



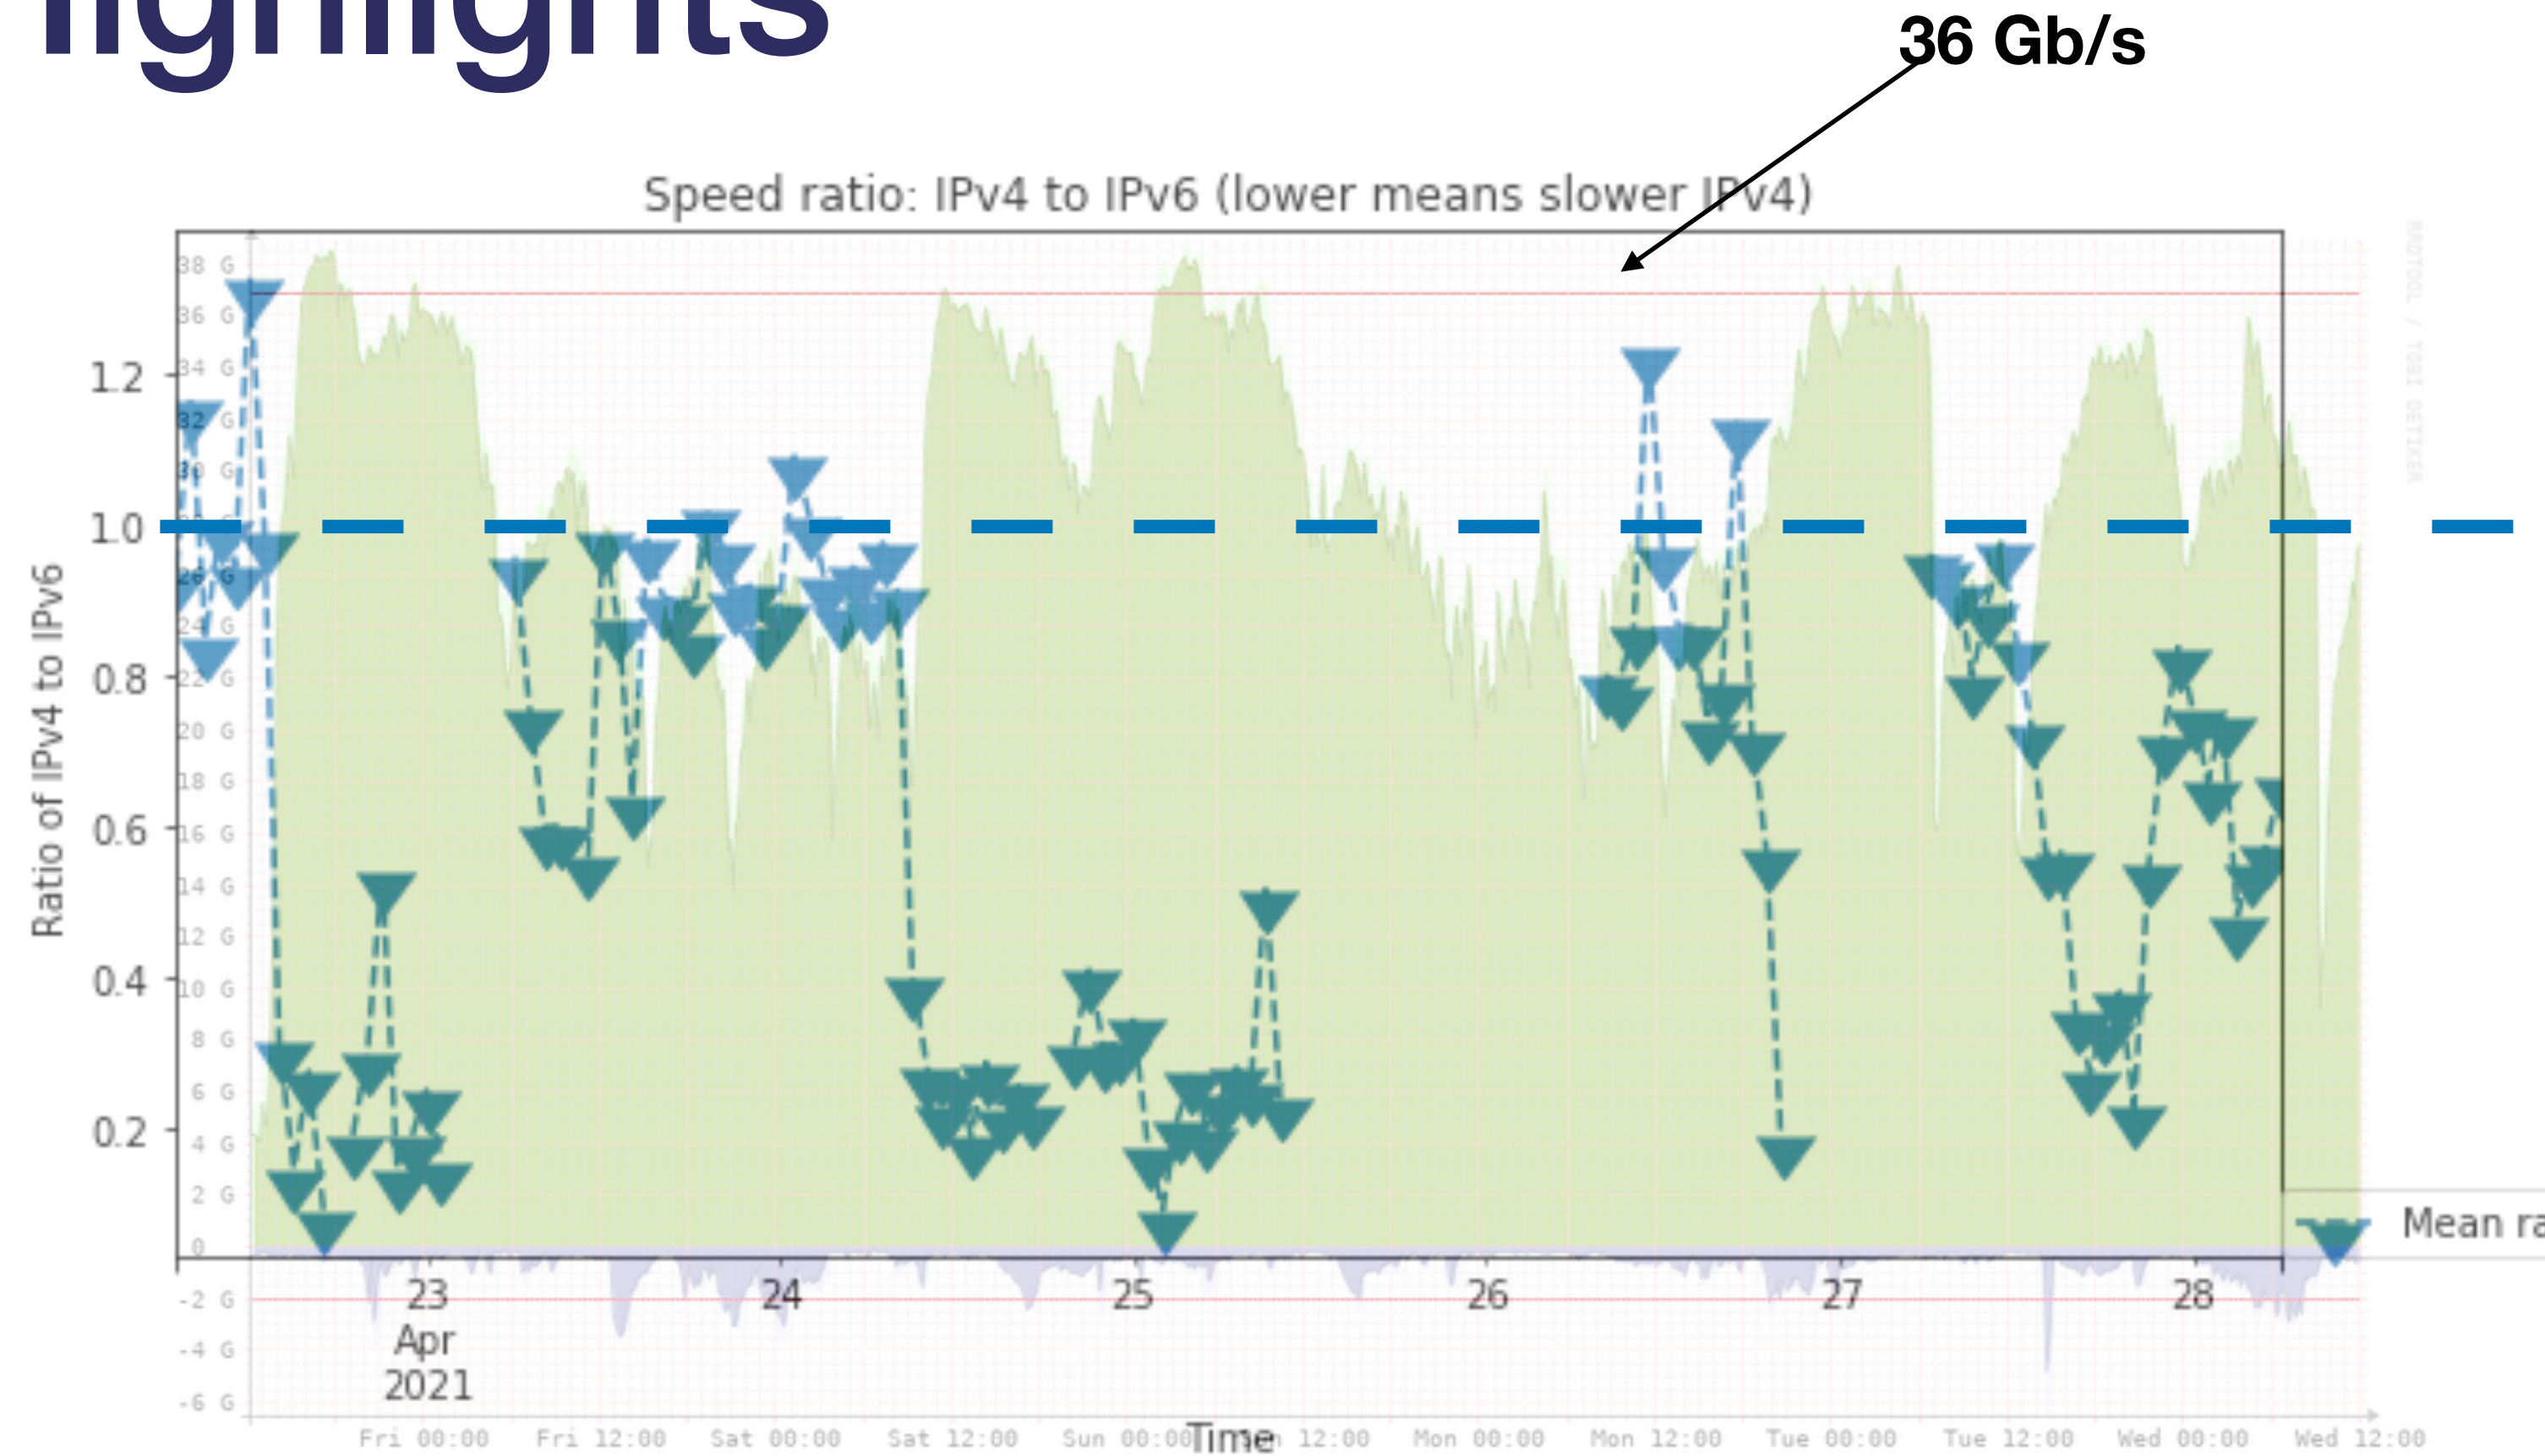
Science and
Technology
Facilities Council

Resources Review Meeting: 2021 Q1 ATLAS RAL

James Walder
28 04 2021

Highlights

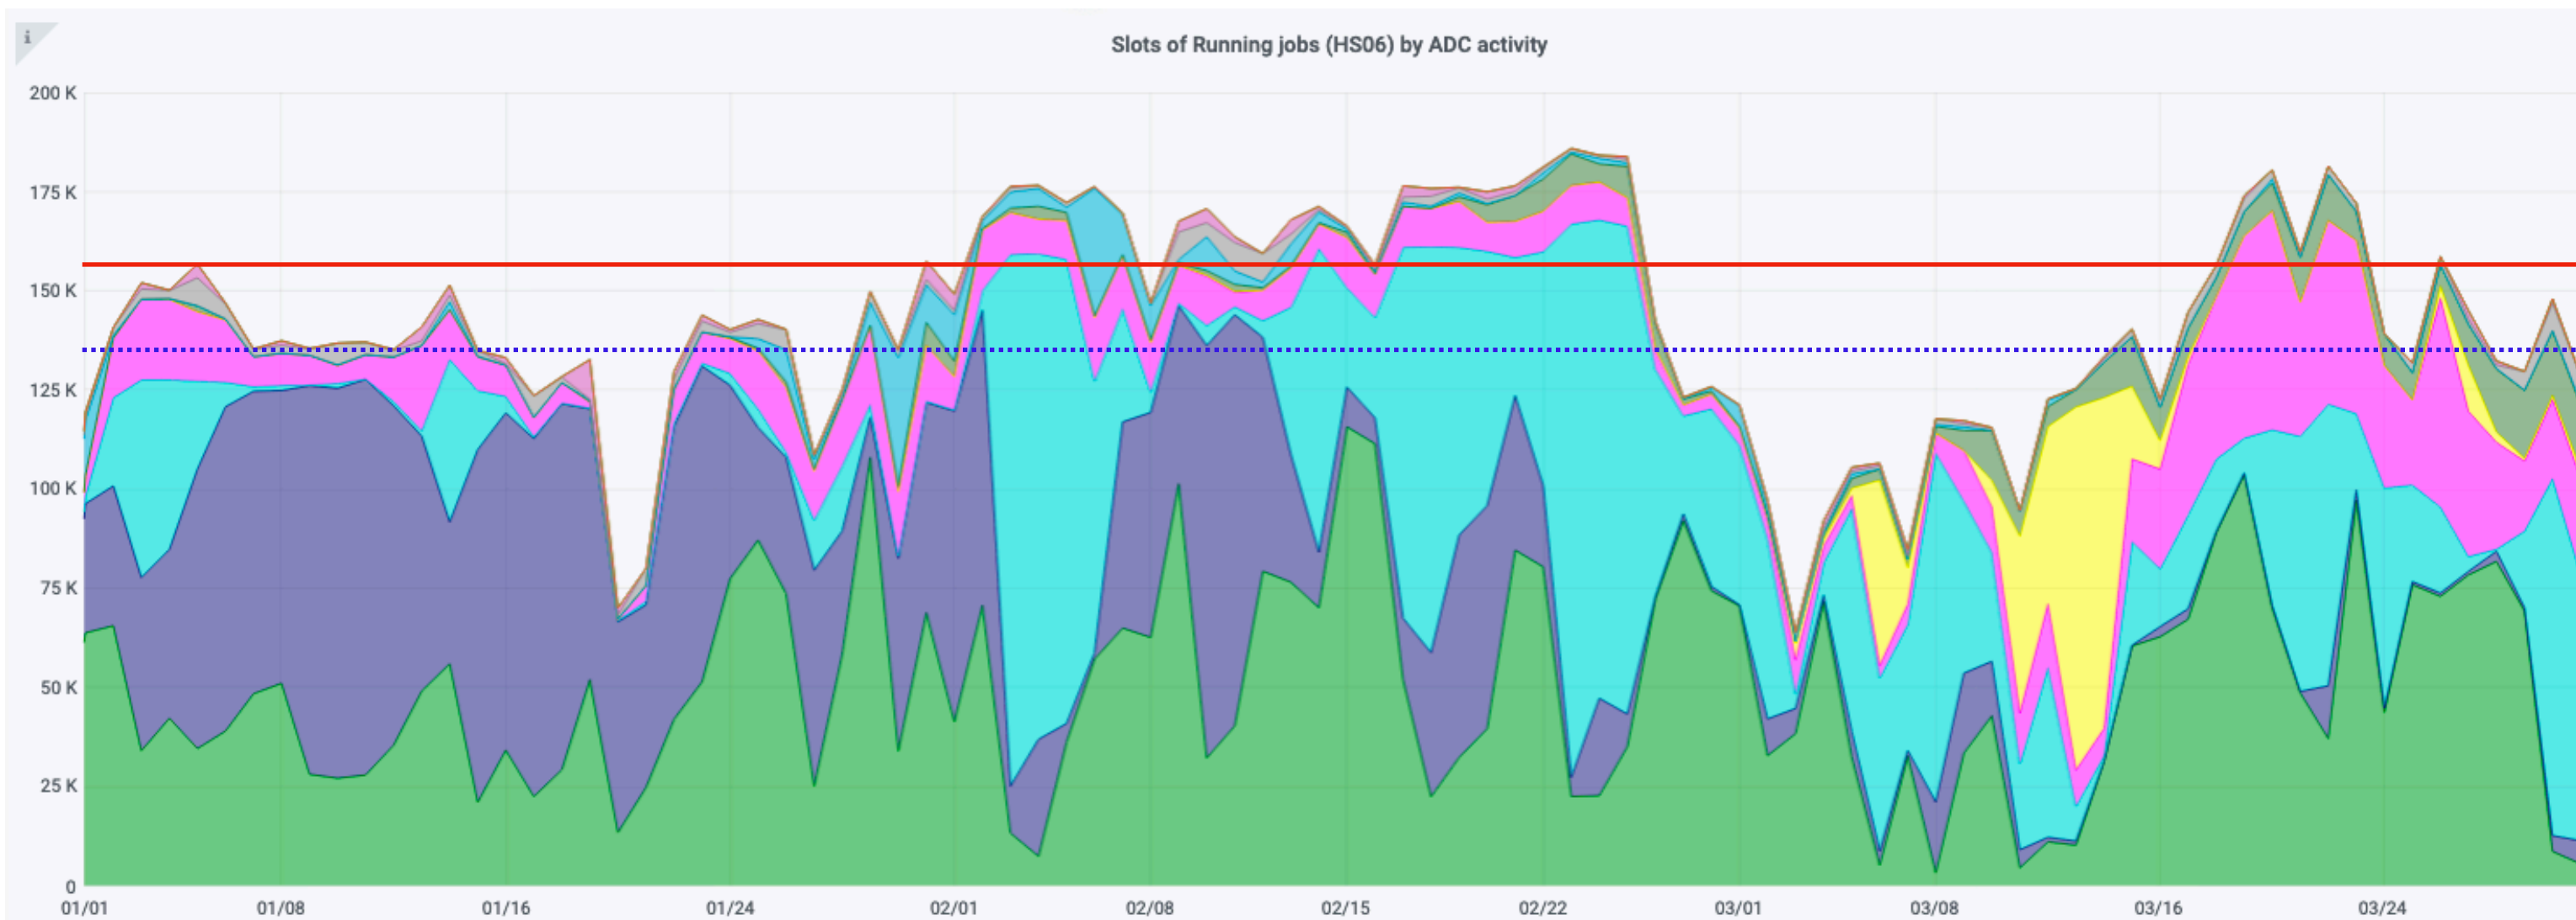
- Reporting period:
 - Q1 2021: 1 Jan - 31 March
- Generally quiet from ATLAS side
 - Largely business-as-usual
 - Effort to push more Analysis jobs to RAL
 - ATLAS currently using direct-IO for analysis work - will benefit from vector read support progress
- RAL-LCG2 provided its CPU, Disk and TAPE:
 - No major issues
 - Tape repack complete at start of Q1
- Anticipated deployments of: Vector Read support, New WNs (with SSDs), future networking, XrootD 5.X (with WebDav TPC) to come.



- Plot on right; recent firewall upgrade at RAL:
 - IPv4 File transfer speeds from UI nodes (similar to the WNs), relative to their ipv6 performance, averaged over a number of different external sites.
- Superimposed (in green):
 - Network load on core router;
 - Various changes/updates over last week, but some association between performance and load?

Resource	2020/21	2021/22
CPU[kHS06]	156.436	173.16
DISK [PB]	13.024	15.54
TAPE [PB]	32.708	34.78

CPU resource




	min	max	avg	current	total
Group Production	3.24 K	116 K	49.8 K	5.50 K	4.53 Mil
MC Reconstruction	0	104 K	32.9 K	5.81 K	3.00 Mil
MC Simulation Full	99.7	140 K	31.7 K	67.1 K	2.89 Mil
User Analysis	231	55.2 K	14.2 K	24.3 K	1.29 Mil
Data Processing	0	91.5 K	4.23 K	2.25 K	385 K
MC Event Generation	0	17.1 K	3.36 K	16.2 K	306 K
MC Simulation Fast	0	32.5 K	2.22 K	634	202 K
Group Analysis	40.3	7.17 K	1.96 K	6.72 K	179 K
MC Merge	0	10.0 K	884	61.2	80.5 K

- 2020–21 pledge: 156.4 kHS06
- Average over period: **164** = **141***(11.7/10) kHS06:
 - Scaling accounts for difference between CRIC Corepower (HS06/core) and farm average.
- Some variation due to various VO utilisations and ATLAS/CERN issues.
- User analysis jobs from ~ 9kHS06 -> 30kHS06
 - Improvements still ongoing; attempting to move to multi-job pilots

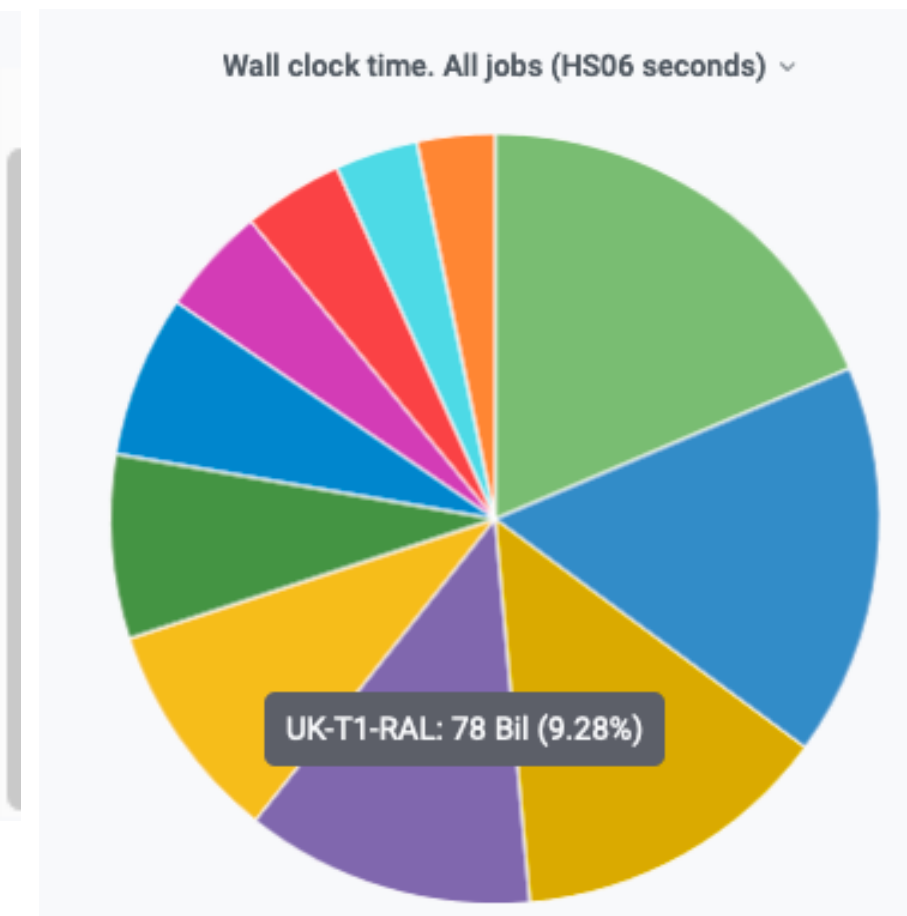
T1 comparison

- ~88% CPU eff. for RAL, cf. ~91% (all T1)

~ 9% of supplied T1 Wall and CPU time



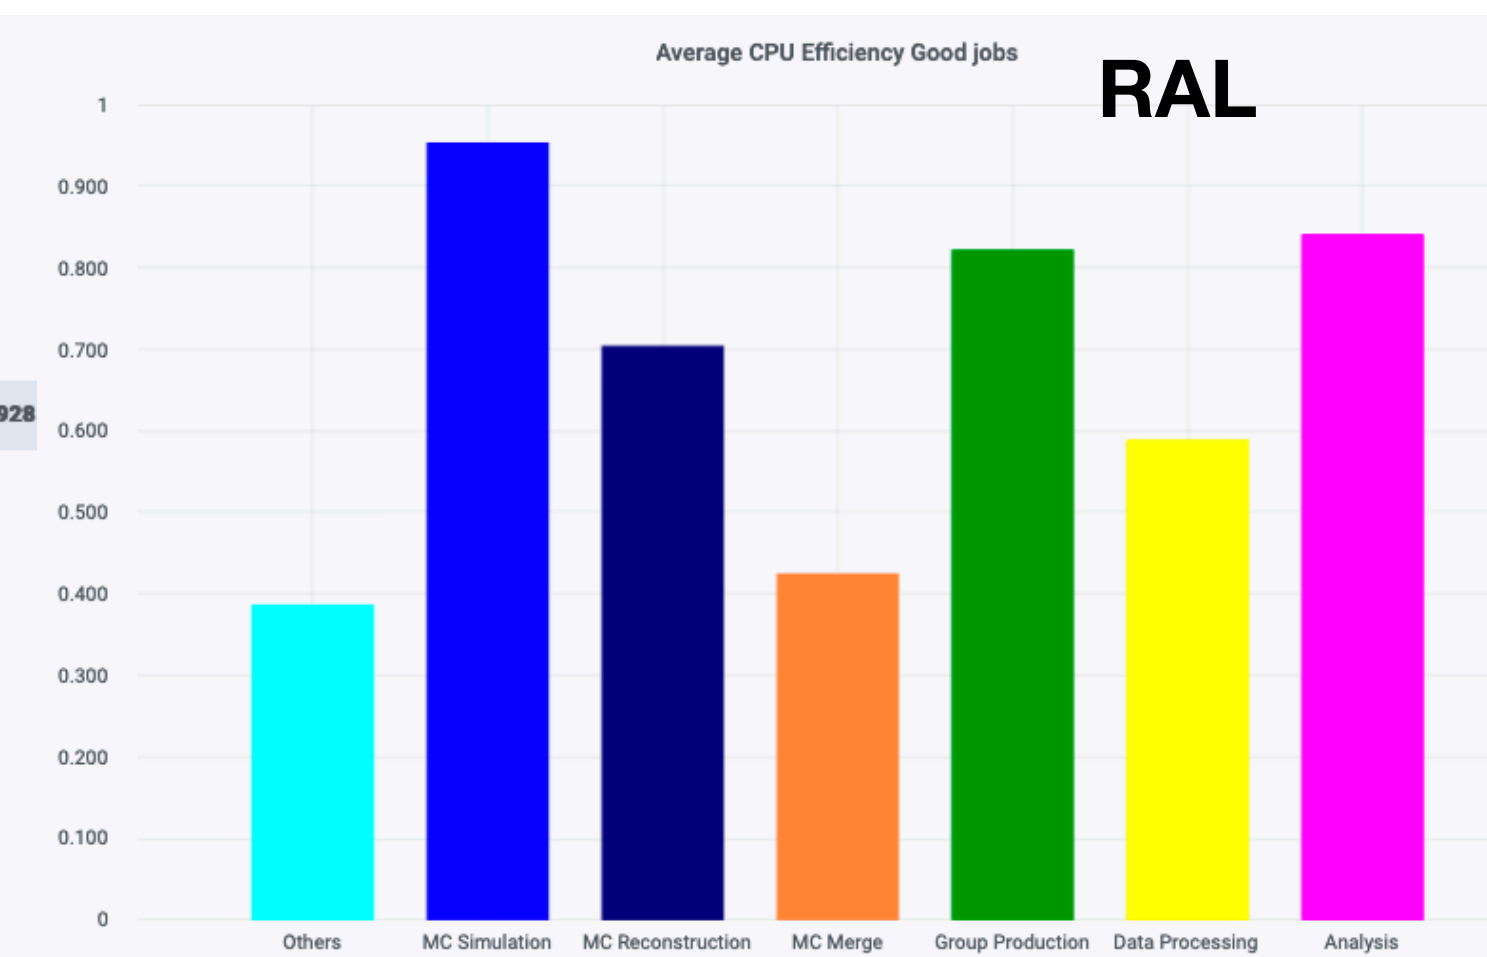
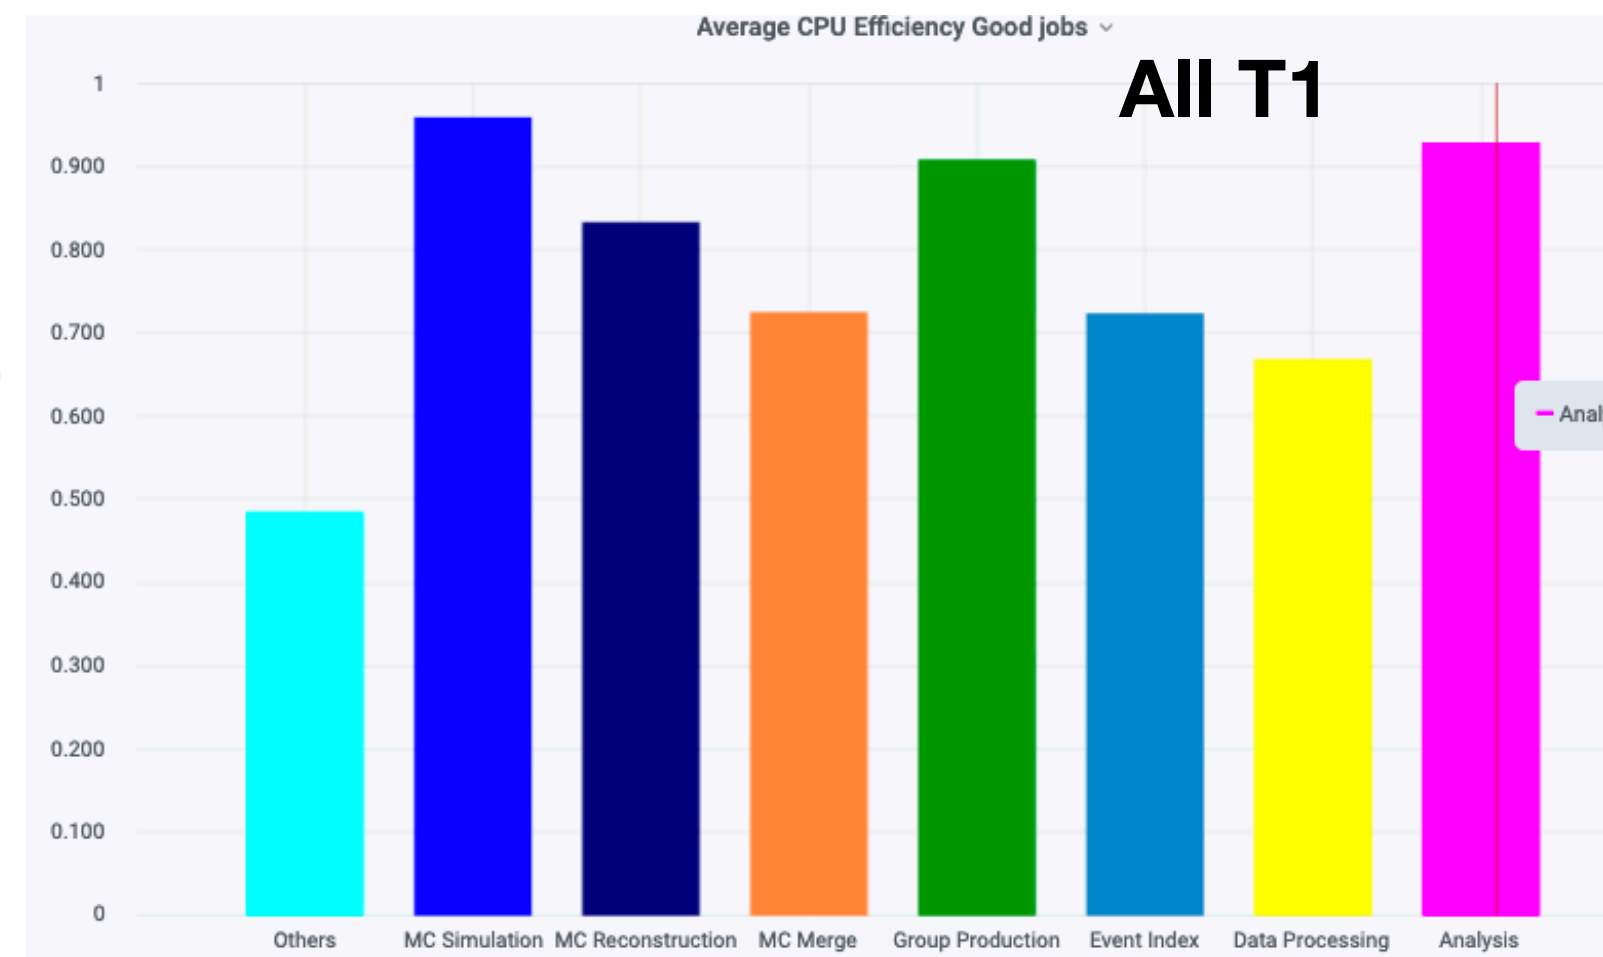
Site	Billions of Jobs	Percentage
US-T1-BNL	157 Bil	19%
CA-TRIUMF	139 Bil	16%
DE-KIT	114 Bil	13%
FR-CCIN2P3	103 Bil	12%
UK-T1-RAL	78 Bil	9%
ES-PIC	65 Bil	8%
NDGF	57 Bil	7%
IT-INFN-CNAF	39 Bil	5%
NL-T1	35 Bil	4%
TW-ASGC	29 Bil	3%
NRC-KI-T1	28 Bil	3%



- As usual; much variability across activity types:

- T1 also share different workloads
- IO intensive tasks (e.g merging) show larger discrepancies for RAL

- Anticipated improvement should appear through vector reads, networking, new WNs.



WN transfer speeds


• Right: Typical staging speeds from Echo to WNs:

- File download, excluding direct-IO
- Separated by tranche and file-size
- Small files suffer large overheads

• Extracted from VO monitoring data



• SSD-based tranche continues to show good performance

- Work currently ongoing to improve WN CPU eff.




WN transfer speeds (2)

- Violin plot of transfer speeds (with median + quantiles shown)
- Extracted from individual transfers on WNs within each tranche using local XrootD proxies (similar to prod. traffic) using ~1GB file.
- General agreement with VO-derived plots
- SSD 2019 tranche not entirely immune from situations of reduced transfer speeds.
- Some association (below) between Pressure on the node and speed.



Disk and Tape

- **Disk:**
 - ~12PB quota in period, fully utilised
 - ~ 6PB of deletions
 - Including within-site transfers (e.g. to WNs), estimate \gg 10 PB of transferred data



- **Tape**
- Limited activity on DATATAPE
- Data policy changes \Rightarrow TAPE-only copies of AOD
- Expect more active deletion cycles for (e.g) MCTAPE than in previous LHC Runs.
- ~ 26 PB data on tape over period
 - ~1PB transfered in.
- No significant deletions over period

