- Performance overview January-September 2017

ATLAS T2 VO REPORT

Gianfranco Sciacca

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

CHIPP-CSCS face 2 face - 04 October 2017

Availability - HammerCloud AFTs and PFTs

Historic view for "panda_queues_all" from 00:00 01.01.2017 to 00:00 30.09.2017

Show 100 entries Search:

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 544 hours	offline		brokeroff		online		NoQueue		test	
					%	count	%	count	%	count	%	count	%	count
ANALY_CSCS	CSCS-LCG2	T2D	DE		3.35	5	0.25	4	94.26	27	0	0	1.47	9
ANALY_CSCS-HPC	CSCS-LCG2	T2D	DE		52.79	14	0.24	4	41.75	19	0	0	4.55	11
ANALY_UNIBE-LHEP	UNIBE-LHEP	T2	ND		1.26	1	0	0	96.9	46	0	0	1.17	31
ANALY_UNIBE-LHEP-UBELIX	UNIBE-LHEP	T2	ND	00:00 01.01.2017	1.26	1	0	0	97.43	19	0	0	0.64	4
CSCS-LCG2	UNDEFINED	UNDEFINED	UNDEFINED		3.84	6	0	0	94.28	26	0	0	1.21	9
CSCS-LCG2-HPC	CSCS-LCG2	T2D	DE		38.41	17	0	0	51.94	16	0	0	8.98	2
CSCS-LCG2-HPC_MCORE	CSCS-LCG2	T2D	DE		38.71	17	0	0	51.72	15	0	0	8.9	2
CSCS-LCG2_MCORE	UNDEFINED	UNDEFINED	UNDEFINED		4.2	7	0	0	93.92	27	0	0	1.21	9
UNIBE-LHEP	UNIBE-LHEP	T2	ND		0	0	0.08	1	98.22	32	0	0	1.02	16
UNIBE-LHEP-UBELIX	UNIBE-LHEP	T2	ND		0	0	0	0	98.65	19	0	0	0.67	4
UNIBE-LHEP-UBELIX_MCORE	UNIBE-LHEP	T2	ND		0	0	0	0	94.69	22	0	0	4.64	7
UNIBE-LHEP_MCORE	UNIBE-LHEP	T2	ND		0	0	0	0	98.12	56	0	0	1.21	41

Showing 1 to 12 of 12 entries First Previous 1 Next Last

ATLAS T2 VO REPORT

Availability

- **CSCS-LCG2**: *93.98% (rank 65 / 76 T2s)*
- **UNIBE-LHEP**: *98.62% (rank 20 / 76 T2s)*

Algorithm:
(CREAM-CE + ARC-CE + HTCONDOR-CE + GLOBUS) * (all SRMv2 + all SRM + all GRIDFTP)

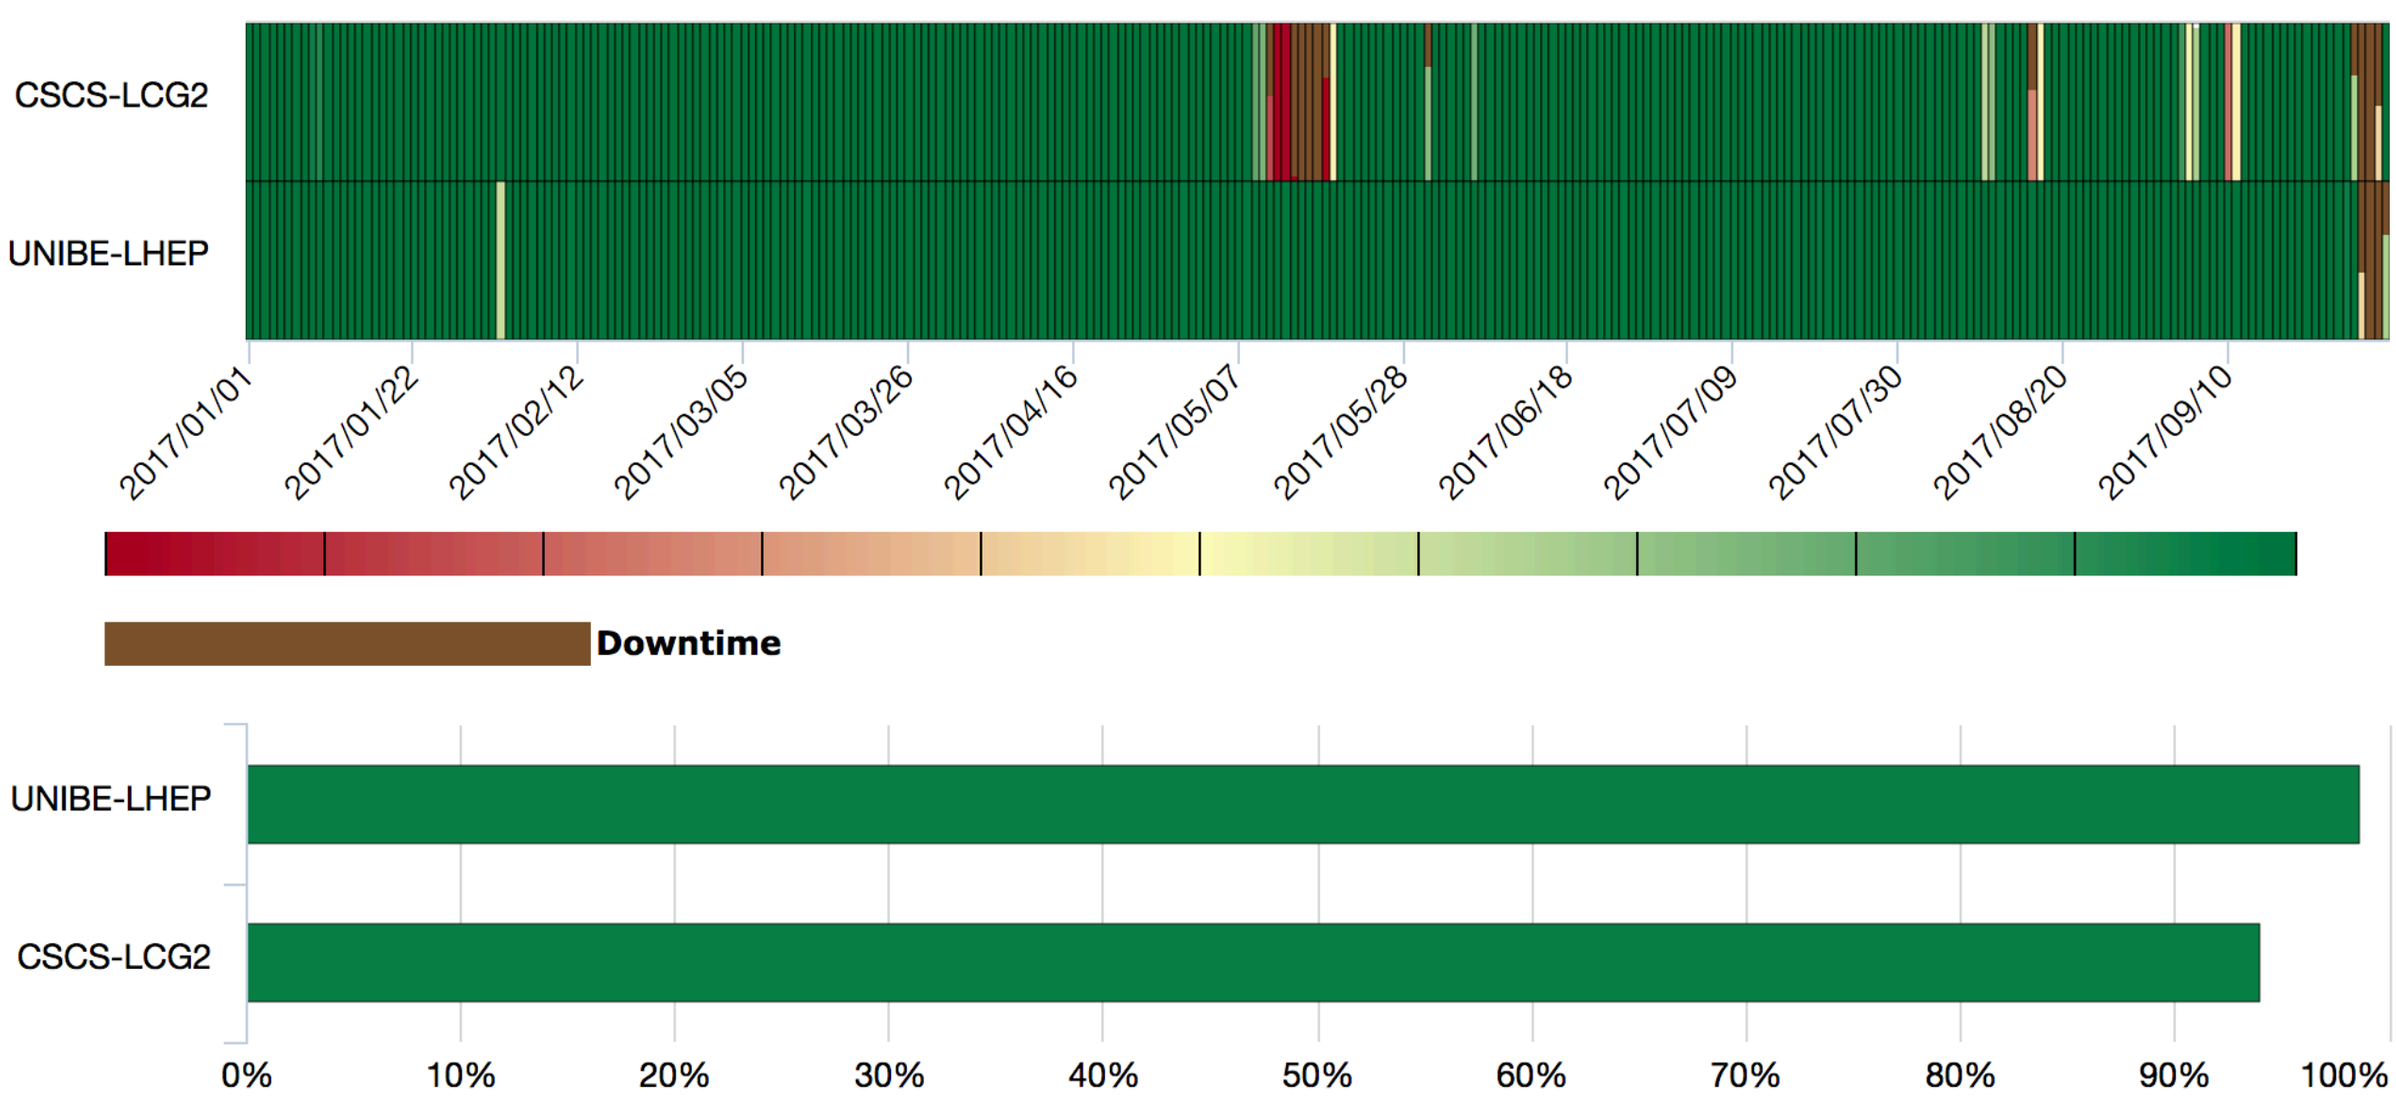

Reliability

- **CSCS-LCG2**: *96.27% (rank 54 / 76 T2s)*
- **UNIBE-LHEP**: *99.69% (rank 8 / 76 T2s)*

Algorithm:
(CREAM-CE + ARC-CE + HTCONDOR-CE + GLOBUS) * (all SRMv2 + all SRM + all GRIDFTP)

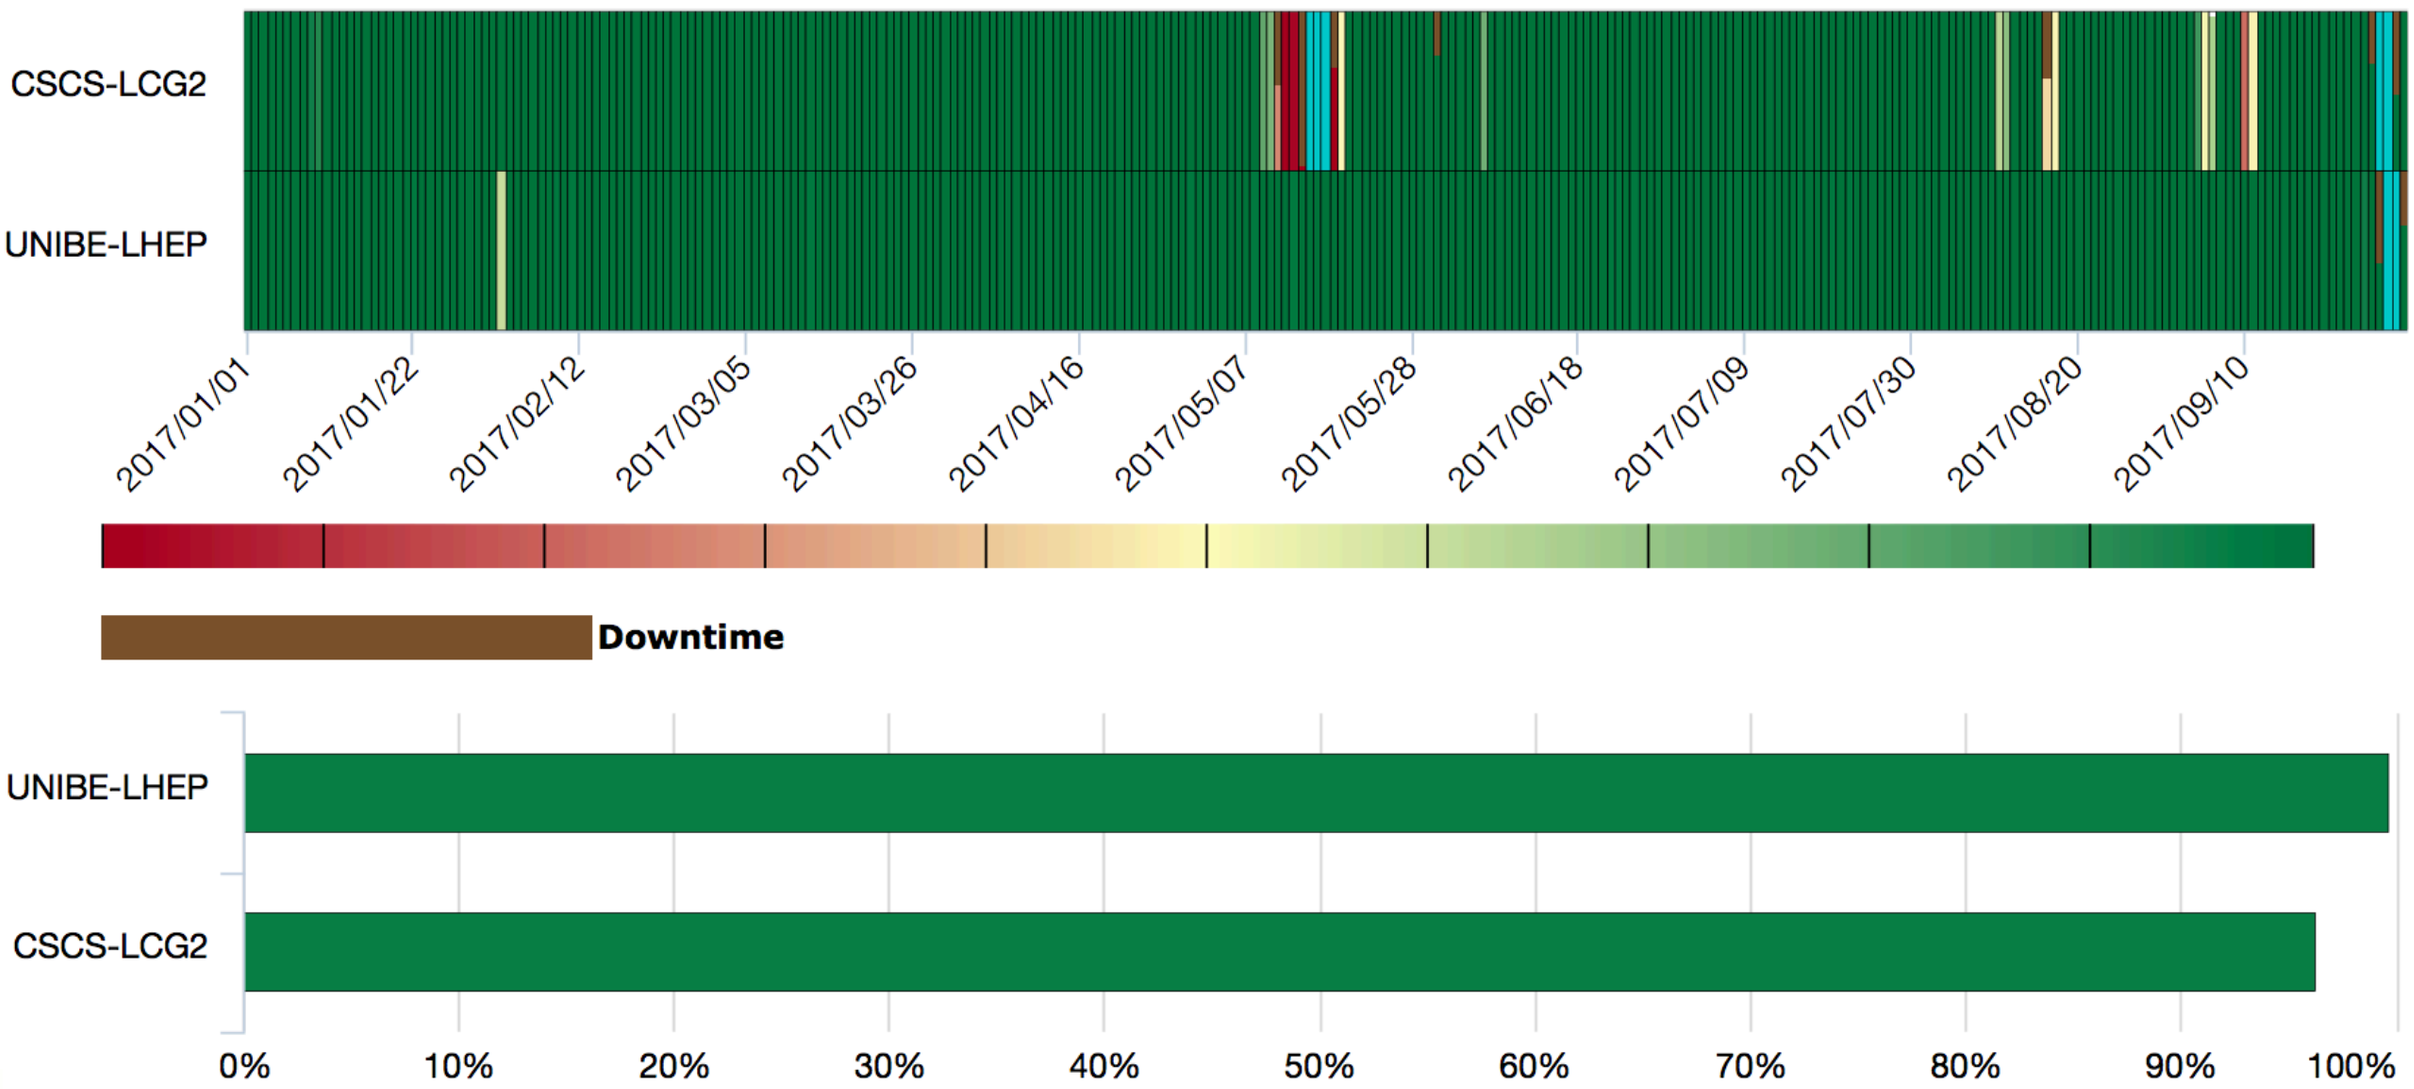

Availability + Reliability last 6 months

This is the period over which performance is evaluated for Nucleus status

Availability

- **CSCS-LCG2**: *94.46% (rank 67 / 76* T2s*)*
- **UNIBE-LHEP***: 99.82% (rank 27 / 76* T2s*)*

Reliability

- **CSCS-LCG2**: *94.46% (rank 60 / 76* T2s*)*
- **UNIBE-LHEP***: 99.82% (rank 5 / 76* T2s*)*

ATLAS T2 VO REPORT

WallClock HS06 - CH-CHIPP-CSCS Federation

- **CSCS-LCG2** *HS06 ATLAS* ***installed (*): 35491*** (59%) ***- pledged: 31200*** (68%) (2017)
- **UNIBE-LHEP** *HS06 ATLAS* ***installed (**): 24601*** (41%) ***- pledged 15000*** (32%) (2017)

total installed

total pledged

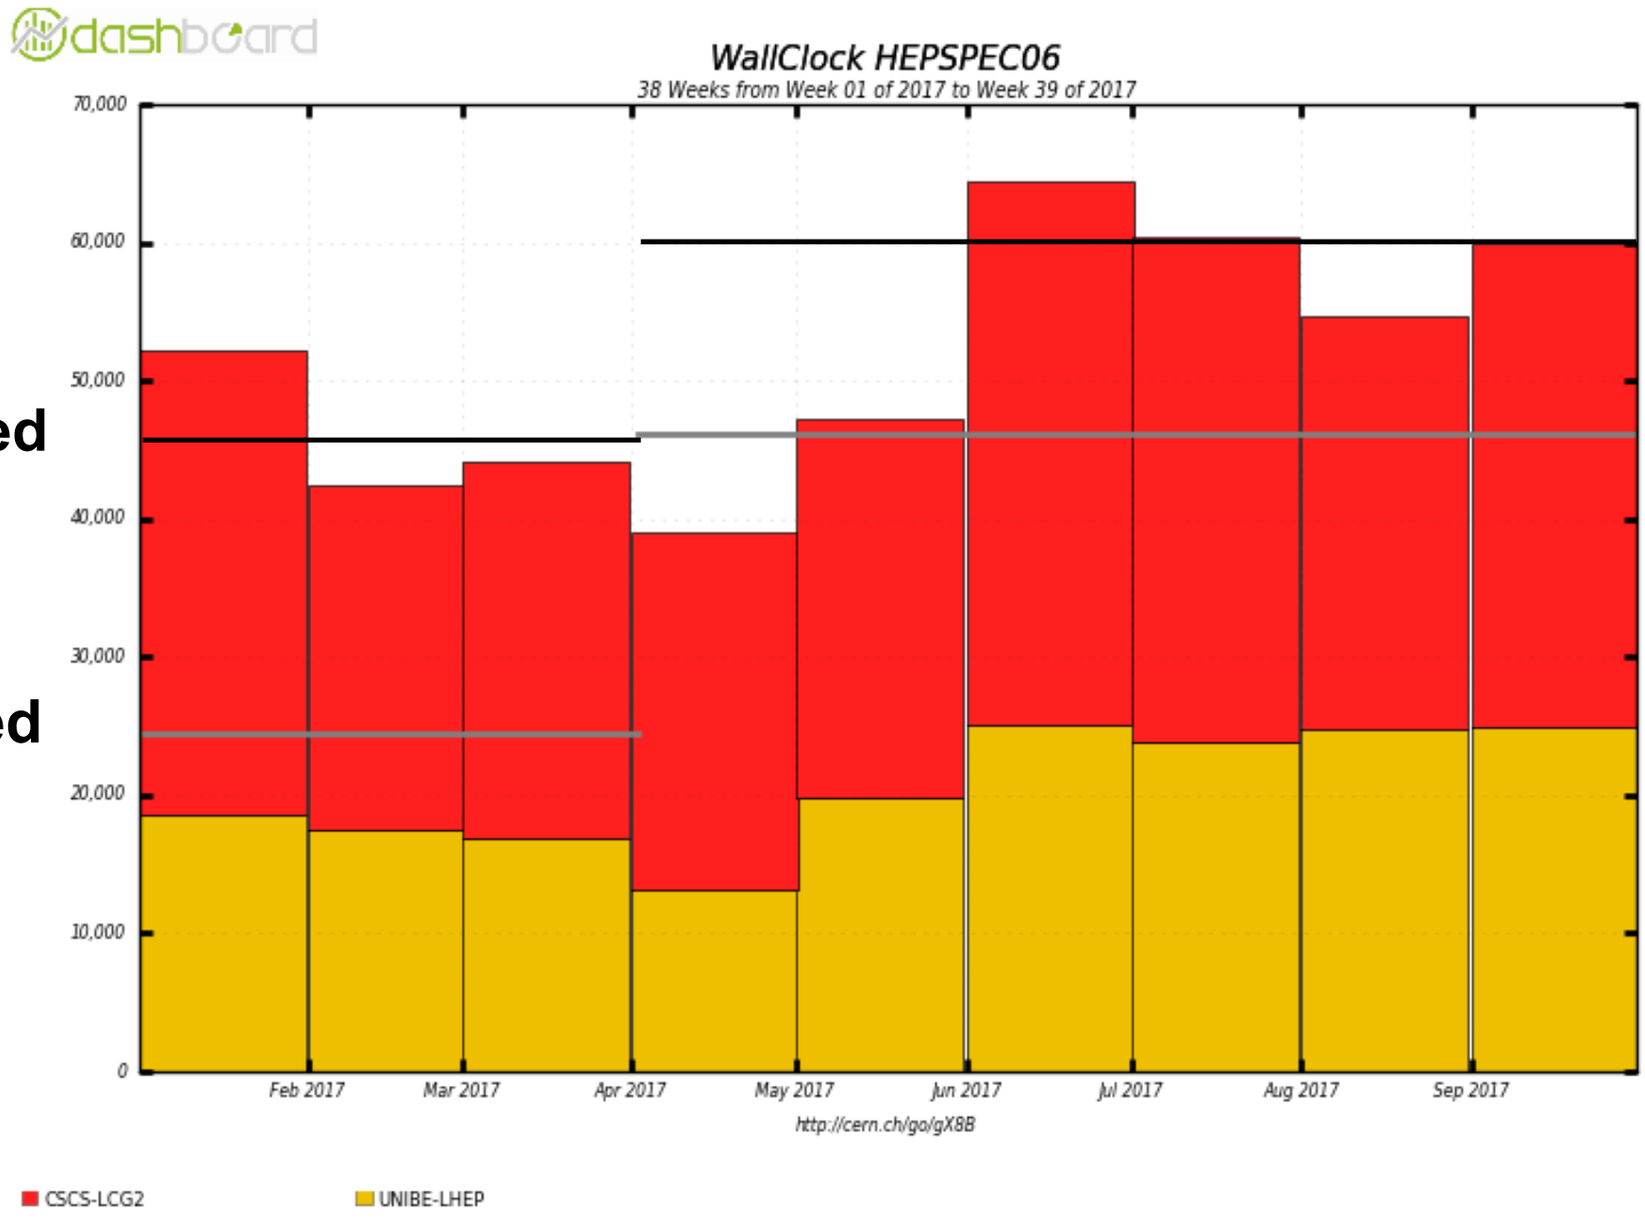

total installed

total pledged

(*) *- ATLAS ahare*

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

WallClock HS06 - CSCS-LCG2

- **CSCS-LCG2** *HS06 ATLAS* ***installed: 35491 - pledged: 31200*** (2017)
- **CSCS-LCG2** *HS06 ATLAS* ***installed: 27291 - pledged: 14500*** (2016)

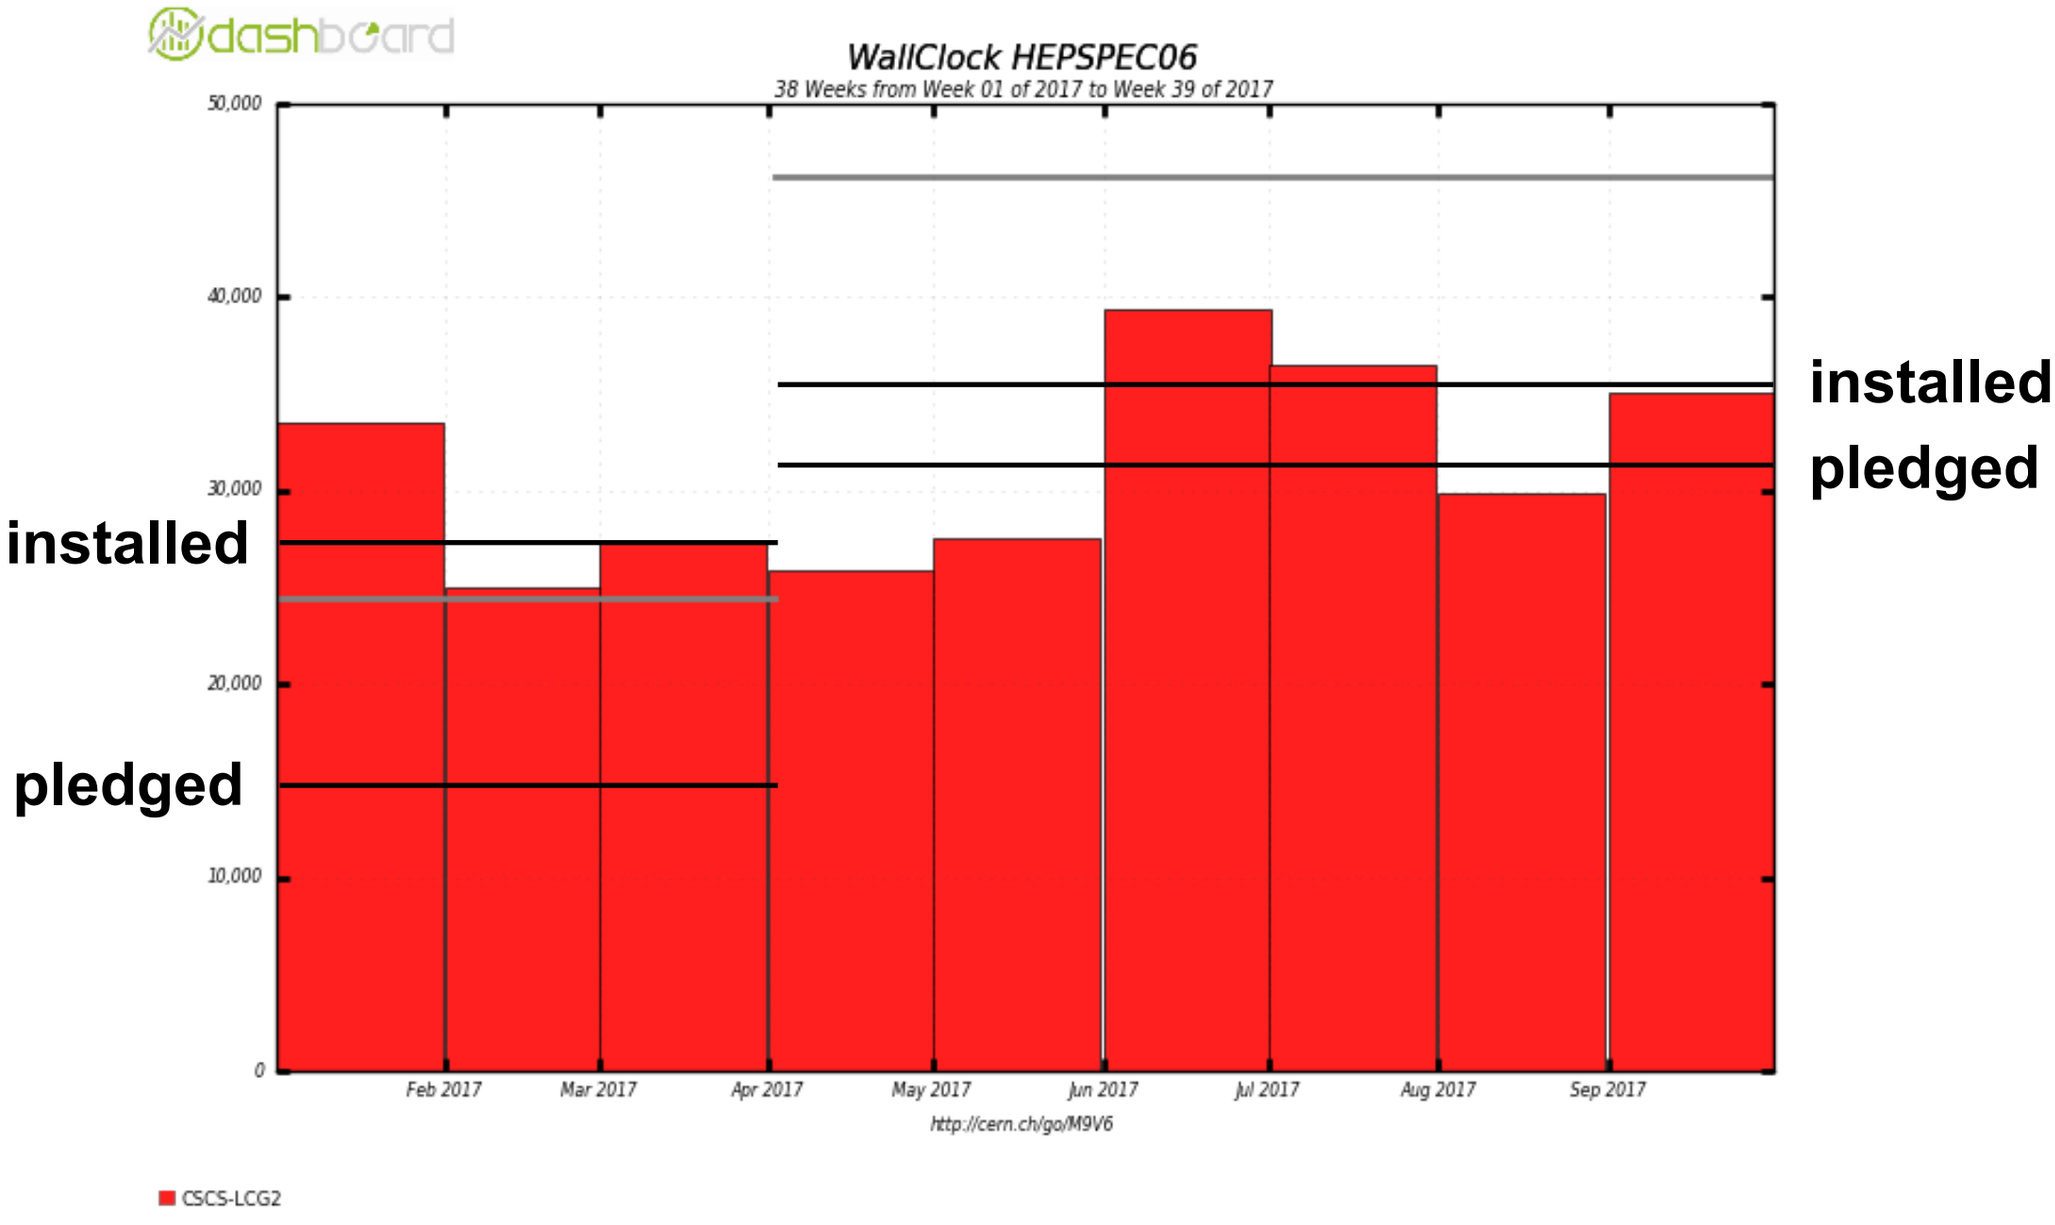

(*) *- ATLAS ahare*

WallClock HS06 - UNIBE-LHEP

- **UNIBE-LHEP** *HS06 ATLAS* ***installed (**): 24601 - pledged: 15000*** (2017)
- **UNIBE-LHEP** *HS06 ATLAS* ***installed (**): 18642 - pledged: 10000*** (2016)

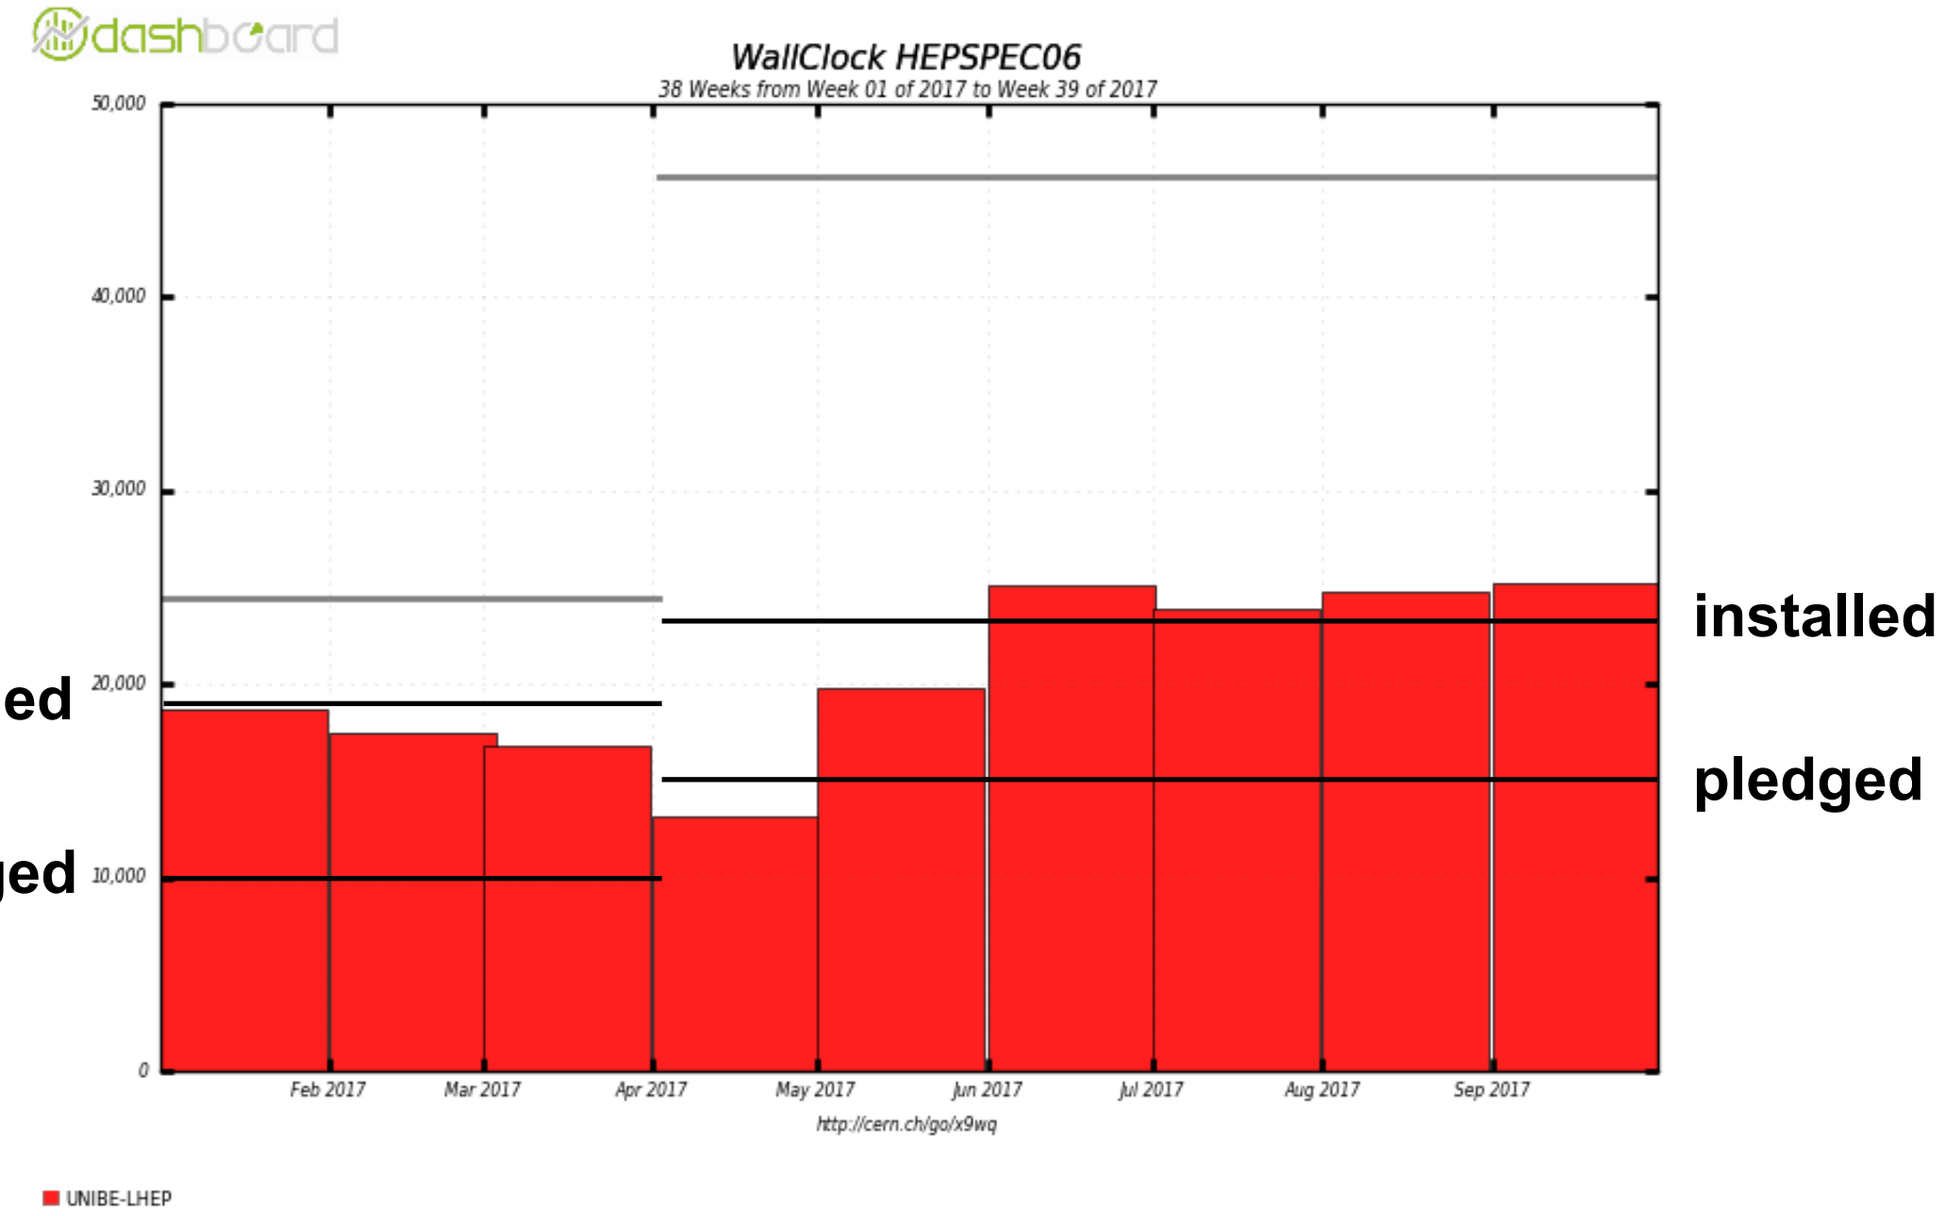

installed

pledged

installed

pledged

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

Relative shares

- **UNIBE-LHEP installed** estimated in about 41***% of the total installed capacity***

WC all jobs

UNIBE - 40.71%

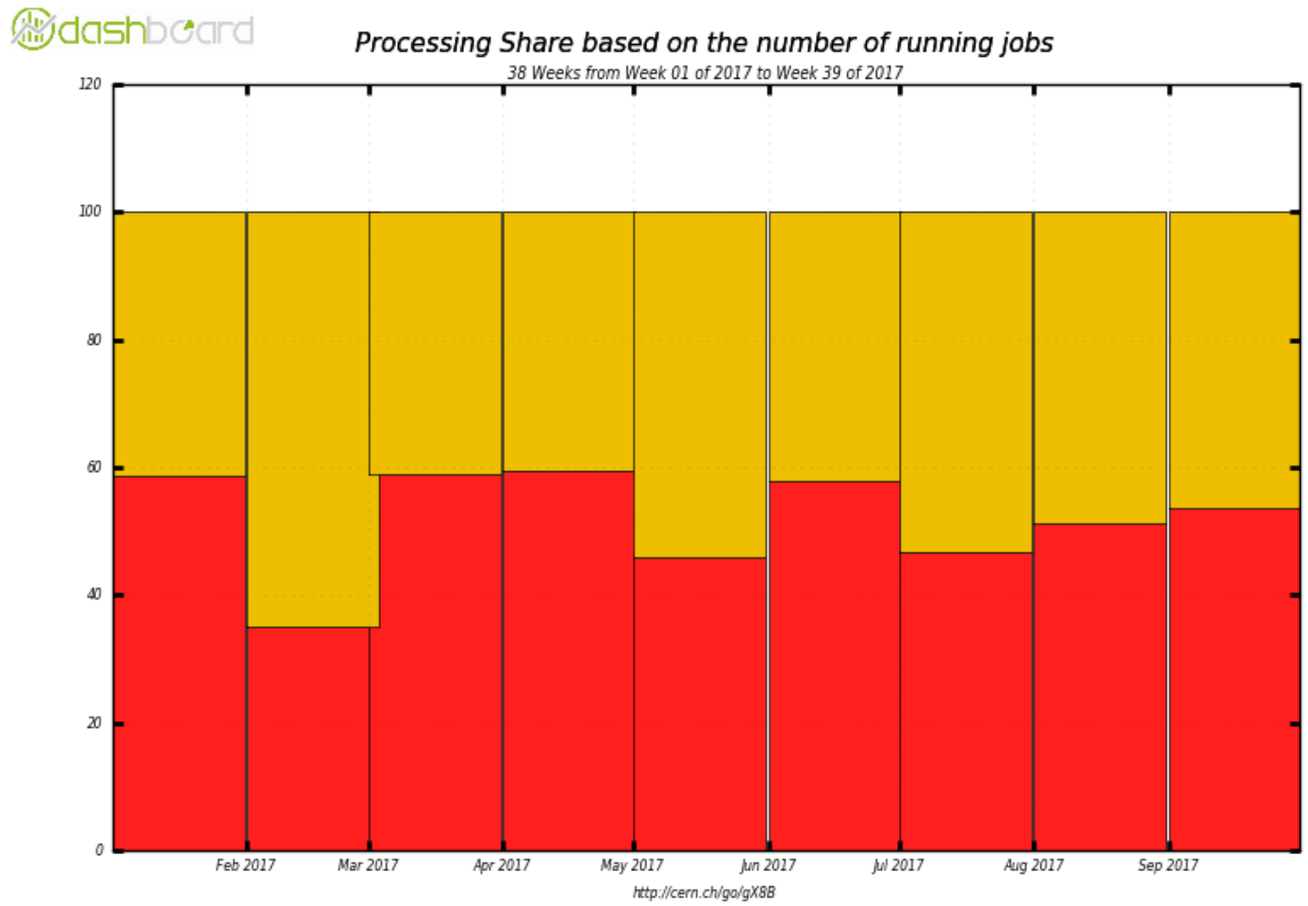

WC good jobs

UNIBE - 39.55%

Success vs fail WallClock efficiency

- **CSCS-LCG2**: *87%*
- **UNIBE-LHEP***: 85%*

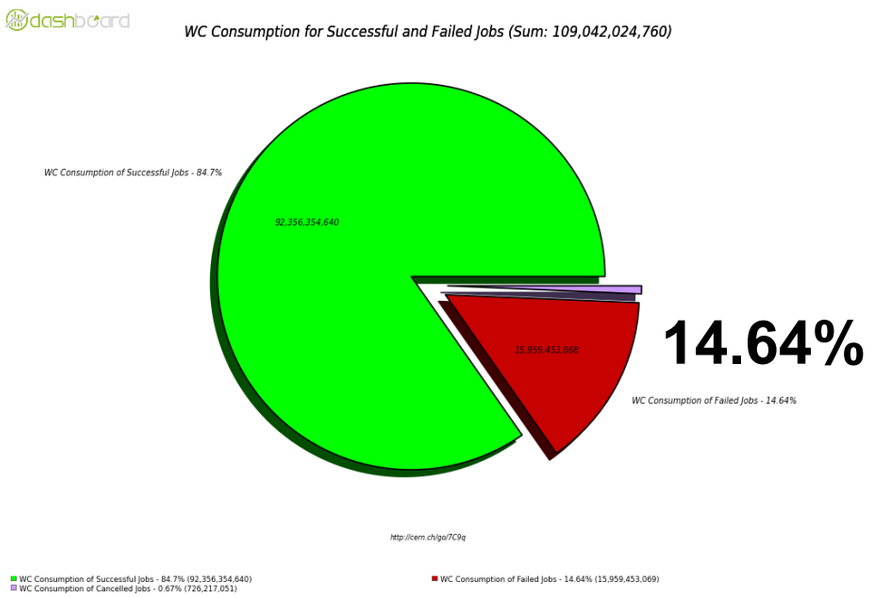

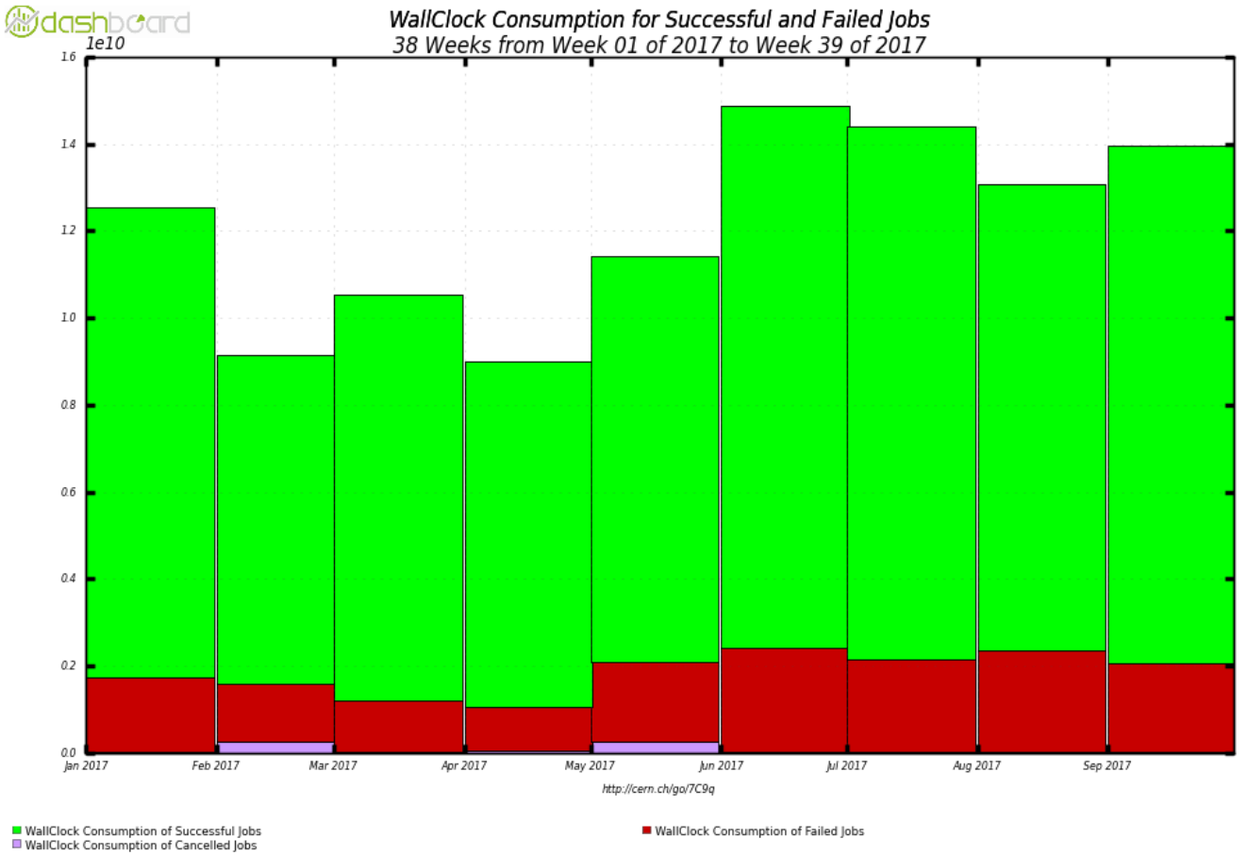

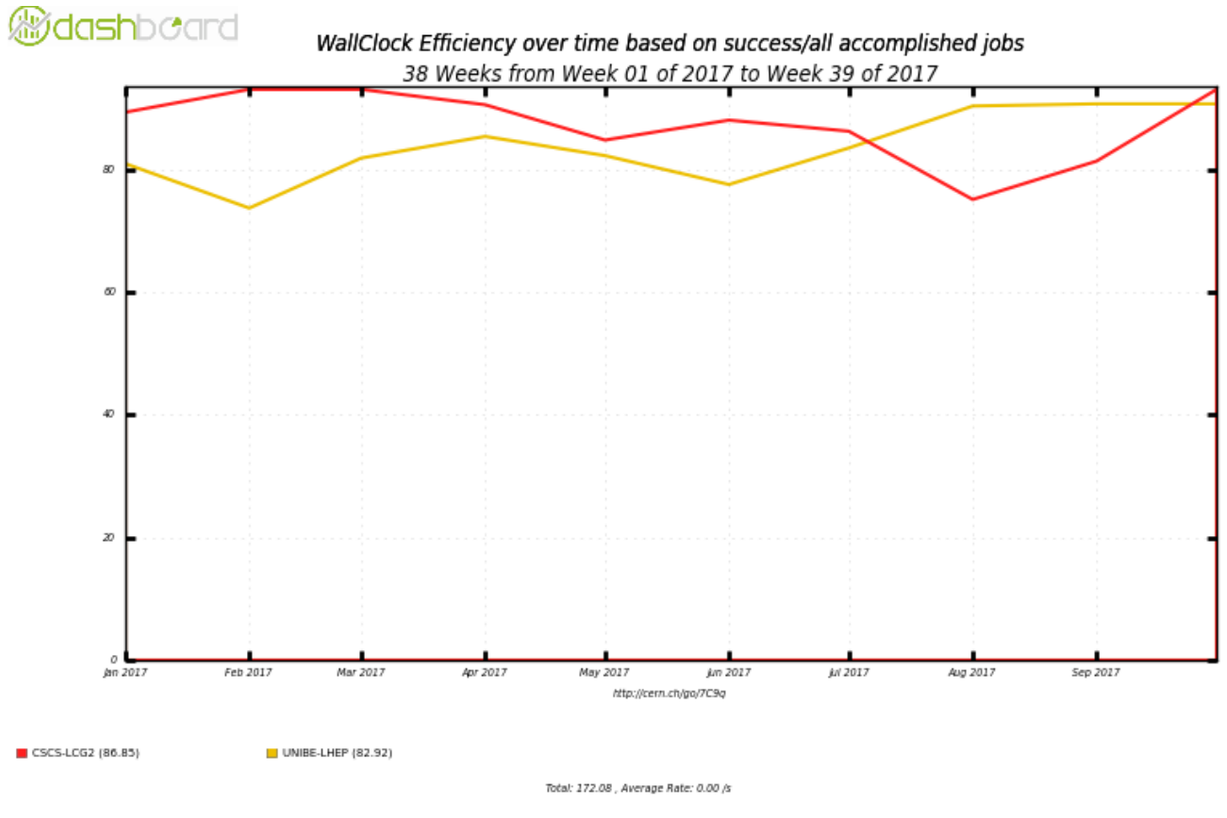

ATLAS T2 VO REPORT

CPU / WallClock efficiency

- **CSCS-LCG2**: *79%*
- **UNIBE-LHEP**: *82%*

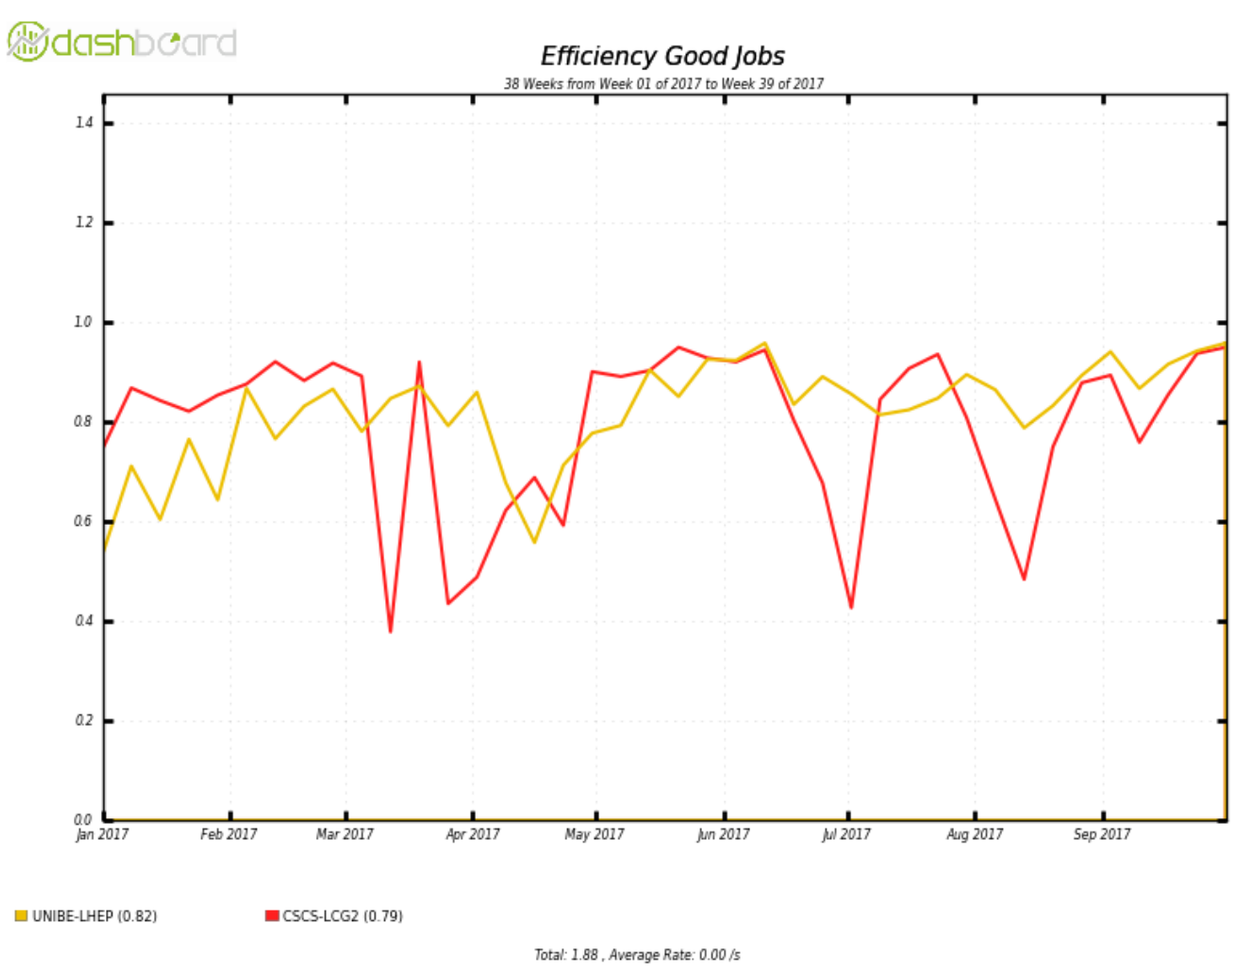

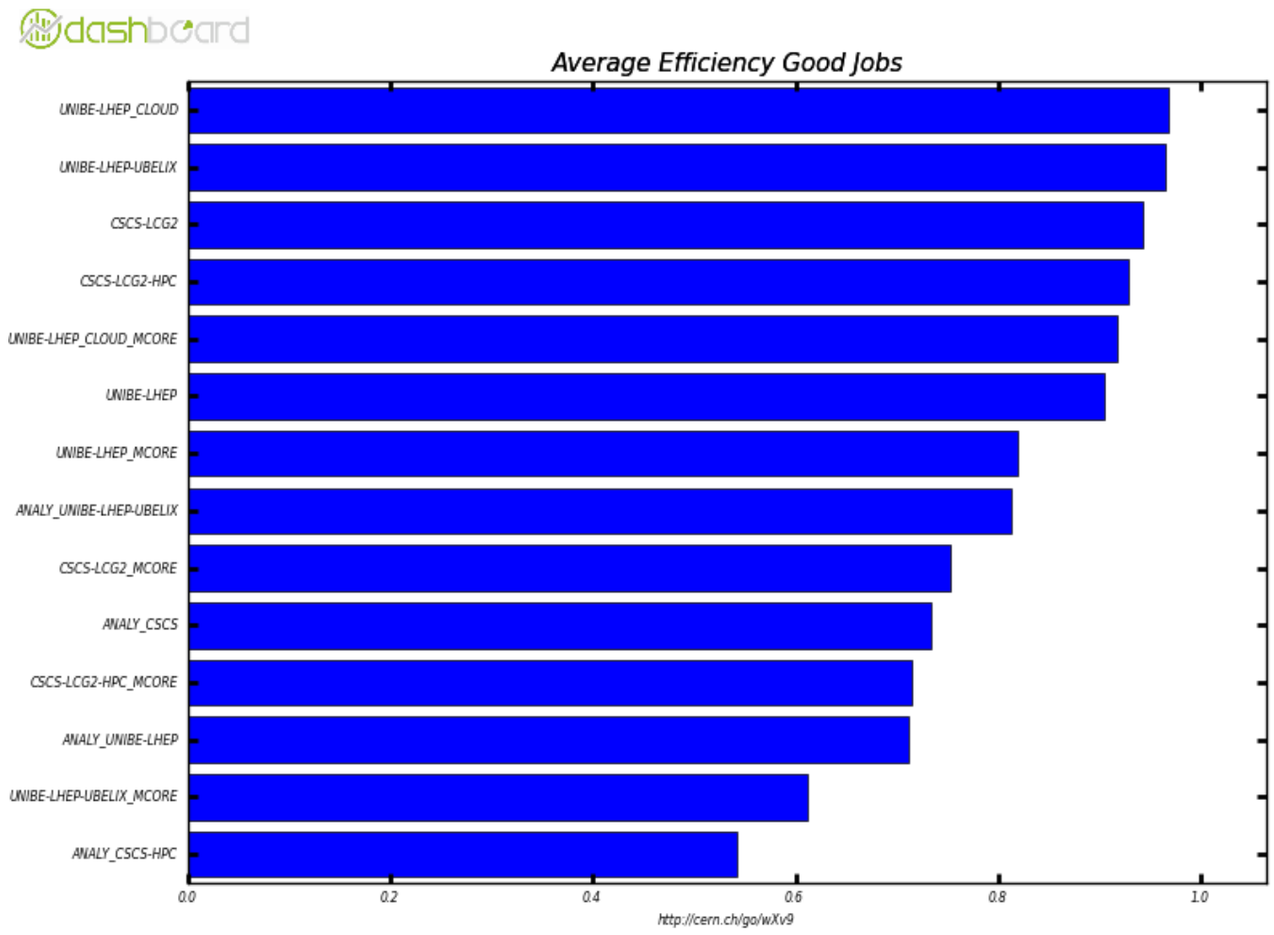

CPU / WallClock efficiency by PanDA queue

CPU / WallClock efficiency by ADC activity

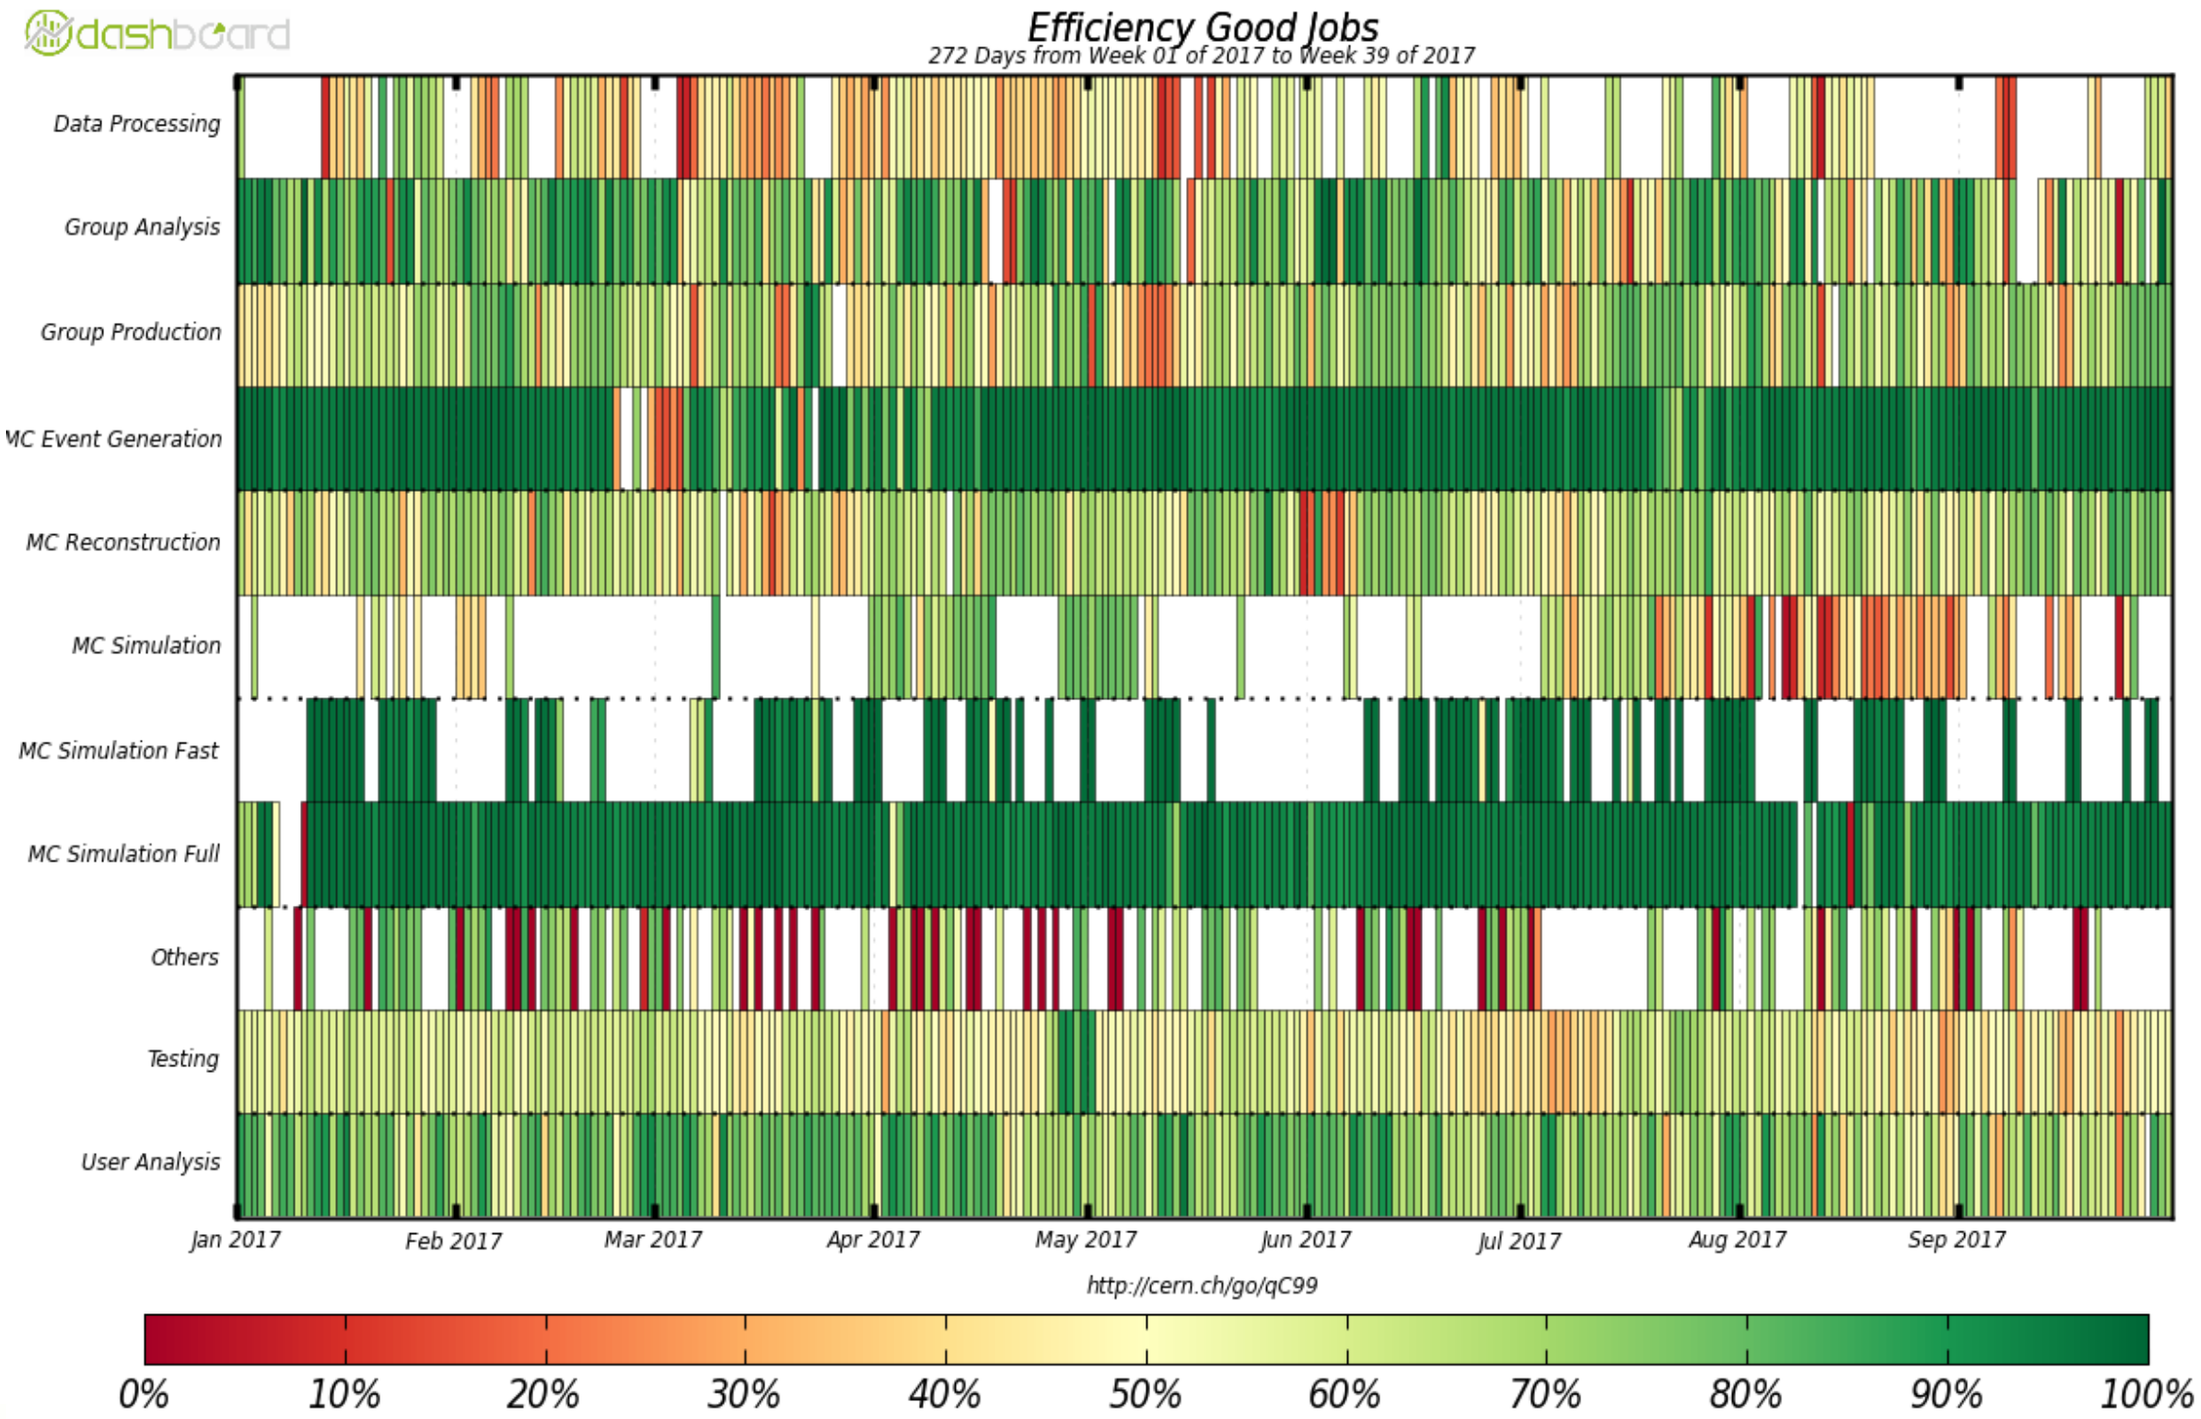

ATLAS T2 VO REPORT

WallClock by ADC activity and by PanDA queue

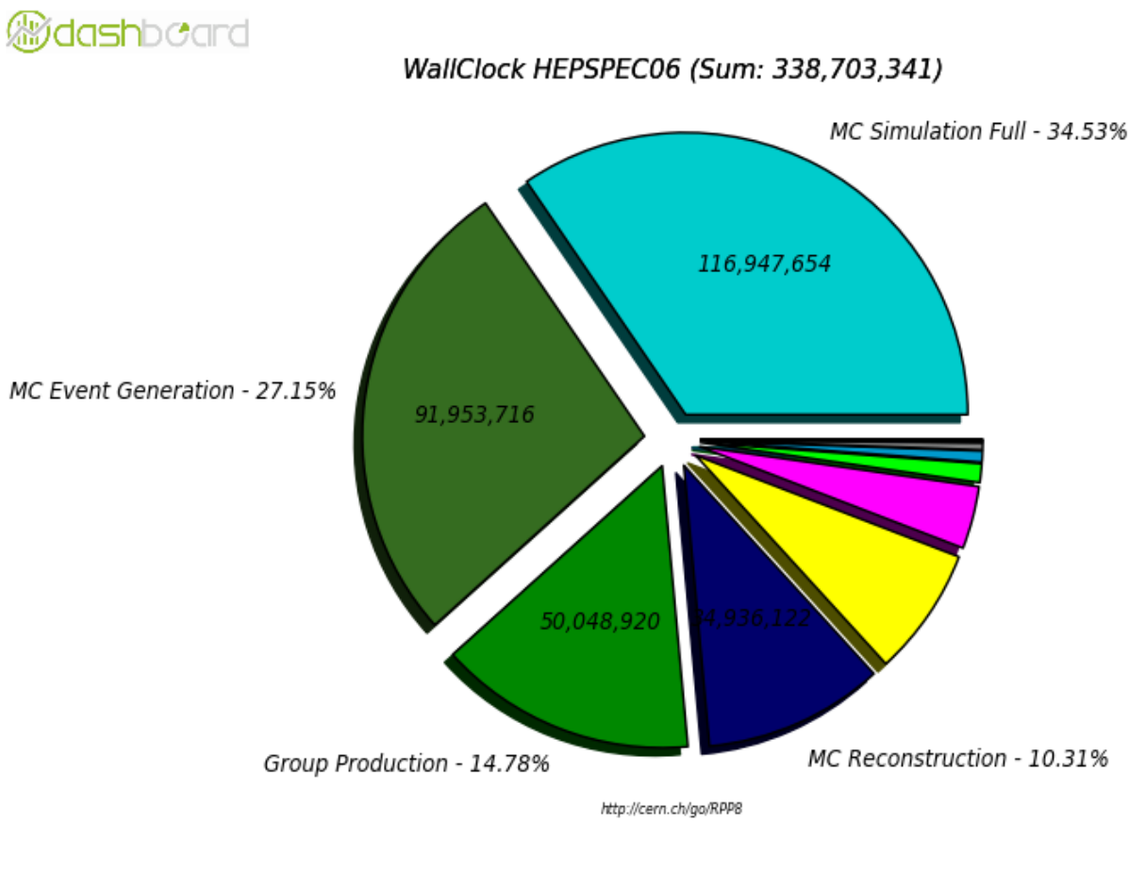

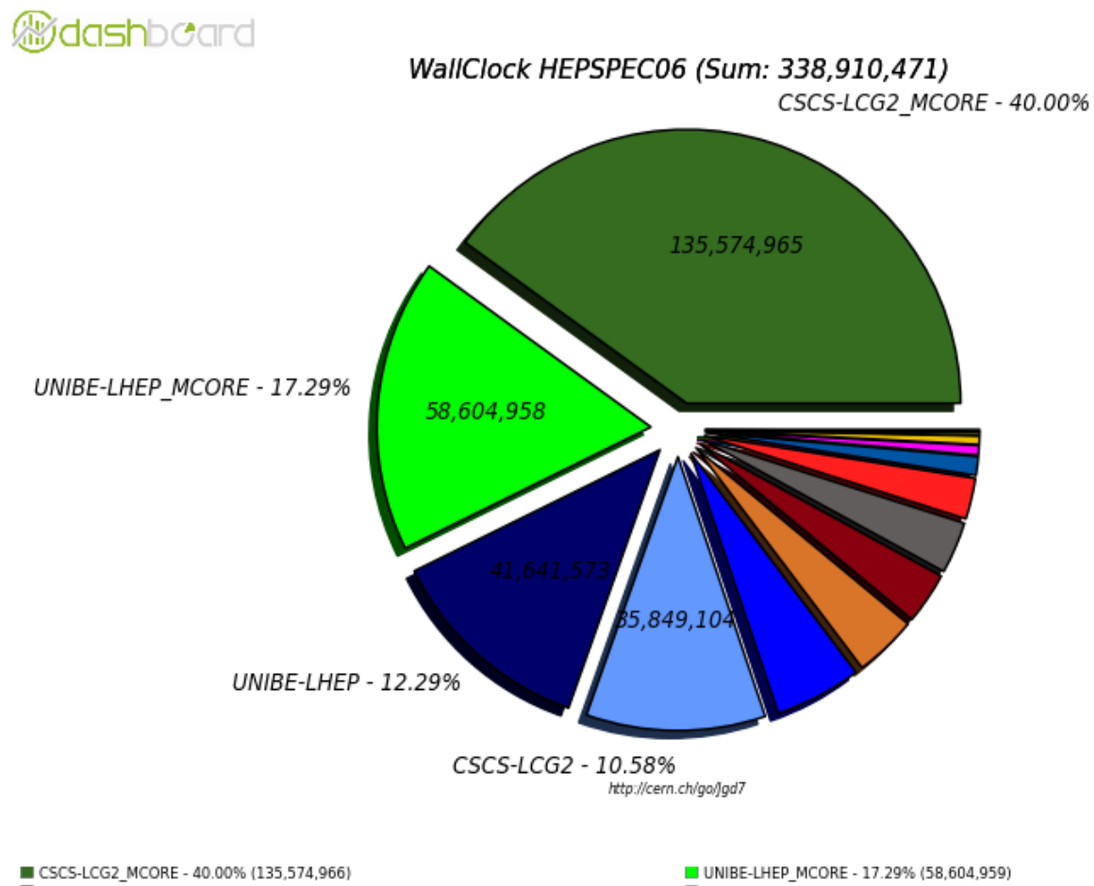

ATLAS T2 VO REPORT

CSCS-LCG2 as a Nucleus

- **Negotiated for quite a long time with the ATLAS management**
  - Initially agreed (April-May), as the 6-month retrospective history looked good
  - Some reservations due to the limited capacity of the site (both storage and CPU)
  - However the first high-profile dCache incident put that on hold
  - A long downtime for the SE is not suited to a centre meant to store data to serve to other sites
  - In the current configuration, also the compute part is taken down when the SE is down
  - More negotiation followed, but the site has since had more instabilities and a second dCache incident just at the time I brought this up again for consideration

- **Current status**
  - *unclear*
  - in principle I got this agreed a couple of weeks ago, after some additional negotiation and pressure, *but with several reservations* (see next slide)
  - however, I have not seen yet any task submitted with CSCS as a Nucleus, and the site is not on the list

CSCS-LCG2 as a Nucleus

- **Summary of objections from the ADC management**
  - Size capacity limited (both in storage and CPU)
  - Site instabilities and lengthy downtimes
  - The feature of Nucleus is an *operational* choice by ATLAS in order to streamline workloads in the most efficient way. The ADC management strongly discourage the use of this status as an argument of pressure on funding agencies
  - Reminder: Nucleus is not a *permanent* status

- **My take on this**
  - we should improve performance and stability of the site
  - reducing both the number of incidents and their duration

ATLAS T2 VO REPORT

Summary

- **We have come a long way since our f2f in CERN 1 year ago**
- **CSCS-LCG2 delivers with better efficiency**
- **However, we are not quite there yet: we need that extra leap to take us up in the ranks**
- **Set permanent remedies to recurring problems**
  - GPFS instabilities and performance
  - SE availability
  - improve reaction time to problems (monitoring, etc)
- **All of this maintaining the improvements since a year ago**
  - e.g. delivery vs. installed, etc.