LEGO train limits

Offline week,

31/03/2016

LB

Incident of 22/03/2016

- A train operator has submitted 107 LEGO trains in one go
  - These resulted in more than 4K master jobs (jobs that contain many sub-jobs, need to be split and optimized)
  - Kept the optimizer busy for 10 hours, blocking other requests (the remaining PWGs, users, production)
- The train operator was not aware of limits
  - Although warned when he became one
- The train were all on ESDs, making them heavier…
- Not much going on these days, limited impact, only one user complained directly

Current train agreement

- LEGO train have absolute priority (over all other jobs)
- There is a gentlemanly agreement that
  - Run max 30K jobs/day (not enforced, often surpassed)
  - Self-imposed limits within the PWG operators on number and content of trains
  - AOD trains have priority over ESD (submitted first, enforced by system) but not if a single PWG ‘grabs’ all available resources before everyone else
- Single Grid user
  - No need of ‘per-PWG’ job quotas (advantage)
  - If one PWG does something ‘bad’, all suffer (disadvantage)
  - Requires communication and respect

Train operation

- No job limits presently, very favorable for periods of high demand

	Series	Last value	Min	Avg	Max
1.	alitrain	18883	0	10366	65562

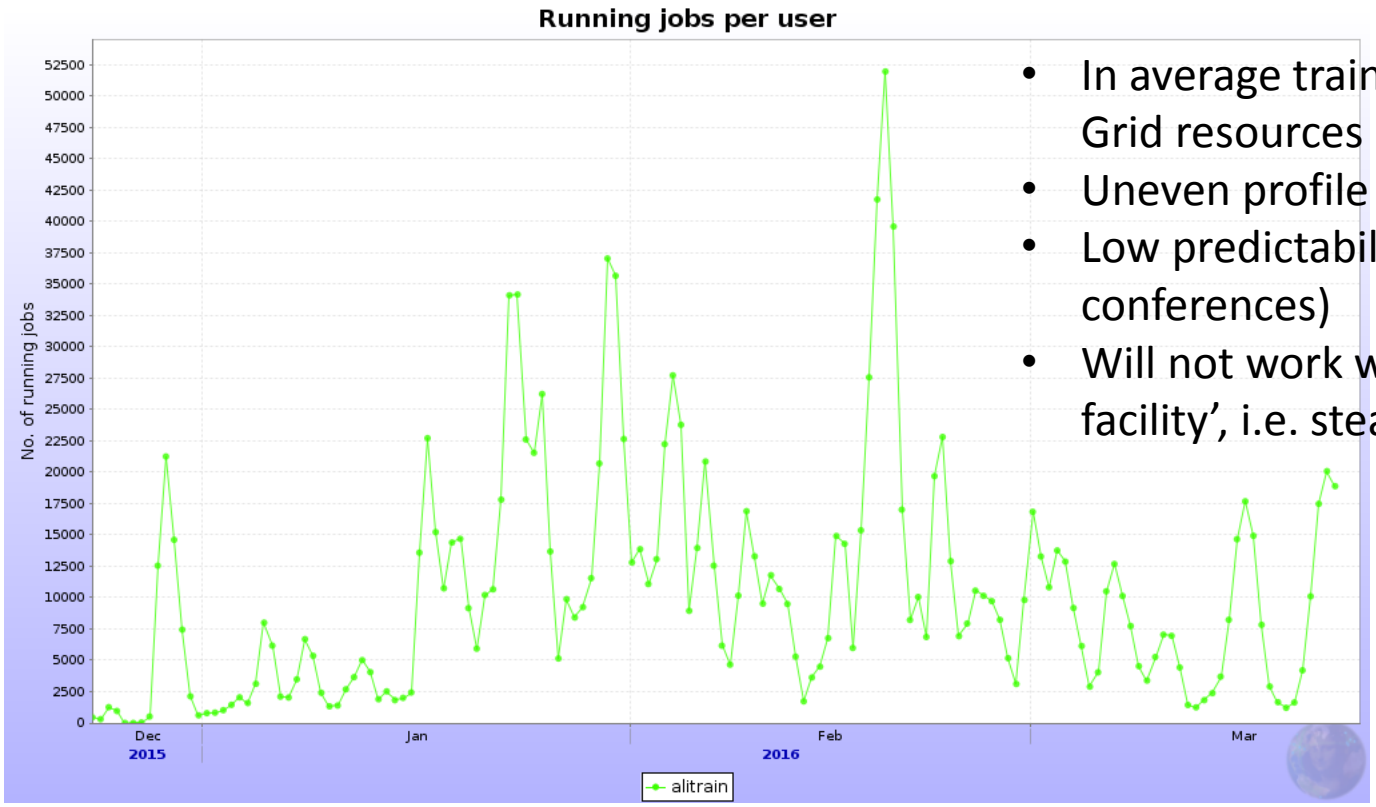

- In average trains use ~16% of Grid resources
- Uneven profile
- Low predictability (except conferences)
- Will not work well with ‘analysis facility’, i.e. steady-state analysis

Train limits - implementation

- Already discussed and agreed last year
  - Per-PWG accounts for train operation
  - Per-PWG limits
- This is mostly implemented but not activated as it has several disadvantages
  - Will actually require to put a cap on running LEGO jobs
  - Will not compensate as well as current ‘single user’ for under/overuse in case of quota not used by PWG(s) – usual principle is ‘use it or lose it’
  - Will require a long and difficult discussion on quotas per PWG

Alternative proposal

- Put caps on submission, not use
- Global (or per-PWG) number of trains per day
  - Advantage: will smooth train running over time - even resources use
  - Advantage: avoid mistakes and system being blocked by a single operator
  - Advantage: allows not to impose limits, keep current analysis priority high
  - Advantage: Favors more compact AOD trains automatically (these are submitted first) + additional optimization (see ‘statistics’ later)
  - Disadvantage: requires planning in case of heavy analysis programme, but only within individual PWGs

Alternative proposal - implementation

- Limit to ***x*** the number of LEGO trains per day
  - Soft limit – allows to submit all trains in one go, the excess over ***x*** is submitted and executed automatically the next day(s)

OR

  - Hard limit – if ***x*** is reached operator receives “You have reached your daily quota” error

Alternative proposal - limits

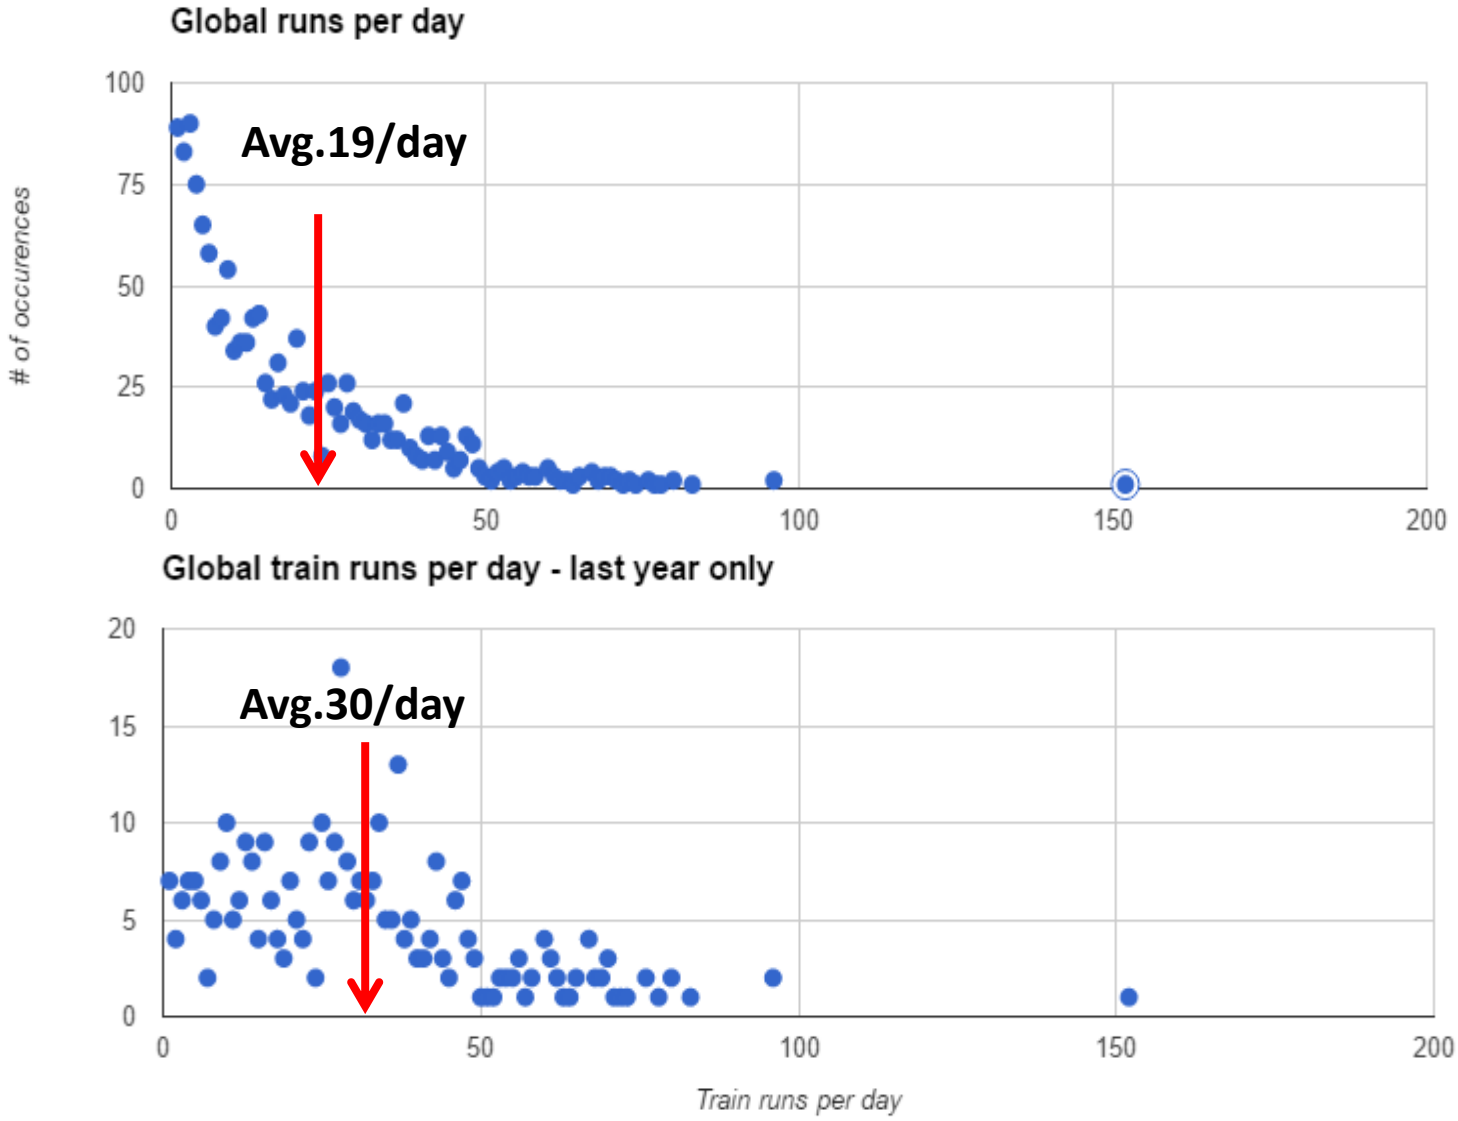

Av. Number of wagons and runs/day

	GA	HF	CF	DQ	JE	LF	UD	Tot.	
2014-2015	3 w 8 r	7.3 w 6 r	3.2 w 5.6 r	1.1 w 2.6 r	9.9 w 5 r	2.6 w 3.5 r	0.9 w 1.4 r	5.1 w 20.3 r	
2015-2016	2.6 w 13 r	6.5 w 7.6 r	2.8 w 7.4 r	1.4 w 4.8 r	9.5 w 5.3 r	10 w* 3.1 r	0.9 w 2.3 r	4.6 w 30.3 r	

Subtracted from trains: tender/centrality/PID wagons

- Overall increase of runs/day by 50%
- Overall decrease of wagons/train by 10%

* biased by few 500 wagon-long tests

Wagons per train global

2014-2015		2015-2016		
wagons	**count**	**wagons**	**count**	
1	1837	1	2938	**27.16%**
2	1300	2	2263	**20.92%**
3	915	3	1891	**17.48%**
4	671	4	1175	**10.86%**
5	423	5	493	**4.56%**
6	345	6	448	**4.14%**
7	200	7	358	**3.31%**
8	244	8	152	**1.40%**
9	118	9	105	**0.97%**
10	211	10	102	**0.94%**
11	180	11	70	**0.65%**
12	239	12	204	**1.89%**
13	55	13	118	**1.09%**
14	54	14	55	**0.51%**
15	67	15	76	**0.70%**
16	110	16	41	**0.38%**
17	19	17	20	**0.18%**
18	40	18	21	**0.19%**
19	24	19	12	**0.11%**
>=20	218	>=20	277	**2.56%**
5.13	Average	4.68	Average	

Efficiency

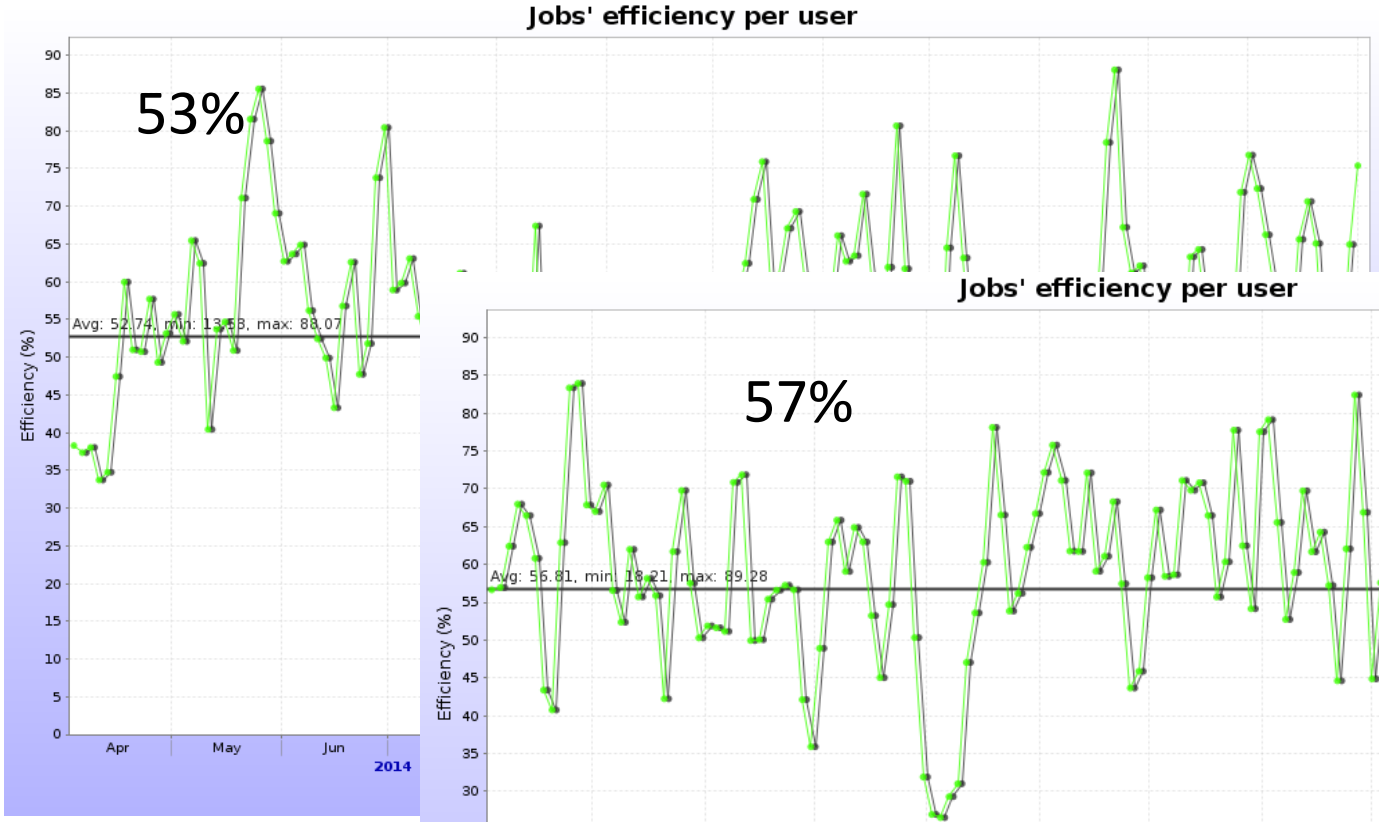
Jobs' efficiency per user
53%
Avg: 52.74, min: 13.58, max: 88.07
Efficiency (%)
Apr
May
Jun
2014

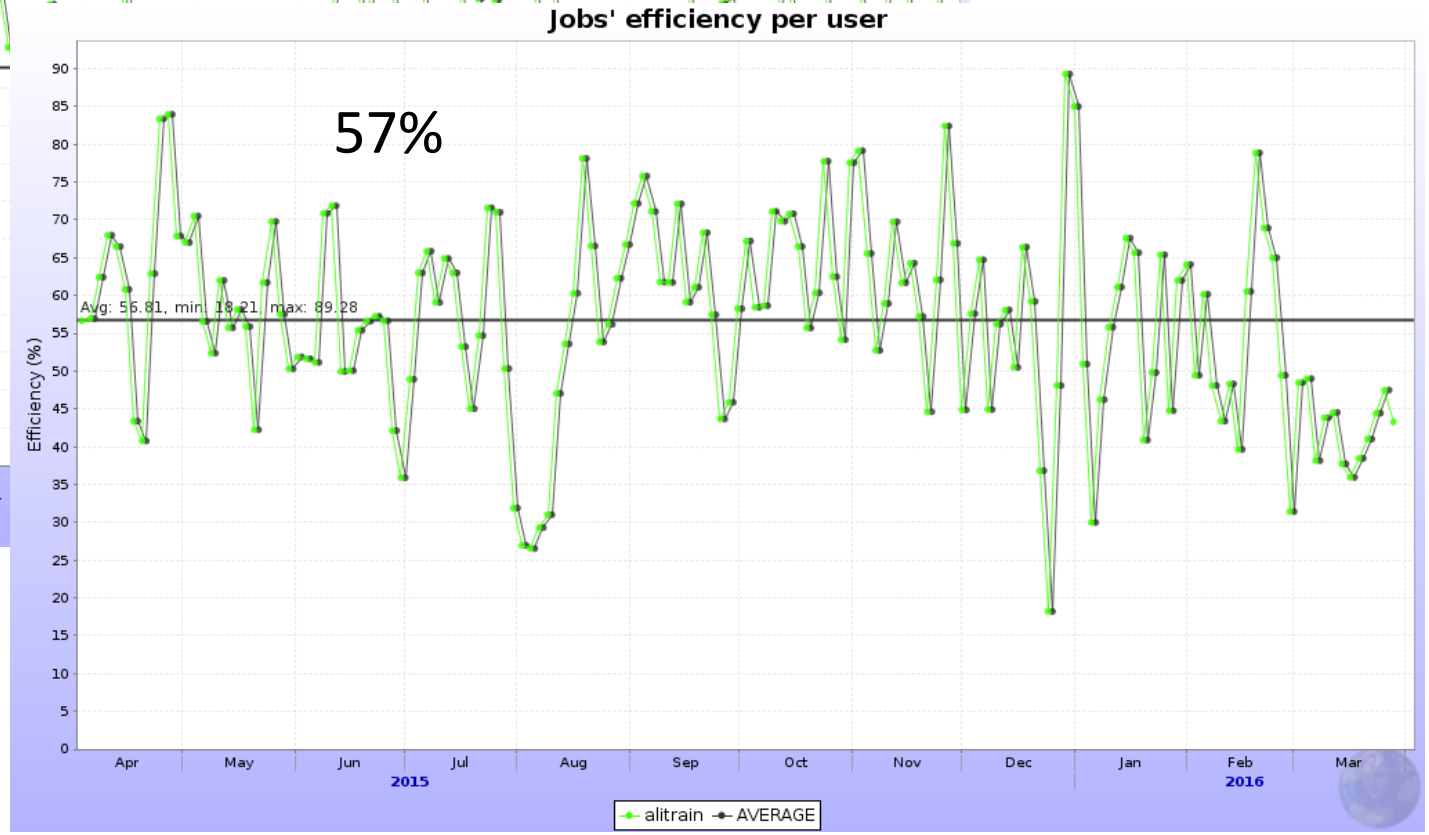
Jobs' efficiency per user
57%
Avg: 56.81, min: 18.21, max: 89.28
Efficiency (%)
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
Jan
Feb
Mar
2015
2016
alitrain
AVERAGE

Resources

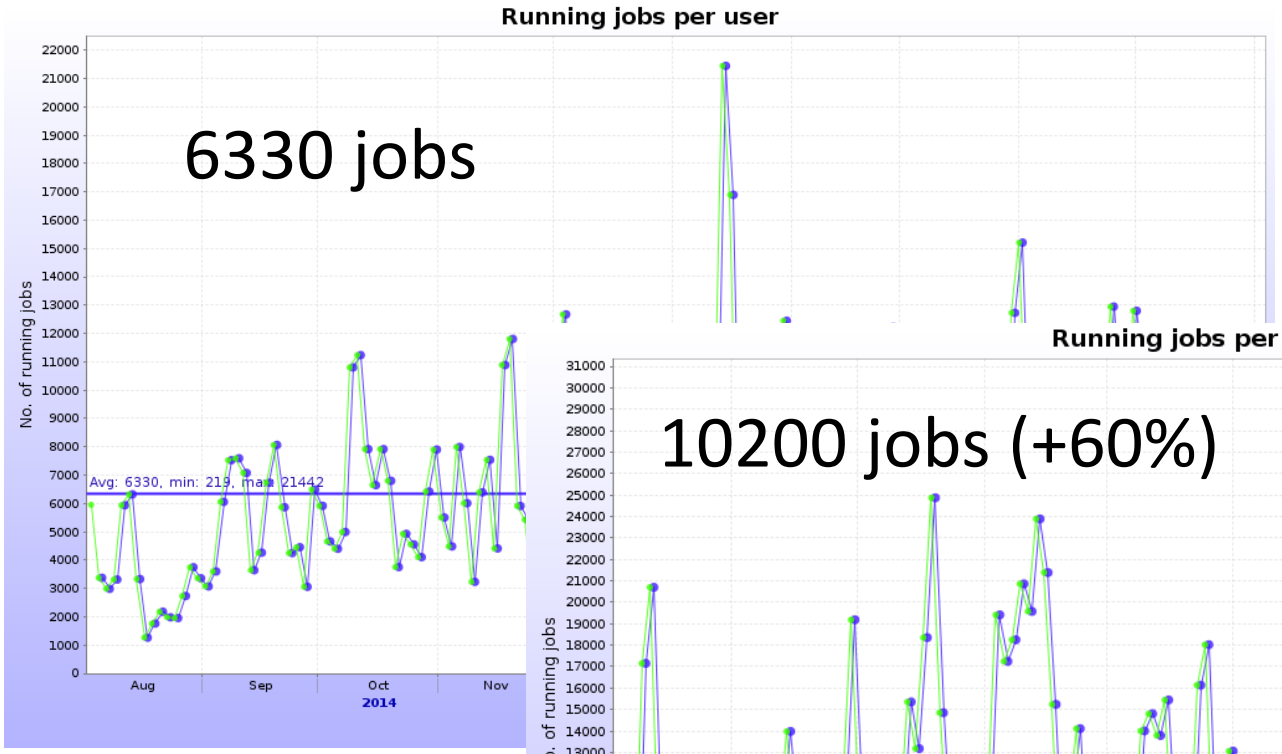

6330 jobs

Growth of Grid resources
+25% per year

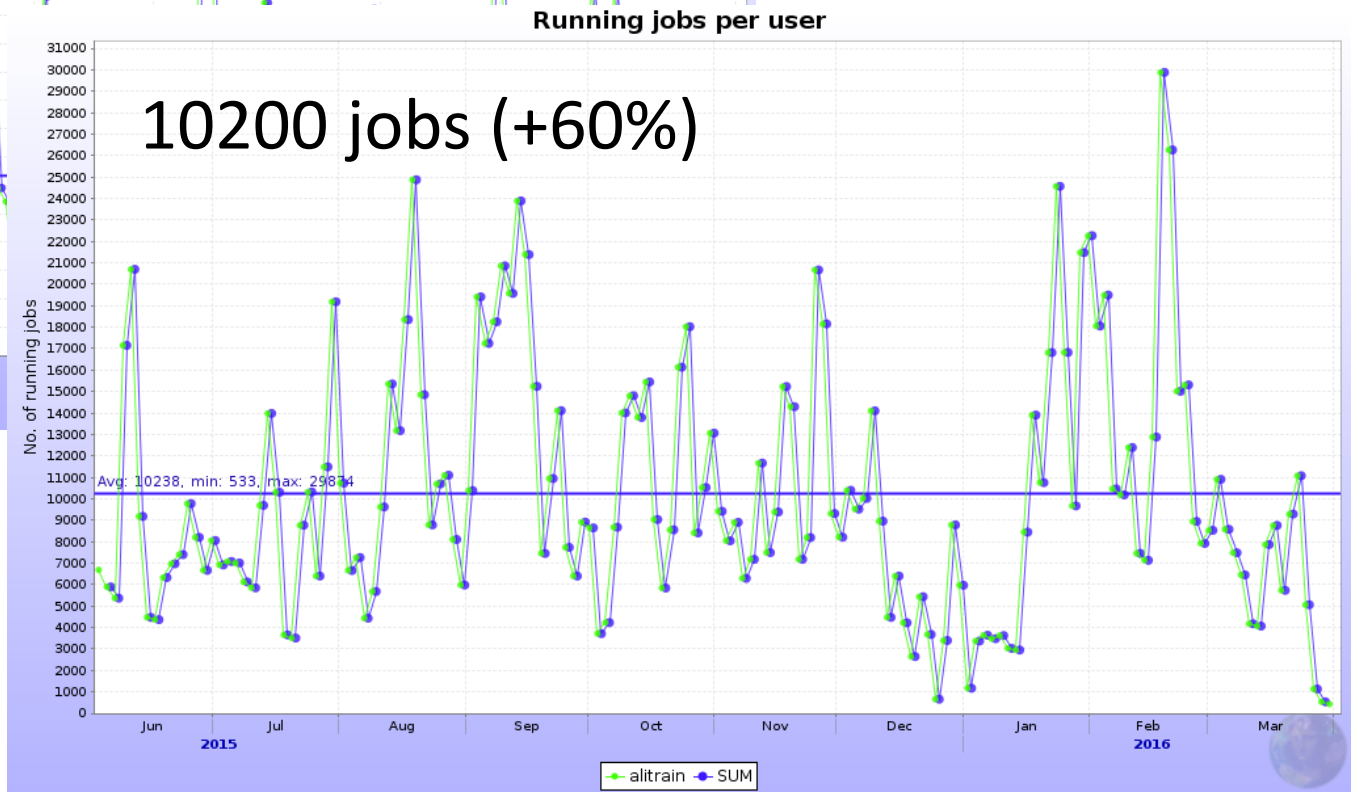

10200 jobs (+60%)

Individual user analysis
9% of total in 2014-2015
6% of total in 2015-2016

Why now

- Period of low activity – we can discuss and test limits without impact on PWG work
  - PWGs can work on trains optimization
- Avoid incidents in periods of high demand
  - The 22/03 type incident just before QM will have larger consequences
- Some limits must be imposed, as the appetite for trains grow
  - Last year 50% more trains than year before, 60% more CPU
  - Resources are limited, +25% CPU per year growth
- Preliminary feedback is positive, but tuning is necessary and will take time

Summary

- We have started (again) the discussion on LEGO train limits
- We favor simple limit(s), easily tunable and PWG-specifics independent
  - Hopefully, can be formalized and agreed upon without too much delay
  - Will be fair and uniform, prevent incidental blocking
  - Leave room for (as presently) cross-PWG agreements
  - Will not limit the overall resources used for analysis
- Will spur some optimization of the train setup and actions from PWG to fit within reasonable constraints
  - Hopefully the ESD to AOD migration will receive a boost
  - As well as train set optimization