News from CMS

A. Sartirana

Plan of the Talk.

- **2017 highlights**
  - few words on **Ops in 2017;**
  - follow up on the **job efficiency** problem;
- something about the future
  - few statements on **Run II-III** and **HL-LHC;**
  - sparse info on computing evolution
    - NanoAOD, DM evol., Dynamic Res., SW evolution
- just a selection of new from some talk/meetings
  - talk at **WLCG-LHCC, Fall17 O&C week** & **DM pre-week, Rio CMS week and few others;**
    - cut & paste (& digest), but of course, I do not master all the details…

References.

WLCG-LHCC Talk

https://indico.cern.ch/event/570975/contributions/2309740/subcontributions/208383/attachments/1521736/2377733/cms_lhcc_wlcg_Sept_2017.pdf

Fall 17 O&C Week

https://indico.cern.ch/event/625239

Rio CMS Week

https://indico.cern.ch/event/611870/timetable

DM pre-week

https://indico.cern.ch/event/676670

Public

CMS « internal »

2017 data.

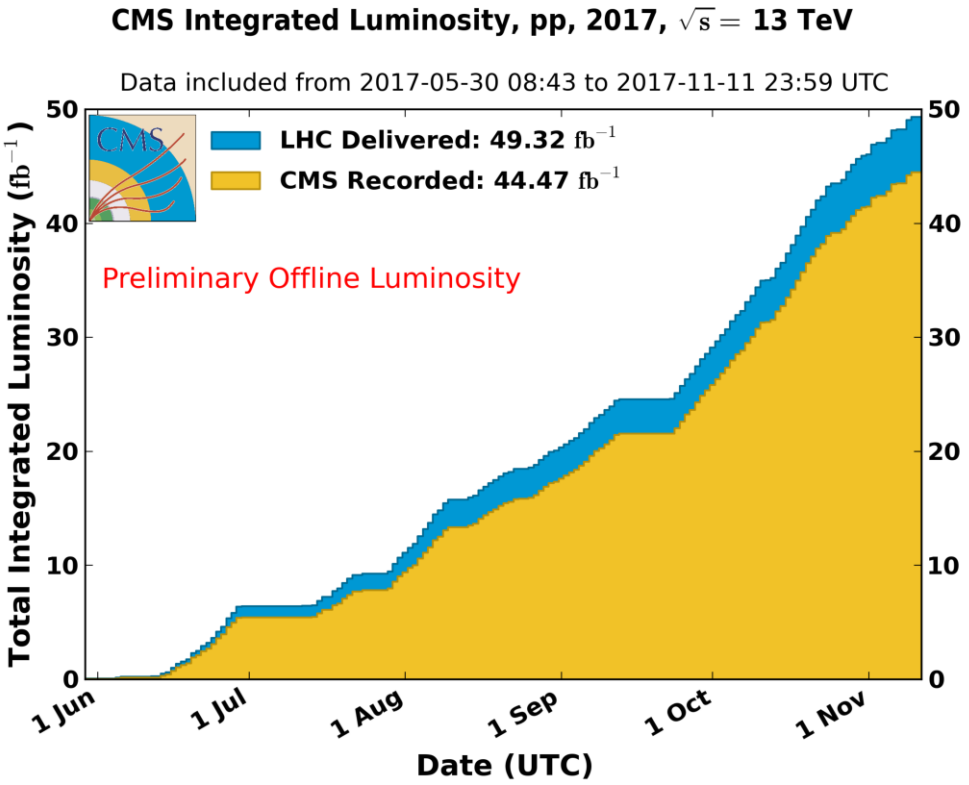

Exceeded 2017 exp

- R2 could **> 150fb^{-1};**
- **smooth** data taking…

… with **few T0 issues**…

- express data kept 3x longer for new (pixel) detector comm.;
- workspace close to quota (6PB);
- few pbs in T0-T1 transf.;

… and **few EOS issues**…

- **scalability** issue: **auth. rate** (GSI) topping at 1 kHz;
- lost files / files becoming 0-sized after some time;
- **solved on Aug 11th** with EOS patch.

2017 activity.

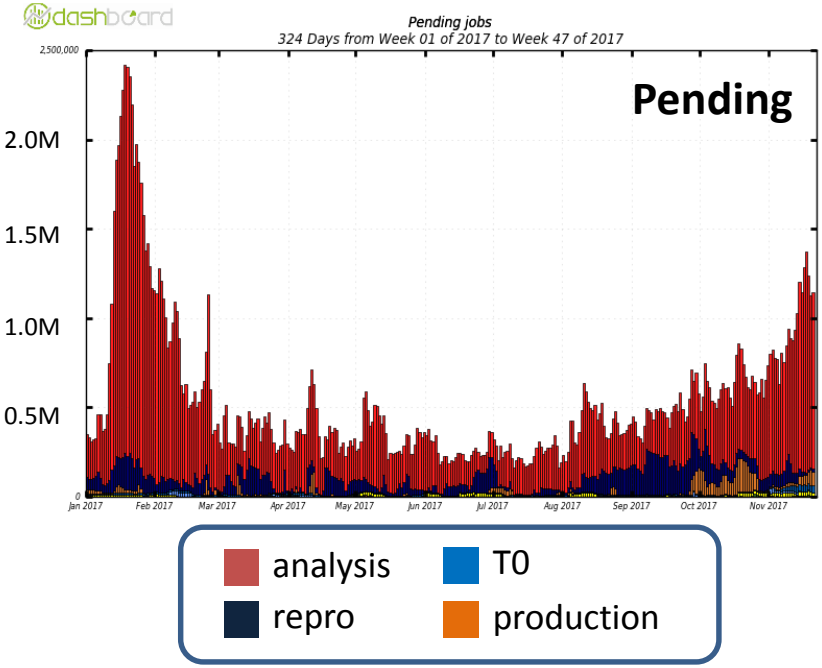

- Pool 150-200k cores;
- analysis O(50k) cores;
- pending O(500k) cores.

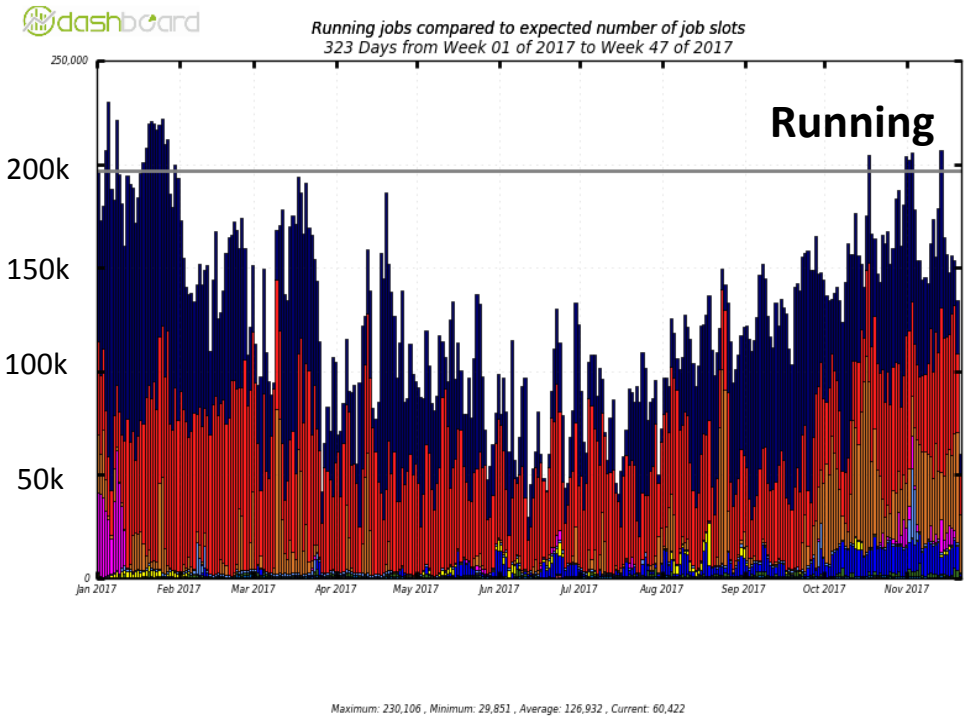

News!

"Important news for sites: we will be **requiring Singularity** at all sites **as of February 2018**. This is because CMS will be **processing data** for 2018 in a **CentOS 7-only release**…" (merci Sébastien!)

Data management.

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%

- **3 disk cleaning** campaigns needed
  - plan: std expir. for data;
- **dynamo** (DDM) keeps usage optimal
  - some side effects: see next slide.

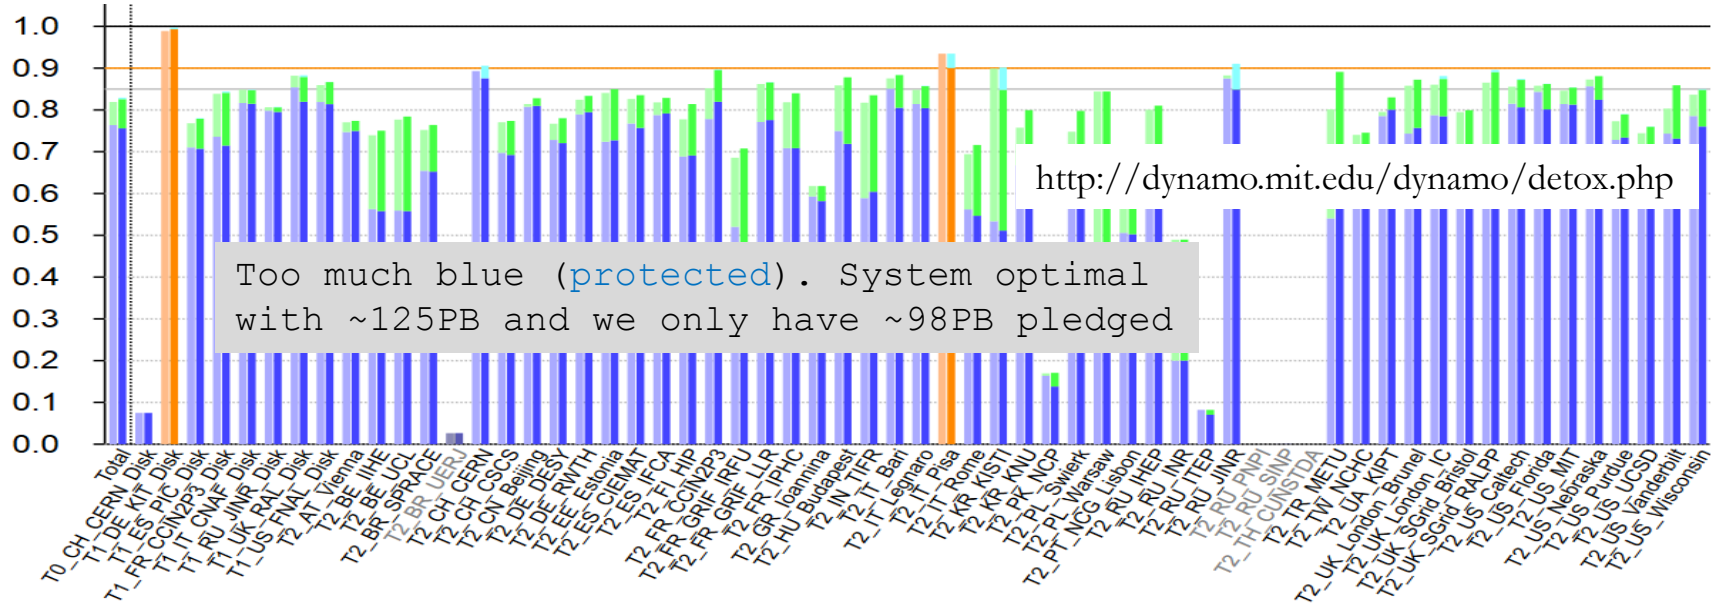

Transfers.

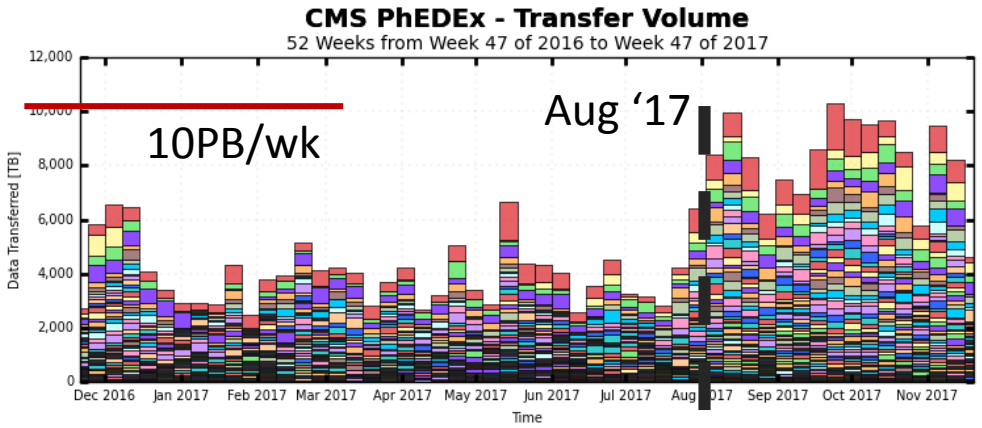

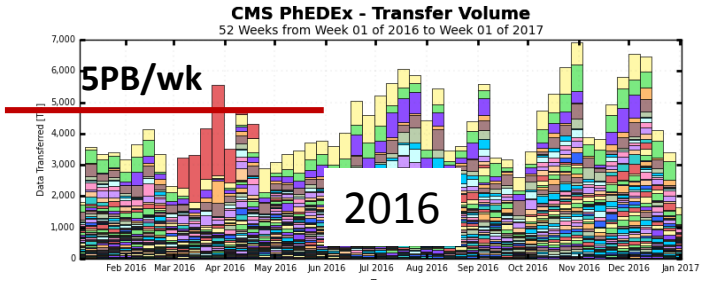

Unprecedented transfer **volume** (10PB/week)…

- larger event size;
- consistent Dynamo (**DDM**) transfers since Aug;

… and the **system** mostly **coped with** it

- … but showed some limits…
- add **4th prio.** queue;
- problem with transfers to/from tapes.

Jobs (in)efficiency.

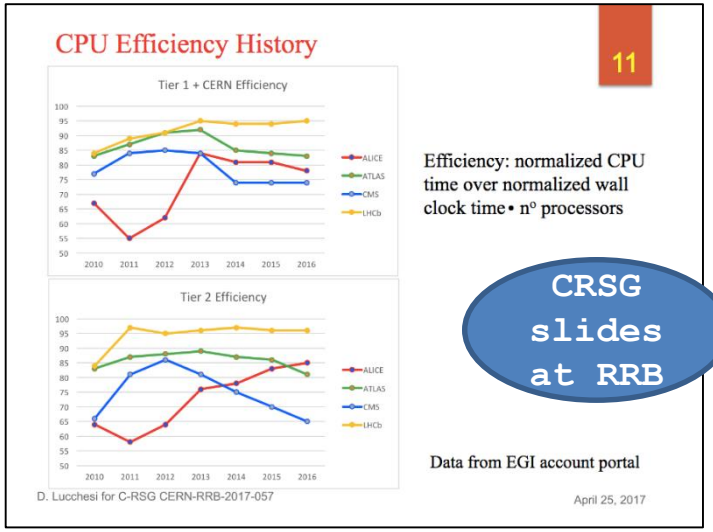

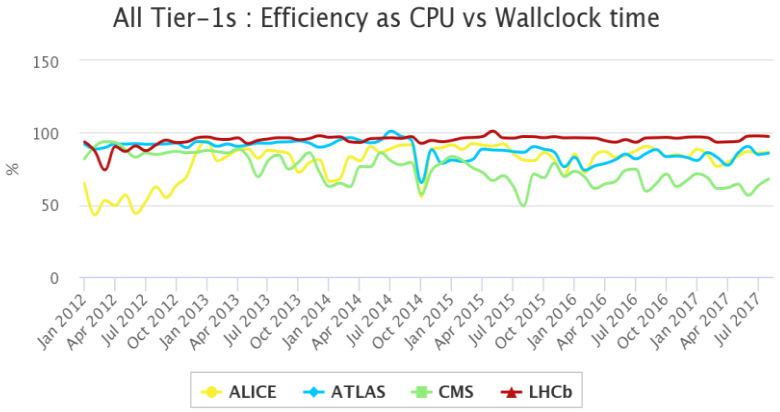

Worse CPU/WC ratio in LHC exp; created (~1st June) a **task force** mandated to

- use monitoring data to **study** all **sources of CPU inefficiency** in depth;
- **suggest** a course of **action** to the management;

in few month of activity the TF already **identified major causes** of inefficiency

- some **actions** have already been taken.

Inefficiency sources.

We can identify **2 CPU inefficiency factors**

- ineff. associated with the **submission infra**
  - CPUs not being utilized as the **pilot infrastructure is not able to direct work** to the pilots resulting in unoccupied and inefficient cores;
  - this is best understood thanks to the great work of the SI group;
- ineff. associated with the **job itself**
  - software not being able to utilize CPU efficiently due to data access latency or **Amdahl's** law in MT apps;
  - **more difficult to study**, but some results achieved also in this sector;

total efficiency is the product of the 2 which can easily lead to a very low overall efficiency.

SI inefficiencies.

MC pilots and Global pool SI **peculiar to CMS**

- often plainly stated as cause of the ineff. but

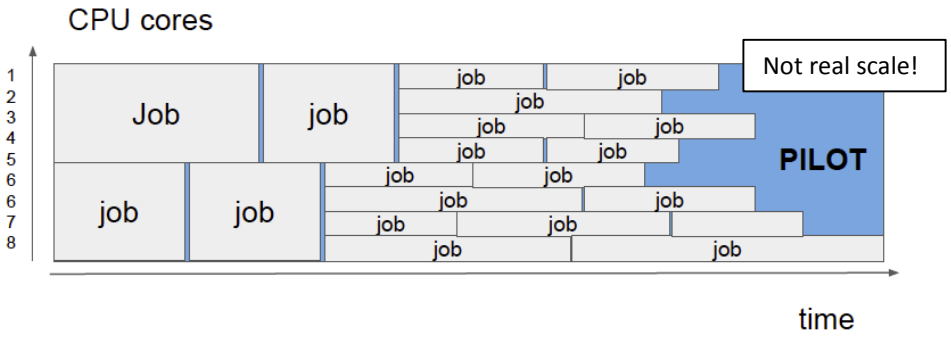

- **ineff. existed long before** we had this model…
- …and, actually the reuse of pilots makes MC sched. easier and saves several mins of WCT per pilot;

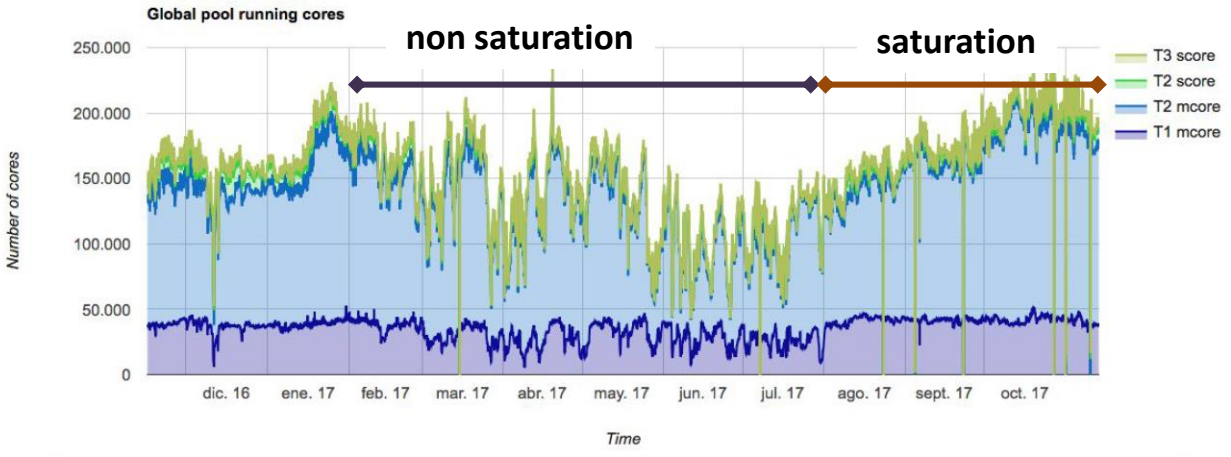

- **Saturation and non-saturation:** 2 types of inefficiency.

Non saturation.

Bursty submission cycles with timescales ~ pilot length

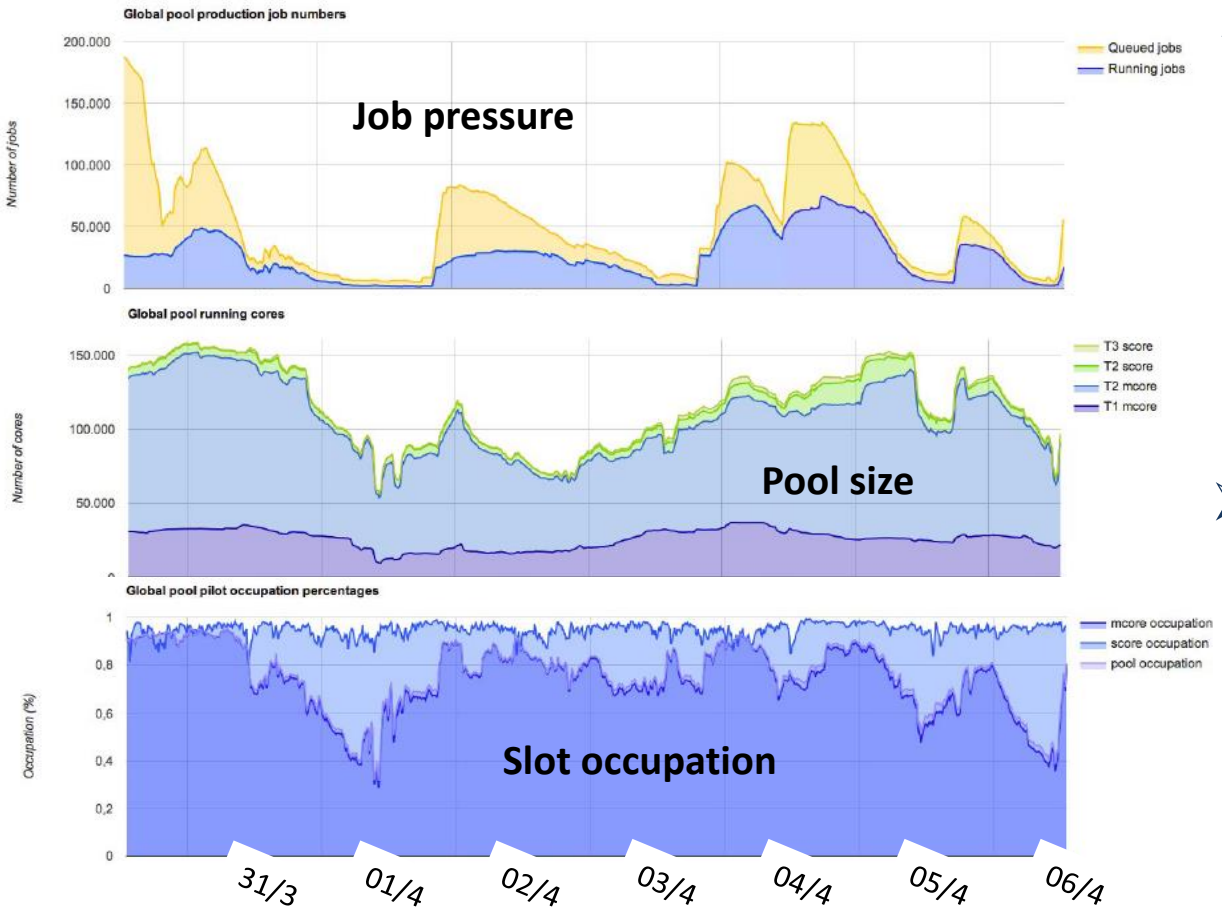

- inefficiency comes from **no longer needed pilots running** after real job pressure drops;
- **up to 60%** of the pool slots **running idle** for hours as the pool drains and contracts.

Pilots de-queuing.

Mitigation: remove pilots before they get to run empty, e.g. retiring them **after 1h in queue**

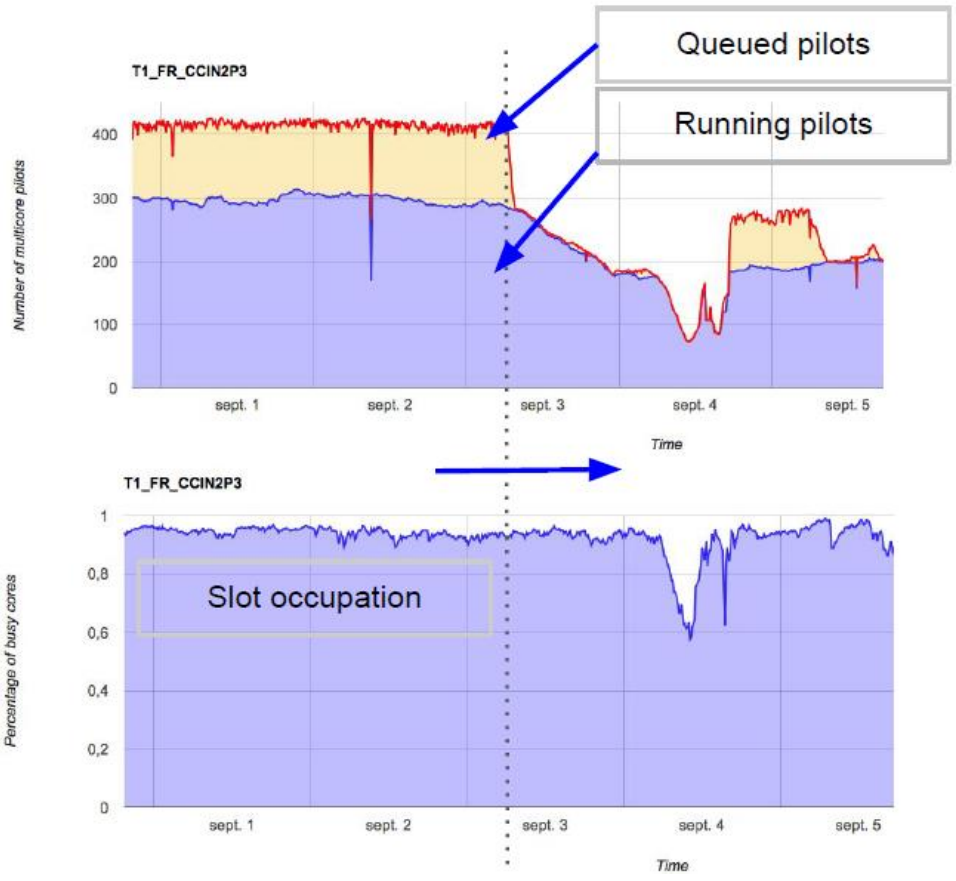

- enforce **tighter coupling** between **pilot pressure and job pressure;**
- **implemented late July.** Soon after the pool entered saturation: difficult asses the real effect on the pool;
- observed on individual sites. E.g. IN2P3 T1: sudden drop of pressure causes much **less dramatic effect;**

Saturation.

Since Aug running in saturation mode. Idle core:

- **mem constrained slots:** jobs with high memory requirement can run on standard resources at the price of keeping some slots idle;

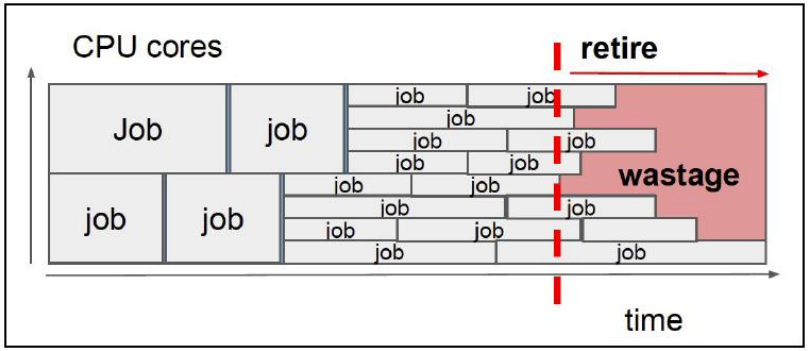

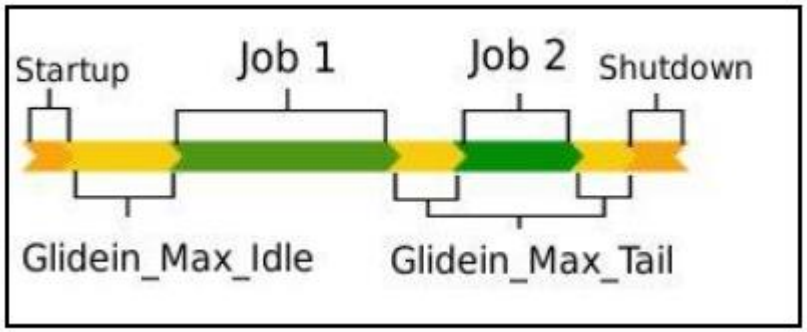

- **idle in retiring pilots:** pilots start draining to prevent payloads run over the max pilot lifetime which leads to wastage. Mitigation: retire time was reduced from 10h to 4h on 9/10;
- **unclaimed idle slots in general:** when idle a pilot wait before retiring to increase chances to be matched. Related to nego cycle (which is kept <300s). Reduced from 20m to 10m on 17/7.

Results.

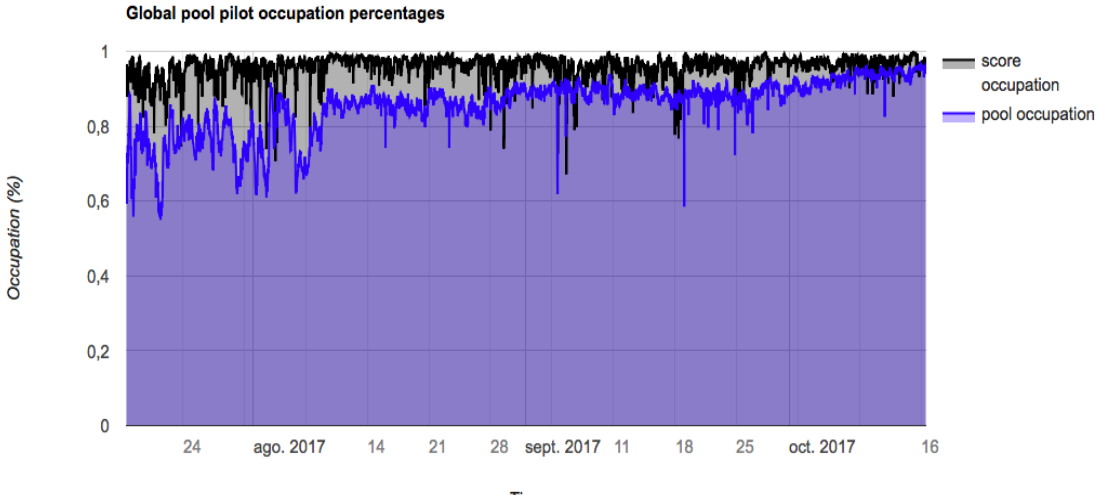

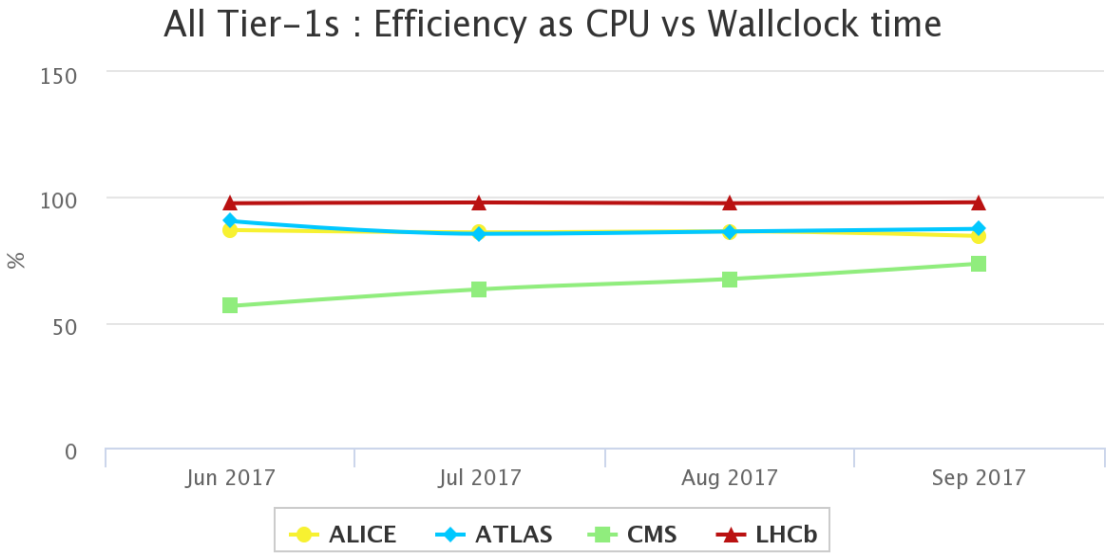

Multiple changes show **improvement.** However partially due to consistent 8 cores jobs pressure.

Clear improvement during the lifetime of the TF. **Still not as high as ATLAS.**

If the core occupancy is ~98% then there **must be ineff. in the payload as well.**

Payload ineff.

CMS only LHC experiment to have **truly MT SW**

- **several advantages**: less memory/core, flexibly scaling from single to multi (and even to many) core;
- may bring **inefficiency**: Amdahl's law.

CMS Apps in controlled env have >95% eff

- but this may not be the case on the Grid.

Study of payload ineff. is **less advanced** than for SI ineff

- uses data collected and stored in ES@CERN;
- seen effects due to **misconf. jobs** (running ST modules with MT conf.) and to **offsite data** access. Other, less obvious, effects are under study

Rest of Run II.

2018: plans unchanged, requests unchanged

- still taking ~ 1 kHz of Trigger at a PU ~ 35, and SB live time not far off;

2019: estim. 4/17 RRB, discussed on 10/17 RRB

- full repro of Run II data + MC;
- phase II studies. Analysis at least at the level of 2018;
- availability of ~100% T0 and 80% of HLT resources for offline processing
  - CERN becomes a 1MHS06 offline processing site;
- +0% for CERN, +20% T1 tape, +10% others (wrt 2018).

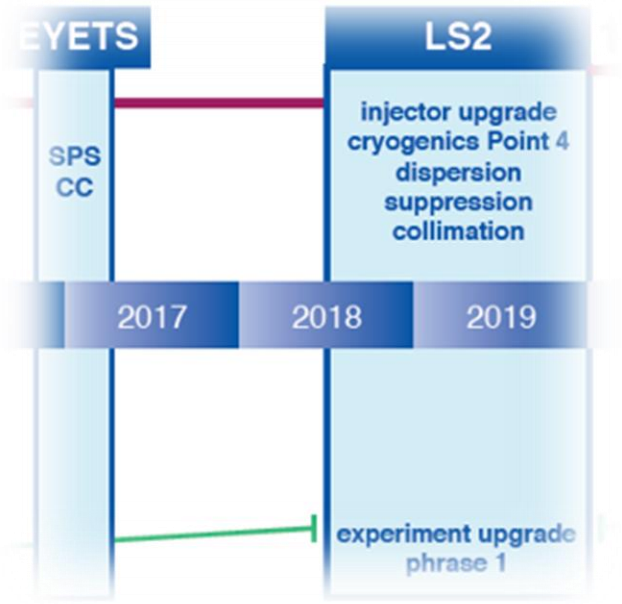

Run III.

No sharp step in resources needs **expected** between end-of-RunII and start-of-RunIII

- target 2.5e34 that is ~130% wrt R2. +0.5-1TeV irrelevant. HLT exp. <1kHz. PhII config not much;
- which does not mean we are not planning changes. Just evolution rather than revolution…

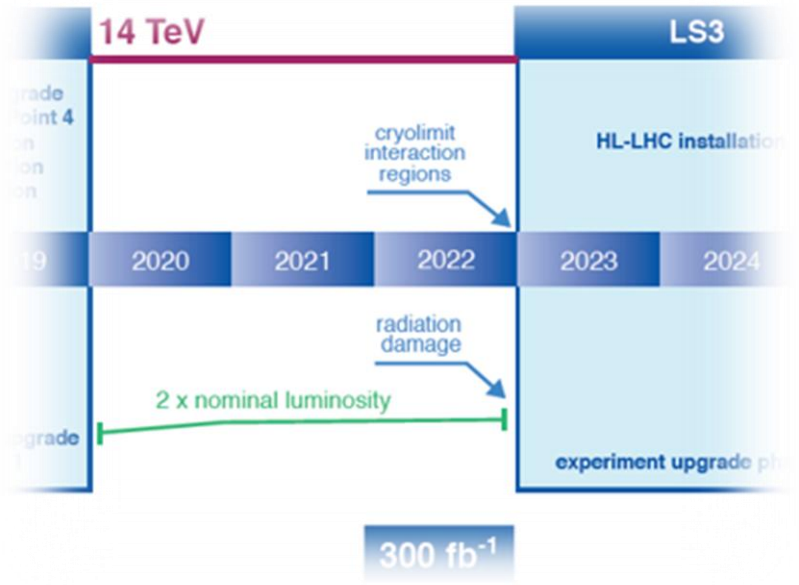

- study of NanoAOD (1-2 kB/evt) shared by ~50% of the analysis
  - if successful, reduce CPU needs for analysis;
- tape/disk data lifetime model;
- all wfs MT by RunII;
- CMS ECoM'17 report with recommendations for the medium term.

HL-LHC.

No officially **blessed** CMS **numbers** for HL-LHC. But **DOE** request to US-CMS for **long time planning** triggered some work

- diff wrt older models are
  - 10%/y code performance improvement exp;
  - largely rely on MiniAOD(SIM) for ops. AOD(SIM) archival thing;
- assuming +20%/y by Moore and friends excess:
  - ~6x CPU;
  - ~4x disk.

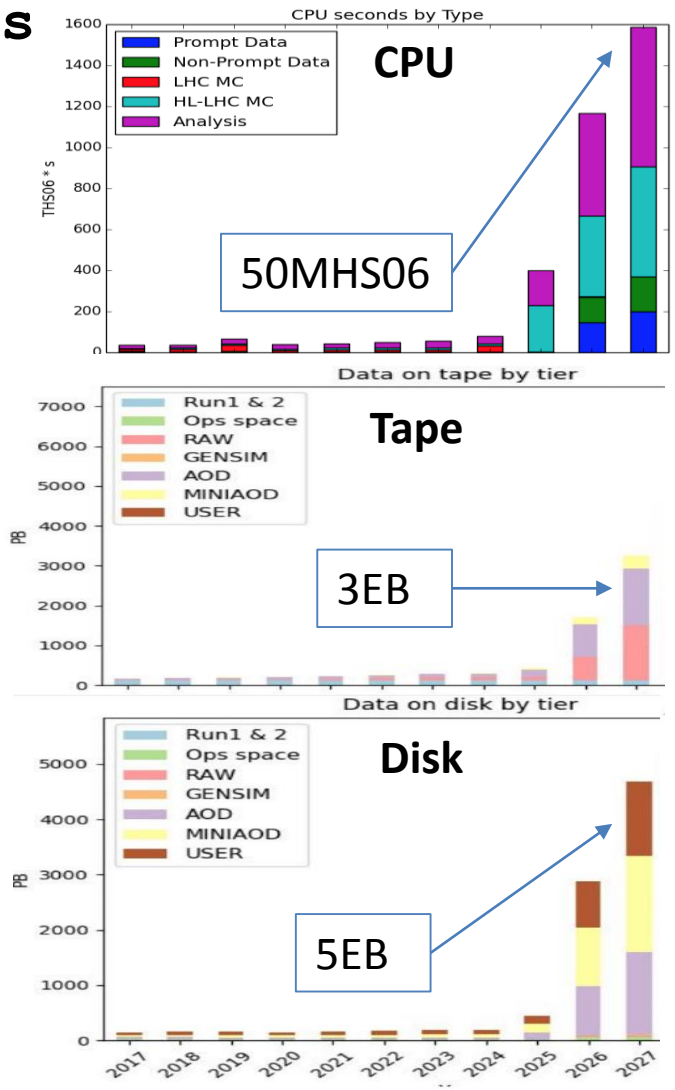

NanoAOD.

New data **format** for analysis **under study**

- at the moment only a prototype (but not only on paper). **Numbers** below are **indicative**. The final performances remain to be seen;
- **smaller** ~1.5kB/event MC, <1kB/event data
  - MiniAOD is ~40-50kB/evt;
- **fast to produce** (from MiniAOD): ~20-30Hz
- **fast to read** O(10kHz)
  - an average for MiniAOD is O(5Hz);
- if we assume that 50% of analysis may be using NanoAOD this may **save up to 15% of CPU.**

Sic transit….

PHEDEx glorious **DM tool** dating start of RunI

- doing its job (and will for all R2);
- **issues emerging** from increased vol…
  - ~1PB/day, O(500k) files. Run III ~ same. Run IV O(50) higher;
  - rate/pressure to/from tape issues;
- … and some function is missing
  - better prio (added 4th queue);
  - no user data handling;
- support heavy (2FTE) and difficult
  - mostly a blackbox to ops = burden on dev;
  - original designer gone. Only 1 dev left;
  - site support by ops is heavy;
  - some Perl libs are orphaned.

Data Mgmt Evol.

1.5 days WS to see **status and possible evolutions**

- **short middle term PhEDEx** dev still there;

Dynamo: CMS DDM, vital for R2 ops. Already taking some PhEDEx tasks (consistency). Planning to incorporate PhEDEx. Will work with the R4 DM or it will be the R4 DM.

Transfer2Go: developed by a Google summer of Code student. De-centralized design (similar to AAA). Rely on Go features. So far 0 cost.

Rucio (Atlas DDM): solid and future proof. Similar use case scale. Ready to collaborate (and CMS willing to test). Some concerns on political implications and openness of the product.

… plus other very interesting stuff: **data lakes, Big CVMFS**

Dynamic resources.

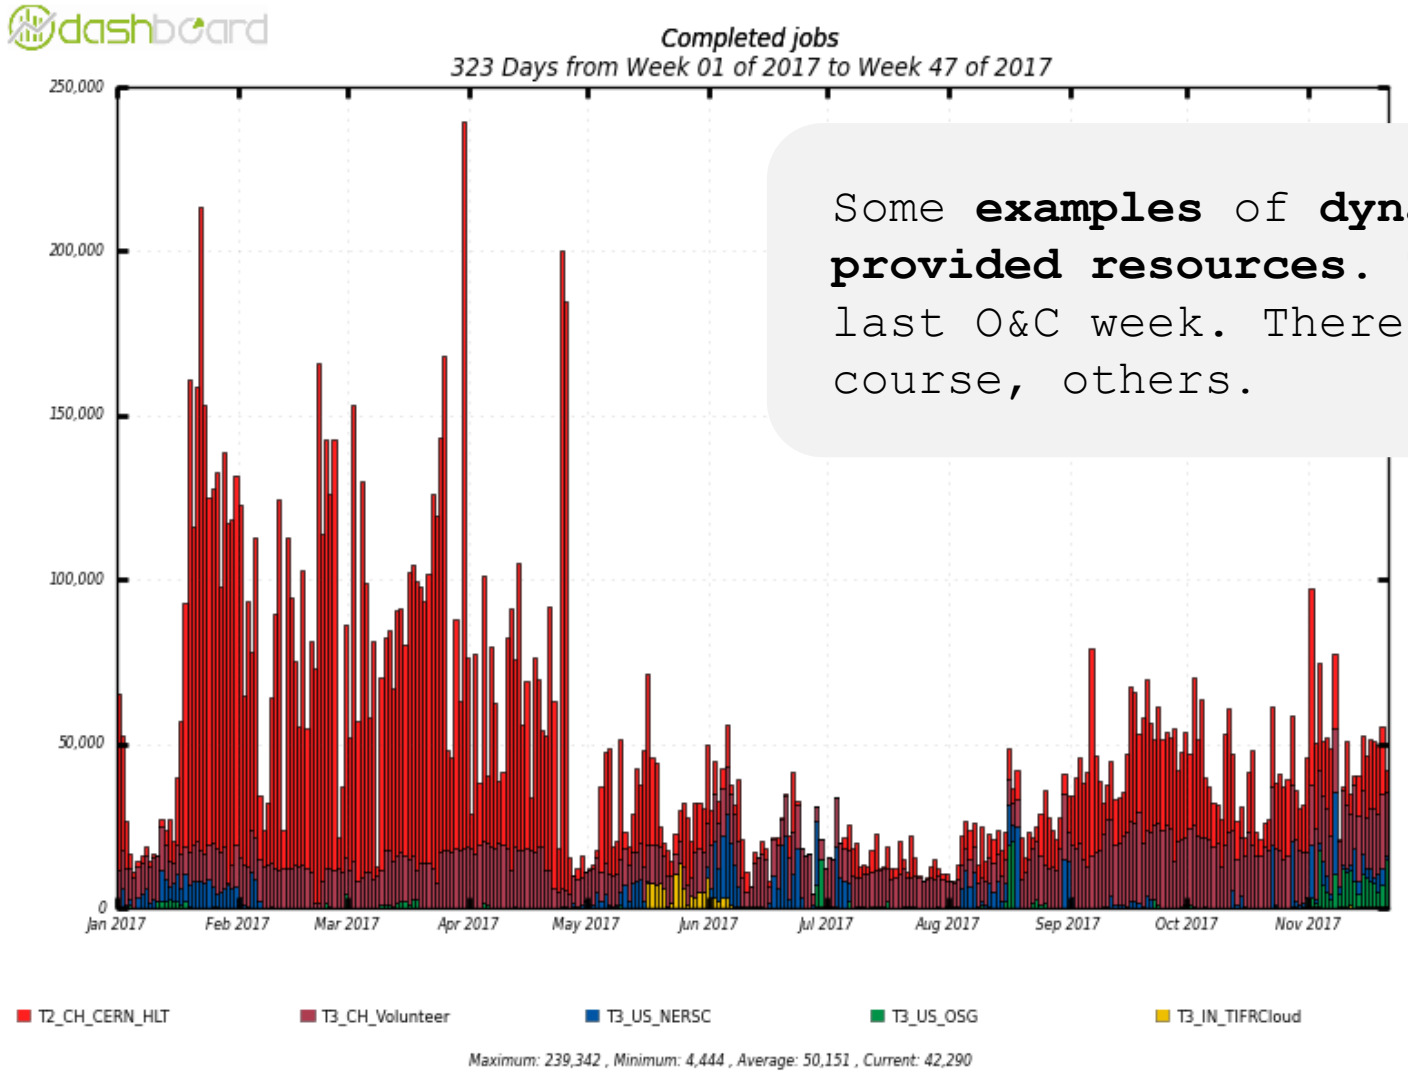

Some **examples** of **dynamically provided resources**. Taken from last O&C week. There are, of course, others.

Dynamic resources.

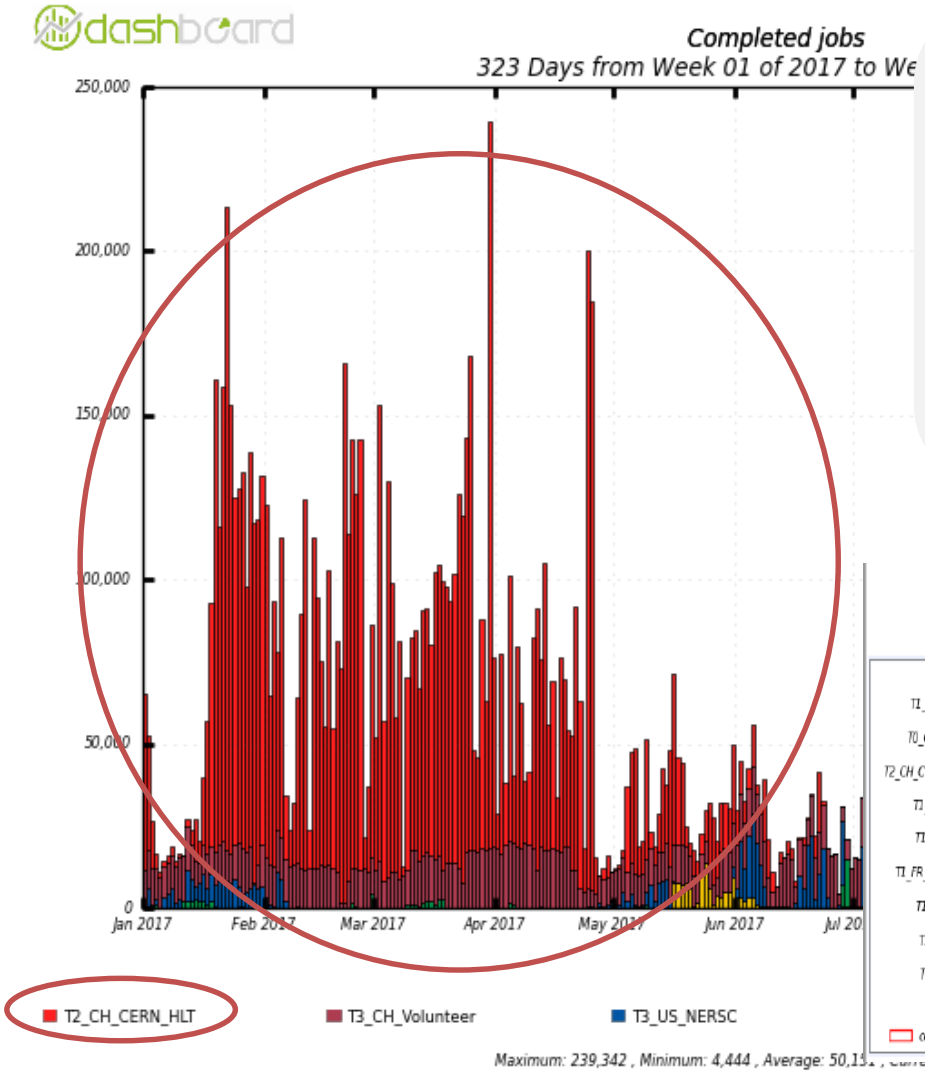

HLT: big. bigger than T1s. Resources not needed for DAQ used for prod. Recently migrated to CC7/Ocata.
In 2018 rump up of cloud usage during fills.

From 11.Nov.2016 to 10.Nov.2017

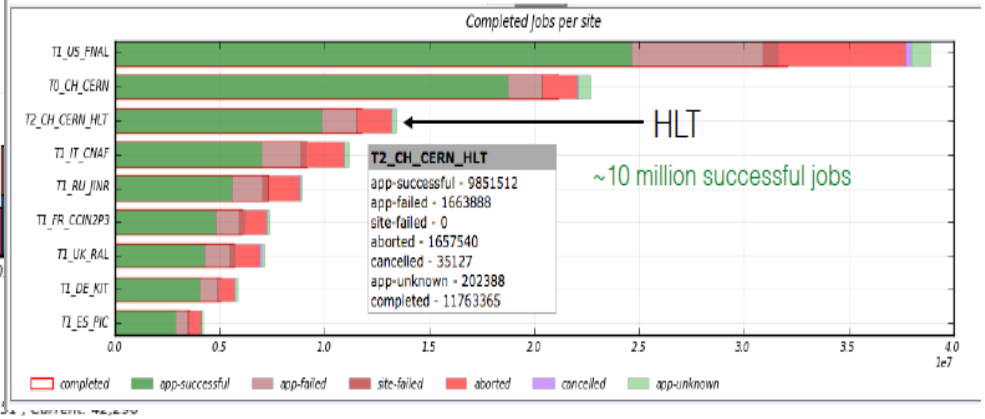

Dynamic resources.

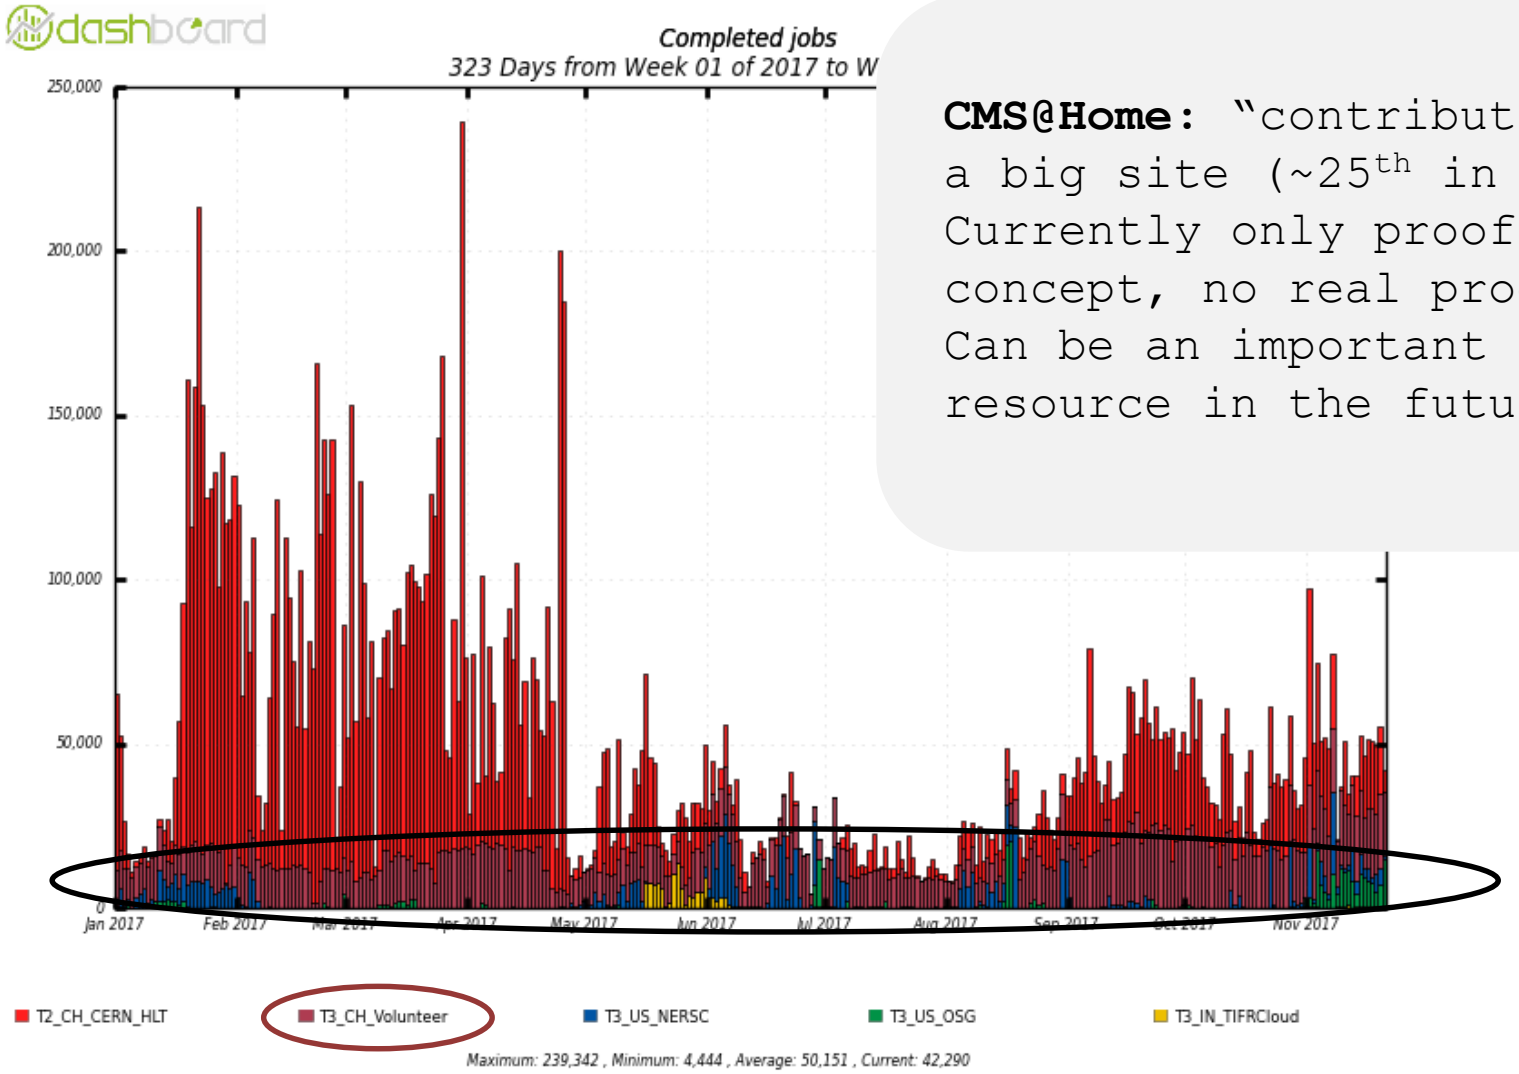

CMS@Home: "contributes" as a big site (~25th in rank). Currently only proof of concept, no real prod jobs. Can be an important resource in the future.

Dynamic resources.

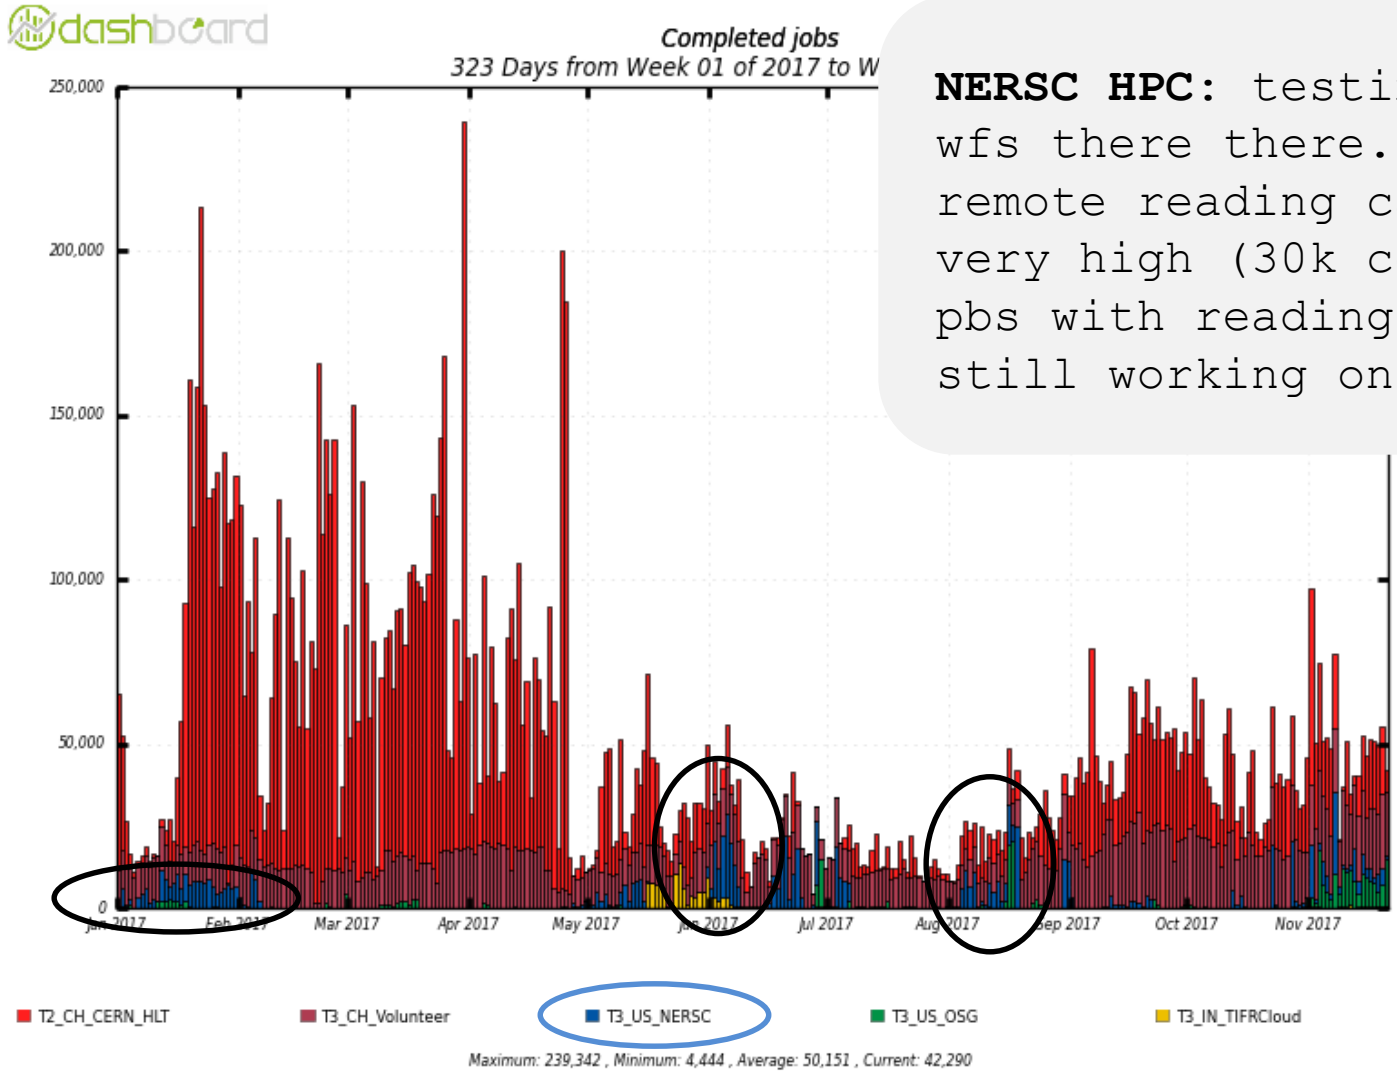

NERSC HPC: testing running wfs there there. if no remote reading can scale very high (30k cores). Some pbs with reading data, still working on it.

Dynamic resources.

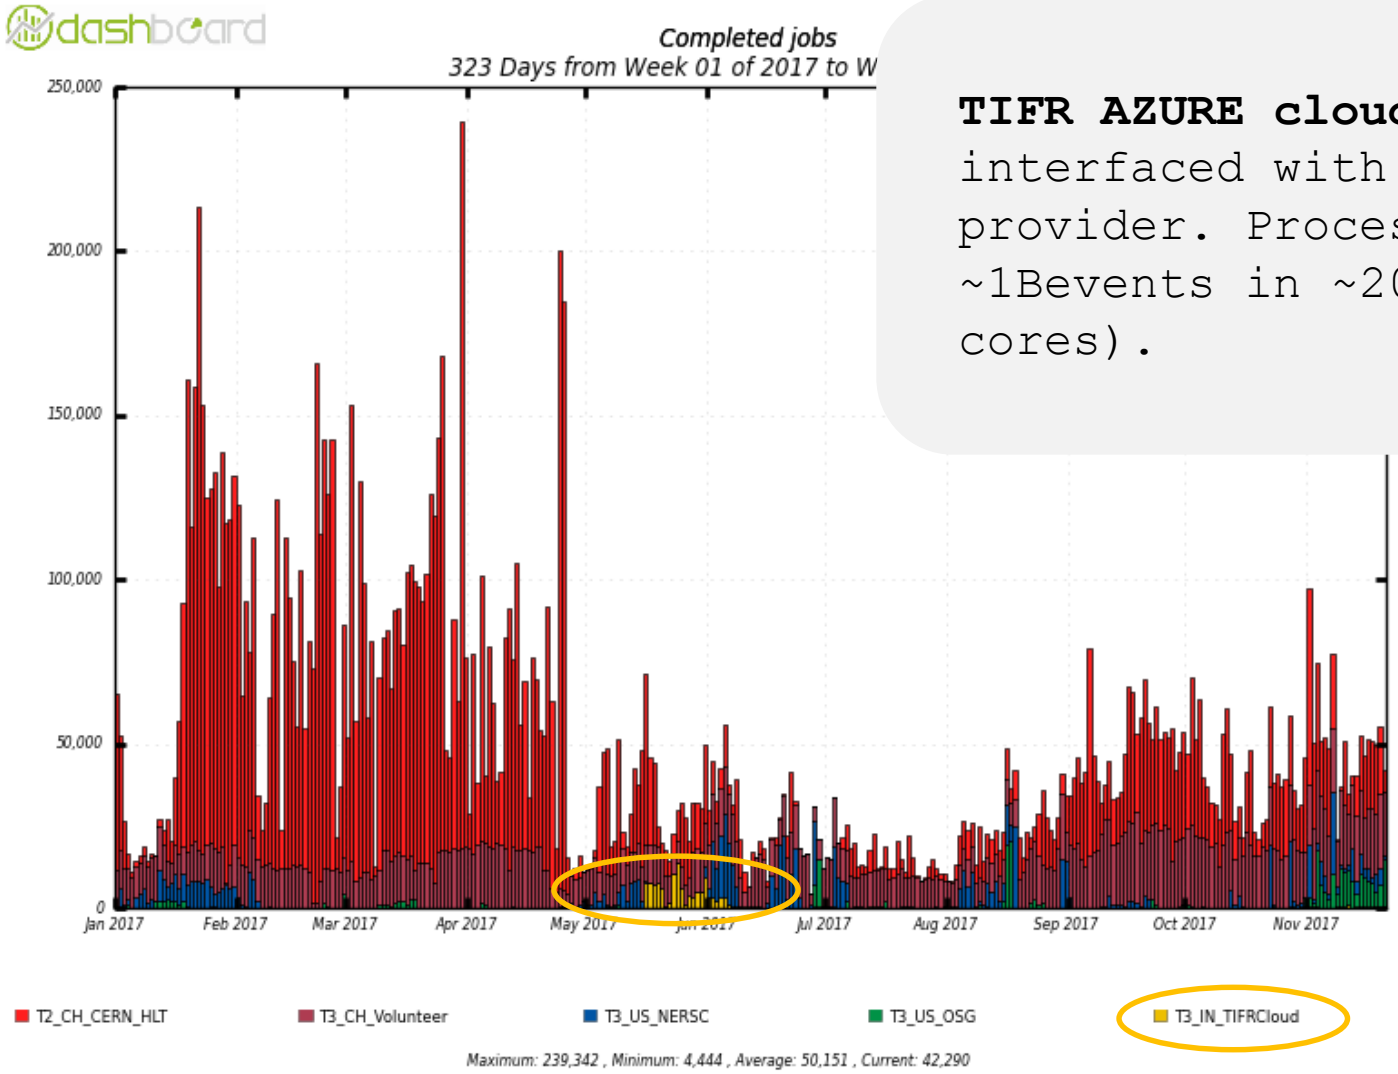

TIFR AZURE cloud: HTCondor interfaced with cloud provider. Processed ~1Bevents in ~20days (10k cores).

Soft evol.

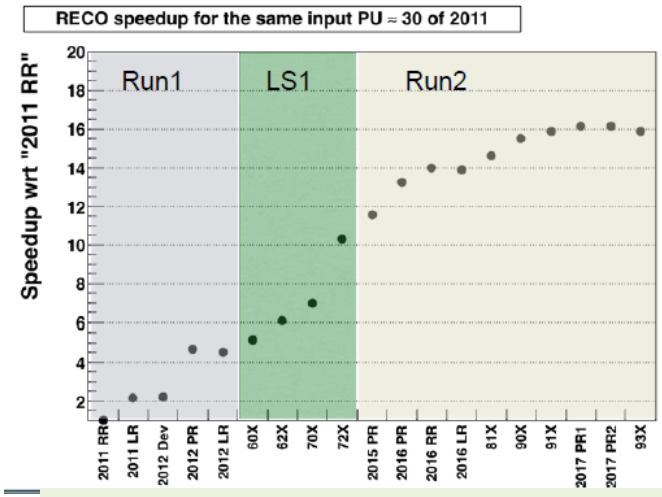

Same run/same HW: x16 in 6y

Need for evol: new detectors, new algos, res not scaling with DT (since 2015), new processors, new platforms…

- constantly modernizing our code;
- active R&D program and fully part of effort for HL-LHC.

- working closely with ROOT/IO team and GeantV;
- already MT soft;
- exploring new arch. GPU, Many-cores (e.g. KNL), FPGA, etc..
- exploring AI algorithms.

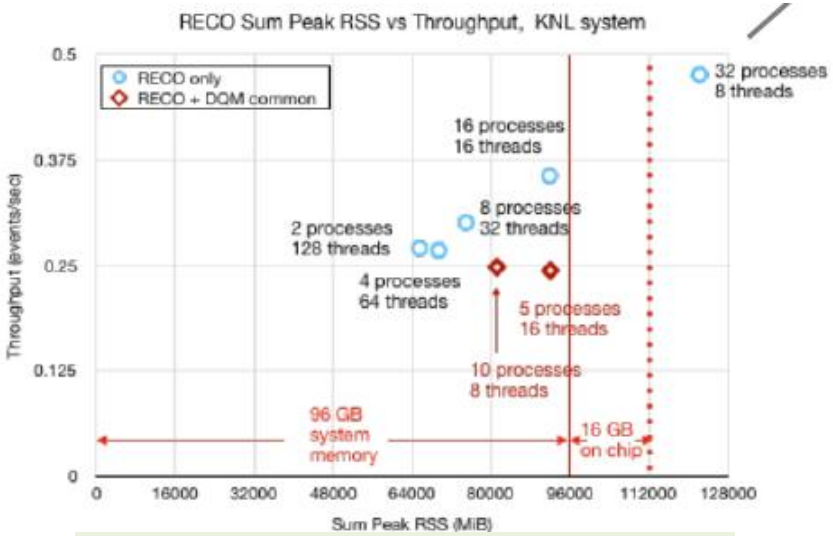

Tests on Cori2 KNL @ NERSC

Soft evol. examples.

Pixel tracking for the HLT using GPUs

https://indico.cern.ch/event/567550/contributions/2627138/

Paral. Kalman-Filter-Based Tracking on Many-cores

https://indico.cern.ch/event/567550/contributions/2629722/attachments/1510614/2355643/KalmanManyCore.pdf

Deep Neural Networks for HL-LHC Tracking

https://indico.cern.ch/event/567550/contributions/2629737/attachments/1511334/2357008/acat_heptrkX.pdf

HSF ACAT

https://indico.cern.ch/event/567550/contributions/2656675/attachments/1512724/2359578/HSF-CWP_ACAT2017.pdf

Public

CMS « internal »

Summing up.

- **2017 smooth** with some "features"
  - EOS auth. rate problems…;
  - **short in disk space**… complex data mgmt;
  - warning: **singularity by Feb'18**
- a lot of effort in studying **CPU eff. issues**
  - major **causes understood, actions taken**;
  - **studies** still **ongoing** (payload ineff.);
- **needs** for (rest-of) **R2** and **R3** defined.
  - and **outlook to HL-LHC** (50MHS06, 5EP disk…);
- **evolving** ops and sw to keep up with challenges
  - **evolution of DM** (phase out PhEDEx in R3 or R4);
  - usage of **dynamic**, opportunistic and non-grid **resources**;
  - **SW evolution**: multi-thread, many-core, deep learning, …

Questions?

Bkp: Res usage

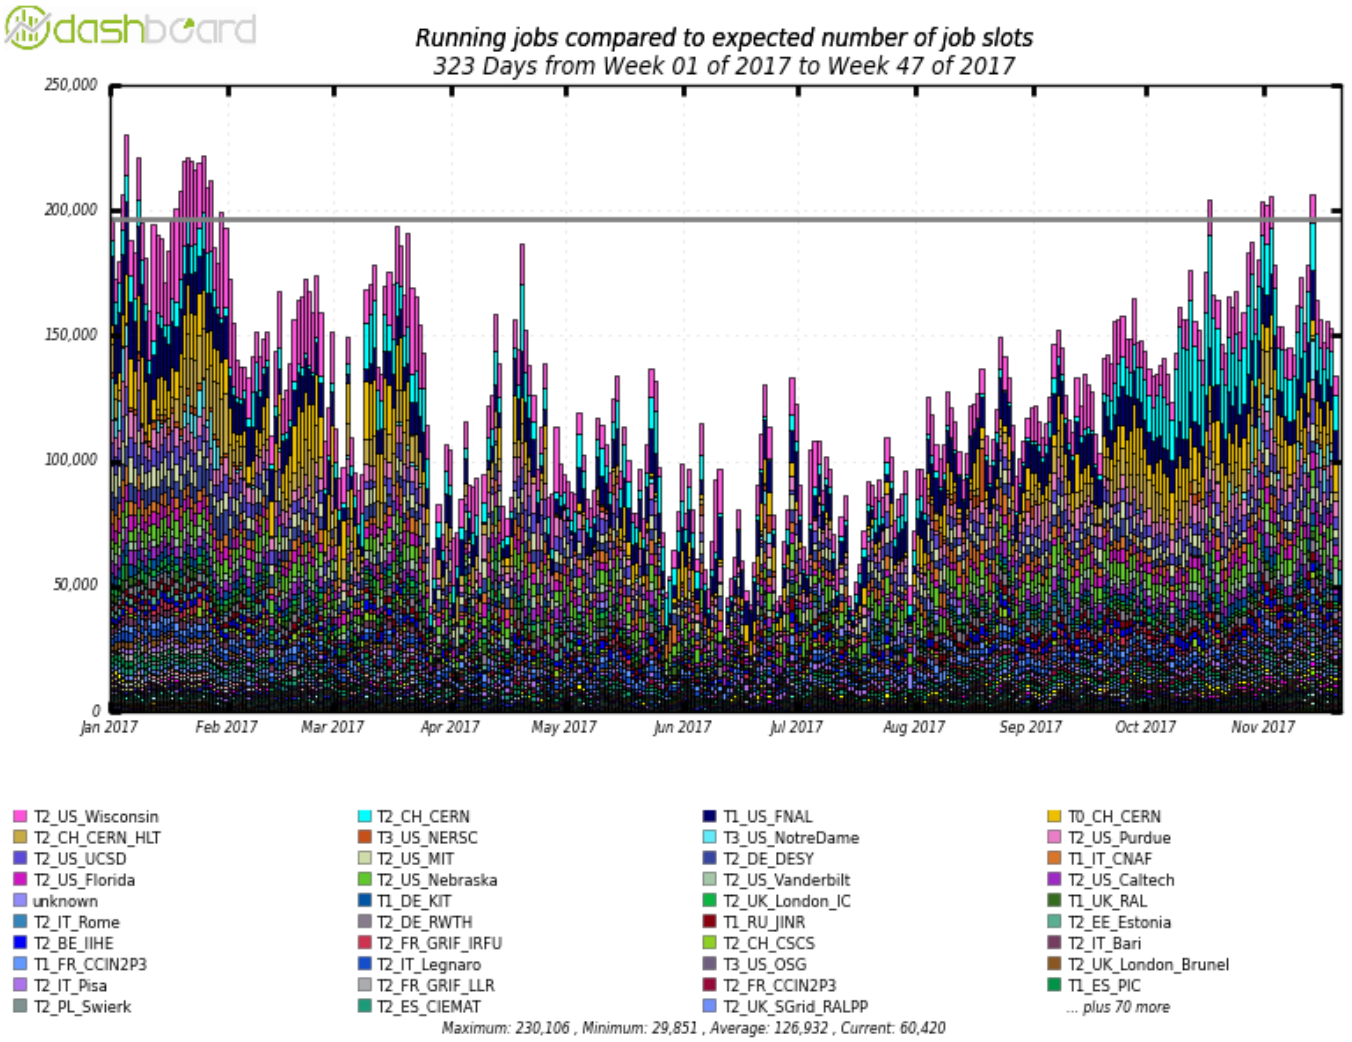

Bkp: DDM transfers.

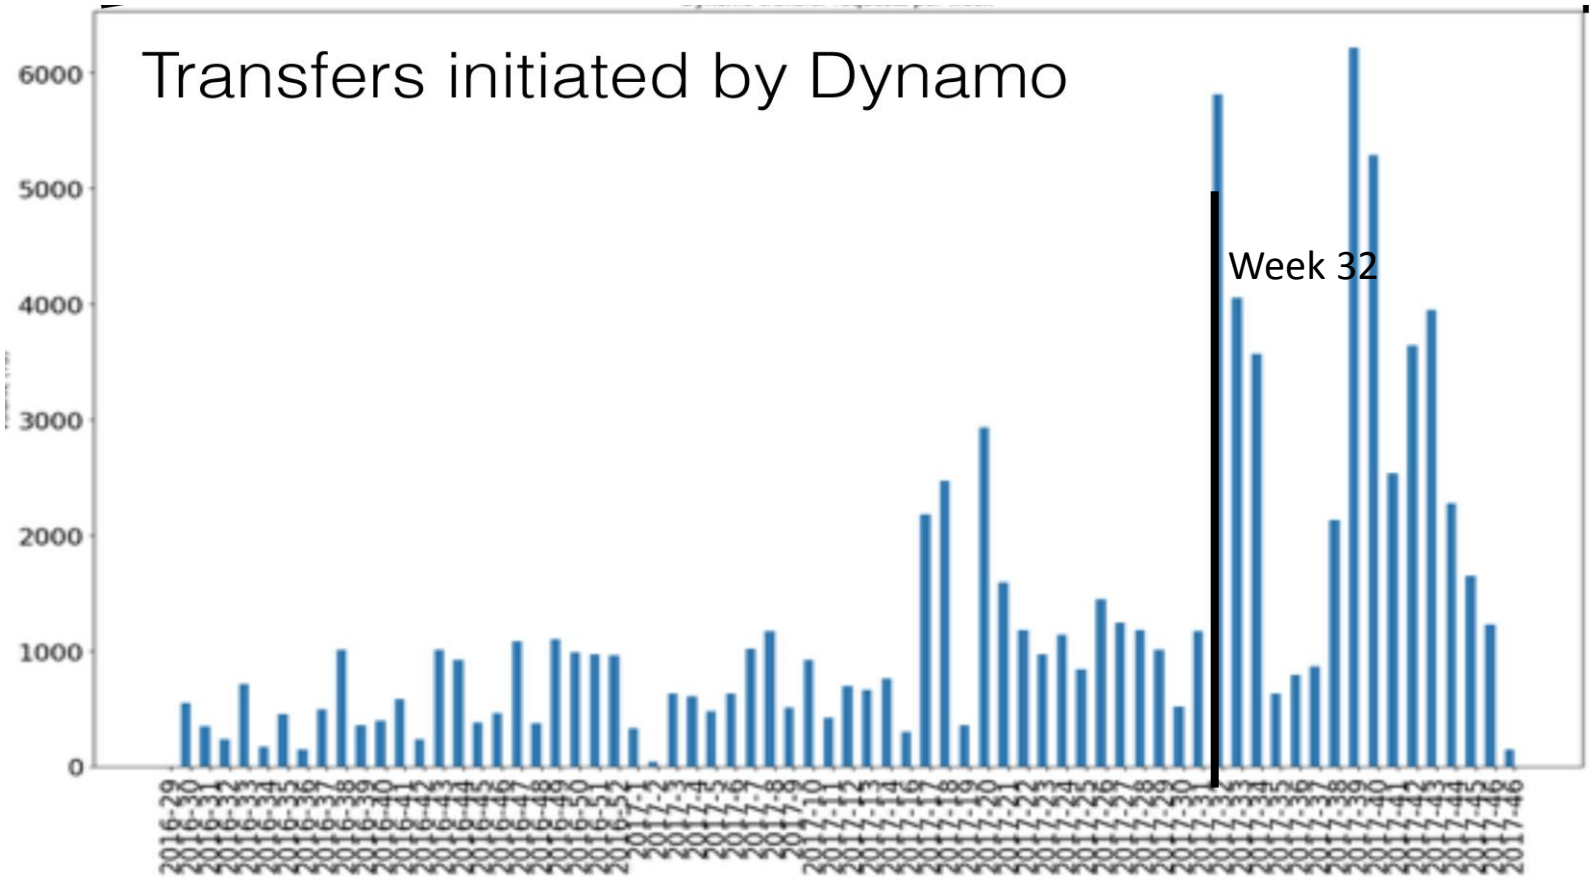

Bkp:Submission infra.

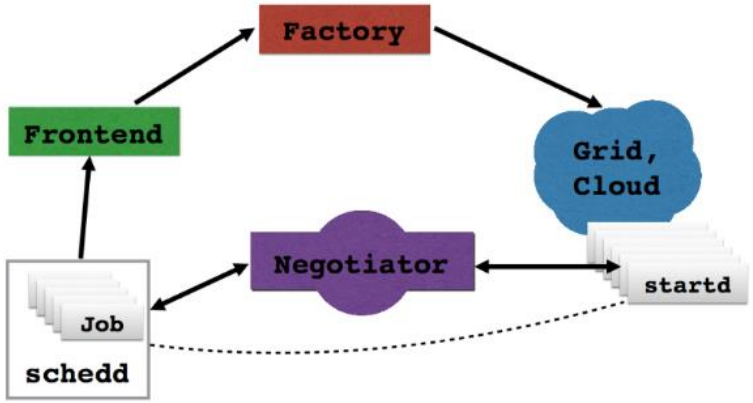

The policies of the Negotiator are configurable by CMS in the glideinWMS frontend. Currently we match jobs to resources based on a list of DESIRED_Sites and requirements for the number of CPUs (RequestCpus), memory (RequestMemory), and disk (RequestDisk). In the near future, we plan also to provision and request I/O capacity.

Bkp: Mem constrained.

Mem constrained slots: flexibility in the pilot model allows for running high mem wfs on standards resources at the cost of matching the remaining cores in the pilot only to low mem jobs… if any, otherwise idle. This is supposed to be more and more a marginal use case.

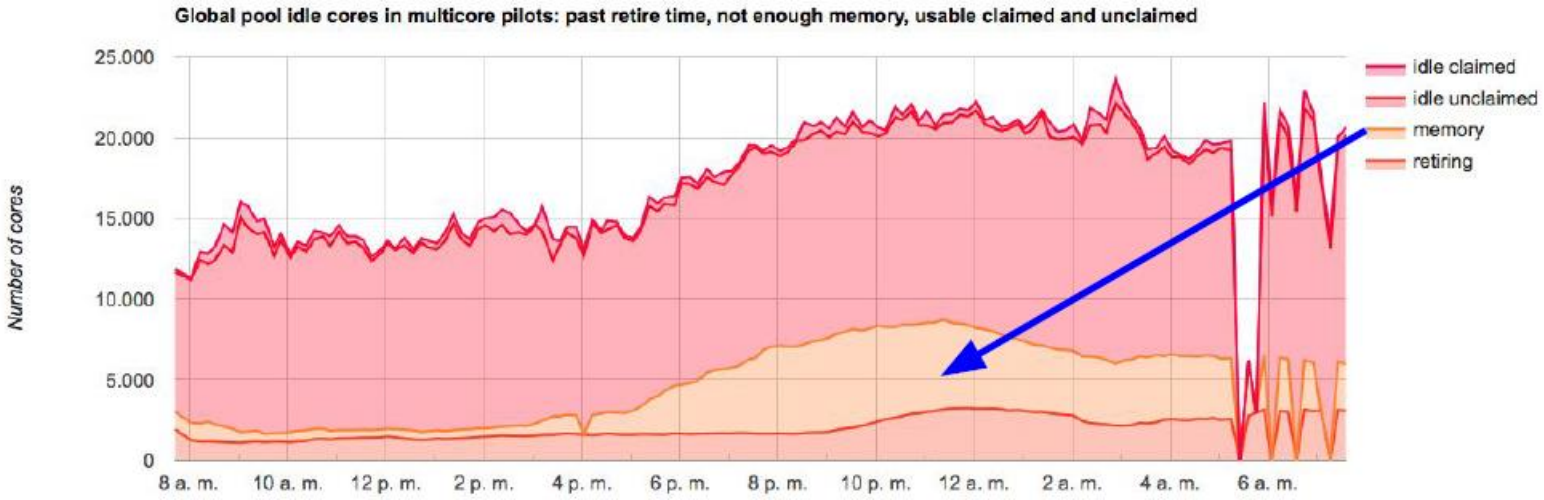

Bkp: Pilots retiring.

Pilots start draining to prevent payloads run over the max pilot lifetime

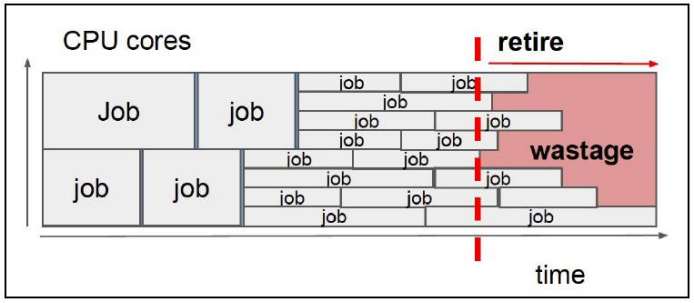

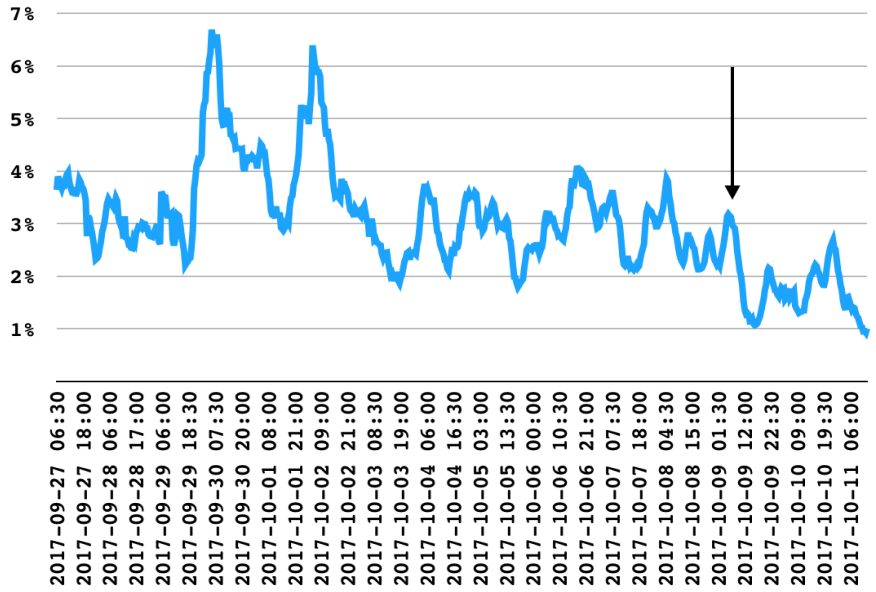

- this may lead to considerable resource wastage as the pilot runs half-drained
  - thanks to better understanding of jobs runtimes the retire time was reduced from 10h to 4h on Oct 9th;
  - immediate detectable effect (however level was already low as we ran 8 cores wfs).

Bkp: Payload ineff.

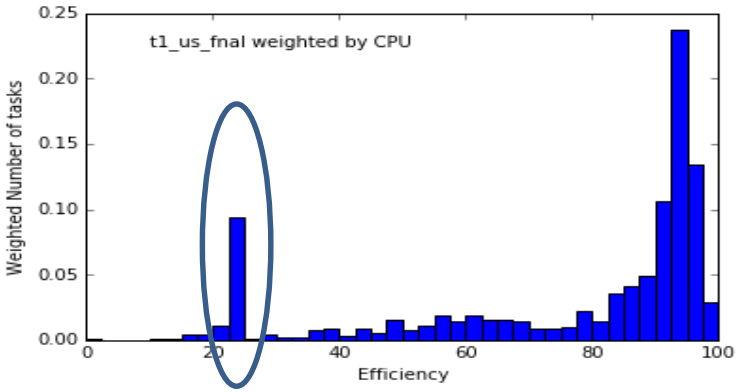

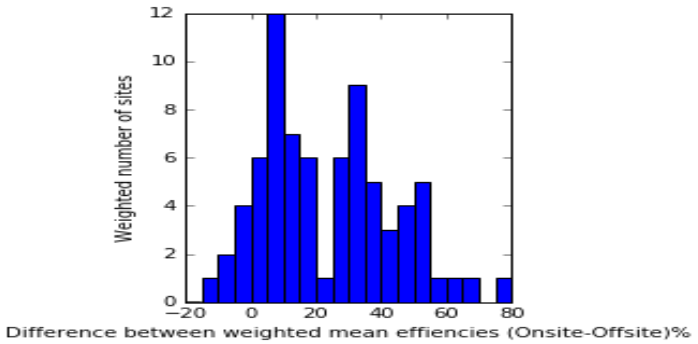

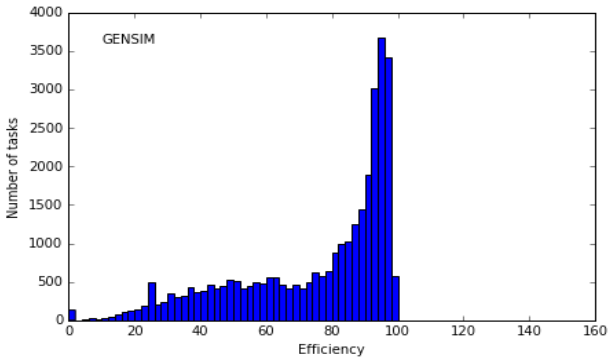

miss-configured jobs: e.g. single FORTRAN module running serially. Implementing a check procedure upon task submission

onsite vs offsite data: plot tells us there is a huge variation. Needs further investigation (not only small sites). However, offsite data is essential to CMS computing.

even then: jobs with little input data, no config pbs, and low eff. We can check their logs with ES: e.g. execute external generators not suitable for MT conf.

Bkp: LHC plan

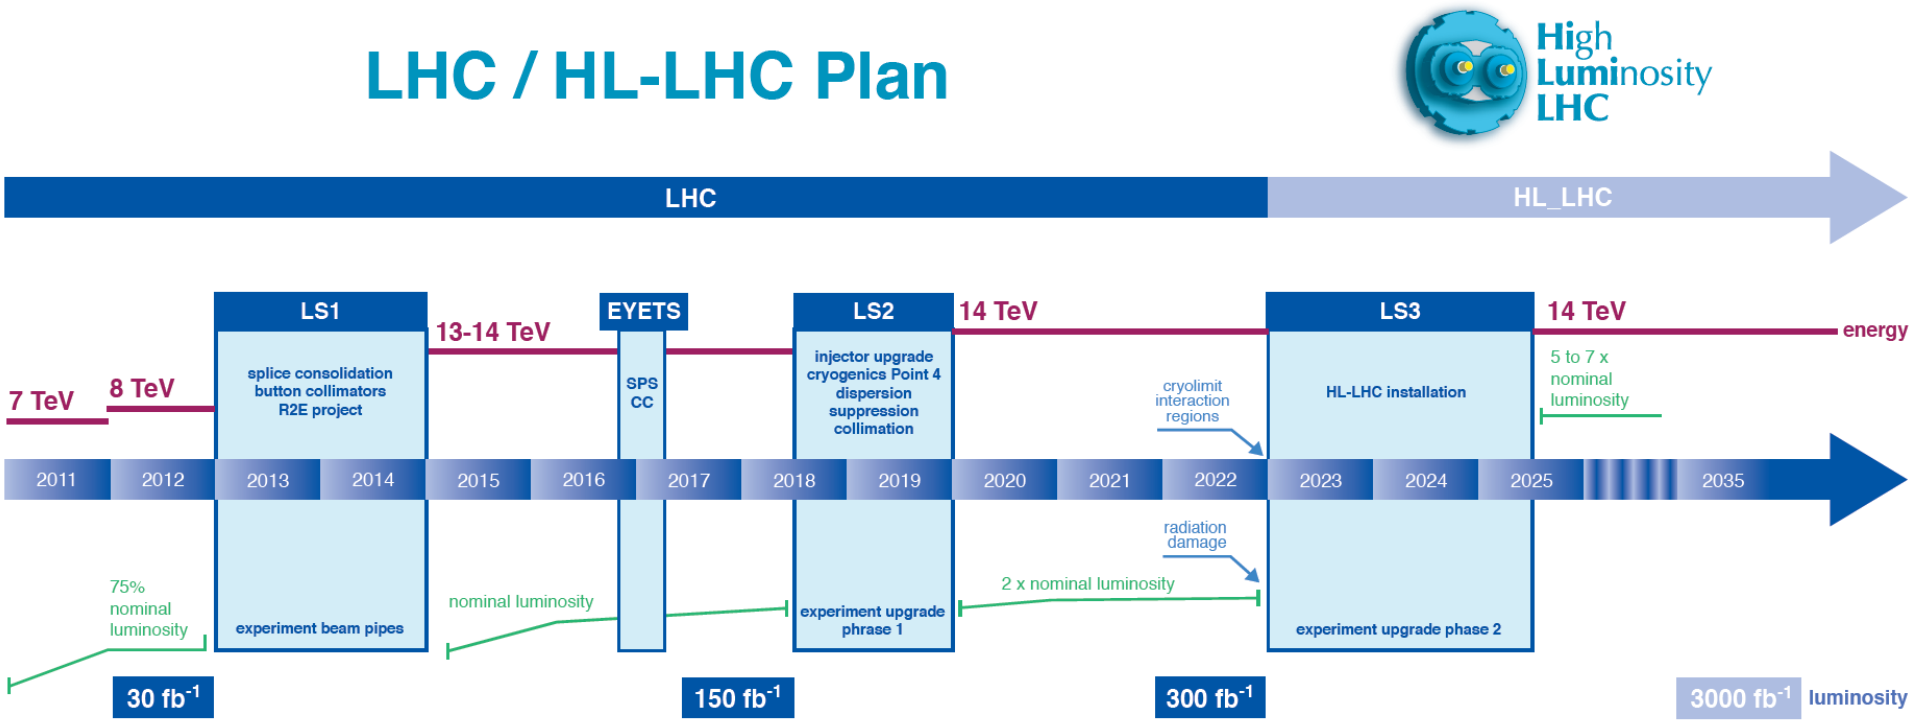
LHC / HL-LHC Plan
High Luminosity LHC
LHC
HL_LHC
LS1
splice consolidation
button collimators
R2E project
EYETS
SPS CC
LS2
injector upgrade
cryogenics Point 4
dispersion suppression
collimation
LS3
HL-LHC installation
7 TeV
8 TeV
13-14 TeV
14 TeV
14 TeV
energy
5 to 7 x nominal luminosity
cryolimit interaction regions
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2035
radiation damage
75% nominal luminosity
nominal luminosity
2 x nominal luminosity
experiment beam pipes
experiment upgrade phrase 1
experiment upgrade phase 2
30 fb-1
150 fb-1
300 fb-1
3000 fb-1
luminosity

Bkp: Resources.

+20%/y from Moore law.

Disk 2018: 156 PB
CPU 2018: 1923 kHS06
Disk 2027 proj.: 5 %B (x32 wrt 2018)
CPU 2027 proj: 9665 kHS06 (x26 wrt 2018)

Tier	Pledge Type	2017	2018
Tier 0	CPU (HEP-SPEC06)	423,000	423,000
Tier 0	Disk (Tbytes)	24,600	26,000
Tier 0	Tape (Tbytes)	70,500	97,000
Tier 1	CPU (HEP-SPEC06)	600,000	600,000
Tier 1	Disk (Tbytes)	57,000	60,000
Tier 1	Tape (Tbytes)	175,000	188,000
Tier 2	CPU (HEP-SPEC06)	850,000	900,000
Tier 2	Disk (Tbytes)	68,000	70,000

Bkp: PhEDEx

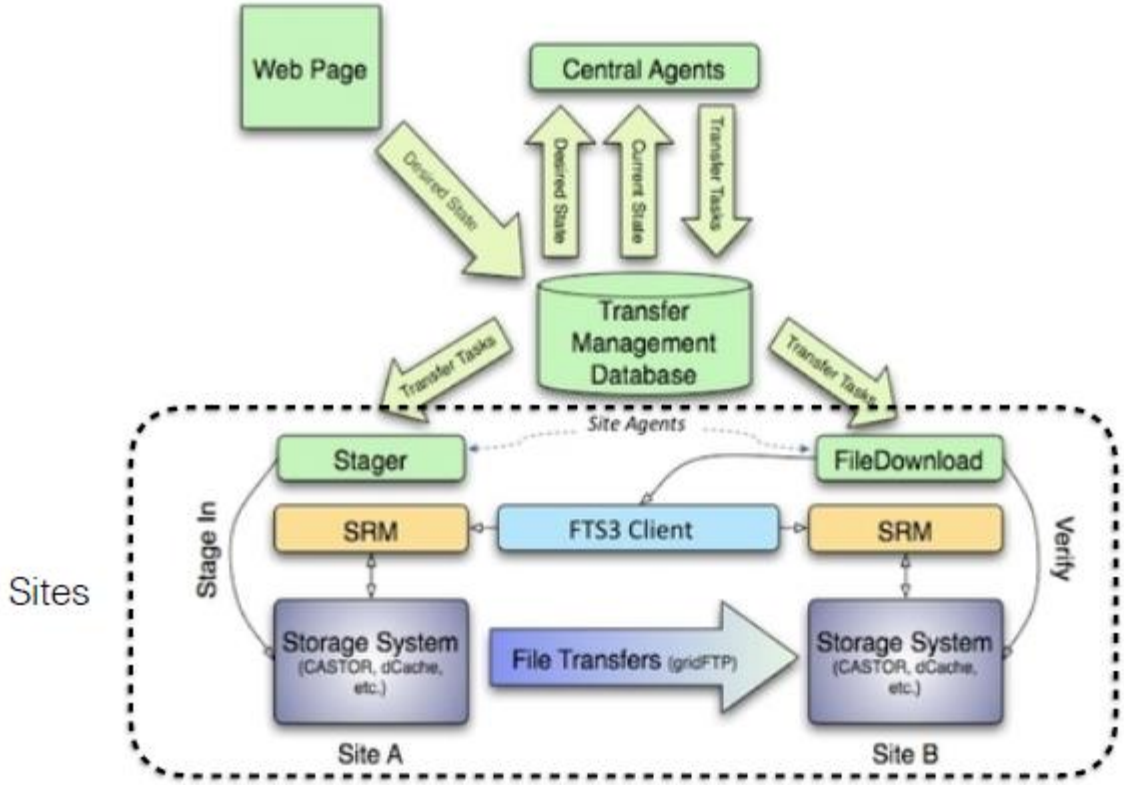

× 3 instances (Prod, Debug, Dev)

Bkp: Data Lakes.

Also discussing **more visionary** solutions: e.g. data lakes

- single (or very **few**) **logical SE** with a significant amount of high-perf storage
  - **DM simplified** with only 3-5 sites;
  - tradeoff DM/storage provider (assuming storage providers are better in this);
- **sites outside** have **no storage**: cached or streamed
  - small/medium site do not need storage, only NW and cache;
  - which ratio? Which latency?
- **like AAA** but need **NS** and **data integrity** concepts.