCS-LCG2_MCORE - 49.84%



HPC capacity: 31% of total

HPC delivered: 35% of total
(after subtracting the T0 activity: 7.22%)

<http://tinyurl.com/yb9cbzjh>


■ CSCS-LCG2_MCORE - 49.84% (33,682,024,194)
■ ANALY_CSCS - 7.69% (5,197,495,162)
■ ANALY_CSCS-HPC - 2.27% (1,533,301,417)

■ CSCS-LCG2-HPC_MCORE - 32.98% (22,289,699,468)
■ CSCS-PIZDAINT - 7.22% (4,882,592,896)

(*) - ATLAS share

WallClock HS06 - UNIBE-LHEP



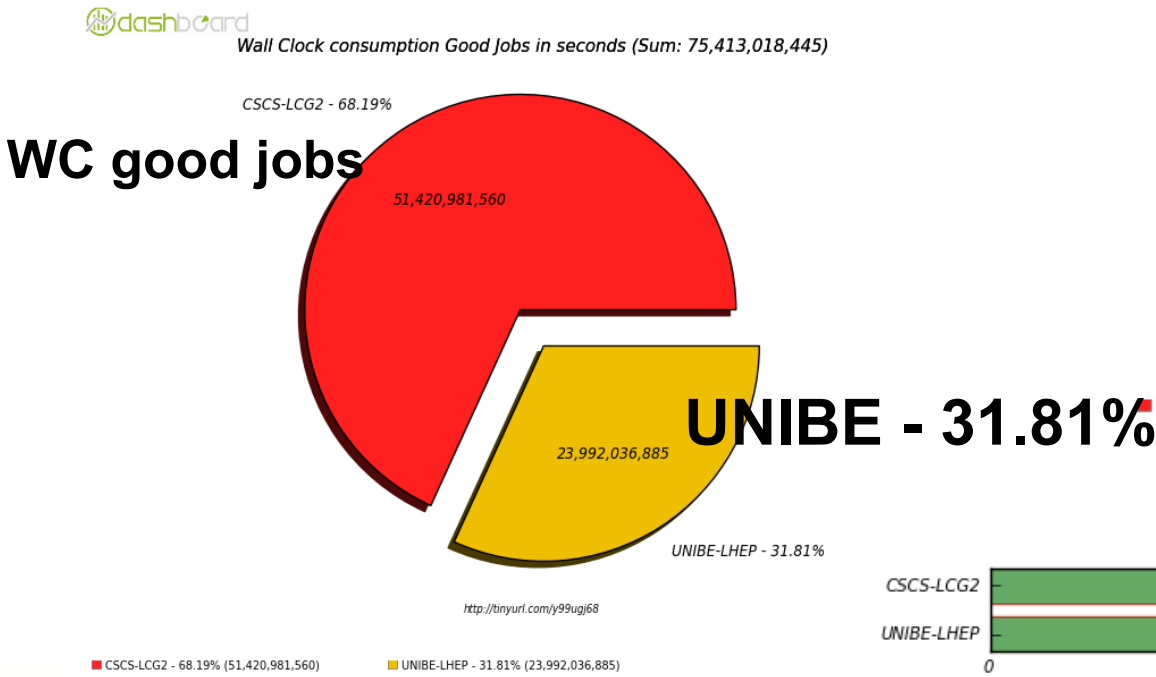
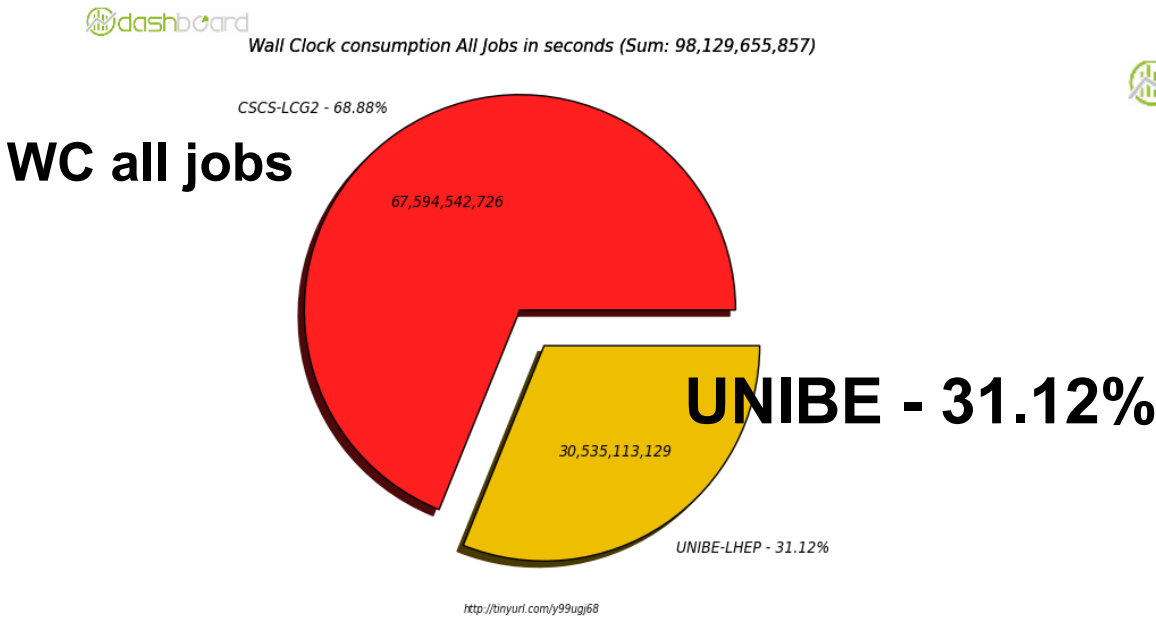
- **UNIBE-LHEP HS06 ATLAS *installed* (**): 22000 - pledged: 18000 (2018)**



(**) - also serving t2k.org and fermilab/uboone
- some opportunistic usage on Ubelix

Relative shares

- **UNIBE-LHEP installed** estimated in about **37% of the total installed capacity**

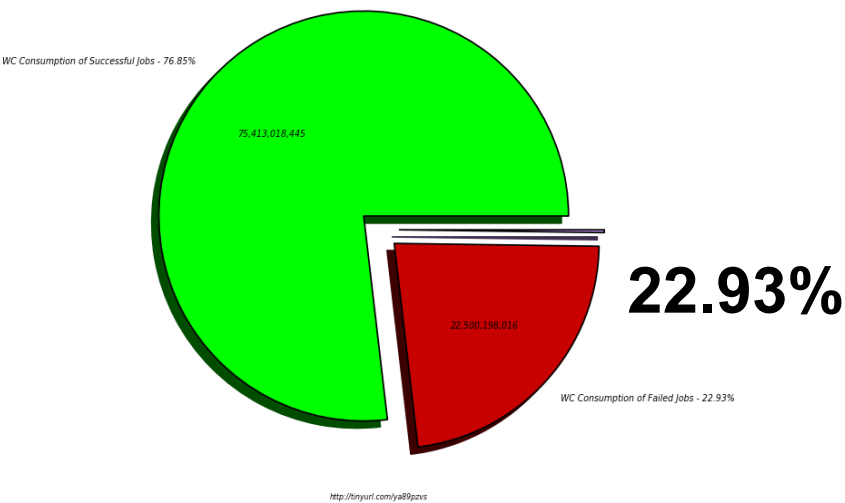


Success vs fail WallClock efficiency

- **CSCS-LCG2: 73% (69% PHOENIX, 87% Daint)**
- **UNIBE-LHEP: 79%**


dashboard

WC Consumption for Successful and Failed Jobs (Sum: 98,129,655,857)



dashboard

WallClock Consumption for Successful and Failed Jobs
29 Weeks from Week 26 of 2018 to Week 03 of 2019



WallClock Consumption of Successful Jobs
WallClock Consumption of Failed Jobs

dashboard

WallClock Efficiency over time based on success/all accomplished jobs
29 Weeks from Week 26 of 2018 to Week 03 of 2019



CSCS-LCG2 (73.35) UNIBE-LHEP (79.20)

Total: 109.44 , Average Rate: 0.00 /s

CPU / WallClock efficiency

- **CSCS-LCG2: 85%**
- **UNIBE-LHEP: 83%**

- **CSCS-LCG2: 71%**
- **UNIBE-LHEP: 76%**

dashboard

Efficiency Good Jobs

29 Weeks from Week 26 of 2018 to Week 03 of 2019



dashboard

Efficiency All Jobs

29 Weeks from Week 26 of 2018 to Week 03 of 2019



■ CSCS-LCG2 (0.85)

■ UNIBE-LHEP (0.83)

Total: 1.66 , Average Rate: 0.00 /s

■ CSCS-LCG2 (0.71)

■ UNIBE-LHEP (0.76)

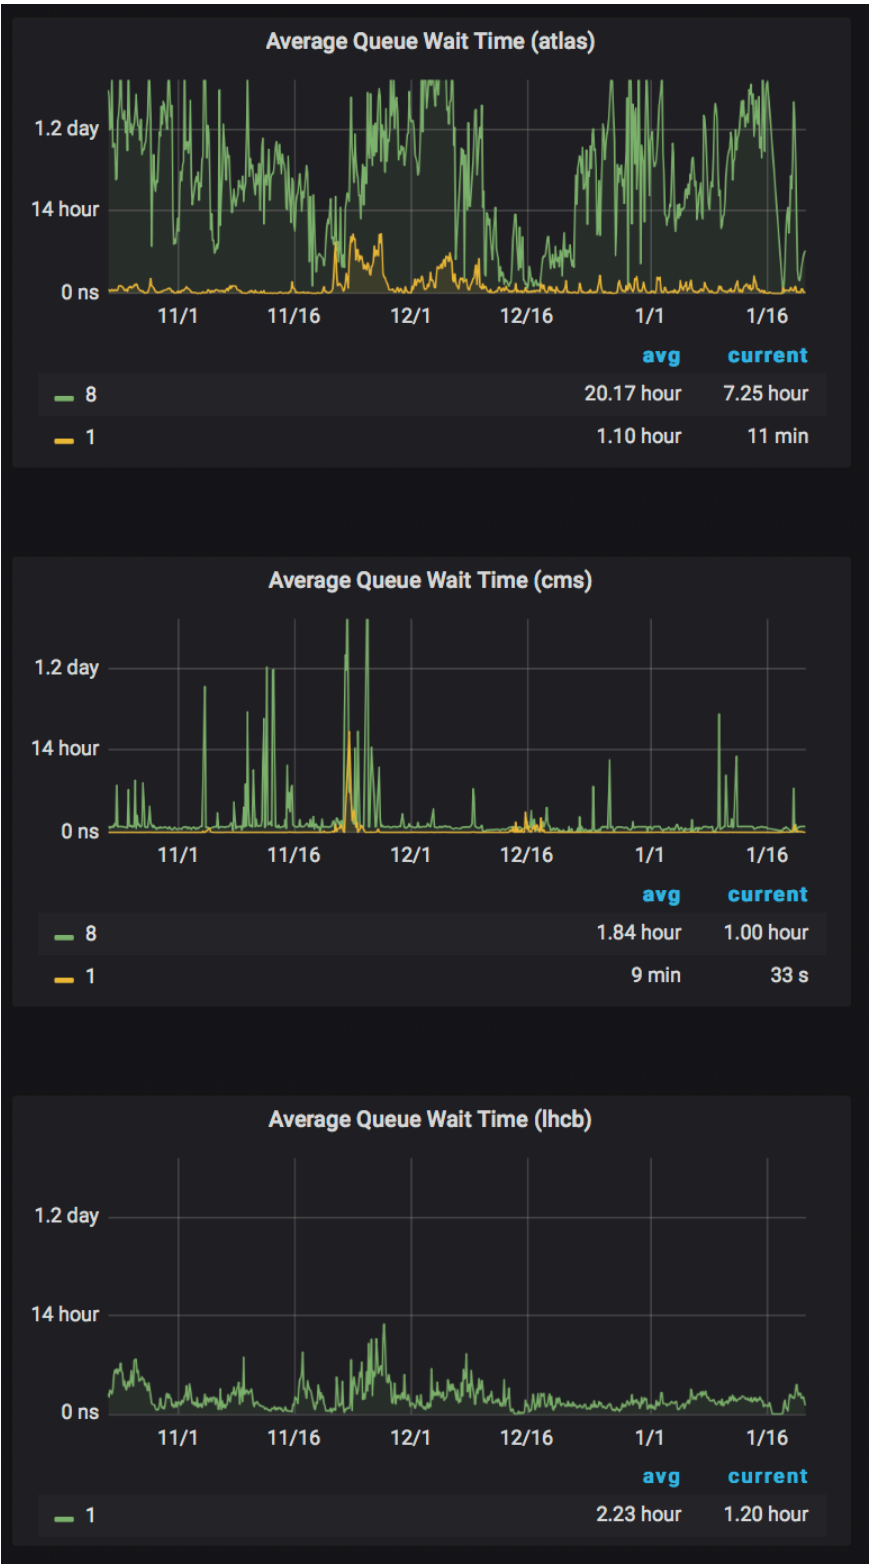


Total: 1.38 , Average Rate: 0.00 /s

u^b

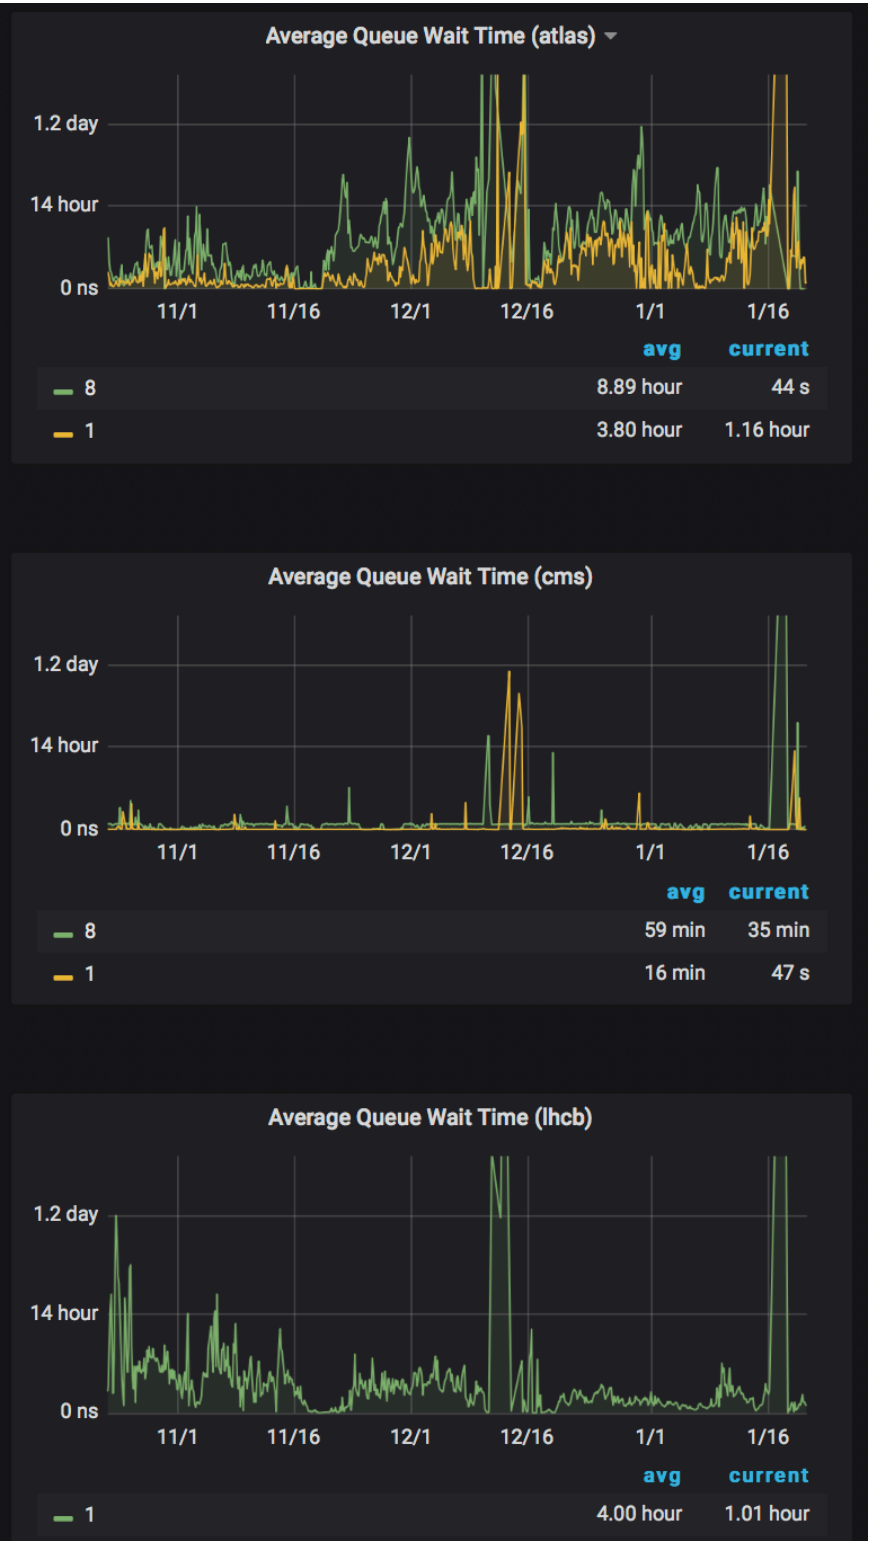


^b UNIVERSITÄT
BERN

AEC
ALBERT EINSTEIN CENTER
FOR FUNDAMENTAL PHYSICS





Fair-share issue - PHOENIX



Fair-share issue - Piz Daint




Fair-share issue



15 Dec =>

<= 16 Dec

Slots used



Dashboard wishes

- ▶ **SLURM dashboard - ATLAS request**
 - ▶ Add a panel for *pending _cores_ per VO*
 - ▶ Display *running vs pledged vs installed (cores)*
- ▶ **Additional monitoring - ATLAS request**
 - ▶ ARC CEs: all gangliarc metrics
 - ▶ HammerCloud alarm status
- ▶ **Public outreach page**
 - ▶ Basic "*propaganda*" plots

Summary and outlook

- ▶ **WC HS06 delivered vs. Pledged on target**
 - ▶ 2018 pledged **lower** than it should have been for CSCS
- ▶ **A bit short in delivered vs. Installed at CSCS**
 - ▶ Seems to have improved since December 2018
- ▶ **Bern ramping down in view of re-installing with PHOENIX hardware and consequential ramp-up**
- ▶ **CPU/WC efficiency stable and ok**
- ▶ **Needs improving:**
 - ▶ Failed WC still high, higher now in Bern as well
- ▶ **VO shares in SLURM**
 - ▶ seems to have slightly improved since December 2018 (VO share rebalancing - still off on PHOENIX)
 - ▶ *are single core jobs being packed to nodes?*
- ▶ **CentOS 7 rollout**
 - ▶ Provisional deadline 1st June
 - ▶ Needs clarifying the use of singularity (wished by ATLAS, is CSCS ready?)
- ▶ **ARC 6 migration**
 - ▶ Still in alpha version
- ▶ **Discussions on storage (tape, data lakes, etc.) are on ice**
 - ▶ Should these be revived?
 - ▶ WLCG and ATLAS storage models are going to evolve considerably, are we going to take part to some extent?

BACKUP

- ▶ **Previous f2f storage slides**

Outlook 1/2

- ▶ **Future** (discussions, r&d, etc):
 - ▶ **Joining the NDGF data-“fjord”?** (*nordic version of a data-lake*)
 - ▶ **Both UNIBE and CSCS**
 - ▶ **dCache based**
 - ▶ **dCache pools under the NDGF head nodes**
 - ▶ **Currently 4 Tier-1 storage sites and 1 Tier-2 storage site (Slovenia)**
 - ▶ ***dCache pools can be managed remotely by an unprivileged user***
 - ▶ **Could open ports to tape usage for CSCS**
 - ▶ *What about politics?*

Outlook 2/2

- ▶ **Future** (discussions, r&d, etc):
 - ▶ **Object stores**
 - ▶ For pure object store, need for ATLAS development for protocols (s3, swift)
 - ▶ First step in that direction would be add a CEPH FS pool to dCache (LHC protocols)
 - ▶ Better QoS for data loss, but it has to be tested vs. offering less disk space
 - ▶ **Tape @ CSCS**
 - ▶ There could be a concrete use case for a new storage model for data lakes
 - ▶ More storage tiers with different QoS
 - ▶ Again, this is new r&d for ATLAS, discussions ongoing
 - ▶ The tape@CSCS case will be brought up next week at the ATLAS s&c week
 - ▶ *Currently tape needs dCache, unless we can provide another interface based on https, root, s3 that is usable by FTS*
 - ▶ **Event service with arc05?**