CMS

Liz Sexton-Kennedy (FNAL)

Tommaso Boccali (INFN-Pisa)

(on behalf of CMS Sw/Comp)

Outline

- Status of Activities in 2017
- Problems / issues
- A preview of 2018-2019 strategy / requests
- News from the project / related

2016 data/MC close-out

- Data
  - Complete ReReco “18Apr” concluded in Mid June (6B events)
  - Sub-optimal ECAL Endcap performance, new Complete ReReco “07Aug” under production (lower priority than 2017 MC)
- MC
  - Production for 2016 configuration had tails much beyond the 10B events used for Moriond17
  - Currently, another low priority / low selection efficiency 10B+ injected
  - Planning a full AODSIM→ MiniAODSIM reprocessing to have consistent data formats with data ReReco (no need/capability for a full reprocessing)

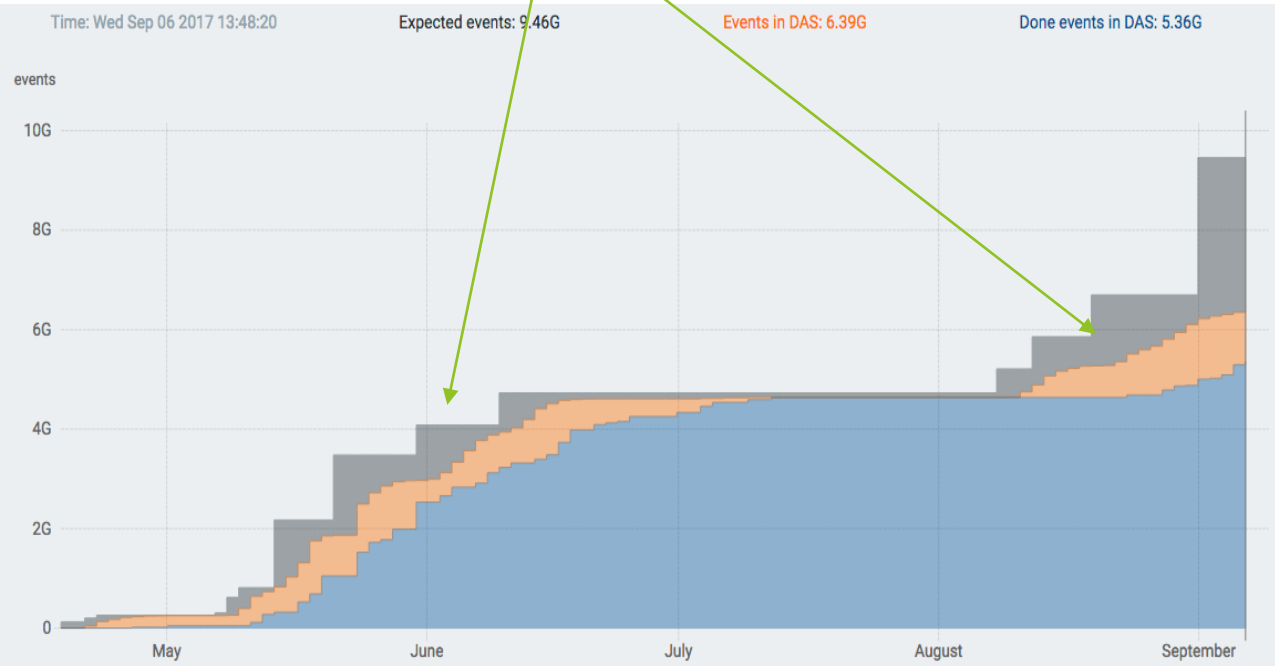

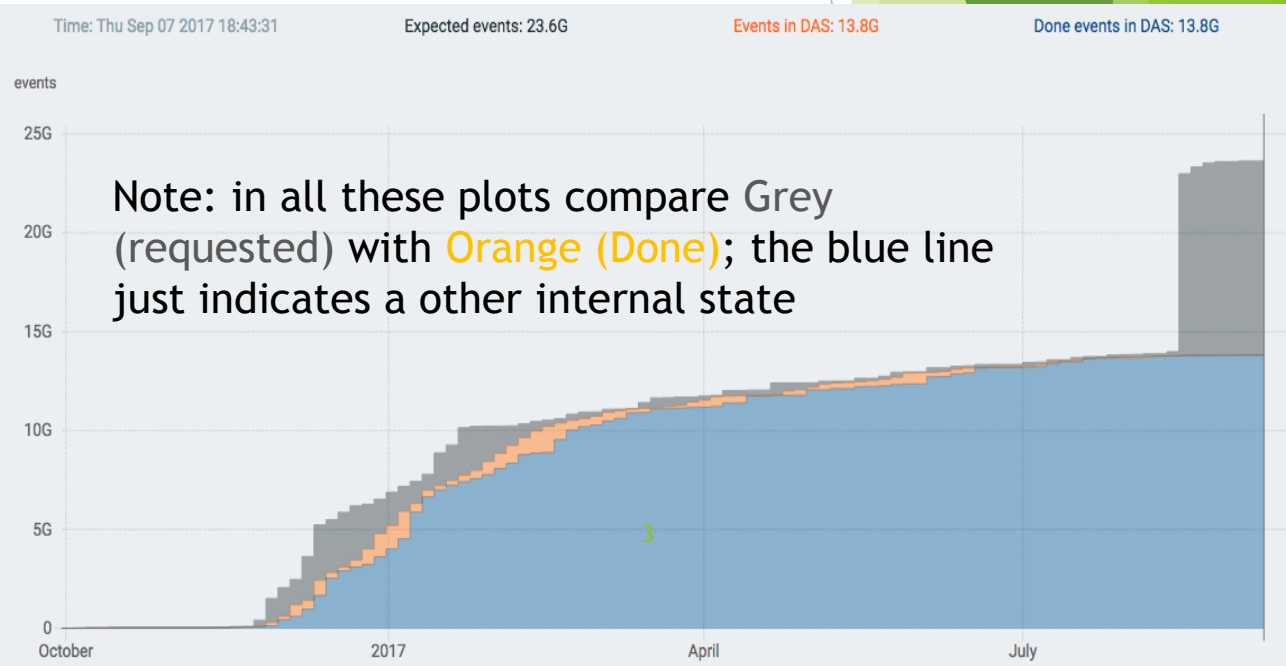

Note: in all these plots compare Grey (requested) with Orange (Done); the blue line just indicates a other internal state

2017 MC operations

- CMS detector underwent important changes wrt to 2016
  - Full upgrade of pixel system
  - Partial upgrade of HE electronics - Late decision (Spring)
- A lot of computing effort for early detector commissioning
  - Up to reprocessing of runs for pixel optimization till ~ July
  - Weekly (or more) release patches
- Consequence on MC17:
  - No possibility to start processing until long term performance reached
  - Two step approach:
    - Late July/Early August: start with a limited MC17v1 production (~1B events) for object validation / calibration optimization – currently ongoing
    - MC17v2 (~10B events) expected later in September – should be on time for Winter Conferences

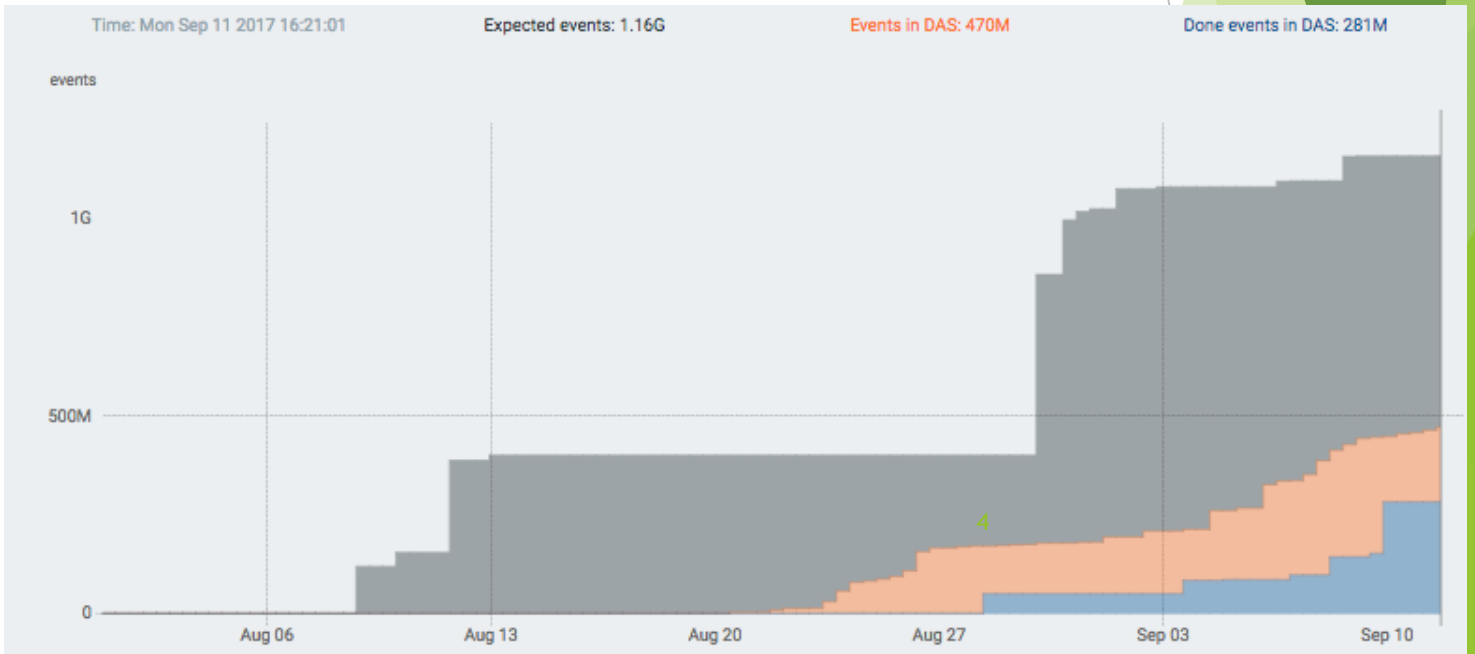

Upgrade TDRs

- In 2017 CMS has scheduled
  - A Tracker TDR (delivered)
  - Barrel+Muon TDRs (nearly delivered)
  - An EndCap TDR (to be delivered)

- Most of the samples produced in preparation are not shared, and are in many versions
  - With different aging settings
  - With different geometries

- Daily effort on small samples, in order to track development and check effects

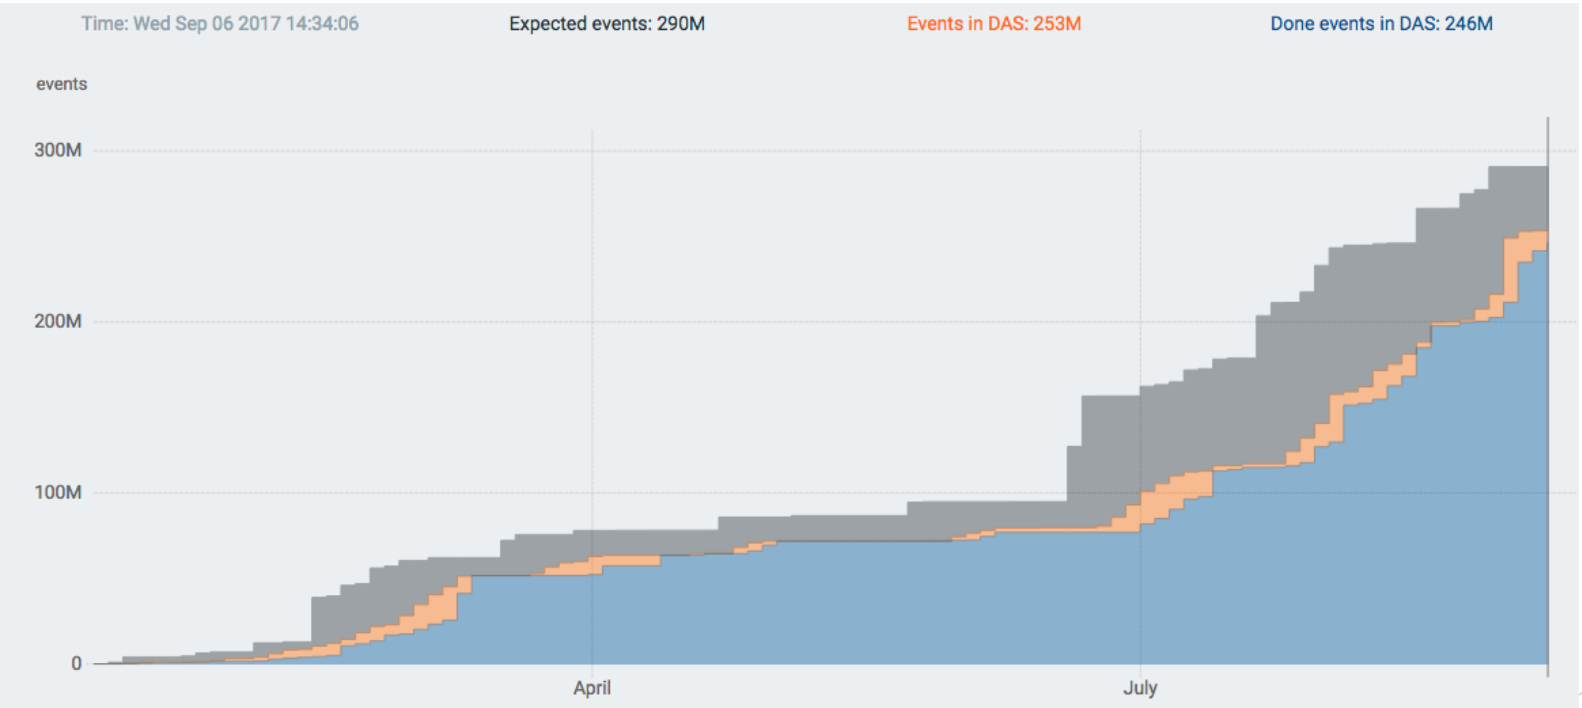

Excluding RelVals, ~300 M events requested, 250 M events delivered. Production for the EndCap TDR starting soon

2017 Data operations

- End-of-Aug Stable Beam time is 800 hours, less than at a comparable time in 2016 (end-of-Jul) – now ~ 950
- Even if instantaneous luminosity is currently not as planned (1 vs 1.9e34), CMS still takes ~1 kHz of data at roughly the same luminosity per-bunch (lower lumi via fewer bunches)
  - Expected increase via beta*=30 cm after TS2

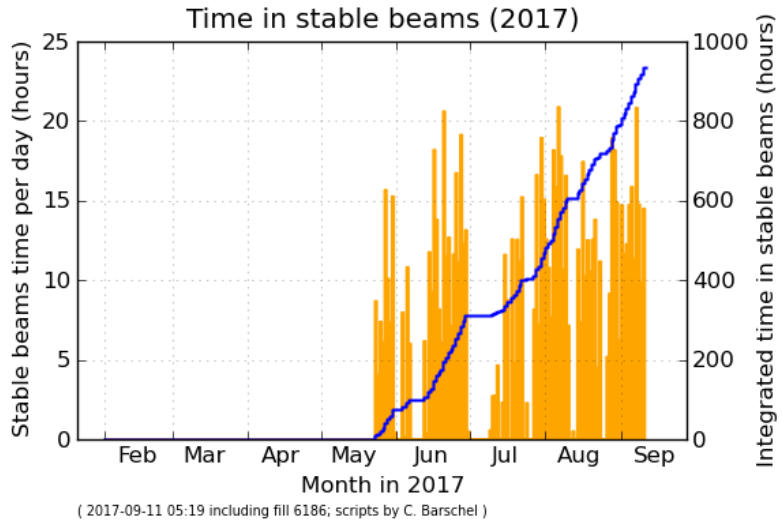

- Integrated luminosity ~ 20/fb by the end of Aug
- Data taking operations smooth, with some “features” (see later)
- Already starting a 2017 Data Rereco in order to get a consistent set with last deployed version of software in data taking
- A second ReReco for the same scope close to the end of data taking

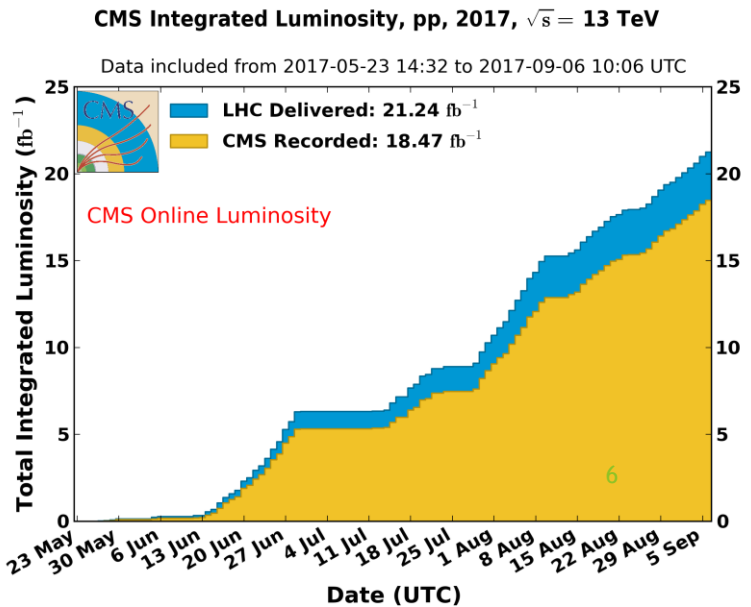

General plans now – start of 2018 (evolving …)

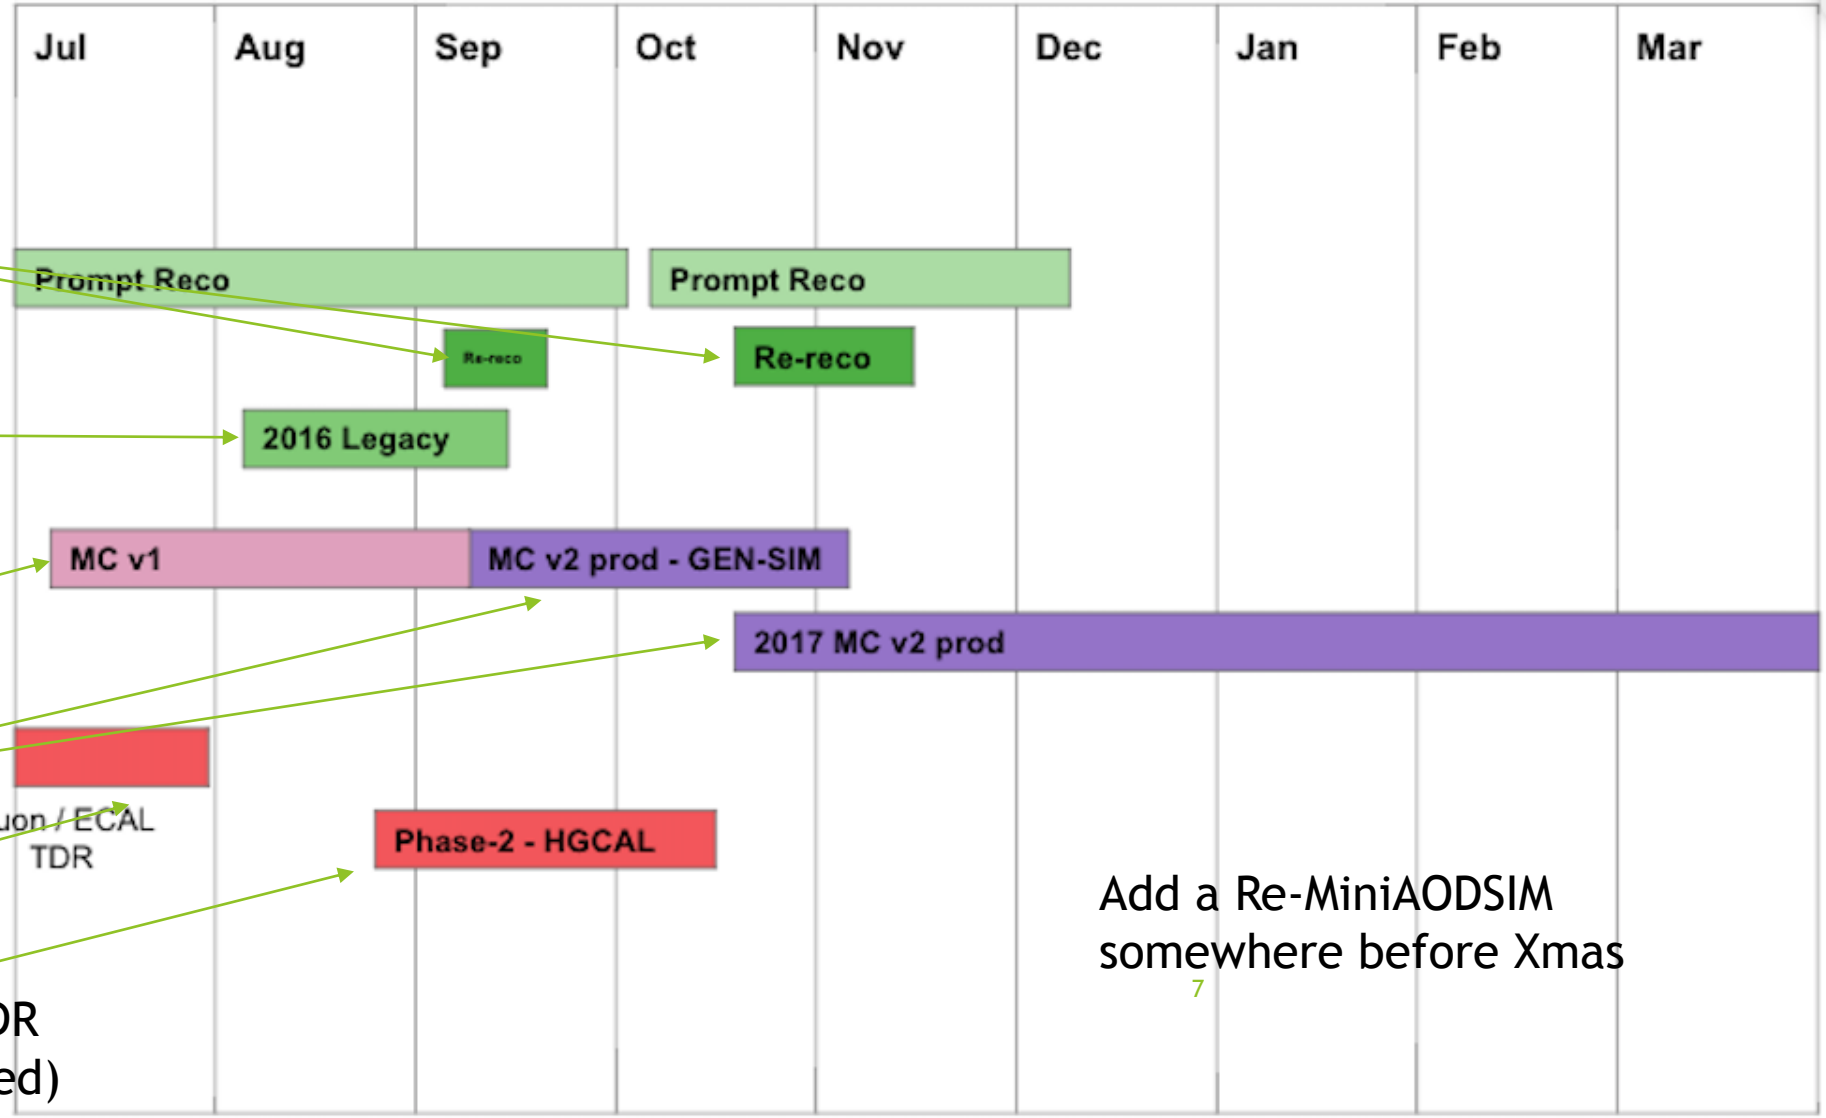

2017 Data ReRecos (not started)

2016 Data ReReco (ongoing)

MC17v1 “object calibration” (ongoing)

MC17v2 “physics” (not started)

Tracker + Barrel TDR (done)

Endcap TDR (not started)

Add a Re-MiniAODSIM somewhere before Xmas

Resource utilization

- Starting from July, full utilization of computing resources
- Global pool 150-200k cores (blue line)
- Pending in the system O(1M) jobcores (red line)

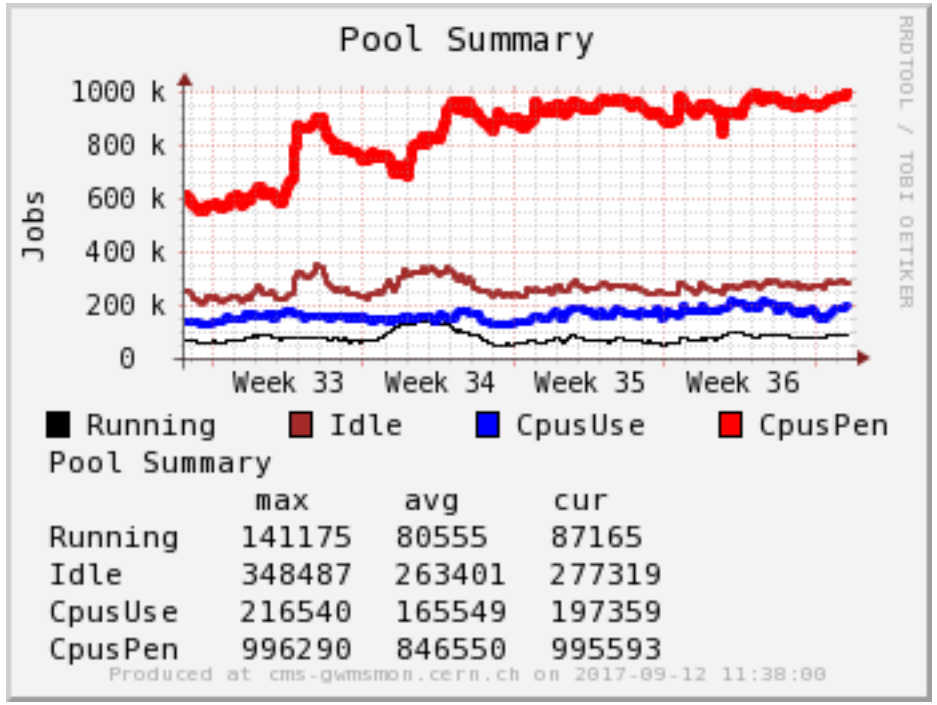

- Analysis utilization at the level of 50k cores or more (following modelling)

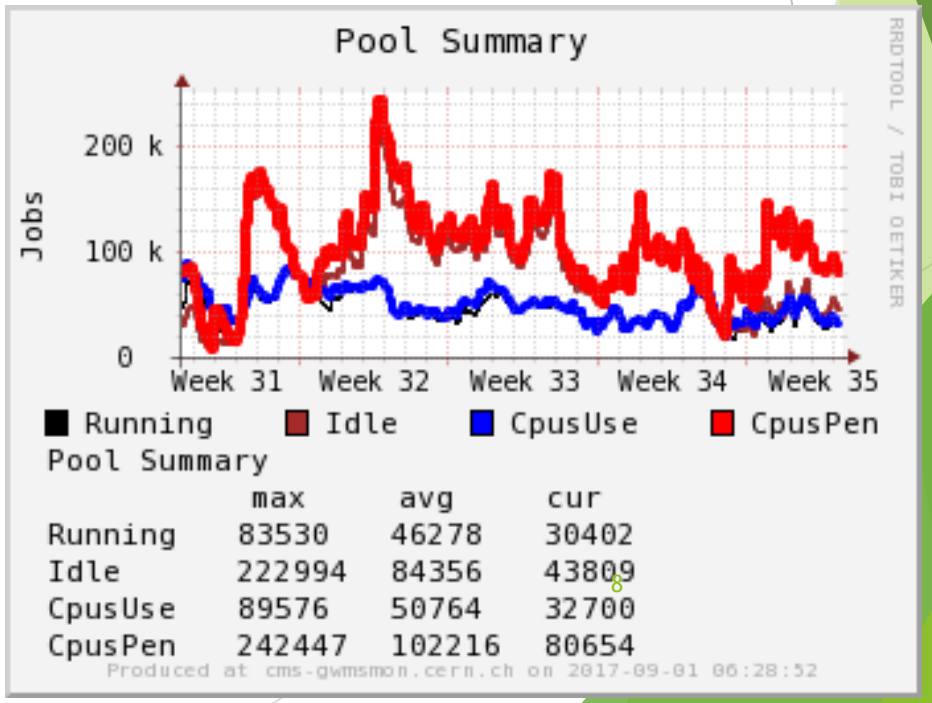

Storage utilization

- As you can remember, CMS is hit by under-recommendations in all storage categories (but Tier-0)
  - Storage more critical
- We will most probably need to perform “unprogrammed cleaning” of storage systems to reach April 2018 (as in late 2016)

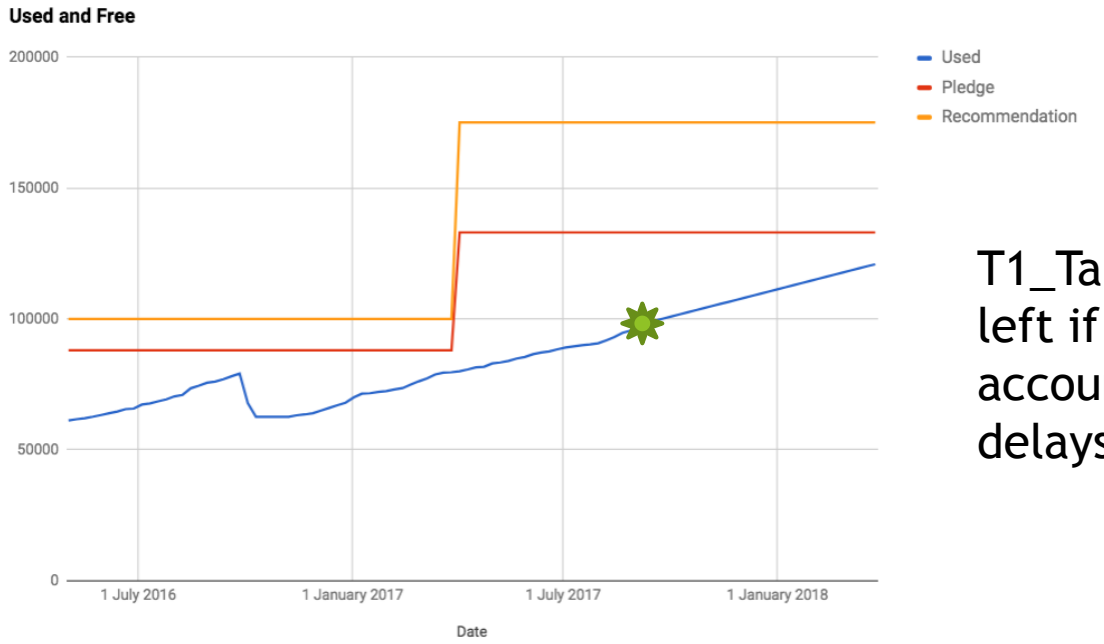

T1_Tape: <20 PB left if you take into account repack delays etc

Pledges Balance

	ALICE	ATLAS	CMS	LHCb
CERN CPU	0%	0%	0%	0%
CERN disk	0%	0%	0%	0%
CERN tape	0%	0%	0%	0%
T1 CPU	-8%	-12%	-14%	-4%
T1 disk	-14%	1%	-21%	-6%
T1 tape	-1%	-7%	-24%	-3%
T2 CPU	-24%	-13%	-7%	27%
T2 disk	-28%	-7%	-22%	-30%

Managed Disk (T1+T2):

- 75% used by un-movable stuff (blue)
- Some Tier-1s over the “attention” threshold (90%)

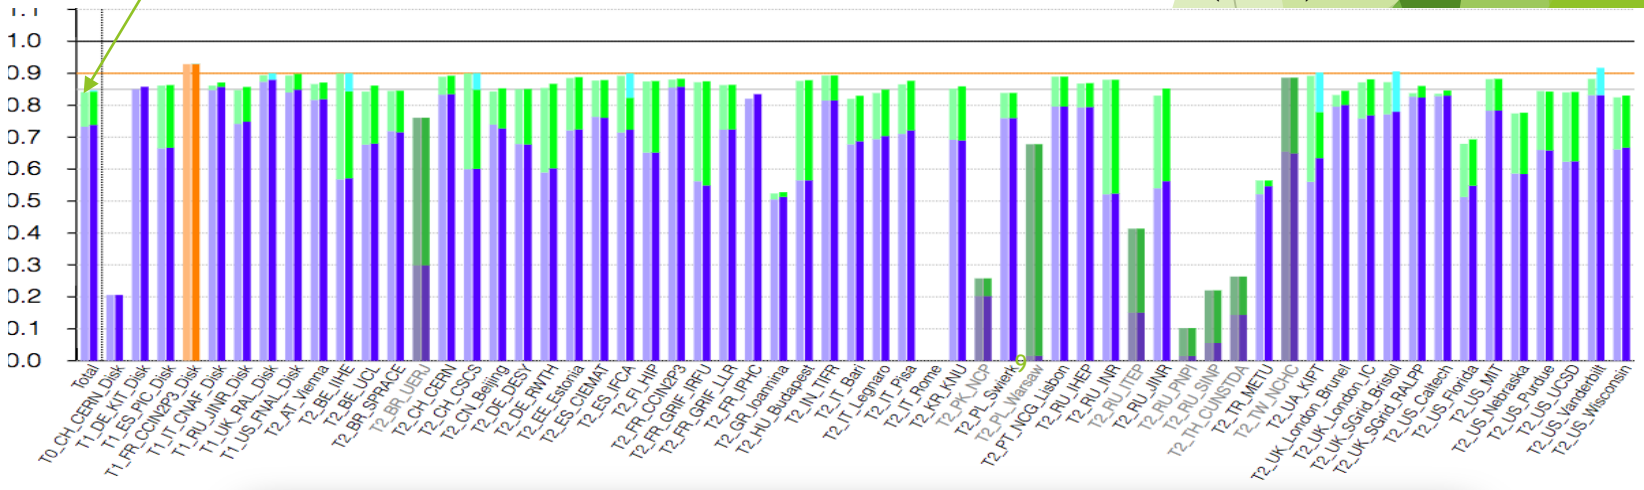

Tra

- In
  tra
  the
- Ma
  -
  -
- Wc
  mc

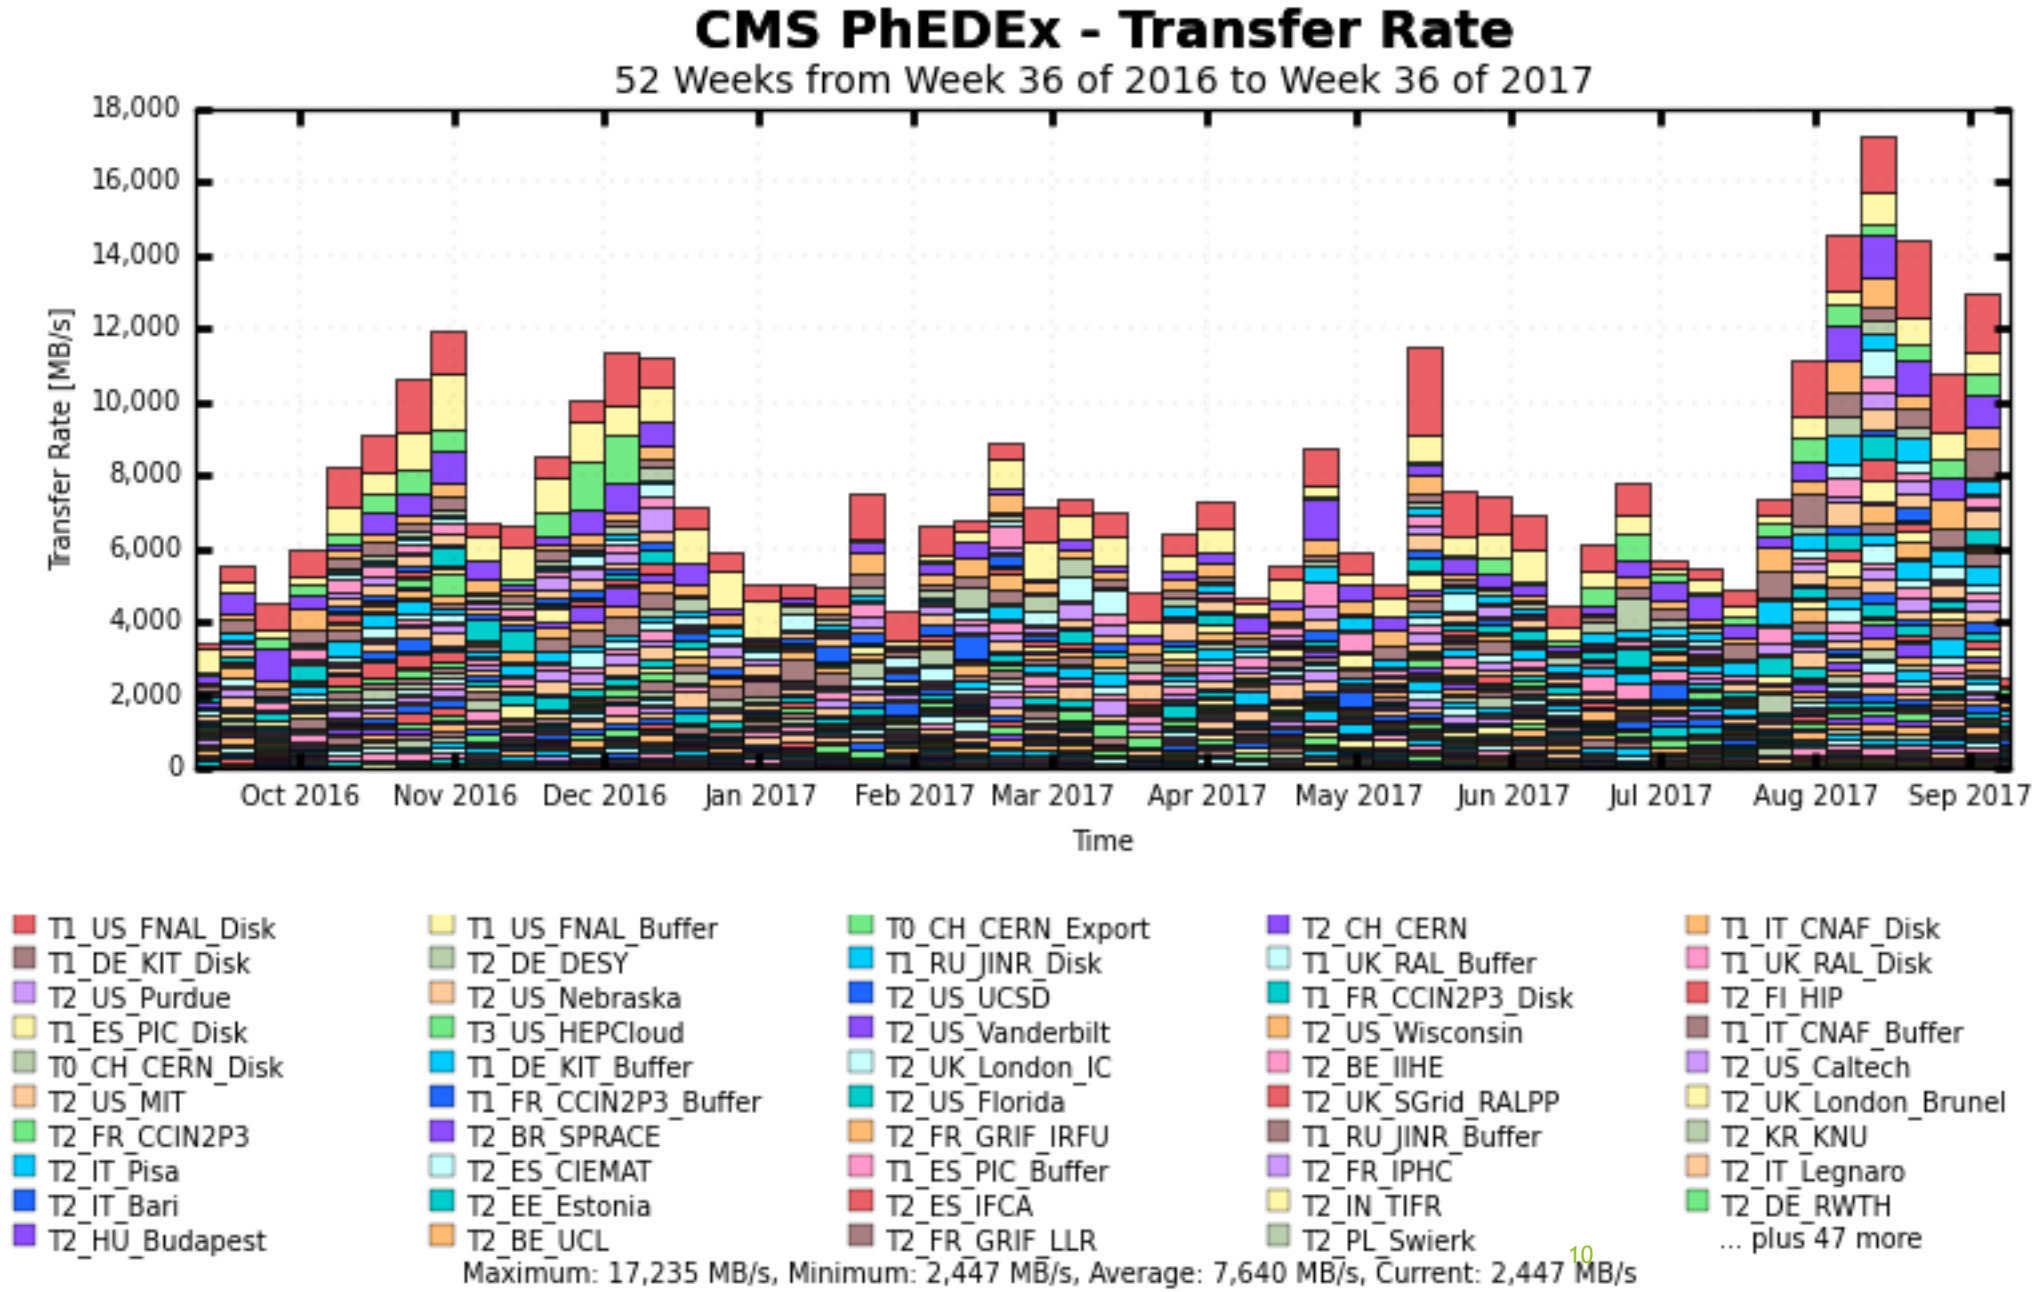

Problems / features

- Tier-0 / EOS struggled a bit up to ~mid August for several reasons:
  - Commissioning of new (pixel) detector required maintaining Express data for 3x longer then expected
  - Tier-0 working space at moments close to its 6 PB quota
    - Problematic transfers to some T1s
    - Problems on source files on EOS (see next bullet)

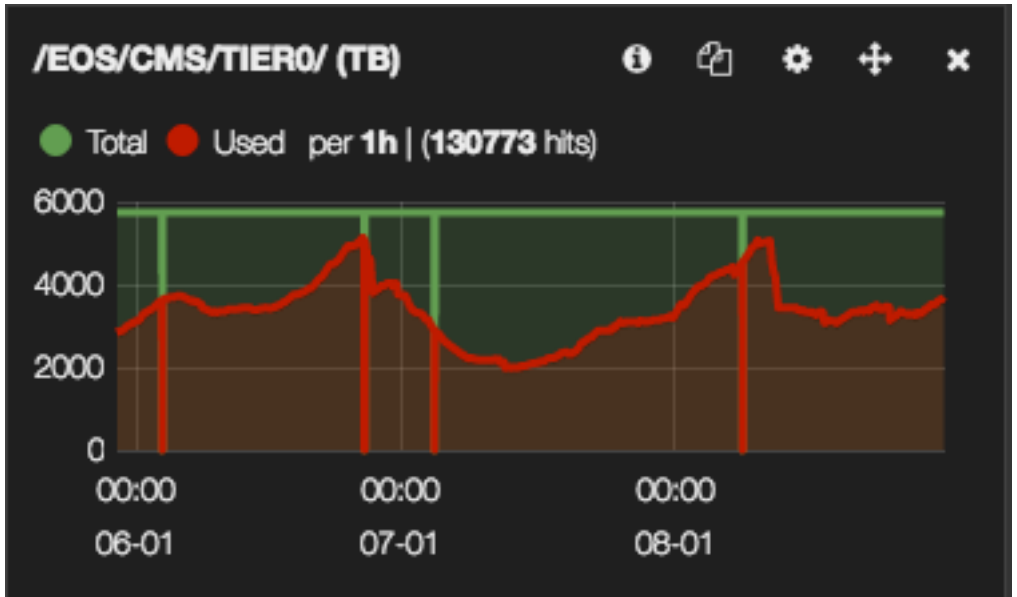

- EOS/CERN: Jun- mid Aug suffered from
  1. Scalability issues: authentication rate (GSI) topping at 1 kHz
  2. Lost Files / files becoming 0-sized after some time
  - The sum of the two effects caused various problems
    - Failures in accessing Streamers and processing prompt calibration loop → Tier-0 processing slowed, with many human recovery actions
    - Problems with transfers to Distributed Computing (continuous interventions to maintain TransferDB / EOS consistency)
  - Problem solved on Aug 11th with an EOS patch
    - No “lost files” problems since then
    - Only recently situation reassuring enough to restart offline processing on Tier-0 resources
  - **We wish to thank the IT/EOS team for the continuous support / recovery procedures!**

Resulted in no permanent data loss, but some derived data was not recovered; it will at next reprocessing

Rest of RunII plans

- 2018:
  - Plans unchanged, requests unchanged → C-RSG
  - Still taking ~ 1 kHz of Trigger at a PU ~ 35, and SB live time not far off
- 2019:
  - After the initial estimate delivered for Apr2017 RRB, detailed needs to be discussed at the Oct17 RRB
  - Guidelines for 2019:
    - Full reprocessing of RunII DT + MC 25+(25-35) B events
    - Continue PhaseII studies
    - Continue analyses at the level of 2018

- Expected availability of **Tier-0 (100%)** and **HLT (80%)** resources for offline processing
  - CERN to become a 1MHS06 Offline processing site for CMS
  - Need to carefully check EOS and P5-IT dedicated link
- Requests at +0% for CERN, O(10%) in all other categories apart from a +20% T1_Tape
  - Wrt 2018 approved requests – 2018 pledges not yet in place in Rebus

CPU efficiency

- Since May, CMS O+C has launched a dedicated task force to investigate the issue
  - A definition problem (what is the correct metric) is being attacked by HSF working groups, in the context of modelling total costs
  - Still, value to understand the sources of “inefficiency” by the current metric – this is what the TF concentrating on since May
  - TF chaired by D.Colling, UK/IC
- No real plot to be shown, since they need > 1 month of accumulation and we are still on a moving ground
- Still, a lot of gained understanding which led to implemented solutions
- Global efficiency can be factorized into infrastructure_efficiency * payload_efficiency

Pilot infrastructure, interaction with batch systems, job matching, pilot draining …

Calibration / code loading IOWait, final stageout, job length …

Infrastructure inefficiency – implemented solutions

- Scheduling inefficiencies, in the form of “idle pilots”, constitute 10-15% of the total inefficiency
- A good part of it we do our internal scheduling, CMS is accounted for inefficiencies which is other cases are attributed to the site (like draining for multi core pilot to run)
- EARLY IMPROVEMENTS:
  - Implementation of depth-wise filling of glideins (pilots) in June (improvement estimated at **~1-2%)**
  - Shortening the amount of time a glidein can sit completely idle from 20m to 10m (July), driven by Negotiator cycle length (improvement estimated at **~1%**)
  - Removing glideins from site and factory batch queues after a certain amount of time (1-3h) in July:
    - Important when the pressure on the system is not constant
    - Avoids having pilots running days after pressure gone

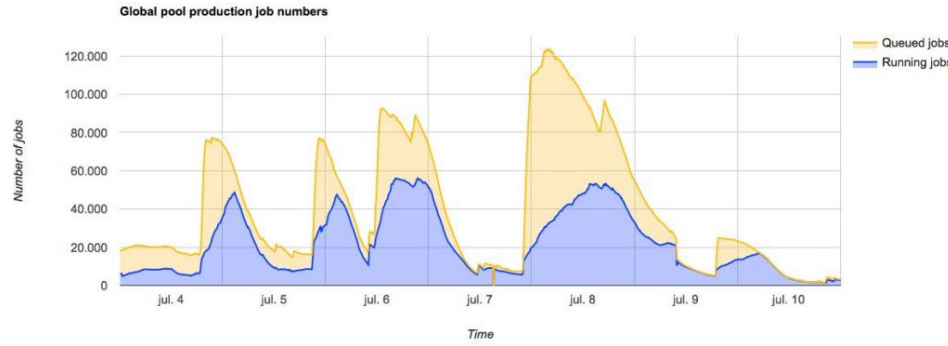

Infrastructure inefficiency – detailed analysis and plans

- Reference study period **after** improvements, with sustained job pressure**: Aug 23rd-29th, 2017**. Total remaining **idle CPU** in multi-core pilots: **11.4%**.
  1. **3.4%**: slots requesting more memory than available (**RED+VIOLET**)
  2. **1.6%**: analysis jobs not starting since they request unreasonable running time (**GREEN**)
     - Automatic job splitting in CRAB3 under test
  3. **4.2%**: jobs not starting since too close to the pilot end-of-life (**ORANGE**)
     - currently optimizing against job failure rates and may consider not draining at all to eliminate this component.
  4. **2.2%**: currently not understood. Possibly irreducible (**BLUE**)
- **2+3 can be recovered in great part – 6%**
- **1 cannot, unless machines with more RAM are acquired**

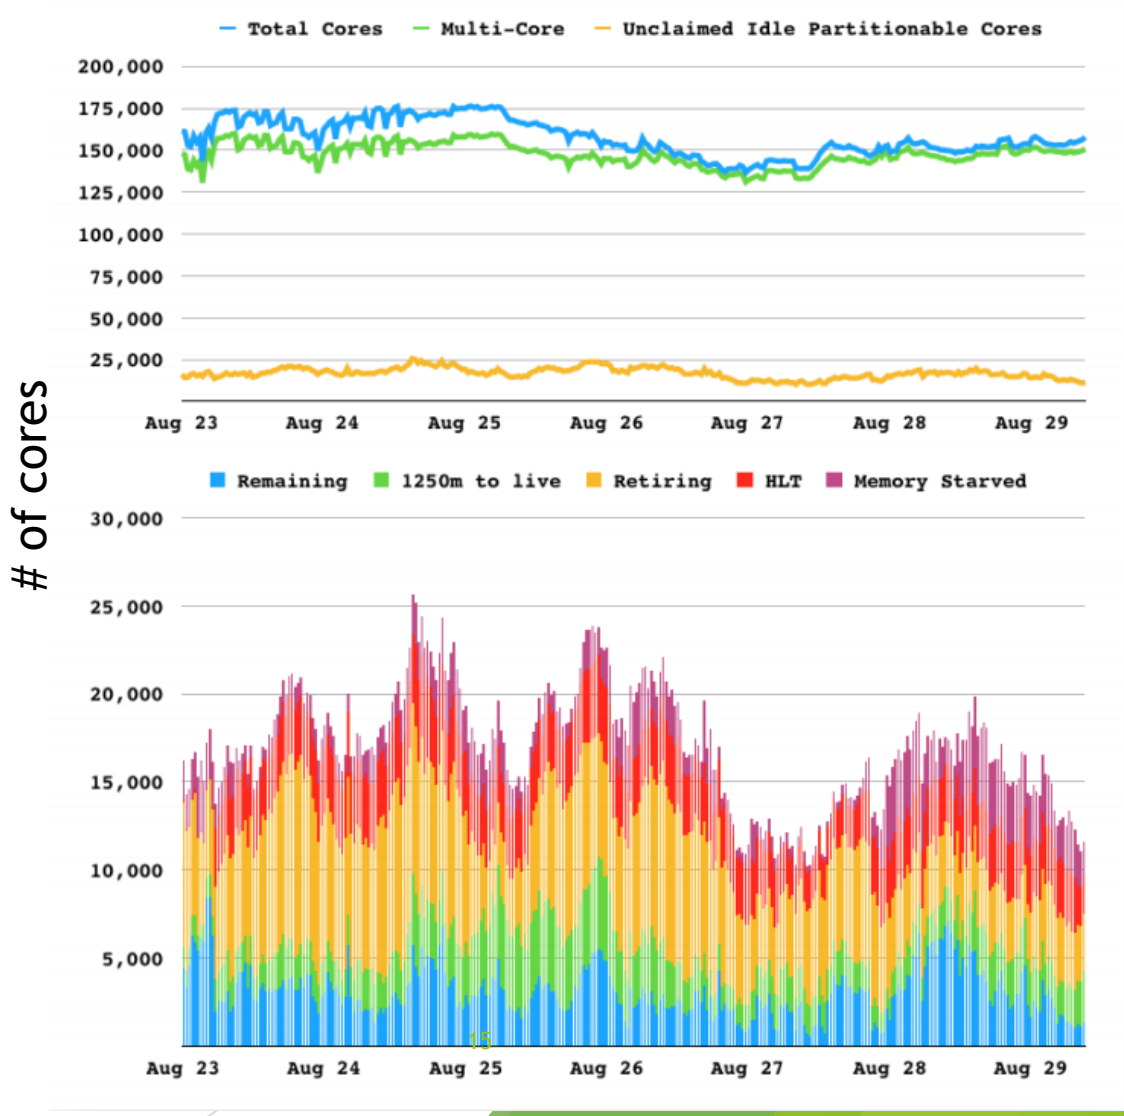

Payload inefficiency

- Complex studies, depend on the many external factors (site load, storage system load, etc) - Extensive use of Elastic Search
- Discoveries up to now:
  1. CMSSW does **not** have an intrinsic efficiency limit. In controlled situations all our workflows can reach nearly 100%
  2. Most processing activities have a O(80%) or plus peak efficiency under load; but we do see tails going very low – studying them
  3. The effect of reading Primary input from remote (XrootD) instead of local seems to create a <**10**% in efficiency
     - Clearly, at the reduction of total storage needed, and in general the ability to be processed sooner
- CMS switched to premixing ~ 1 year ago, and premix is read from 2 copies at CERN and FNAL
  - Switching to it divided by 2x the DigiReco time (and resources) + reduced the storage needed for MinBias Distribution
  - From our analysis we think that the impact on efficiency is at the few % level at most
- Job length found sub-optimal in some categories (analysis, some production workflows…)
  - Worked on obtaining better estimates from submitters + WMAgent
  - CRAB3 auto splitting (under test now) should make sure analysis jobs are not too short/long

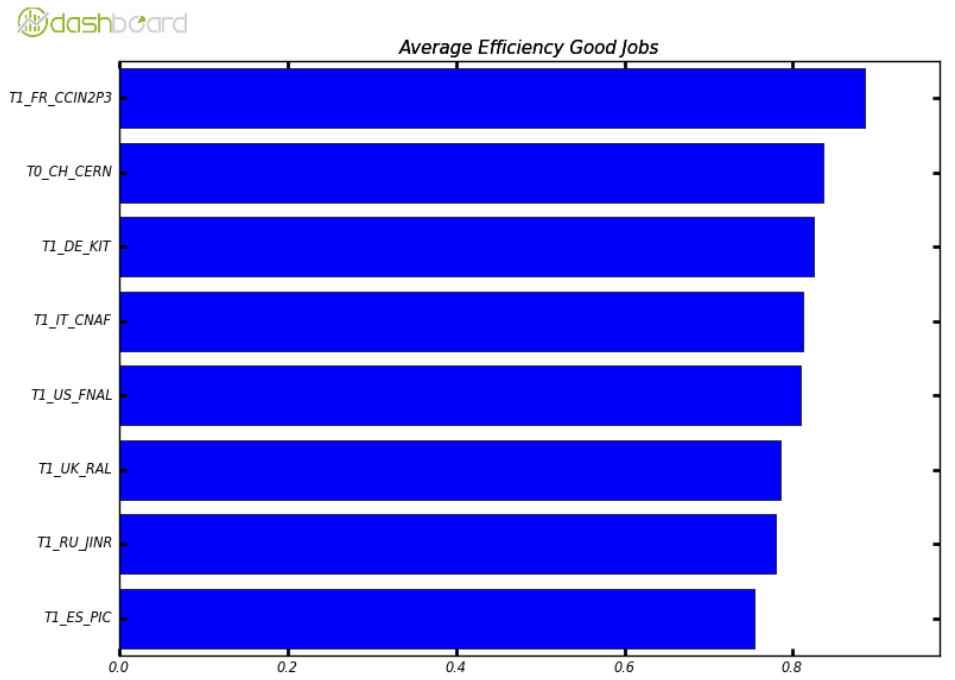

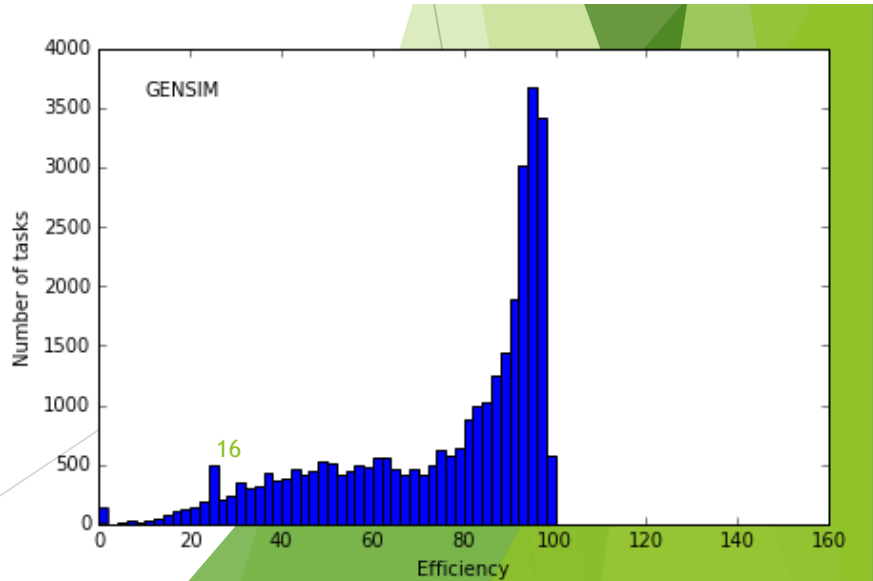

RunIII expectations

- 2018 can currently be though as a “16L2-free implementation of 2017”
  - Up to 2.3e34 instantaneous luminosity if beta* = 30 cm
  - Official numbers for 2018, on which we model, still 1.7-1.9e34
- Difficult to see a sharp resource need step between end-of-RunII and start-of-RunIII
  - RunIII target seems still to up to 2.5e34; so if you want a factor 2.5/1.9 = 130%, but three years later
  - The +(0.5-1) TeV is ~ irrelevant
  - HLT total rates as we know them are not expected to increase above 1 kHz
  - Detector PhaseI configuration does not seem to imply any excessive increase in needs
- This does not mean we will not plan changes / activities!
  - Just, a simple evolution scenario wrt RunII seems at the moment sufficient
- Some major changes on the table:
  1. Study a **NanoAOD** data format, at 1-2 kB/event level per event, shared by ~ 50% of the analysis
     - If successful, expect a decrease of CPU needs for analysis in the medium-long term
  2. Implement a lifetime model, with end-to-end interfaces, for tape and disk
  3. All our workflow should be multi-threaded by RunII, nothing to be expected there
- CMS is finalizing the report from a Evolution of Computing Model task force, with recommendations for the medium term. 1+2 are for example included there

HL-LHC new working numbers

- CMS does not have newer officially blessed numbers for HL-LHC
- Still, work has been ongoing also due to the DOE request to US-CMS for long time planning
- Main changes wrt to older models (see for example ECFA presentation by S.Campana) are
  - Expectation of 10%/y code performance improvement
  - Rely largely on MiniAOD(SIM) for operations; AOD(SIM) an archival thing

Take home messages for 2027:

- **50 MHS06 CPU**
- **5EB disk**
- **3EB tape**

- **Wrt to 2017, assuming a +20%/y by Moore and friends, the excesses are ~6x for CPU, ~4x for storage**

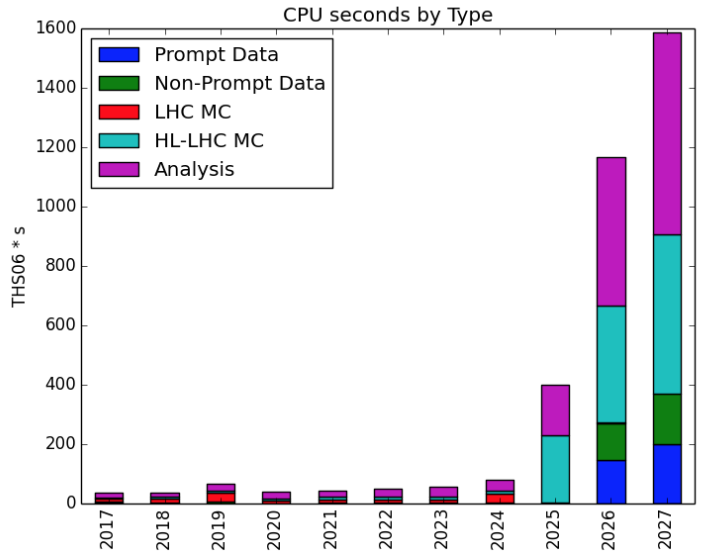

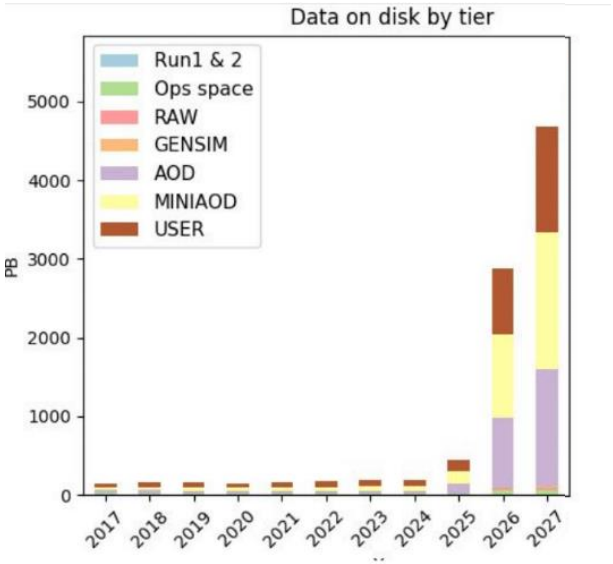

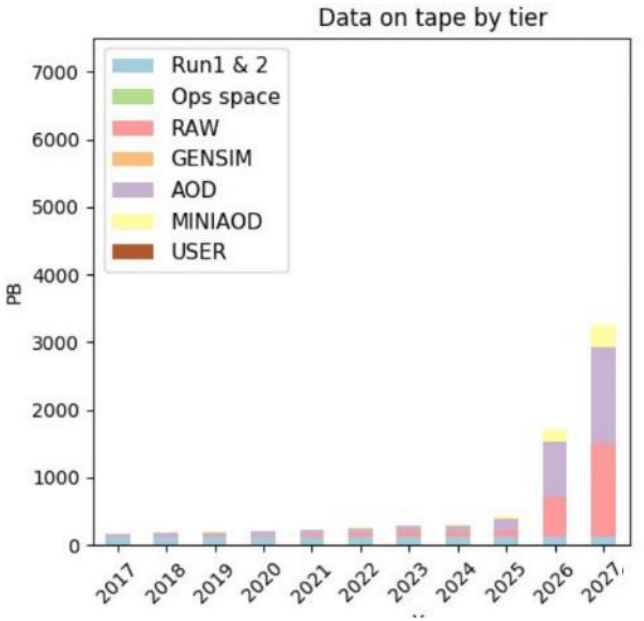

Conclusions

- Many activities happening at the same time – most considered urgent; CMS computing able to cope
  - 2016 MC and Data
  - 2017 MC and Data
  - PhaseII samples
- The lower-than-recommended storage pledges has indeed materialized as a real problem in operations – as we were expecting
  - Storage lower than requests; balancing between sites difficult
  - More pressure on tape and transfer systems
  - Some sites / FAs are doing special efforts in helping over official pledges
- 2018+2019 RRB requests submitted
  - **2018**: confirmed (approved) RRB/April17 requests
  - **2019**: plan for Legacy RunII ReReco (MC+DT), with large utilization of CERN HLT + Tier0. Requests (well) below standard increases
- CPU Efficiency task force at work
  - First solutions implemented
  - Expect more to come in the next few weeks

BACKUP

CMS (Payload) Workload Studies with Elastic Search

- Over the last 1B of core hours, here's the approximate usage:

Job Type	Percent CPU resources	Payload CPU efficiency	Percent of idle *payload* CPU
Analysis	37%	64%	41%
Simulation	27%	74%	22%
MC reconstruction	27%	68%	29%
Data reconstruction	7%	73%	5%
Other	2%	56%	8%
	Overall:	68%	

- For each row:
  - Different source of inefficiency. Some inefficiencies may be irreducible.
  - Different approach (and cost!) for improvements.
- Goal: given finite effort, maximize payoff.

WLCG CPU efficiency =
(pilot scheduling eff) *
(payload CPU eff)
WLCG measurements roughly match collected CMS monitoring!

Use of our Global Facility

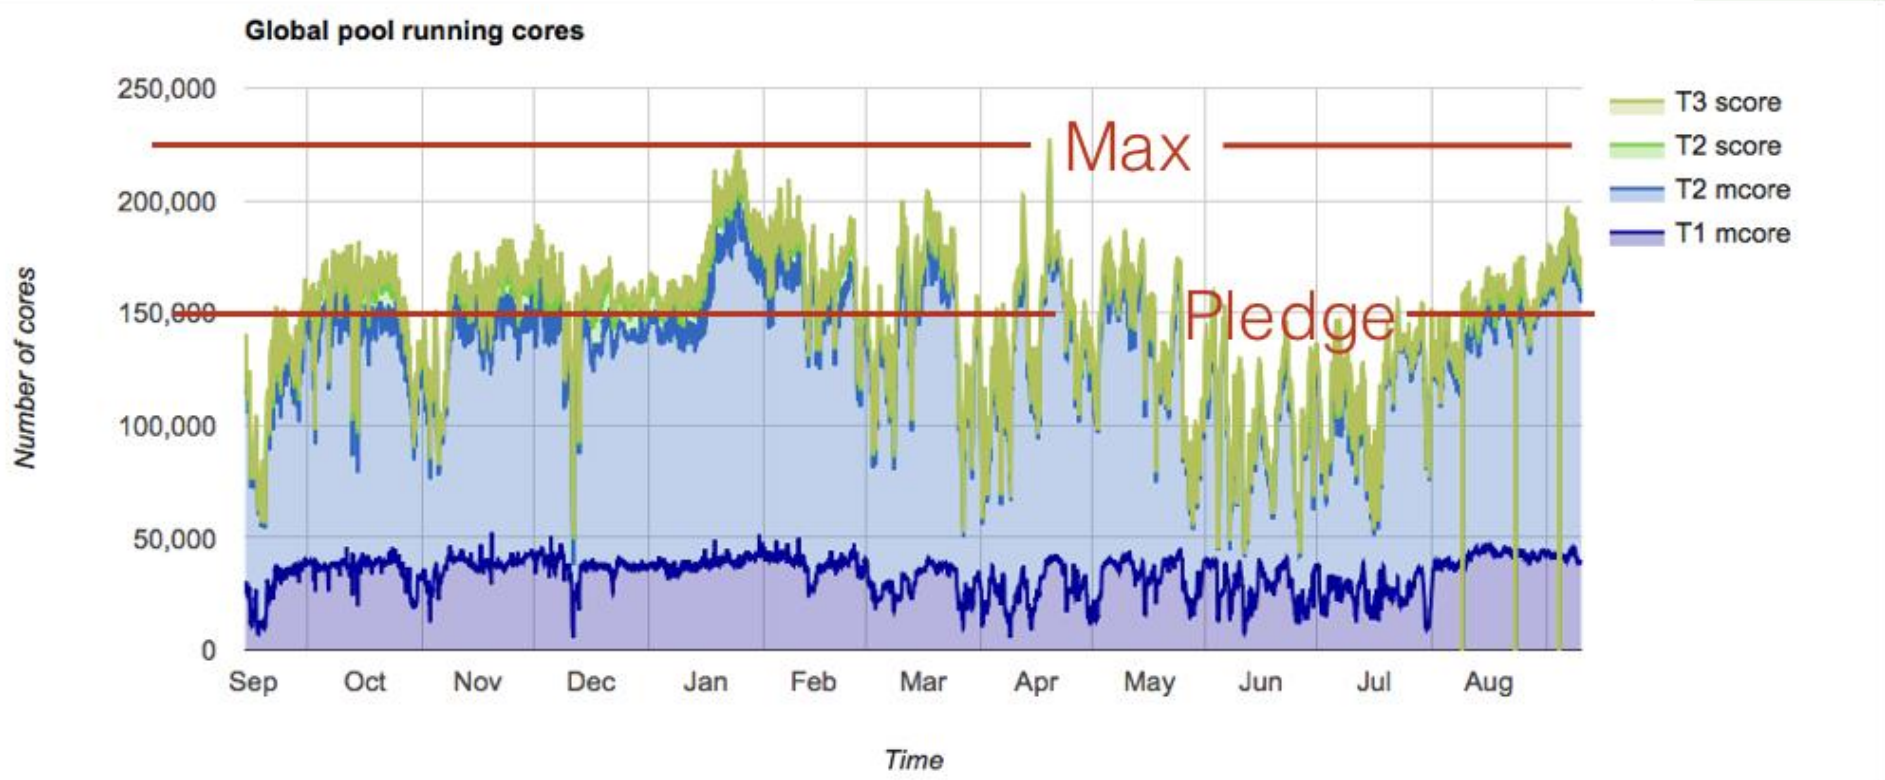

- The scale up of our facilities to deliver for Moriond17 was a huge success.
- Large fluctuations in demand are a challenge for our provisioning systems.
- With the submissions of the MC1 campaign resource usage is adequate again.

X

Schedule in Spring'17

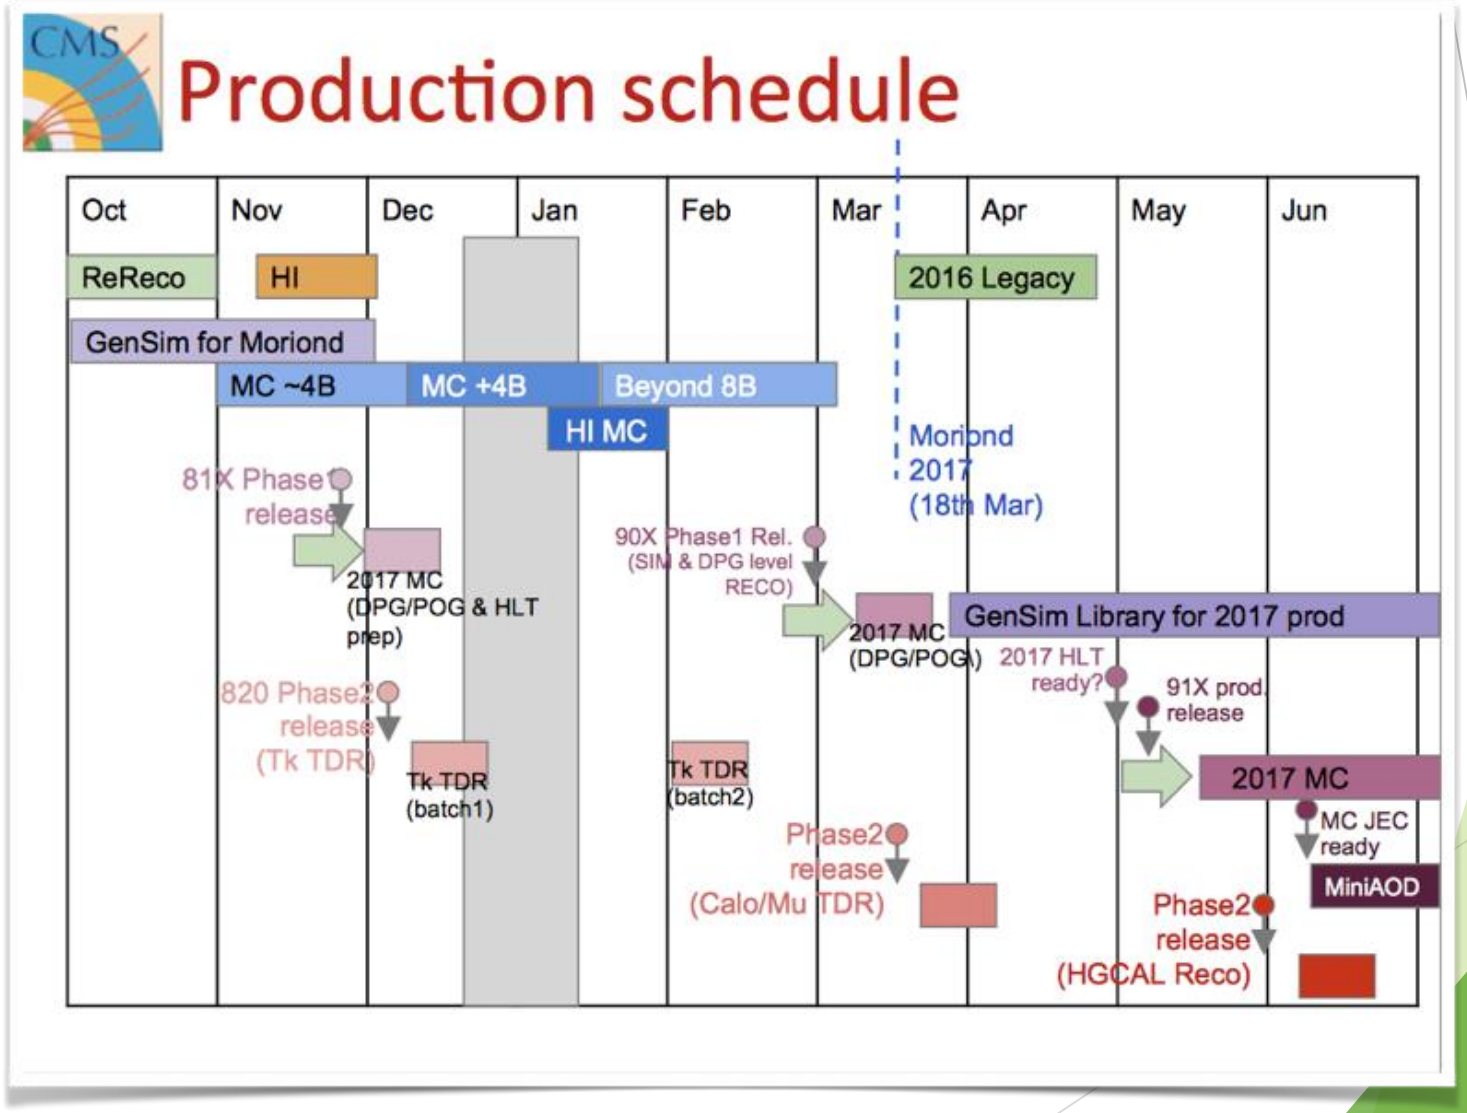

X