



Science and
Technology
Facilities Council

Batch farm rejigging and ATLAS HS06

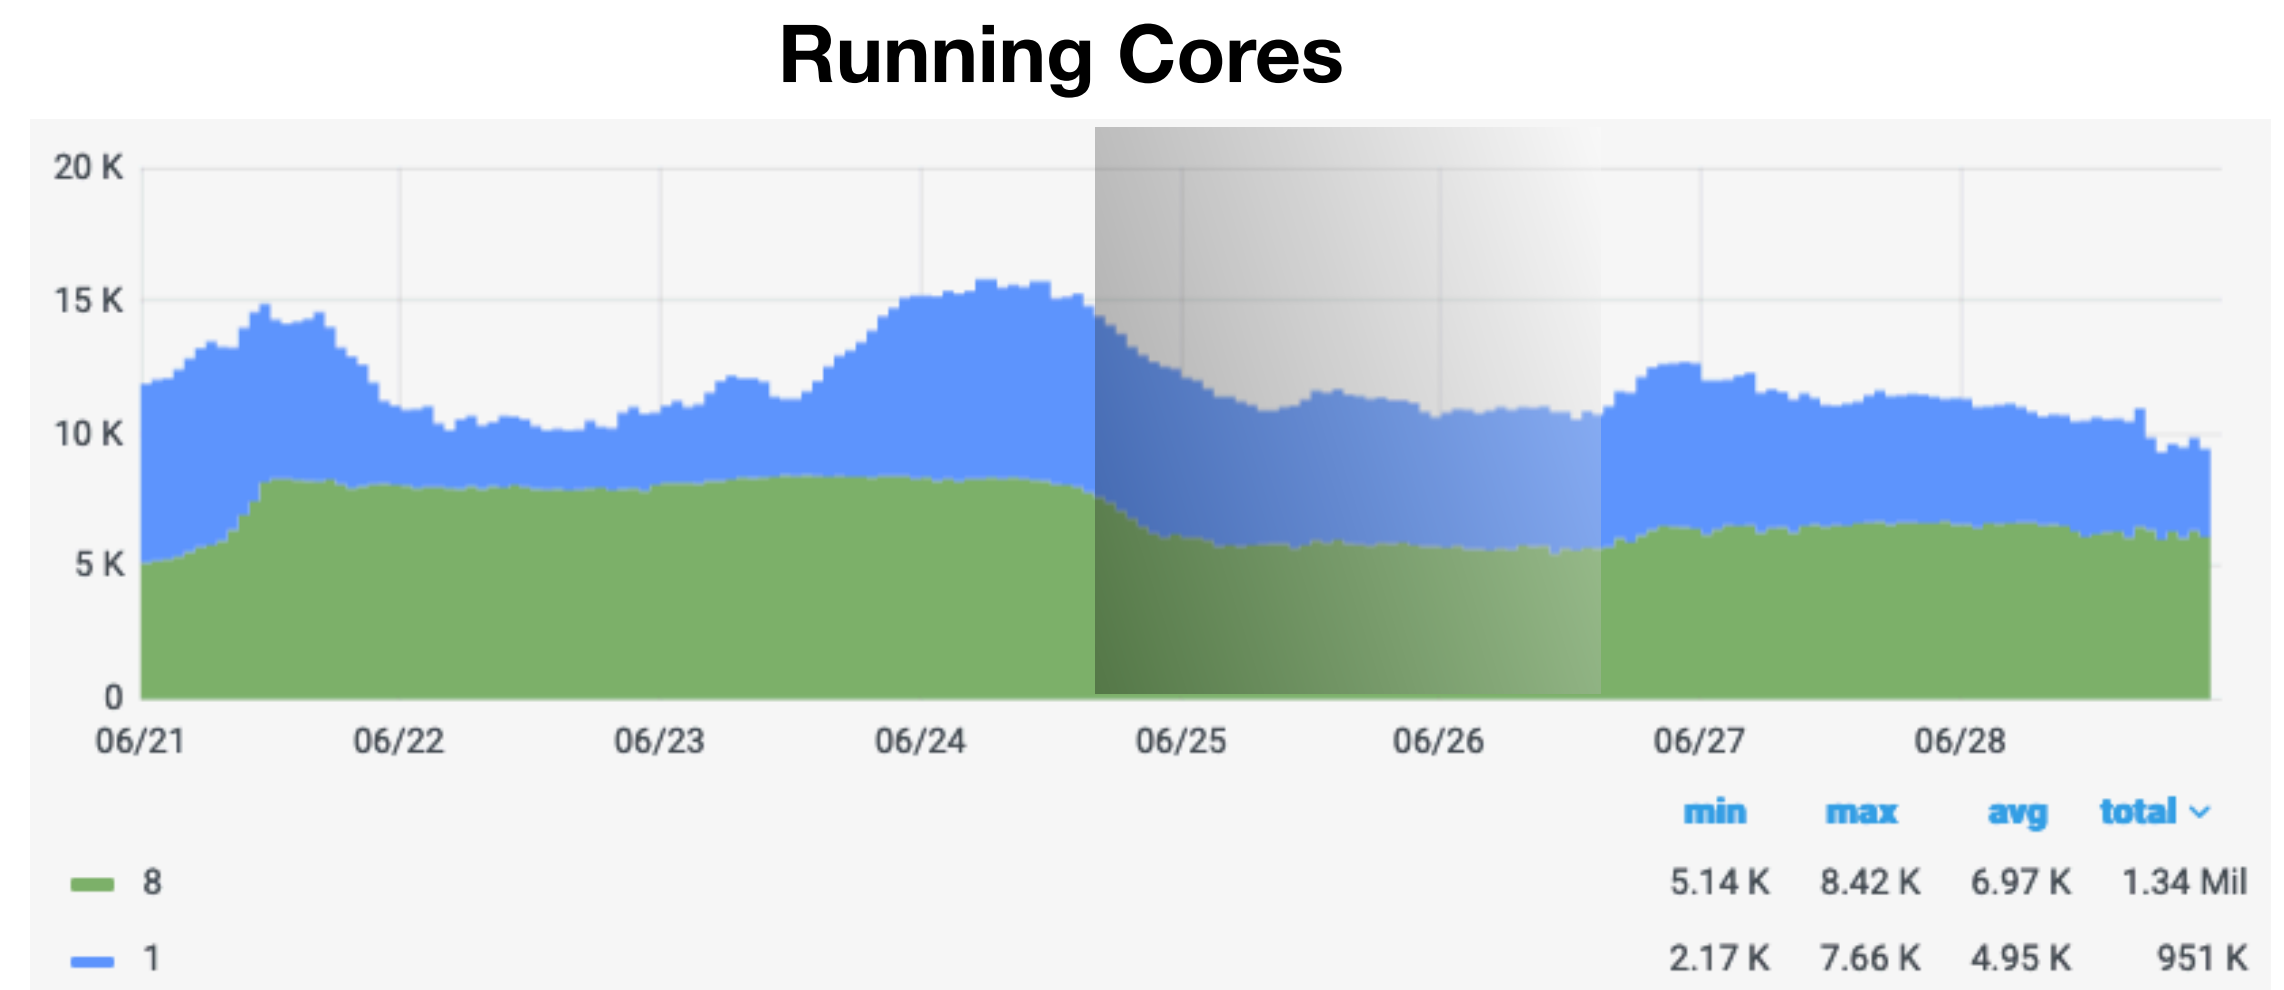
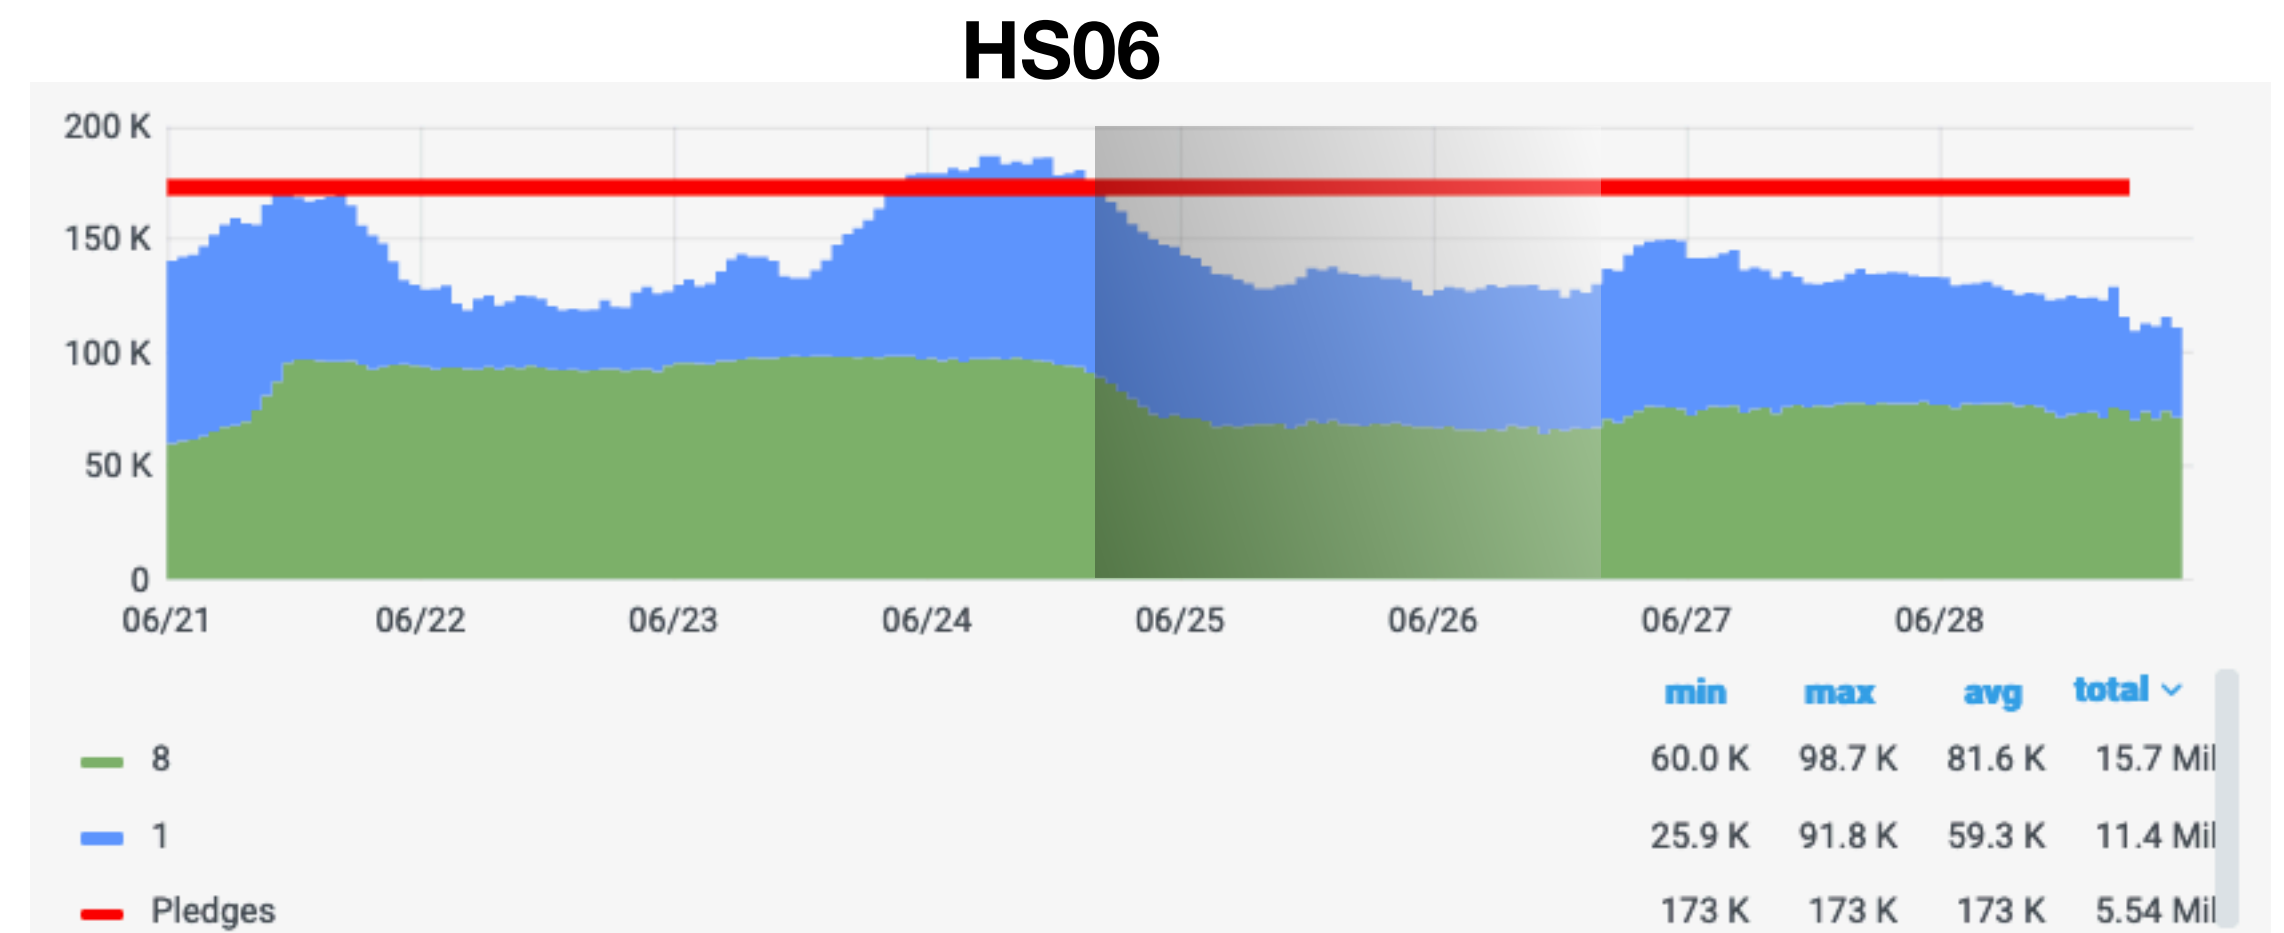
30/06/2021

Overview

- Started to drain older tranches evening Thurs, 24th June 2021:
 - Setting of 50% job-slot capacity applied:
 - (which should correspond now to the number of physical cores?)
 - 2014, 2015 and 2016
 - New personalities created: multicore-reduced-jobs, ml-reduced-jobs
- Some discrepancy between ATLAS monitoring and Vande, and ATLAS missing pledge (from its perspective).
 - Number of running cores in reasonable agreement; extracted HS06 (or Corepower scaling) appears to be the issue


ATLAS view

- Between 21 – 29 June
- Average HS06 ~ 140 kHS06
- Average number of running cores: 12k
 - Observe ‘capping’ (?) of mcore slots
- As of 29th; ~9.6k cores and ~110 kHS06.
- HS06 is determined from the number of running cores and the site-averaged Corepower value
- Historically Corepower was 10,
 - More recently, set to 11.7
 - Current values is ... ?
- In a non-homogeneous / fragmented cluster average value is potentially biased ...
 - For comparison, max host corepower value 13.8 for 2019 hosts (almost 20% difference to current value)




Vande View




- For same period, Vande shows movement around 100% values.
- Number of cores appears in good agreement.
- What is the definition of the “use / allocation” ?
 - i.e numerator and denominator?
 - Capping of ATLAS MCore slots?



atlas @ arc-ce-test02 Current: 675 atlas @ arc-ce01 Current: 502 atlas @ arc-ce02 Current: 739
 atlas @ arc-ce03 Current: 338 atlas @ arc-ce04 Current: 274 atlas @ arc-ce05 Current: 826

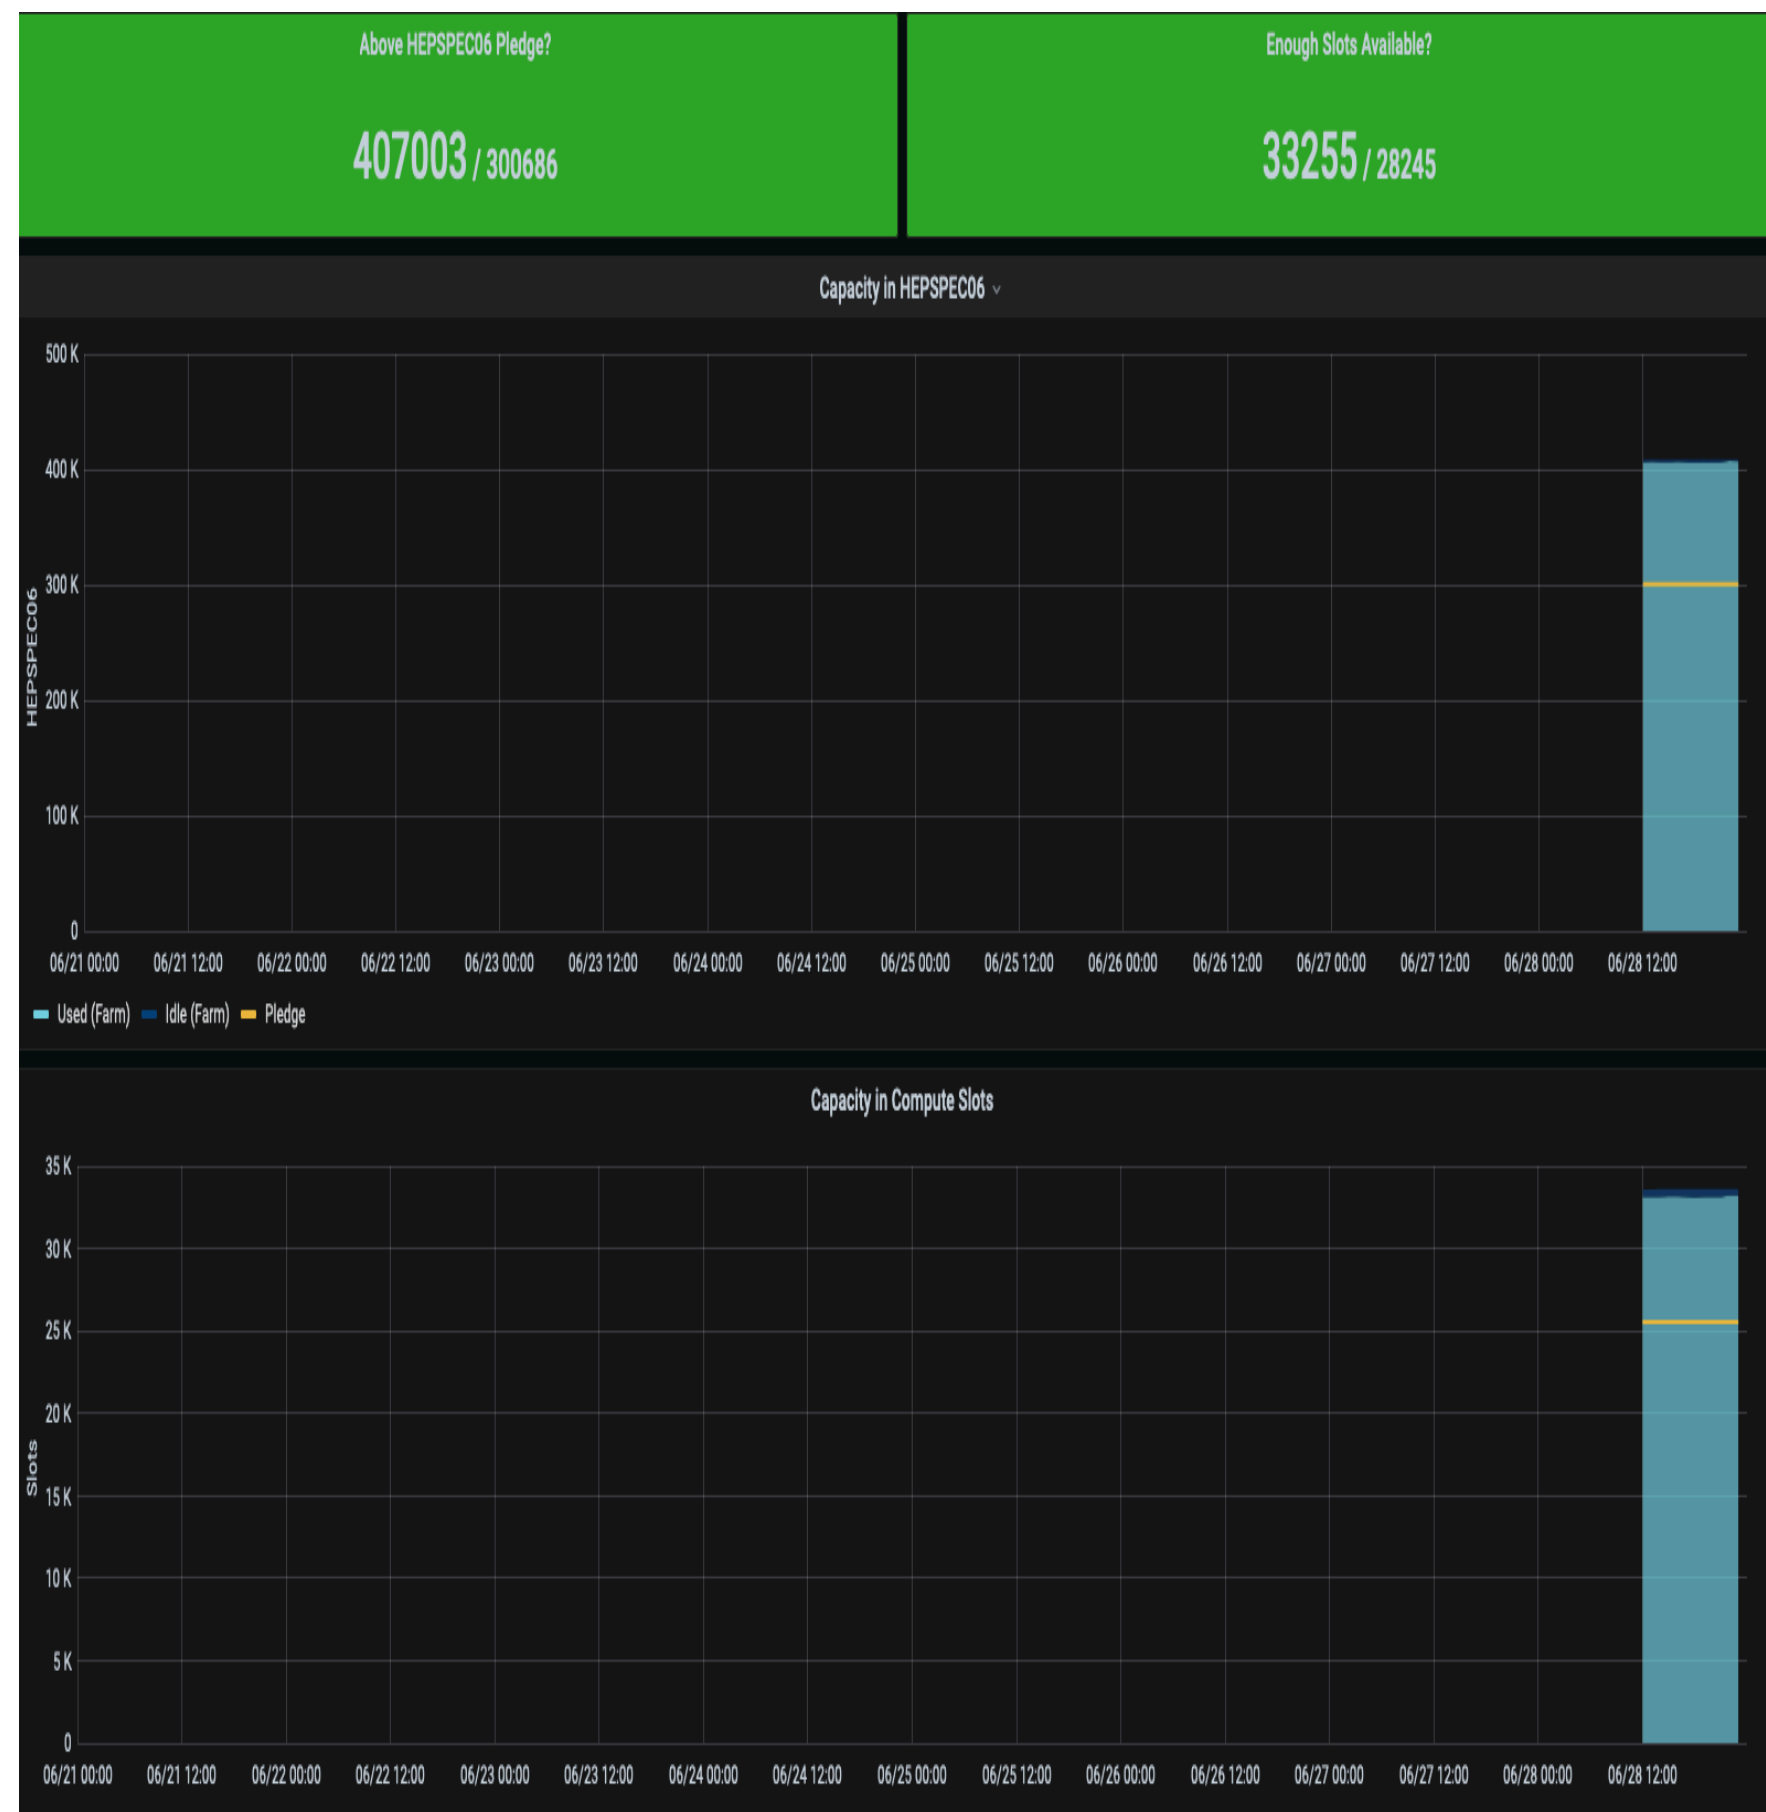


atlas @ arc-ce-test02 × 8 Current: 253 atlas @ arc-ce01 × 8 Current: 98.4 atlas @ arc-ce02 × 8 Current: 69
 atlas @ arc-ce03 × 8 Current: 139 atlas @ arc-ce04 × 8 Current: 184 atlas @ arc-ce05 × 8 Current: 102




Batch system capacity

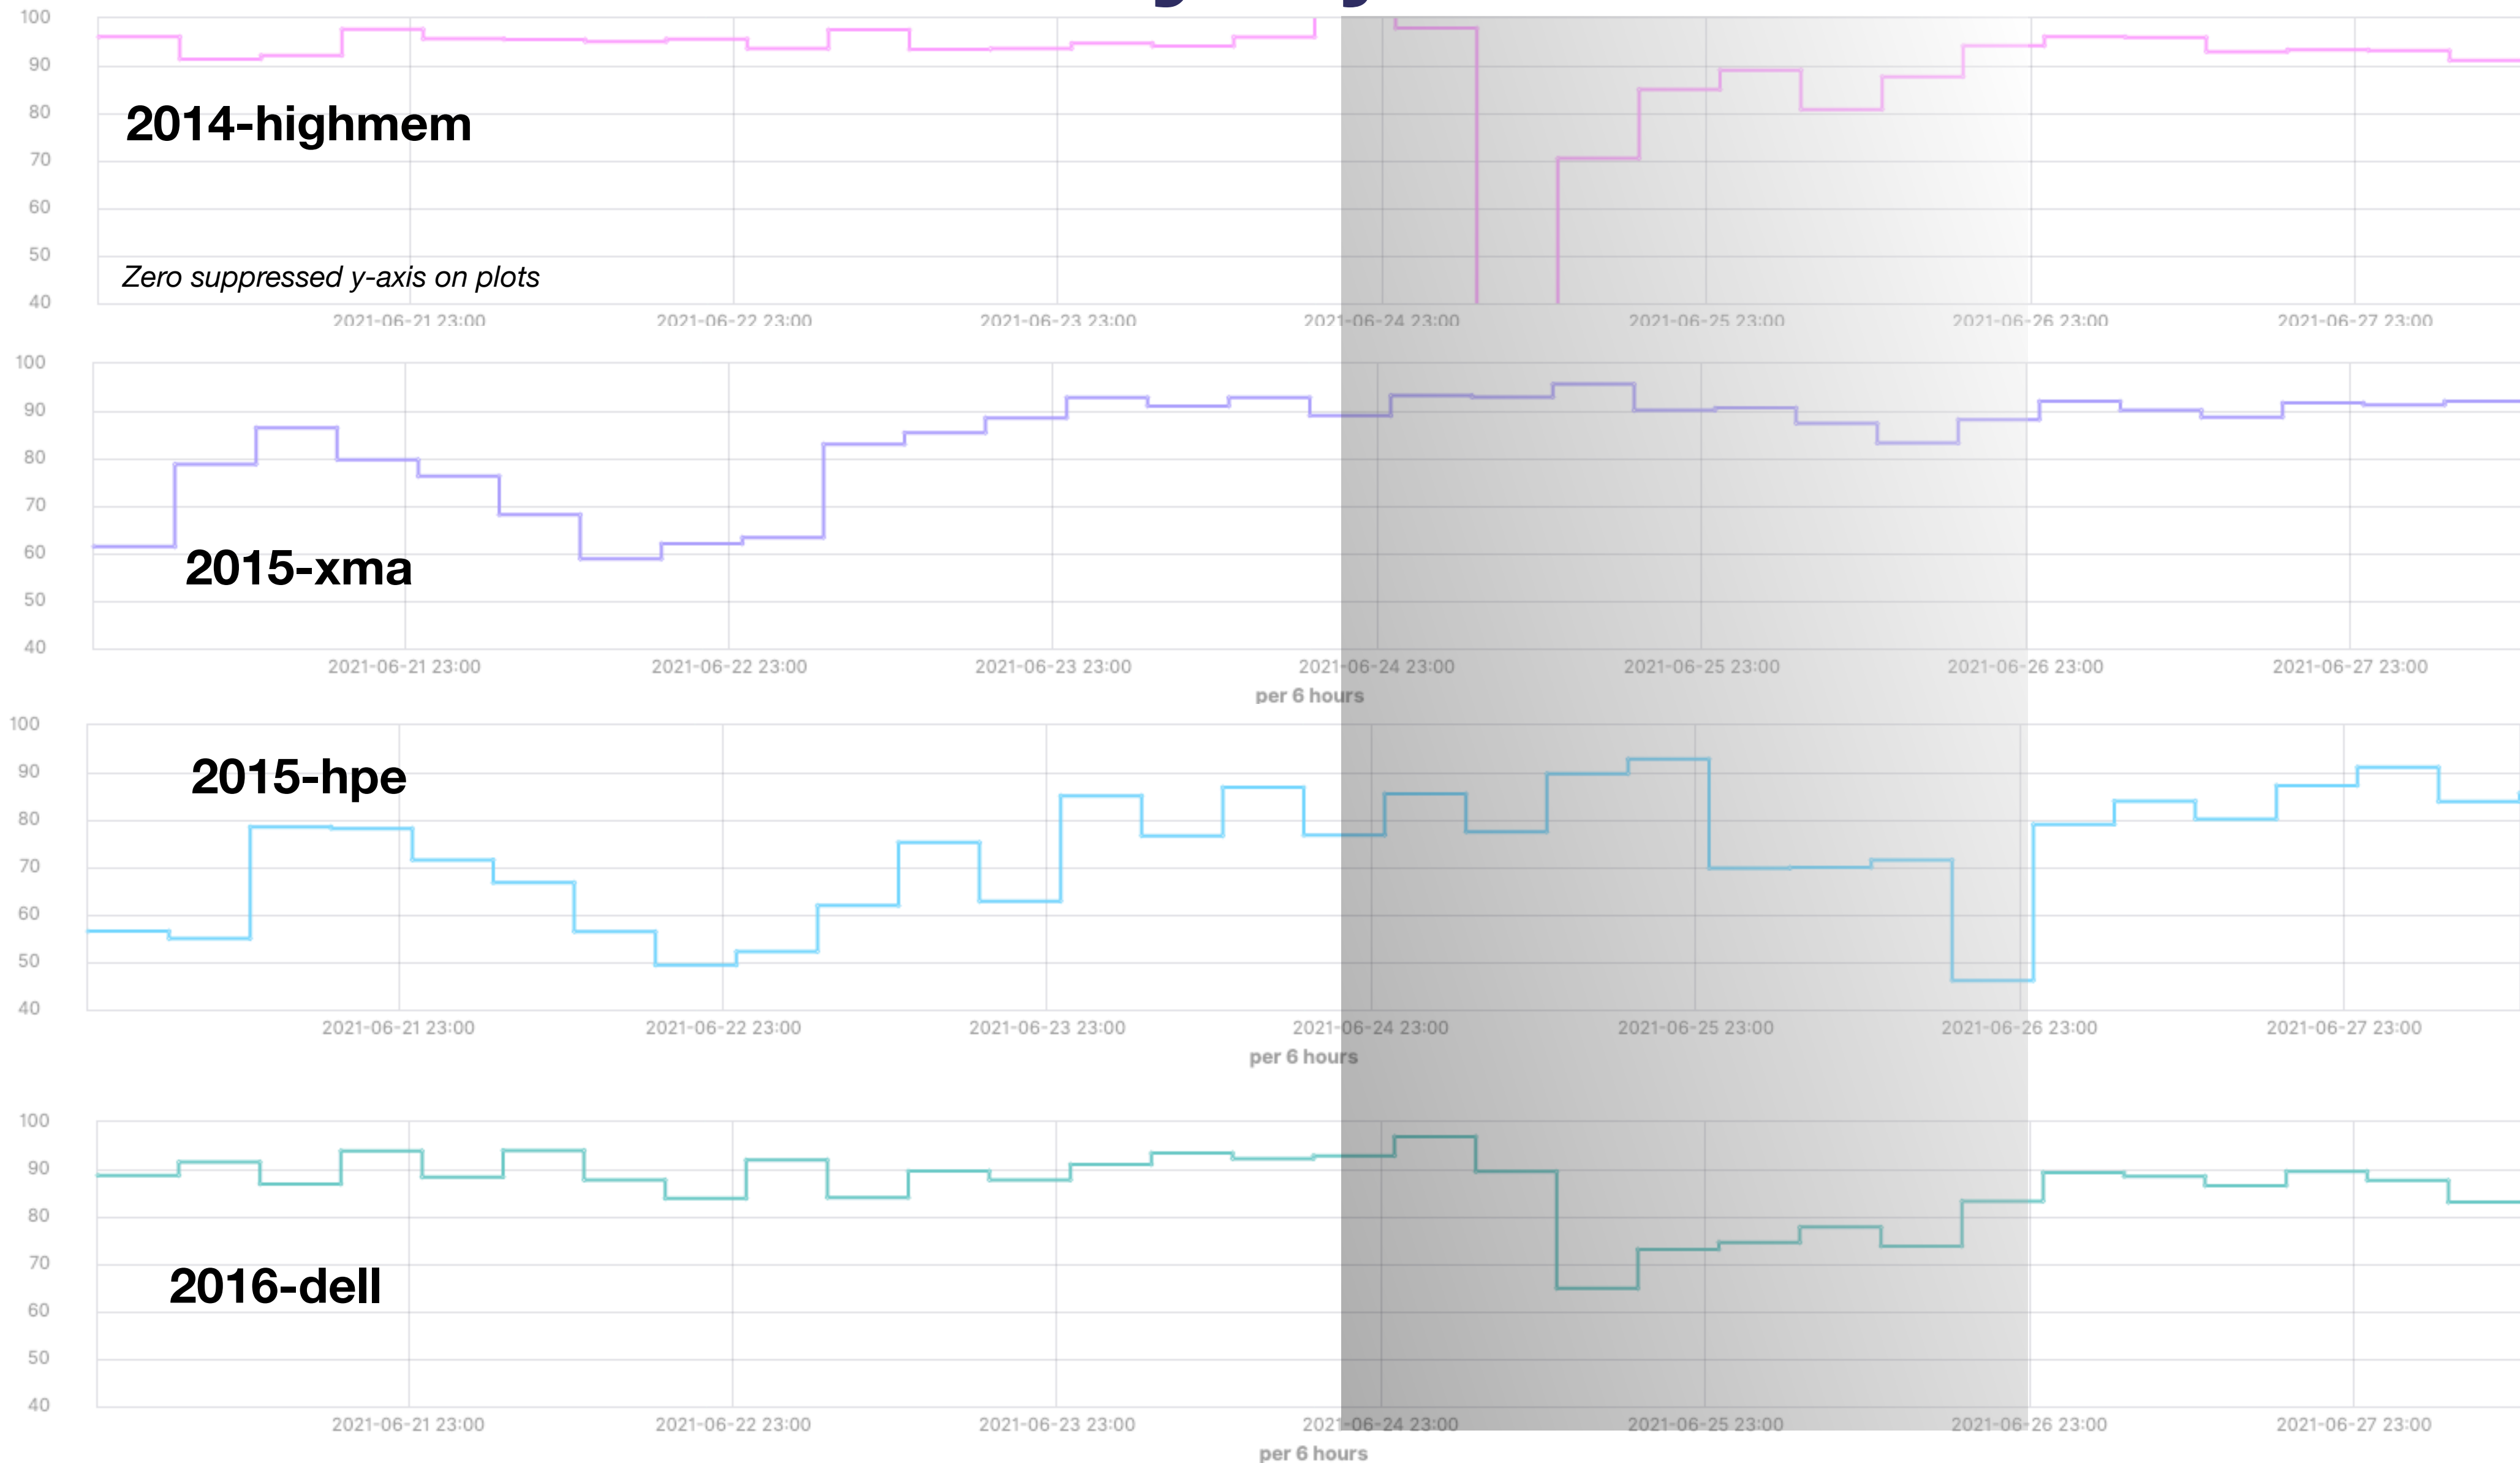
- 400k HS06 used vs 300k pledged ?
 - LHC VO pledge is 369k for 2021-22 ?!?



**Current reported HS06 per tranche
Totals to around O(400k) HS06**







CPU Efficiency / job waltime ratios (1)




CPU Efficiency / job waltime ratios (2)


No Change to job slots




ATLAS Monitoring plots







21-24th June



27-29th June



CMS and ATLAS



	min	max	avg	current	total
TOTAL	0	14.0 K	8.67 K	0	6.23 Mil
Production	0	11.5 K	7.00 K	0	5.02 Mil
Processing	0	7.32 K	1.46 K	0	1.05 Mil
Analysis	0	679	169	0	121 K
Merge	0	118	43.3	0	31.1 K
LogCollect	0	25.4	1.75	0	1.25 K
Cleanup	0	6.60	0.520	0	373

	min	max	avg	total
User Analysis	0	5.82 K	3.30 K	2.29 Mil
Group Production	0	8.80 K	3.20 K	2.23 Mil
MC Resimulation	0	7.44 K	2.34 K	1.63 Mil
MC Simulation Full	0	8.73 K	1.57 K	1.09 Mil
MC Event Generation	0	2.84 K	884	615 K
Group Analysis	0	2.85 K	745	518 K
MC Reconstruction	0	3.73 K	606	422 K
Data Processing	0	6.26 K	194	135 K
MC Simulation Fast	0	2.20 K	50.8	35.3 K
MC Merge	0	934	48.2	33.5 K
Testing	0	109	39.3	27.4 K

IO-wait comparisons


- 21-24th June



- 27-29th June; arguably higher IO-wait now; but more job-mix related ??



Vande Job Walitime ratios



atlas Current: 79.3% cms Current: 37.3% lhcb Current: 83.6%

True Mean Current: 60.7%

Summary

- Discrepancy between ATLAS and Vande;
 - On ATLAS, only have a site-averaged Corepower available
 - On Vande, how is it defined? Is it set to current values?
 - Might be useful to get plots of HS06 (and cores) groupby tranche, VO if possible ?
- For reduced-slot efficiencies:
 - 2014-highmem tranche:
 - Originally had as good an efficiency as the 2019's;
 - no need to have reduced these;
 - How does the slot reduction affect single-core and multicore workloads and distributions
 - Affected all VO's equally ?
- For ATLAS, Difficult to quantify any real change in efficiency; and differences more likely due to external (other VO?) or job mix factors.
 - Number of cores used has decreased from around 15k to 10k over June.

Backup



Finished Jobs

	Analysis	MC Simulation	Group Production	Others	MC Merge	Data Processing	MC Reconstruction
All	181,967	62,667	16,034	13,060	1,266	274	1,217
wn-2019-dell	62,500	26,591	8,699	3,647	314	179	621
wn-2018-xma	26,373	7,906	233	1,800	217	6	7
wn-2017-xma	34,920	9,510	491	3,090	321	10	26
wn-2017-dell	40,337	11,282	985	2,455	283	22	63
wn-2016-dell	12,486	2,947	666	660	95	9	54
wn-2015-hpe	2,000	1,950	2,221	688	18	24	234
wn-2015-xma	41	1,503	2,402	561	4	20	196
wn-2014-viglen-highmem	3,000	845	292	134	13	4	14

- 1 - 100,001
- 100,001 - 200,001
- 200,001 - 300,001
- 300,001 - 400,001
- 400,001 - 500,001
- 500,001 - 600,001
- 600,001 - 700,001
- 700,001 - 800,001
- 800,001 - 900,001
- 900,001 - 1,000,000



Finished Jobs

21-24 June

27-29 June

	Analysis	MC Simulation	Group Production	Others	MC Merge	Data Processing	MC Reconstruction
All	99,355	26,562	6,067	5,883	677	81	374
wn-2019-dell	31,374	10,075	2,950	1,588	164	32	178
wn-2018-xma	13,376	3,021	158	704	102	4	2
wn-2017-xma	17,927	3,848	260	1,340	152	5	4
wn-2017-dell	20,354	4,486	398	928	134	7	18
wn-2016-dell	11,320	2,229	118	529	90	6	1
wn-2015-hpe	1,963	1,244	967	399	18	12	92
wn-2015-xma	32	875	1,182	284	4	13	79
2014-viglen-highmem	2,973	743	12	99	13	2	

Top 12 unusual terms in data.resourcesreporting

	Analysis	MC Simulation	Group Production	Others	MC Merge	MC Reconstruction
All	33,526	31,694	6,034	4,403	727	609
wn-2019-dell	10,484	14,914	2,983	1,206	172	315
wn-2018-xma	6,015	4,274	40	624	151	4
wn-2017-xma	7,847	4,844	107	1,056	192	16
wn-2017-dell	8,545	5,883	282	939	203	26
wn-2016-dell	518	656	322	87	8	42
wn-2015-hpe	11	510	1,068	236		118
wn-2015-xma	8	448	1,059	224		77
wn-2014-viglen-highmem	1	79	166	21		10

Top 12 unusual terms in data.resourcesreporting

