CMS Report

ELIZABETH SEXTON-KENNEDY

TOMMASO BOCCALI

Outline

- Final part of 2017 Run
- Plans now-Winter Conferences (and a sketch beyond that)
- Resource request / status
  - 2018– (or: “how to survive the run”)
  - 2019 plans
  - RunIII: expectations (not yet a model)
- CPU efficiency – a rolling status report
- AOB

Final part of 2017 Run

- At last meeting (Sep 12th):
  - We had reported on the main problems we had (mostly the CERN/EOS) – confirmed solved, no problem since then;
- With that, we were expecting a very smooth end-of-run period; but then Sept 20th: 8b4e
  - Pl
    bu
  - CM
    (e
    Pl
- Al
  ra
  "e

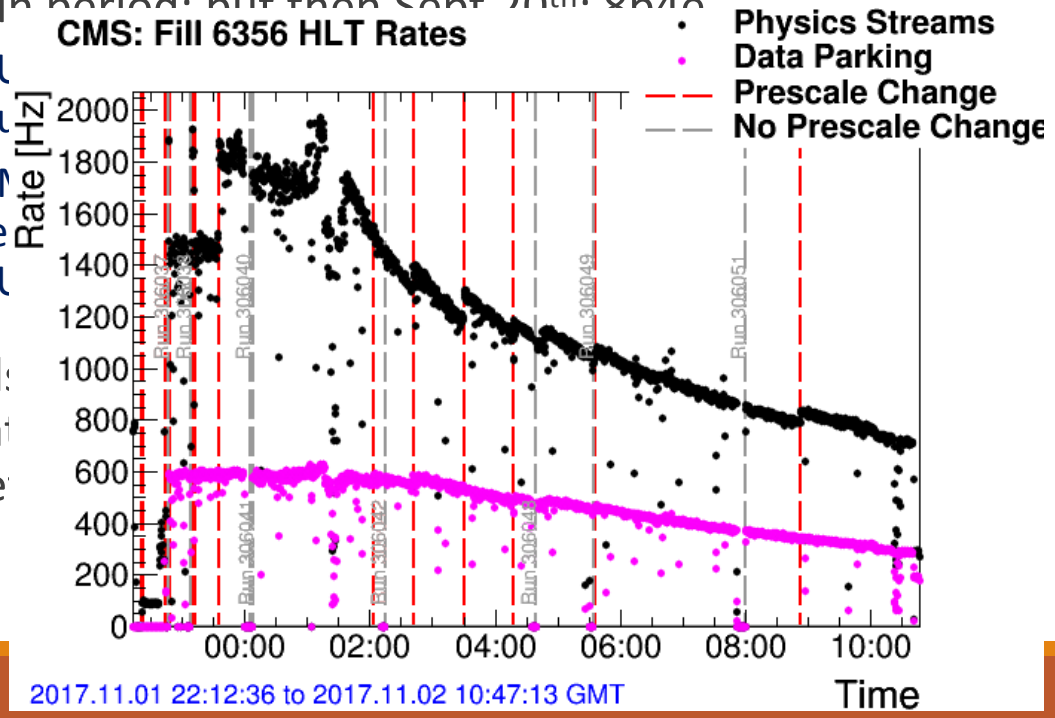

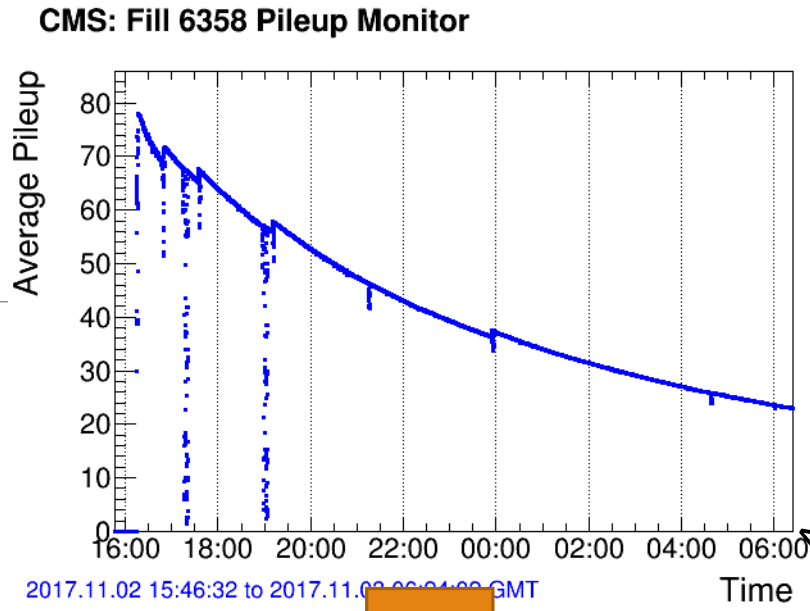

A single fill taken w/o levelling

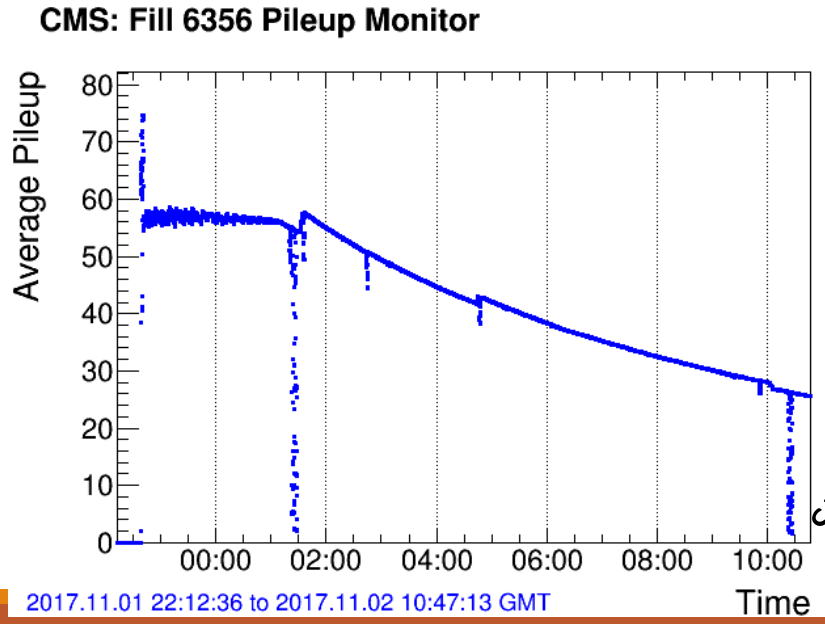

Standard fill after Sep 20th

All in all – Tier-0

- Processing > 1 kHz of events with <PU> exceeding 45 - quite over specifications (1 kHz @ 35)
  - Interesting test. It might end up as a standard operation mode for 2018 (no new Tier-0 resources, given that the “official model” is still @ 1kHz,<PU>=35
  - As expected, **not** smooth operations
- CPU: backlog building up, with Prompt Reco delay going 2 → 5 days (so 3 more than design)
  - Up to 60000 “late jobs” after XeXe run
  - CMS has commissioned solutions for this: run on distributed sites. In principle at T1s, but this time, given the (opportunistic?) availability of resources on T2_CH_CERN, we used CERN T2 as part of the T0
    - Slope inverted
- Disk Buffers:
  - T0streamer Under pressure from streamer files from P5 → increased
  - Output Buffer under pressure from backlog to T1s (mostly KIT)
  - Express not particularly suffering

T2_CH_CERN in the game

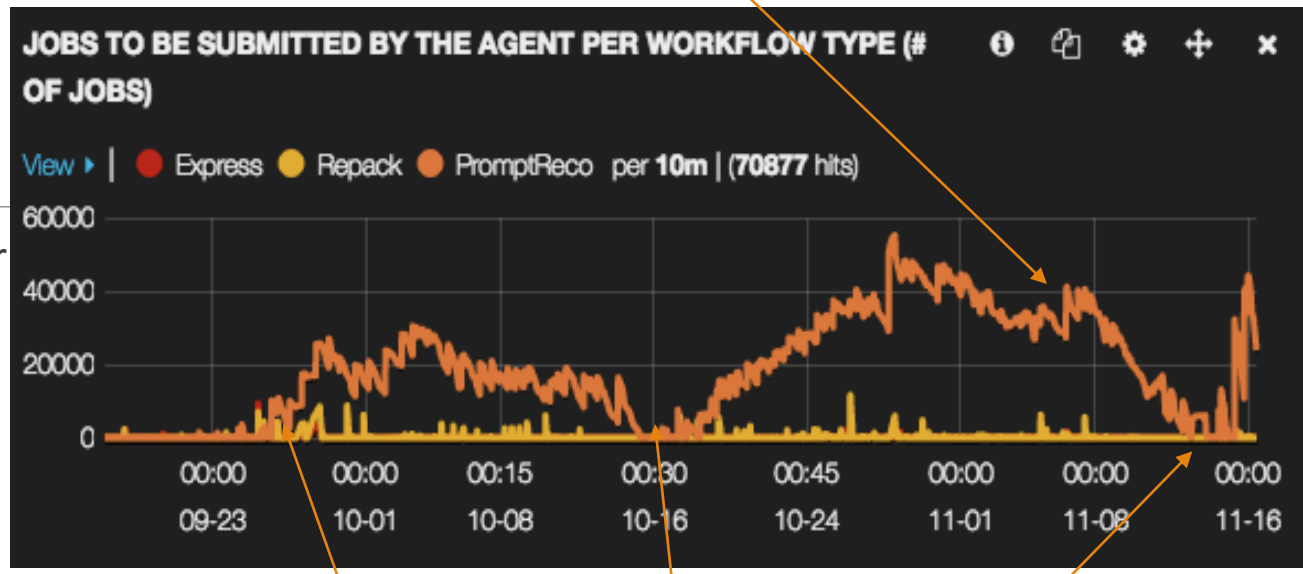

Levelling start XeXe Run ppRef Run

# of late Tier-0- jobs (“backlog”)

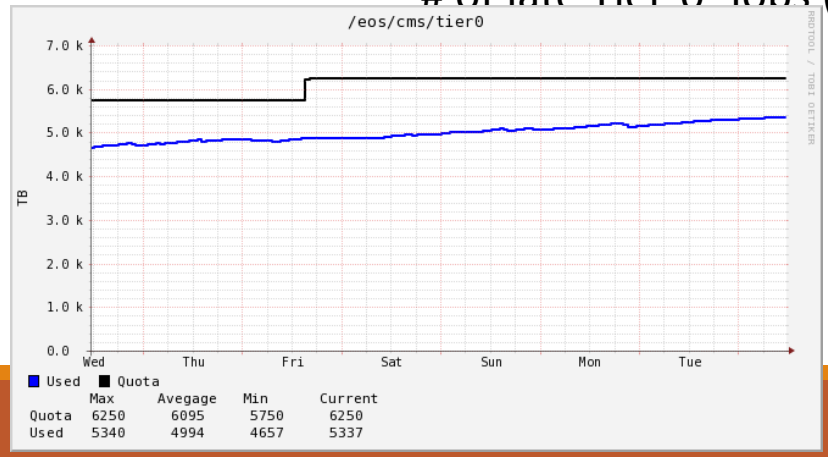

XeXe, ppRef

- CMS took two interesting special runs since September, which temporarily tested the system well above specifications
  - Fill 6295: XeXe:
    - Trigger rate @ 4 kHz, 40M events with complexity ~ PU 70
    - Digested by Tier-0 w/o major problems, AOD at Vanderbilt within 2 days, analysis could start the next week
  - ppRef: 1 week of reference pp run @ 5 TeV, PU ~3
    - Needed mostly by HI guys as reference for 2018 HI Run
    - Extreme conditions: **20-30 kHz HLT output (not a typo), mostly ZeroBias**
    - Extremely heavy for Tier-0: **input from P5 up to 5 GB/s (2-3x is the norm)**, all events undergoing prompt reco (~3 sec/ev)
    - All in all Tier-0 running 3x over specifications
    - Tier-0 buffers suffering and needing full attention; had to decrease some measures (streamer file retention, …)
    - CPU backlog reabsorbed in 3 days
    - Backlog to Tier-1s to 5 PB
    - **Lesson learnt: we can work like this, but not more than for 3-5 days**

End of ppRef

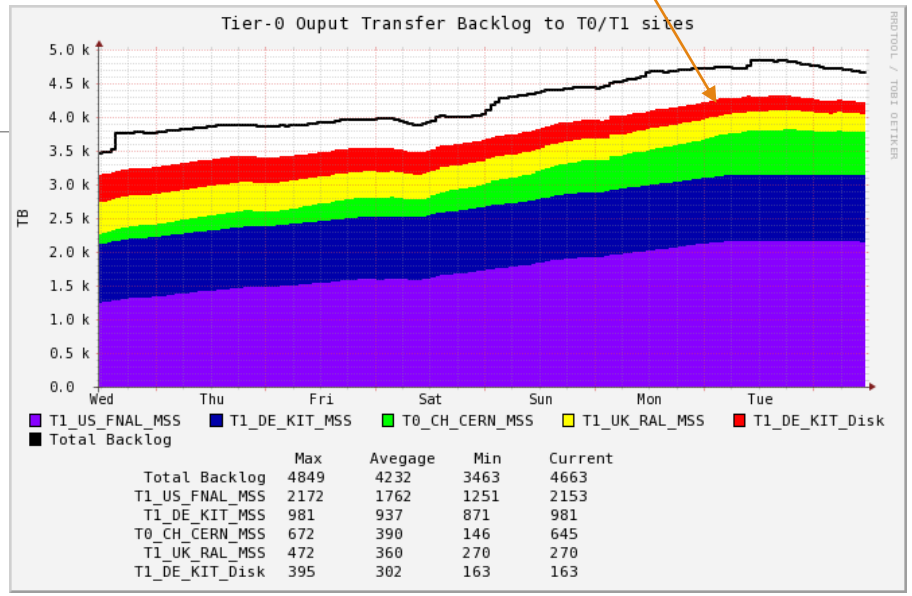

4.5 PB backlog to Tier1s (mostly KIT and FNAL)

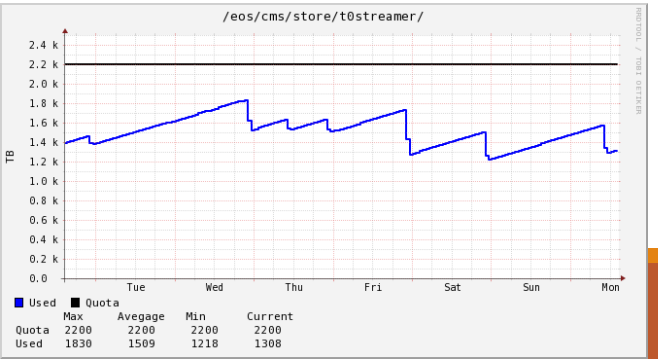

T0streamer area: removed 7 days safety

Highest HLT output rate seen in CMS so far!

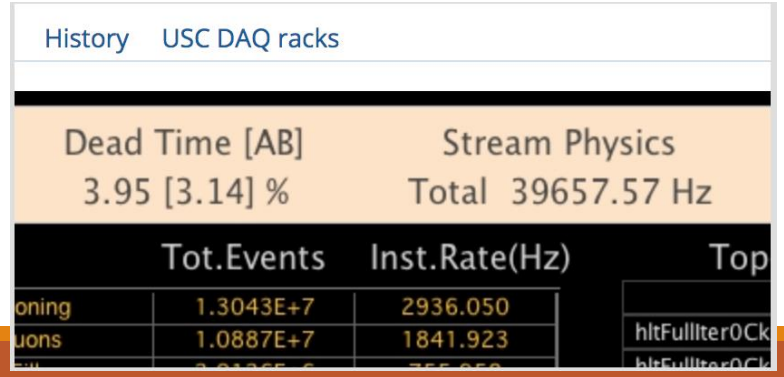

PhaseII samples

- PhaseII production for the HGCAL TDR submission
  - Arrived very late (started 2nd week of October) – should have started in August
  - Due to still not finalized local reconstruction, forced to be at GEN-SIM-RECO level (and partially at RAW level) in order to fix algorithms at analysis time
    - Very disk heavy, 30 MB/ev minimum
    - Due to that, PhaseII uses on CMS disks 13 PB, 3x what expected at the time of the modelling
  - Total PhaseII production in 2017 is at the level of 500M events requested (PU 0,140,200)

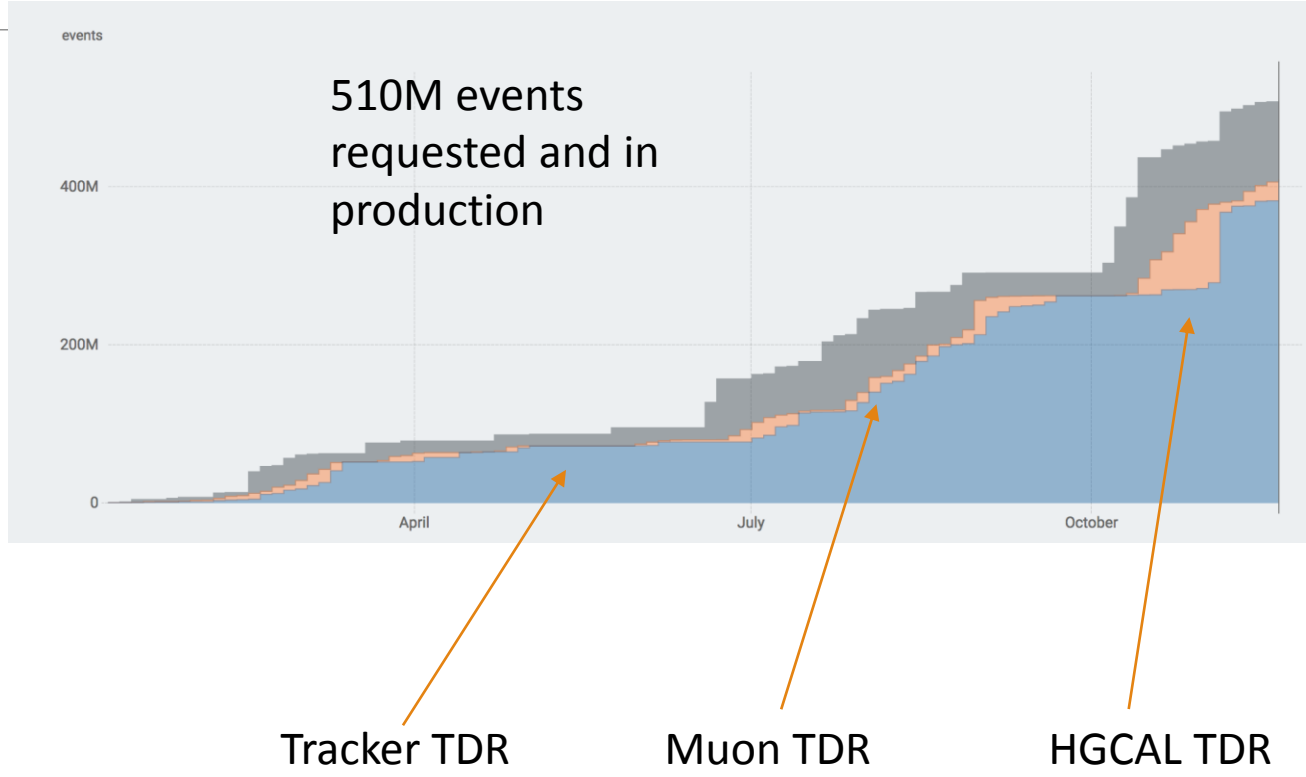

- 2018: CMS just introduced a Timing Detector TDR
  - (Not in the planning, will have to fit standard resources)

DT+MC processing

- 2016 ReReco:
  - 18Apr and 07Aug, both finalized
- 2017 ReReco:
  - 12Sep, partial, finalized
- EOY2017Reco:
  - Just started (expected to start Nov 27, ok)
  - 2B events needed by Mid Dec (to extract Jet calibration, mostly)
  - All the rest by December 2017 (!)
- Total is 14B DT events reprocessed in major campaigns

- MC17v1 (for algorithm development)
  - Finalized, ~1.5 B events
- MC17v2 (for Physics analyses → Winter Conferences)
  - Late by ~ 2 weeks (calibrations and algorithm development)
  - Now:
    - 2.5 B Geant level events done
    - 0.5 + 0.5 B MinBias events done (to be used as PileUp)
    - 250M Premixed events done (to be accessed remotely from CERN and FNAL)
  - Submission of ~10B events started
    - Out of which, 2B by Mid December, for calibration extraction

Status of resources

Central space utilization: “blue” protected on disk (last copy); “green” is dynamic buffer; white is 10-15%, tentatively kept free

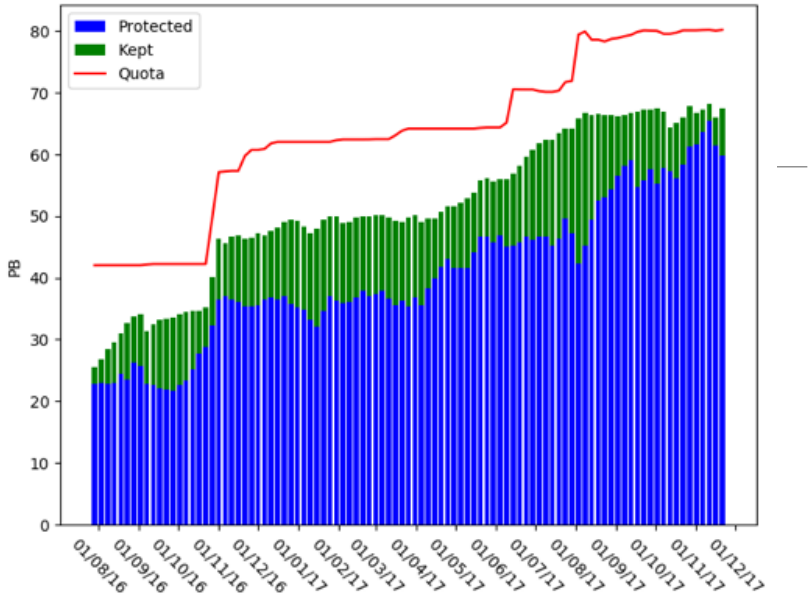

- CPU: full utilization since Summer.
  - We are entering a phase in which speed of processing is essential in order to present results with 2017 data at Winter Conferences
- Disk: We reached at the end of the data taking 85% of Central Space declared “unmovable”
  - 5% still usable
  - 10% is what we try and keep for transfers
  - Getting better after data taking (RECO disappearing, etc), but large reprocessings ahead of us
- Tape: after CNAF’s incident, less than 20 PB free @ T1s. But additional 20PB at Tier0. Seems under control

Weekly transfers 2005/2017

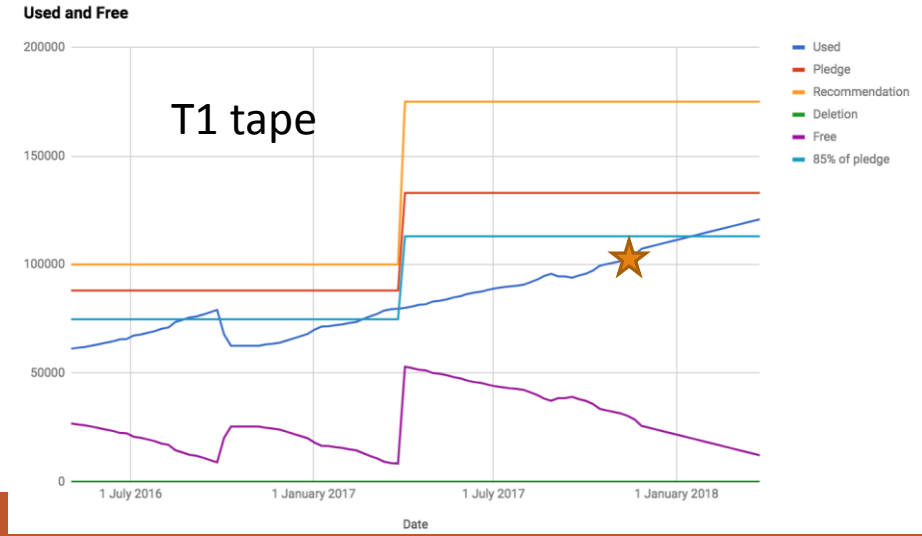

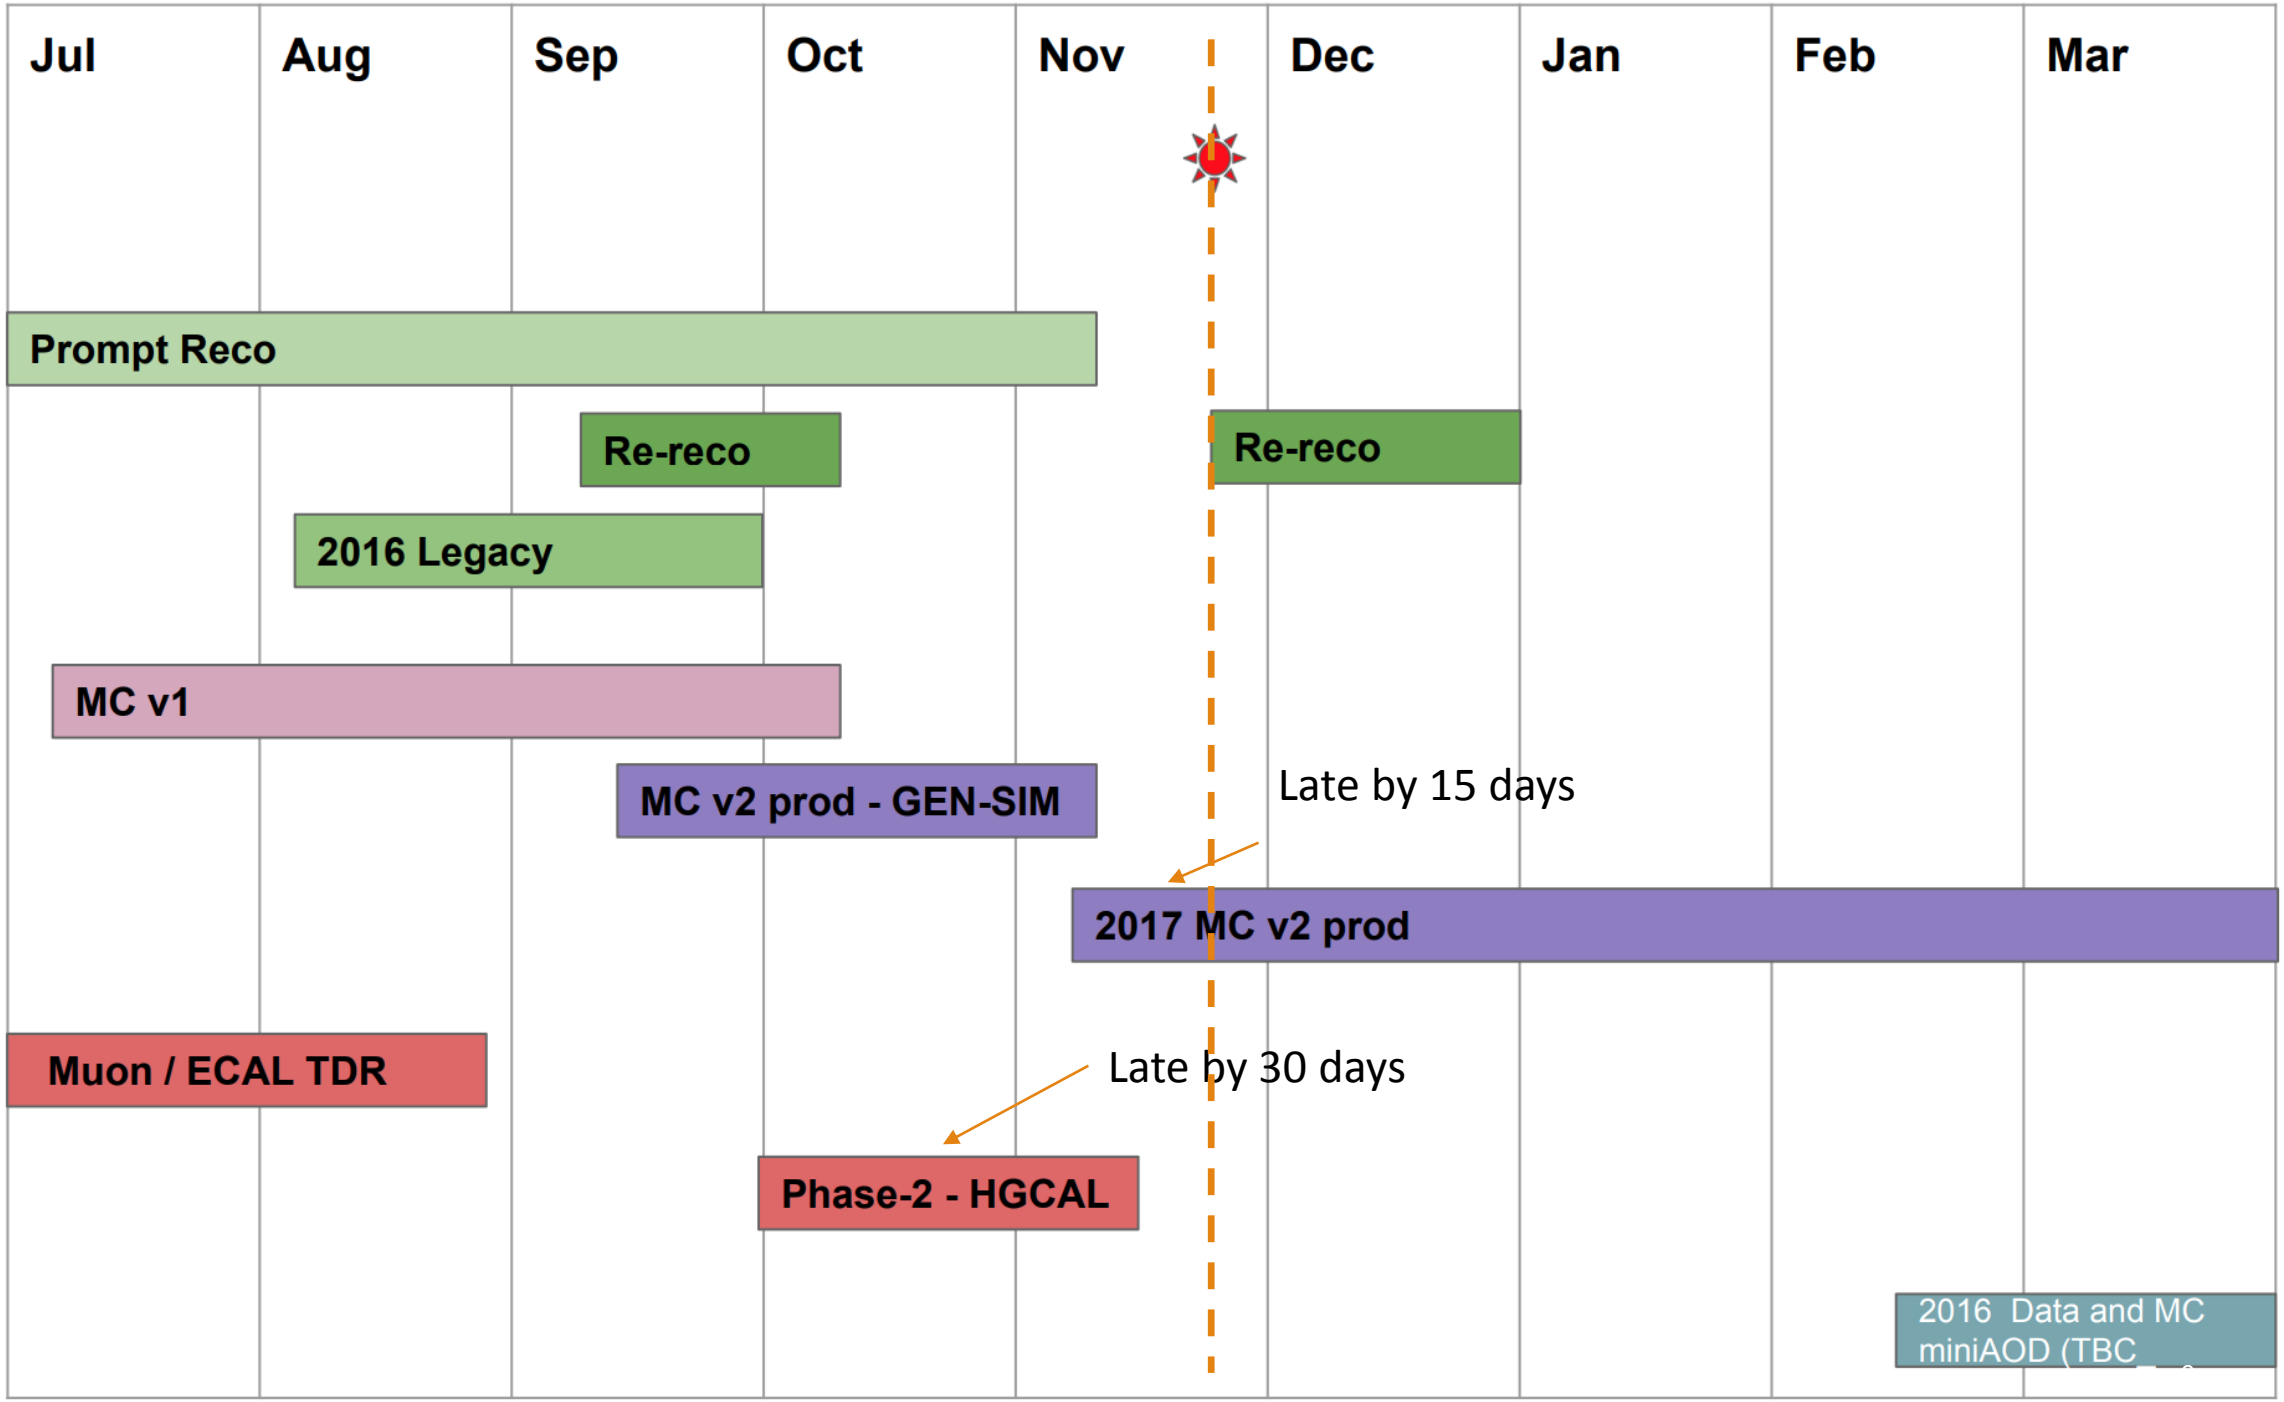
Jul
Aug
Sep
Oct
Nov
Dec
Jan
Feb
Mar
Prompt Reco
Re-reco
Re-reco
2016 Legacy
MC v1
MC v2 prod - GEN-SIM
Late by 15 days
2017 MC v2 prod
Muon / ECAL TDR
Late by 30 days
Phase-2 - HGCAL
2016 Data and MC miniAOD (TBC

Resources – 2018 and 2019

- 2018 request fully confirmed at the Oct RRB
  - With a +2 PB Tier-0 tape due to a machine rescaling of HI 2018 live time predictions
  - How to mitigate 2018 if needed (<PU> > 45 in a high availability year)?
    - CPU: extend Tier-0 processing to distributed computing
    - Disk: limit RECO, limit number of MC events
    - Tape: stop parking
    - (just ideas, no real plan for the moment)
- 2019 request, drivers
  - Full RunII legacy reprocessing in 2019-2020
    - 25B (DT) + 25-35B (MC) events to be reprocessed/generated
    - Starting April-May (critical dependence on calibration availability)
  - Trigger TDR
    - O(50 Mevents), many in "expensive format" (need to simulate trigger at analysis level)
  - No major change for the moment – expect refinement but no revolution for C-RSG submission in February

Resource	Site	2018 CMS Approved Request (Apr17)	2019 CMS Request (Oct17)	2019 CMS Request Increase (Oct17)
CPU (kHS06)	T0+CAF	423	423	+0%
	T1	600	650	+8%
	T2	900	1000	+11%
Disk (PB)	T0+CAF	26.1	26.1	+0%
	T1	60	68	+13%
	T2	70	78	+11%
Tape (PB)	T0+CAF	97	97	+0%
	T1	188	230	+22%

Table 5: *CMS resource request for 2019. The first column shows CRSG Spring'17 recommendations for 2018. The second shows the current requests for 2019, and the third the relative increase.*

Not endorsed by C-RSG, but positive feedback

Resources 2018: requests vs pledges

- 2017: the “change in LHC parameters” crisis; requests updated in November 2016 – too late for most FAs
- This led to pledges under recommendations (up to -24%)
  - This is Rebus view, some FAs tried to help with more resources on the table
- 2018: having been the request process much more linear, recovery expected
  - At the level of 2016 and before: up to 10% @ T1s, compatible with 0 otherwise

2017	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%

2018	ALICE	ATLAS	CMS	LHCb
CERN CPU	0%	0%	0%	0%
CERN disk	0%	0%	0%	0%
CERN tape	0%	0%	0%	0%
Tier-1 CPU	-9%	2%	-6%	-1%
Tier-1 disk	0%	11%	-8%	7%
Tier-1 tape	3%	1%	-12%	12%
Tier-2 CPU	-21%	-2%	3%	16%
Tier-2 disk	-18%	-2%	-6%	-36%

CNAF crisis

- For CMS:
  - CNAF the second biggest T1, with
    - 21 PB of tape (15 used)
      - 6 PB free are also the 2nd biggest reserve we have (had)
    - 3.9 PB of disk (all used) – still not at 2017 pledge, which can be even be a lucky aspect
    - Some 7k cores
  - At the moment considering the 40 ”drown tapes” for CMS
    - 6 are RAW data, the rest is derived data
    - CMS for the moment is finishing replication of those 6 tapes CERN→IN2P3; for the rest of RAW (~2 PB) we are waiting for status assessment from CNAF
      - If no news asap, we will eventually replicate all of the 2 PB
  - Note: excluding CNAF, < 20 PB free overall in the Tier-1 CMS federation
  - Acquiring information from the other T1s + CERN (is there excess availability / early procurement?)
    - Anyhow, the situation needs to be followed, but does not appear as critical
    - CRB next week will follow up

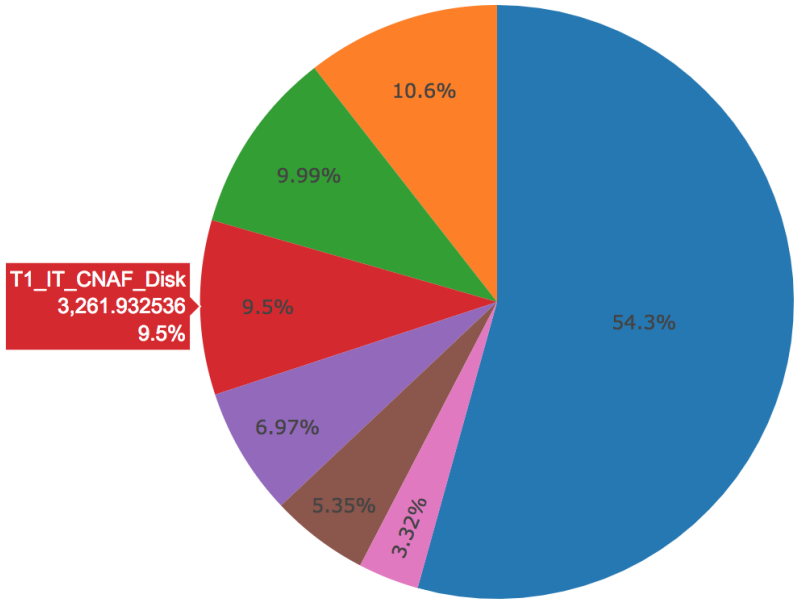

Disk

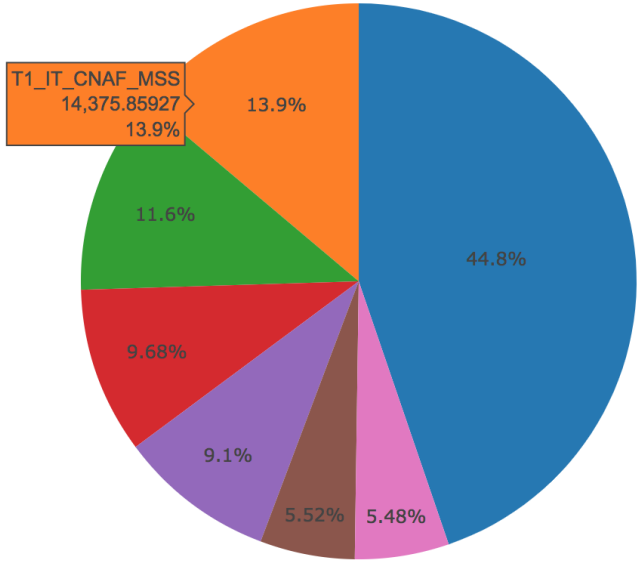

Tape

CPU efficiency

- Work still ongoing, first results can be seen (and were shown at the RRB)
- Infrastructure inefficiency (the pilot model):
  - Down to 3% in optimal conditions (not too many 1-core payloads around, …) – it was up to 20%
- Payload inefficiency
  - Identified many problems with generator fragments, causing multi threaded call stalls
    - Code not our under control (theoreticians, …)
  - Implementing eviction protocols for jobs which did not use enough CPU in the last N hours
    - Should prevent jobs stuck / IO events and such

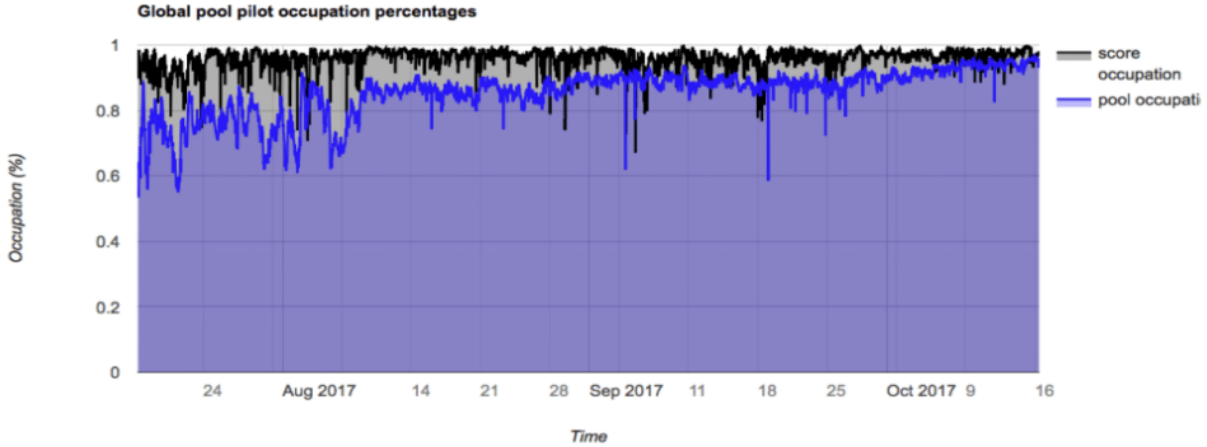

Figure 1: Occupation of multi core (mcore) and single core (score) pilots from July to October 2017. The inefficiency due to the association of payload to pilots has decreased from greater then 20% to less than 5%.

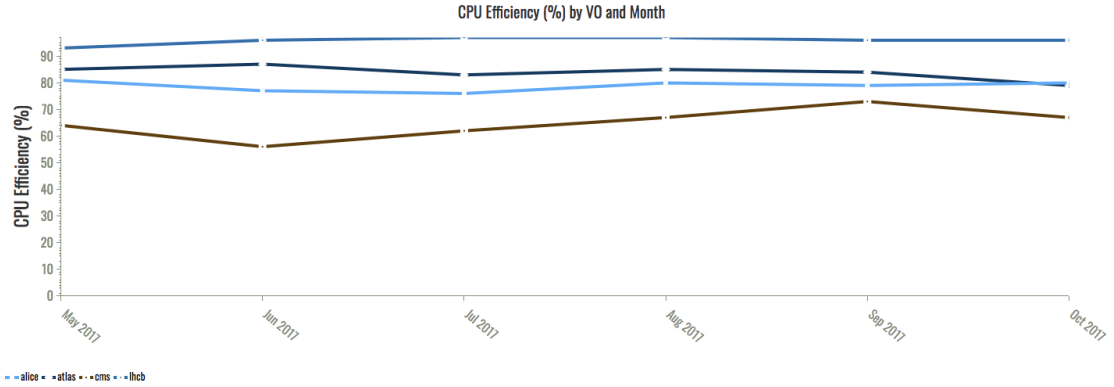

T1 plots: recovery Jun-sept

Oct: many MC16 single core payloads

CPU efficiency – Tier0

- Needs clarification!
- For Tier-0, efficiency is not computed as CPU_time/WALL_time, but as CPU_time / UPTIME of the machines
- So they appear always as 100% used, with a CPU efficiency that can be 0% if they are not used
- So, what happened?
  - May-Aug: EOS problems, Tier-0 was not able to sustain additional load wrt to data taking
  - Sept on: started to climb, expect the problem to be gone
  - 2018: going back to fairshare?
    - Under discussion with IT; Tier-0 systems will be again uniform between experiments and not allocated a priori

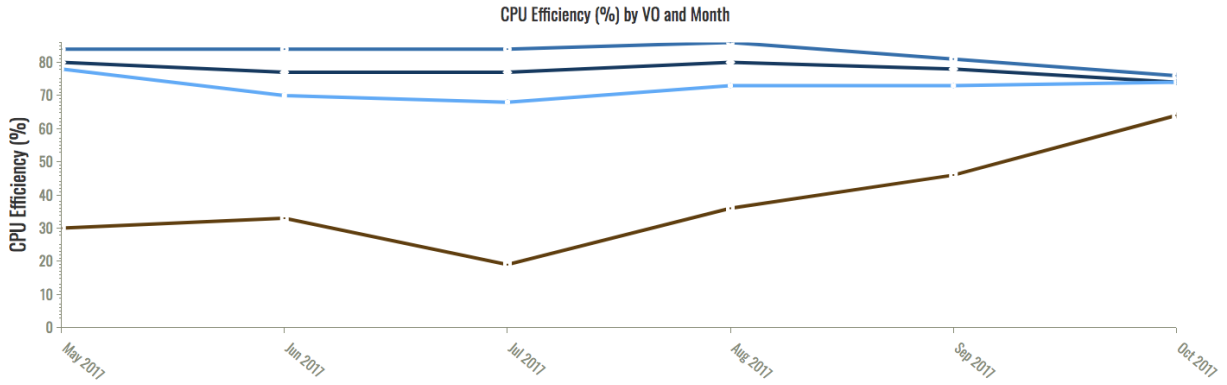

Another example:Tier-0 had 2 days of pause last week, and restarted lower due to the start of tests for the new CERN HTCondor deployment. It will result as “inefficient”

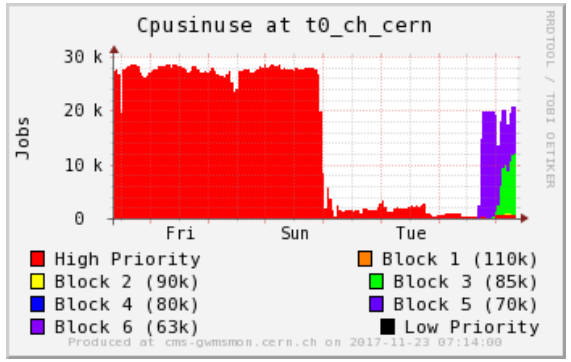

RunIII

- Not yet a real modelling
  - 2021 is 4 years from now
  - 2022 should be what counts, since 2021 a startup year (is it ? – we need modelling from the machine!)
- According to the last available machine parameters (2.5e34 on 2500-2800 bunches), it should not be far from the 2018 scenario. No evidence of HLT rate increases on CMS side.
- Expect 150/fb on 2022 and 2023 (again, is 2021 a startup year)
  - Storage: total RunIII = 150/fb as RunI+RunII. Expect a +100% needed
    - Strictly true for tape, could be less for disk depending on how fast we close the RunII analyses
  - CPU: less dependent on past history, +50-100%
- Correction factors?
  - NanoAOD is are up to the promises (reduce need for analysis CPU – 50%?)
  - NLO/NNLO if needed on large scales (Madgraph → Sherpa could be a huge increase – see ATLAS vs CMS today)

For the moment an increase of

- 100% storage
- 50-100% CPU

seems reasonable

AoB - NanoAOD

- NanoAOD preparation going fast
  - Current prototype smaller than expected:
    - MC: 1000 bytes/ev
    - DT: 600 bytes/ev
    - Producing at @ 15 Hz from MiniAOD
    - Reading up to 5 kHz (analysis dependent)
  - (the expected Legacy RunII, 60 B events, would fit in 50 TB)
  - Tests with the production system positive: our WM can handle the new “flat” format
  - Idea is to try and produce NanoAOD on a regular basis starting from Spring18
  - Still ironing out details like
    - Release blessing – who?
    - Fraction of analyses for which it will satisfy (50% is the bare minimum)
    - The fact that this is smaller than expected can allow for easier inclusion of more analyses

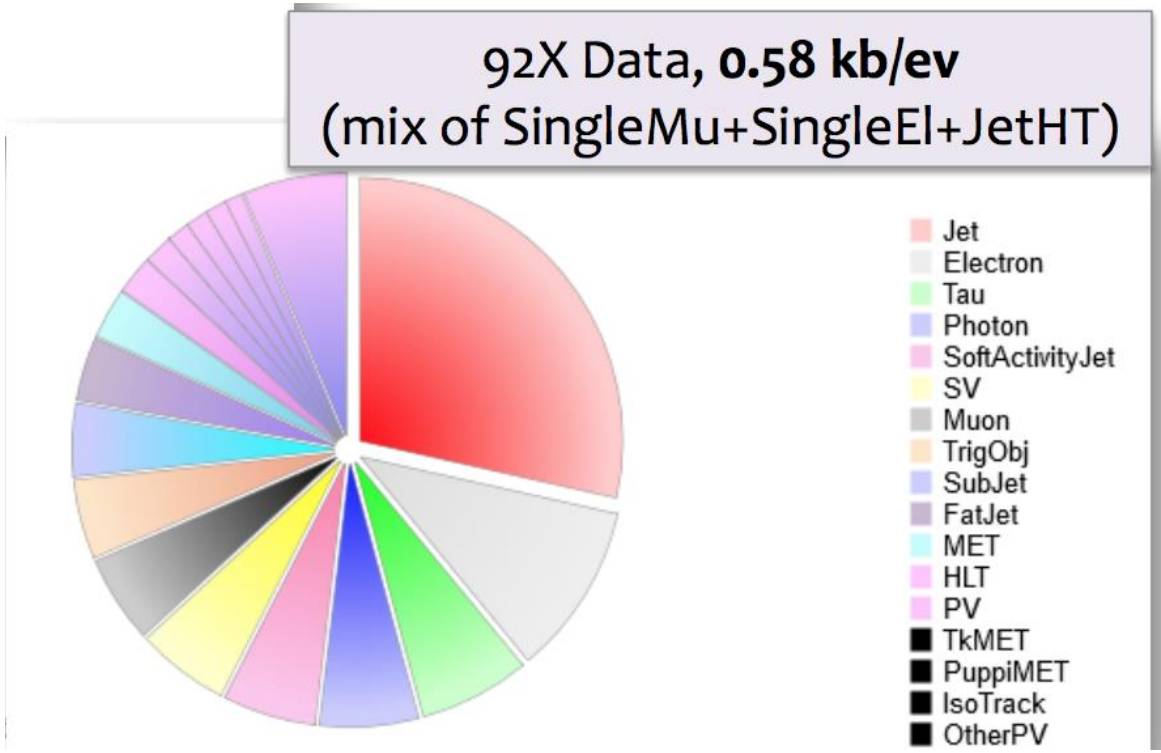

AoB #2

- SW releases: 4 cycles open
  - CMSSW_9_2_X: used for Data taking (Tier-0) and for MC17v1
  - CMSSW_9_3_X: used for HI, and for Geant4 simulations
  - CMSSW_9_4_X: used for Winter 18 DT and MC reprocessing
  - CMSSW_10_0_X: used as devel release for 2018 Run
- Data Management
  - CMS held a workshop, we are trying to finalize a plan.
  - At the moment, on the RunIII scale and beyond, two viable solutions
    - Extend the development of Dynamo (MIT)
    - Try Rucio (ATLAS) – in case there is a manifest interest in common solutions from all the involved parties
  - Evaluation to be concluded in Fall18, after two parallel reviews

Conclusions

- Last 2 months of the Run more complicated than expected – LHC surprised us with smart solutions to the 16L2 problem
  - This is potentially a “problem” for 2018: resources requested with old parameters (which are still the baseline on the LHC page we use for planning)
  - It is pretty clear that we can expect <PU> can be quite larger than 35
    - Working on a number of mitigations – CPU, disk and tape
- Preparation of 2017 DT and MC for the Winter Conferences late (reflecting the difficulties CMS detector experienced in 2017) – still trying to meet the goals
- CNAF issue:
  - For the moment operating recovery only on RAW on the “underwater” cassettes (6) – already replicated to IN2P3
  - Waiting for more precise info on the situation of the rest of the tapes (~15 PB there, out of which ~ 2 PB custodial RAW)