CMS Computing evolution(s)

Daniele Bonacorsi
[University of Bologna, Italy - CMS Software/Computing coordination]
on behalf of CMS Sw/Comp

LHC machine is evolving

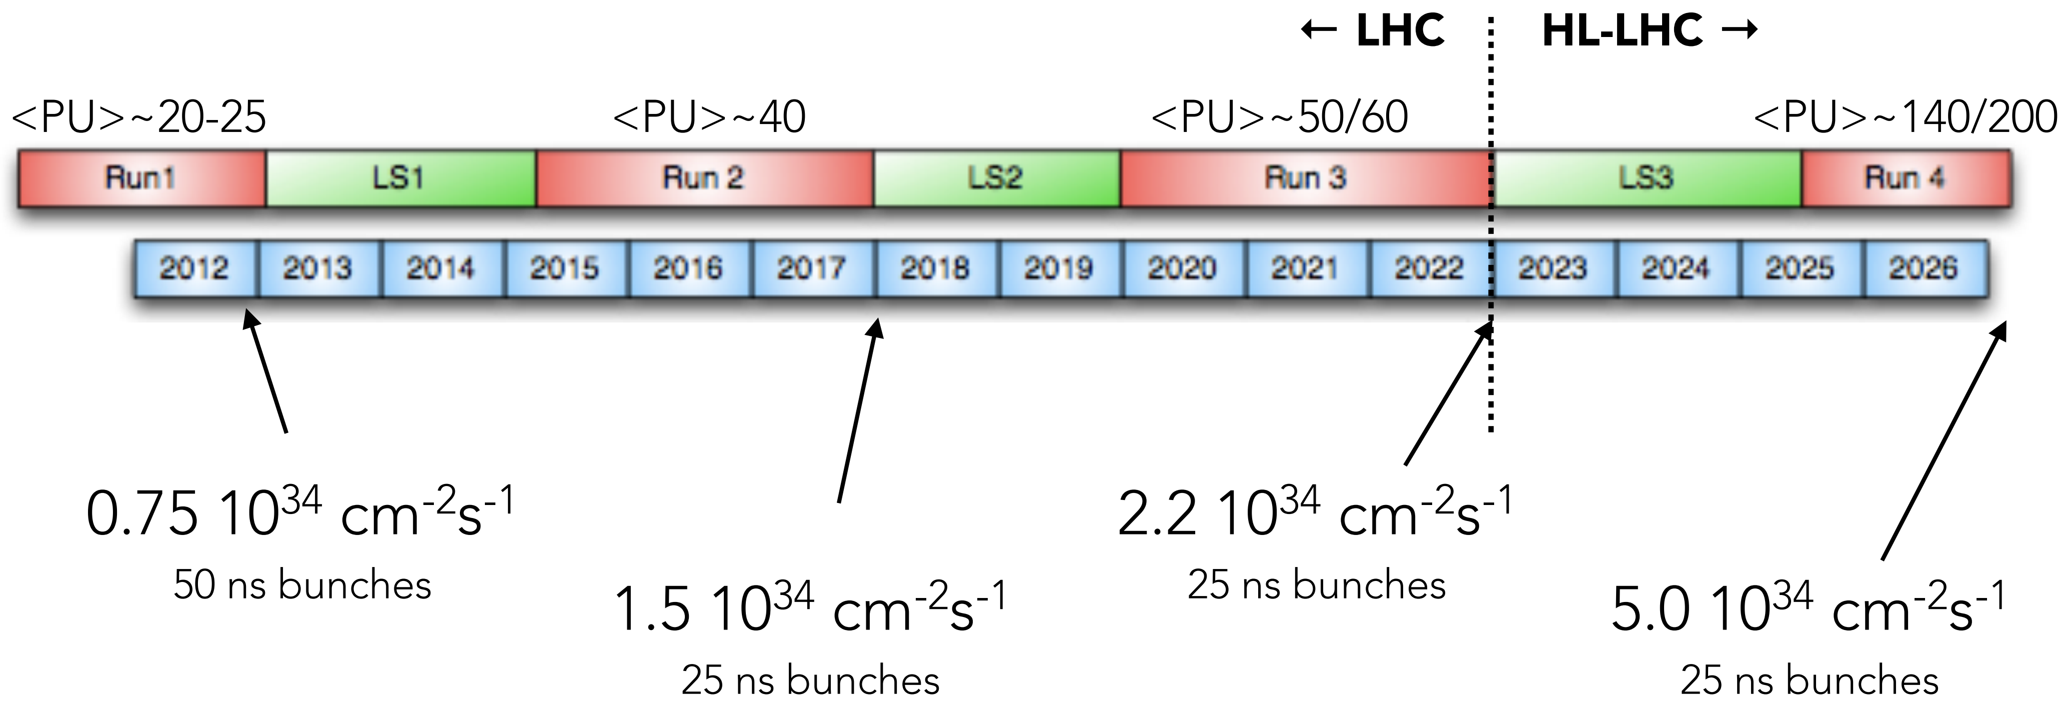

LHC machine lumi is steadily increasing

- between LHC Runs, as well as within Runs
- Far horizon is a ultimate goal of **3k fb^{-1}** in 10 yrs of HL-LHC

The most relevant metric for event complexity for reconstruction software is **<PU>**

Detectors are evolving

ALICE

- major upgrades to inner tracking and TPC
- Pb-Pb readout at 50kHz trigger-less, online-DAQ integration, data compression

ATLAS

- major upgrades to inner tracker, calorimeter readout improvements
- track trigger operations at Level-0 rates

CMS

- **new silicon tracker with trigger (@ interaction rate) capabilities**
- **new calo endcap, more read-out improvements**

LHCb

- fully sw-based trigger @ interaction rate
- upgrades in readouts and front-end electronics

LS2 in view of Run-3

LS3 in view of Run-4

Computing models are evolving

CMS Computing model under continuous evolution

- since Computing TDR in 2005
- adiabatic (or not) evolutions to the model in Run vs LS periods

LHC **Run-1**

- already less hierarchical than MONARC
  - workflows generally performed where they were expected as from the model
  - but e.g. full transfer mesh being commissioned while taking data in Run-I

LHC **LS1**

- explore flexibility, optimisations and efficiency gains
  - workflows executions: ran where expected (Run-1), then started not to (≥LS1), including also PR; data federation based on xrootd; more dynamism and automation in e.g. data placement and replicas removal to make more efficient use of the storage

LHC **Run-2**

- in principle, a tough game in 2015 already
  - HLT raising to ~1 kHz in 2015, i.e. a factor x2.5 wrt Run-I; expected increase in the <PU>, implied a factor 2.5 increase in reco time; already a factor ~6 increase to maintain current activities
- Tactics: formerly independent sites started to evolve towards a coherent worldwide pool of resources which brings more power to compute faster and more efficiently

Blurring the Tiers boundaries

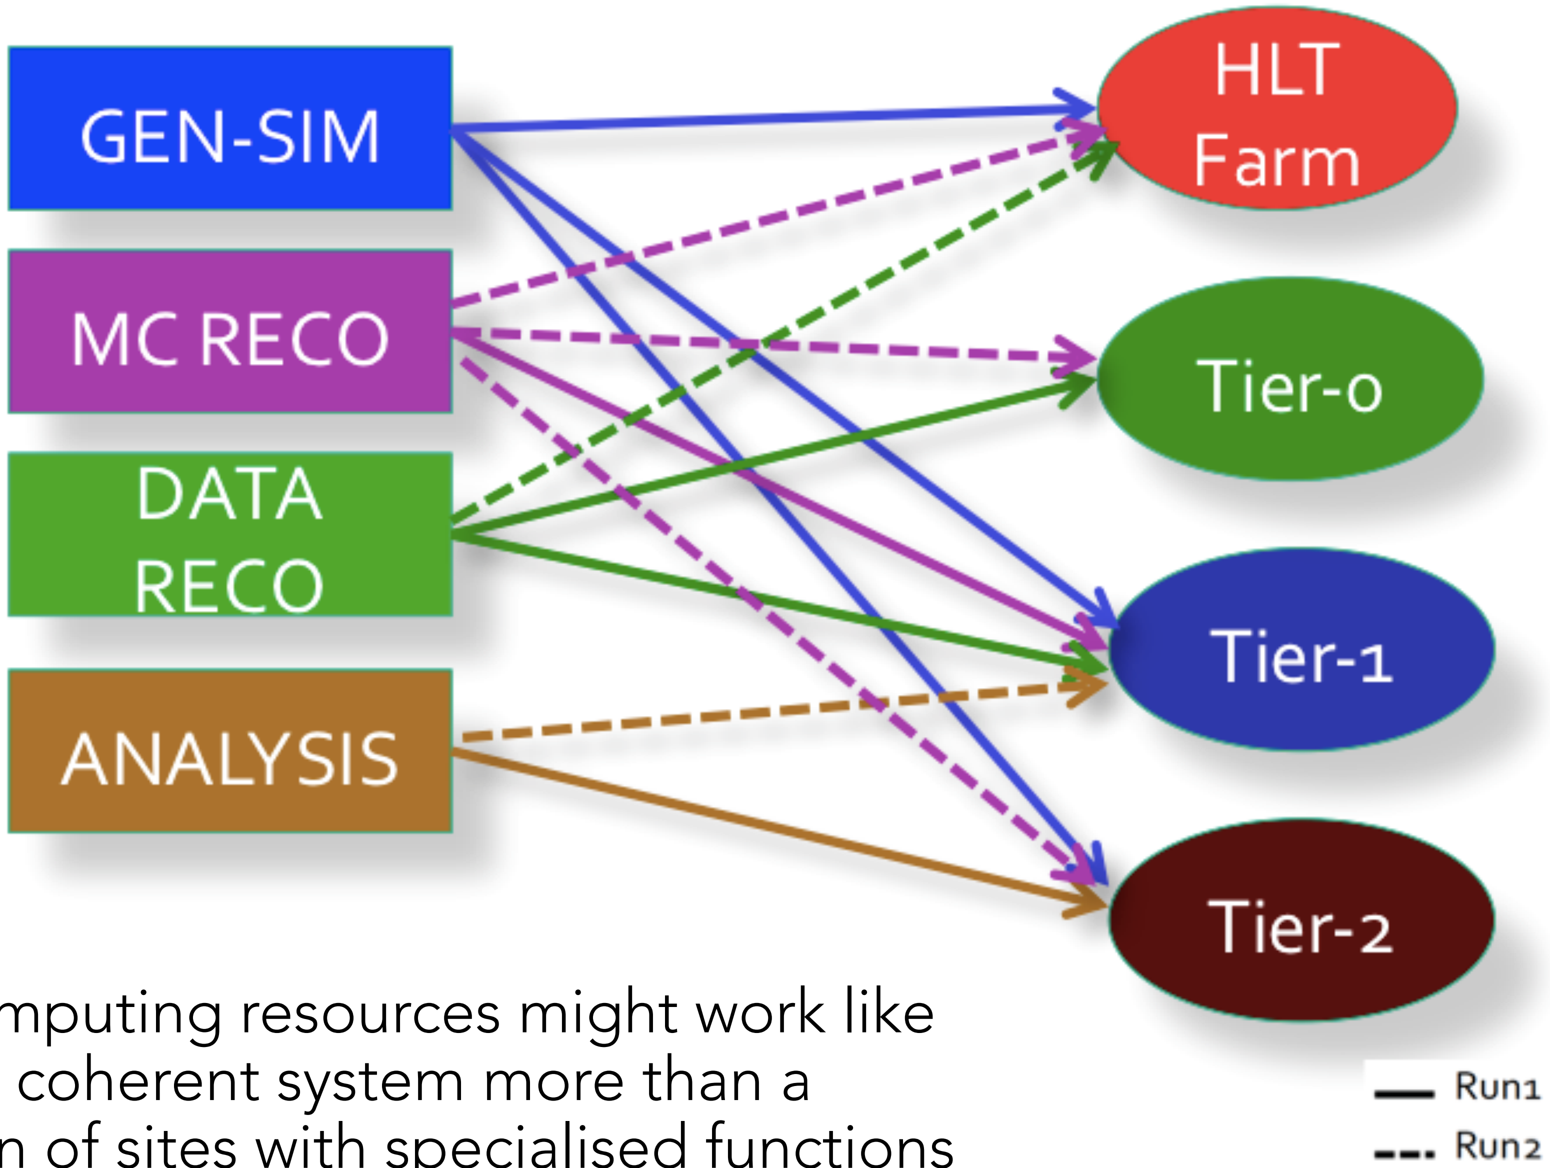

CMS Computing resources might work like a unique coherent system more than a collection of sites with specialised functions

Changes in LS1 towards Run-2

(not a complete list)

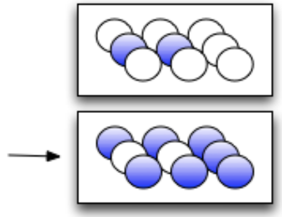

CMS-wide xrootd-based **data federation**

- ✦ rethink about the need for data locality
- ✦ allow CMS applications to read data over the WAN from remote storage, so to exploit processing capacities at multiple sites to execute a workflow which utilises the storage at one location
  - AAA project: commissioned and stress-tested in 2013/2014, now in production at full steam
- ✦ Target: add transparency in data access, add flexibility to wf assignment to Tiers

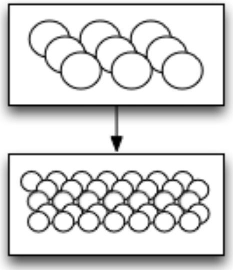

MiniAOD as a **new** (thin) **data tier**

- ✦ ~30/50 kB/evt (~10% size of the Run-1 AOD)
- ✦ Can serve roughly 75% of the analysis needs
  - Validated and included in production in 2014, now stably produced and accessed
- ✦ Target: lower the "cost" of hosting meaningful samples on CMS storage

Changes in LS1 towards Run-2 (cont'd)

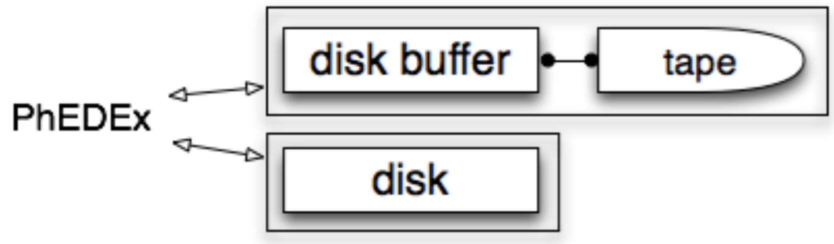

Disk/Tape separation at T1s

- Implementation of a logical separation T1 disk vs T1 tape in the CMS DM
  - It follows a WLCG Storage and Data Management TEG recommendations
- Data can be PhEDEx-subscribed to either as an independent resource
  - a T1D0 system with a relatively small disk buffer
  - a bigger T0D1 system that serves data to running jobs
- Target: less functional differences between T1s and T2s; allow MSS protections; T1s can be used for analysis (and remote access)

DM → **D**DM: complementing CMS DM with **Dynamic** features

- Popularity-based automated system for dataset replication and removal
- Target: improve the storage utilisation efficiency

Changes in LS1 towards Run-2 (cont'd)

New CRAB major version: **CRAB 3**

- Latest version of the distributed analysis toolkit, commissioned in 2014
- Good user experience,
  reached 20k simultaneous jobs, 200k daily subs
- Target: ease the analysis workflows and user experience, improve the analysis throughput at any Tier level open to analysis

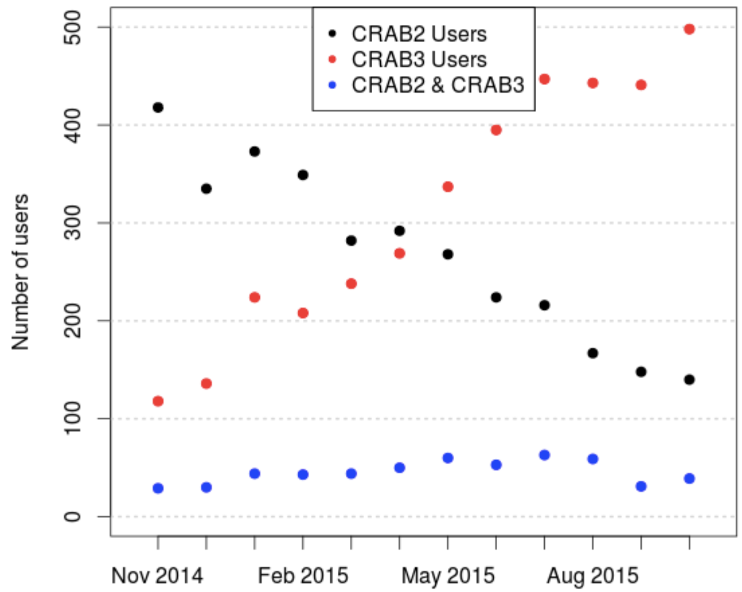

Global HTCondor pool

- a unique pool for production and analysis, with priority management inside the pool itself
- Target: equip yourself with a wheel to globally steer the resources on most critical tasks upon need

HLT integration in CMS processing

- OpenStack-based cloud infrastructure over HLT hardware
- Launched in early 2013, commissioning in 2014, in production since 2015
- Target: exploit HLT offline usage, use it flexibly during CMS idle times (technical stops, inter-run periods, also interfills), relieve the T1 work load

Changes in LS1 towards Run-2 (cont'd)

Ability to provision **cloud(s)**

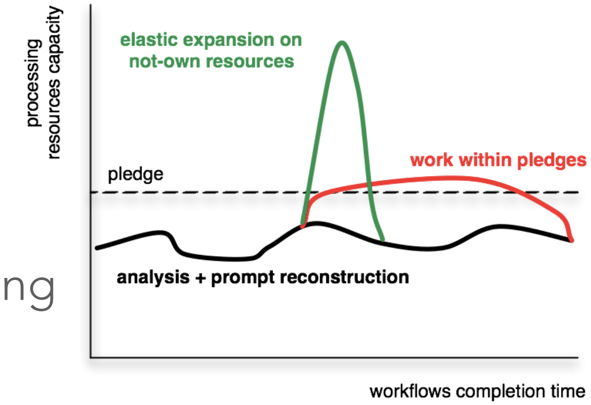

- Run-1: we enjoyed a "WLCG consistency" in accessing resources
- from Run-2 on: we need a service architecture with a provisioning layer on top of a diverse set of grid(s), cloud(s) interfaces, opportunistic resources, specialised resources, local traditional farms, whatever-will-be-next.. which would make this profitable in a coherent manner
- work done on HLT directly applicable to resources provisioned with OpenStack
  - CERN Agile Infrastructure resources, academic clouds, potentially even commercial clouds (once cost model is demonstrated to be competitive)..
- the pilot infrastructure through the glideIn-WMS offers a clean interface to the CMS production service
- Target: commission CMS ability to use cloud-provisioned resources to increase the simulation capacity for a defined period of time, and smoothly absorb peaks upon need

Volunteer computing

- HEP simulations on home computers was an untapped opportunistic resource for CMS
- Now: CMS@home. Probably useful only for specific workflows, but still..
- Target: run it, exploit the potential of being large and free resources potentially, plus the outreach benefits

Changes in LS1 towards Run-2 (cont'd)

More (operational) choices to gain more improvement factors:

e.g. make **RECO** a **transient data format**

- from 2015 on, considered not cost-effective to keep systematically
- proposed to be transient and stored on disk temporarily - or not written at all
- scaling down to few months only of RECO store, from ~now to 2017

Reduction of data replicas on T2 disks

- replication factor was 3 in Run-I
- scaling down to reach ~1.7 by the end of 2017

And so on..

Run-1/2 show-off

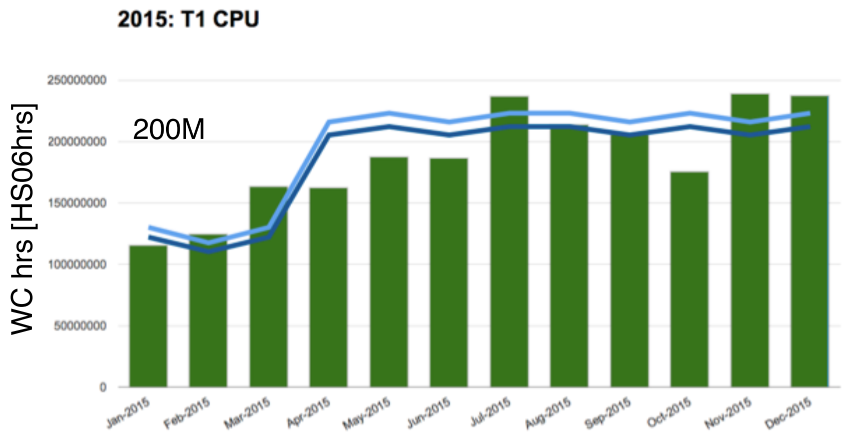

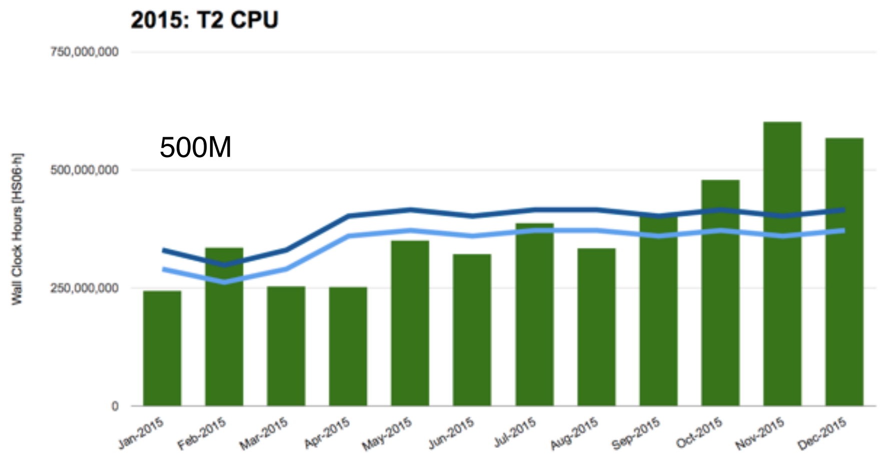

average close to **~100%** usage wrt the pledges

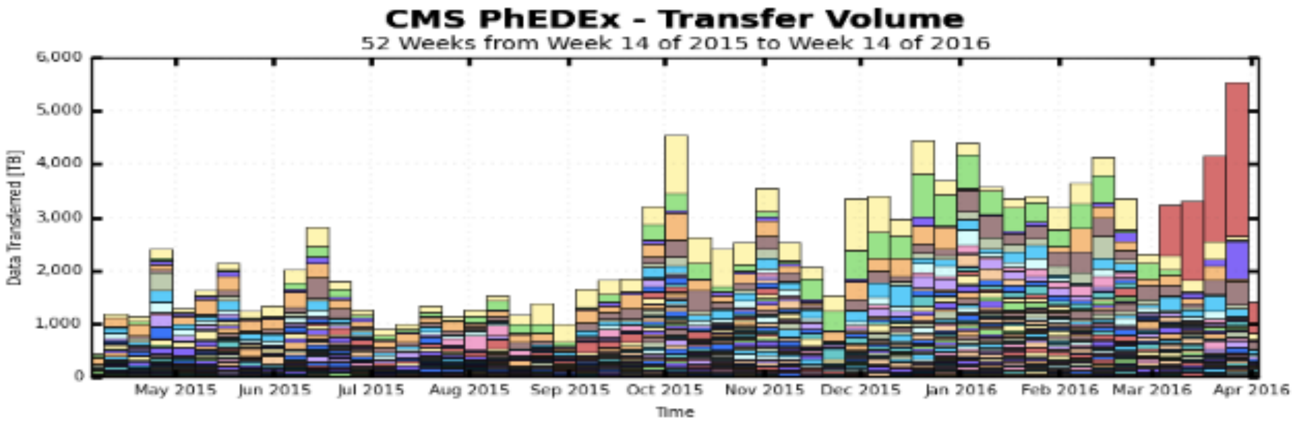

4 PB/week (2M files/week)
Increased 4x since end of 2015

Plus: ~15-20% use WAN access.
Global traffic ~1/4 PhEDEx over most time periods

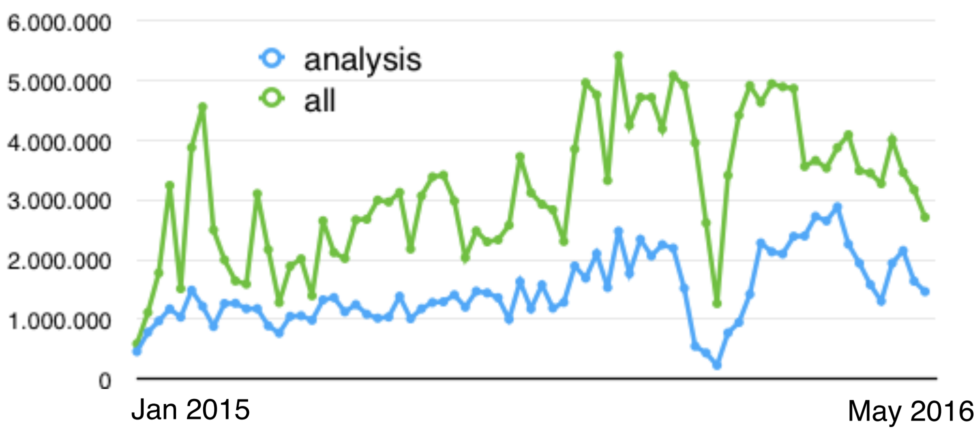

typical seasonal trends, plus an evident increase in average # jobs getting into 2016

The Computing challenge

Improvements of existing features and optimisations, addition of new features, changes in the model..

All this allowed to face **Run-2**, successfully (so far), within a ~flat budget.

Where does the challenge become really tough?

Computing Challenge = Event Complexity x Rate

What's ahead of us

Run-1/2 Run-3 … Run-4

CPU

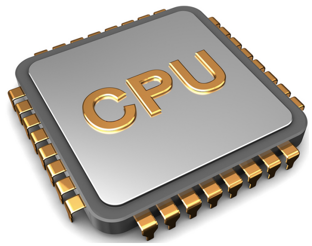

Event reconstruction

Resource usage in event reconstruction driven by:

- output rate from HLT
- event complexity (i.e. <PU>)
- choices in the acceptable level of physics algorithm performance (e.g. p_T cut)
- detector characteristics (e.g. channel count, complexity of geometries)

Assumptions (first attempt), based on extrapolation of the present:

- continue to fully reconstruct all HLT-saved events into a std event data format
- algorithm performance for physics objects remains unchanged wrt Run-2
- include all RECO application developments for Run-2 up to today
- use <PU>={25, 50, 140/200} as {Run-2, Phase-I, PhaseII} conditions
- Phase-II detector contains only tracker and muon upgrades

reco time and AOD size: resource requirements per data event (wrt Run-2)

	<PU>	reco time (ratio wrt Run-II)	AOD size (ratio wrt Run-II)
Phase-I	50	4	1.4
Phase -II	140	20	3.7
Phase -II	200	45	5.4

AOD size is somehow less a concern that reco time:

- work on algorithms, integration of C++11 standards, advanced compilation options, …
- a further **x2** improvements is expected

Event simulation

Detector simulation: resource requirements do not change with detector conditions

- CMS chose to simulate PU evts in the digi algorithms by overlaying evts
- As for the reco algorithms, work in progress to reduce the resource needs /evt
  - technical improvements (e.g. novel sampling techniques for shower development simulations)
  - gain **x2** in CPU needs already in Run-II simulation application. More work in progress with Geant4

Digitization (of Geant4 output to create the simulated raw data)

- due to PU simulations, it scales linearly with the <# PU evts per BX> * (# BX that must be simulated)
- on the former: CMS software already been redesigned to be able to scale to as many PU evts as necessary within an essentially fixed memory budget
- on the latter: assumption that we continue to simulate 16 BX (before/after the one of the simulated hard interaction, to fully capture the calorimeter electronics response)
- CPU time within the digitization dominated by the detector-specific digi algorithms and thus can depend significantly on the final Phase-II configuration of the detector

Detector sim and digi: CPU requirements per simulated event (wrt Run-II)

	<PU>	detector simul (ratio wrt Run-II)	digitization (ratio wrt Run-II)
Phase -I	50	1	4
Phase -II	140	1	9
Phase -II	200	1	13

Reco and sim resource needs

So, if we assume:

- no radical changes in the model (just extend what we do today)
- the accelerator operation time is the same as in Run-2
- the size of the simulated samples continues to increase according to the HLT output rate

.. we get:

Estimated CPU needs, including HLT rate and processing time (wrt Run-II)

		Data	Simulation			tot
	HLT rate (kHz)	reco	det. sim	digi	reco	
Phase -I	**1**	4	1	4	4	~3
Phase -II	**5**	100	5	45	100	~65
Phase -II	**7.5**	340	7.5	100	340	~200

Last column gives you the sense of the challenge:

- roughly computed as weighted averages of the processing steps based on how frequently each step is run

CPU evolution(s)

Flat spending profile from early Run-2 to early Phase-II gives you a **x8** in processing capacity

- effective growth in CPU power for fixed cost is ~25% /yr [*], i.e. x2 every 3 yrs
- ~3 doubling cycles between early Run-2 and early Phase-II

Assuming a **x8** growth in processing capacity and a **x2** improvement in pure technical changes to algorithm performance, this leaves **a deficit of a factor 4 (13)** to be gained from other progresses in computing capabilities for operation at PU 140 (200) respectively

Note: CPU requirements for analysis (user) jobs not included

- not standard application configurations, their scaling into Phase-II is more difficult to estimate

Historical increase in LHC CPU capacity has been ~360 kHS06 /yr [*]

- **"OK PANIC!": not even close to what we will need..**

[*] WLCG evolution doc

Disk

CMS

Disk evolution(s)

Disk space needs in CMS scale with HLT output rate and event complexity

- Phase-II at PU 140(200): data volume will increase by x5(x7.5) in # evts, with an increase in AOD evt size of x4(x6)
- AOD total size increase on disk is roughly a factor 20 (40) at PU 140 (200)

	HLT rate (kHz)	AOD size evt (ratio wrt Run-II)	tot
Phase -I	**1**	~1.4	1.4
Phase -II	**5**	~3.7	~19
Phase -II	**7.5**	~5.4	~41

Flat spending profile from early Run-2 to early Phase-II gives you approximately **x4** in disk capacity

- effective growth in disk capacity for fixed cost is ~20% /yr [*], i.e. x2 only every 3.8 yrs
- only ~2 doubling cycles between early Run-2 and early Phase-II

This leaves **a deficit of a factor O(10) at least** to be gained from other progresses for operation in Phase-II

- **"OK PANIC!": not even close to what we will need..**

This indicates that CMS will need to develop techniques that need smaller data formats (or more innovative approaches) to accomplish the same processing and analysis goals

[*] WLCG evolution doc

Expected time for the (storage) panic?

Timing of "the big computing challenge" varies per experiment:

- ALICE/LHCb: Run-3; ATLAS/**CMS**: Run-4

Very rough estimate of collected RAW data per year of running, grouped by Run, done using a simple extrapolation of current data volume scaled by the HLT output rates foreseen in each LHC Run

Note: this is only RAW.. no derived data (ESD, AOD), no simulation, no user data…

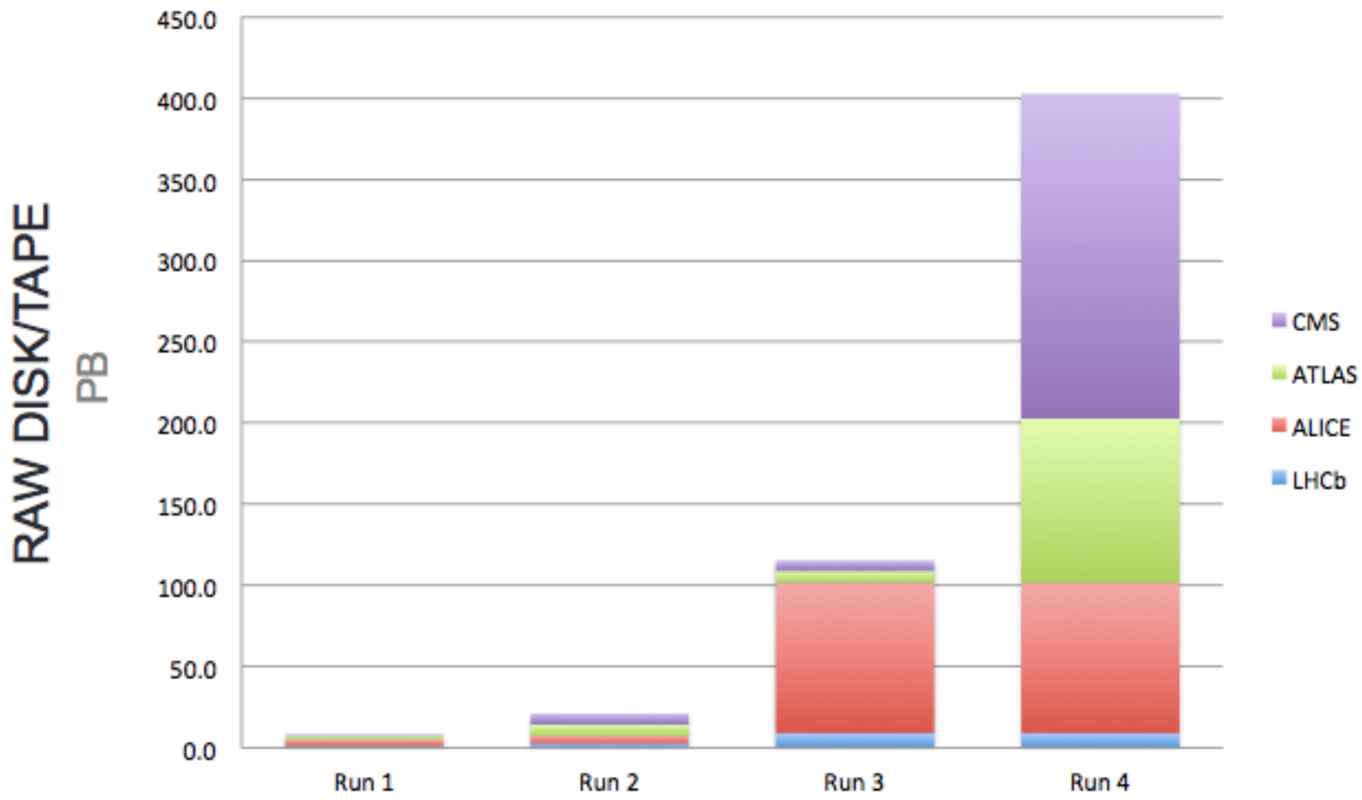

Next question is: "<u>will the scaling law(s) save me?</u>"

Technology evolution(s)

We are too small to use anything that is _not_ off-the-shelf hardware

We have enjoyed huge and steady advances (so far). Will it continue?

All scaling laws are empirical, and we see slow-downs..

- many technology we use are mature enough to exclude rapid growths..

E.g. Moore's law is not dead

- but more uncertainties in density doubling times
- which make a huge difference in a 10 yrs time frame

Clock speed stalled ~ 2005

More recent estimates are the ones used in the previous slides:

- CPU +25%/y
- Storage +20%/y

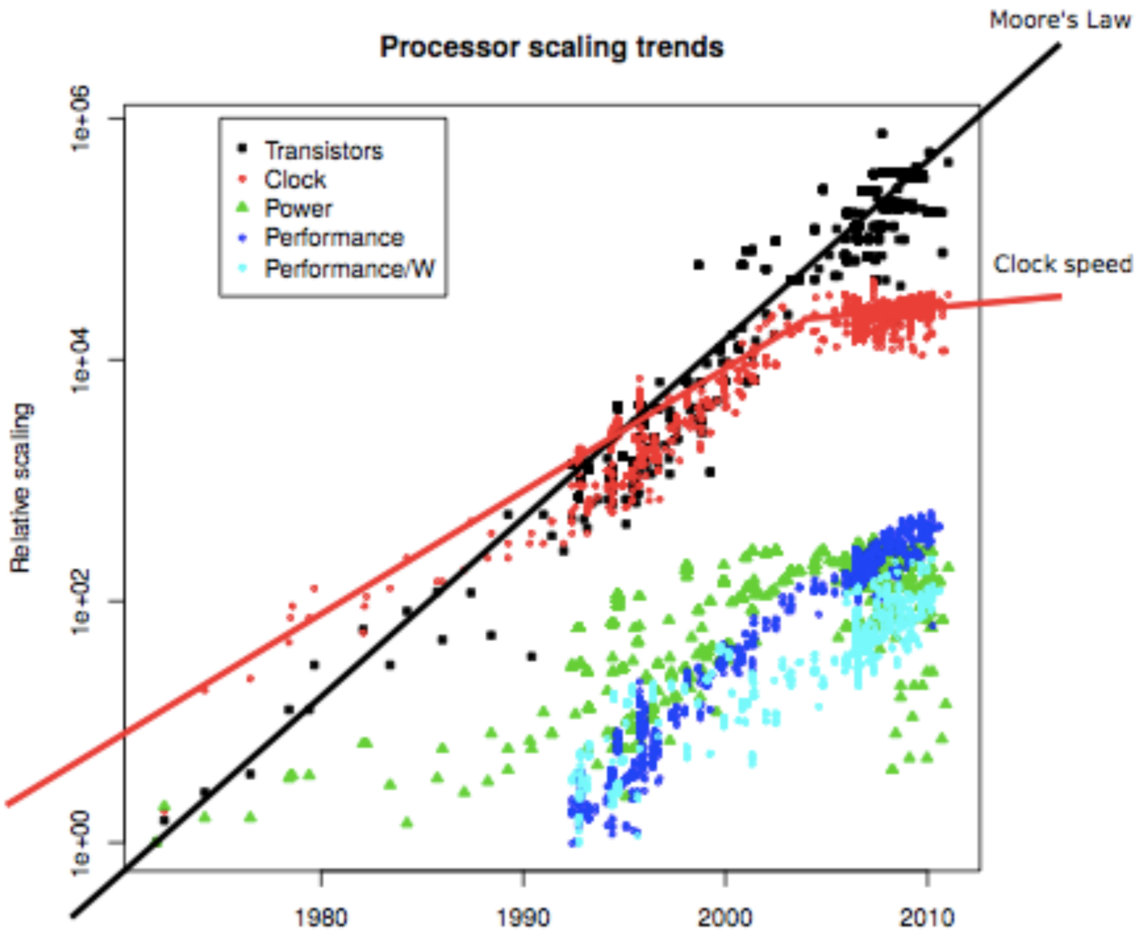

Crossing the river without drowning

Use more efficiently the **resources you own**

- ~DONE. It will continue, but do not expect miracles
- you can exploit e.g. "trains" concept in analysis, but it has pros and cons

Grab and use additional **resources you do not own**

- e.g. use HLT for offline processing. DONE and more BEING DONE.
  - Drawback: limited availability, partial control
- e.g. use opportunistic resources (AmazonWS for CMS, HPC, ..). DONE and more BEING DONE.
  - Drawback: no predictability. Plus, need to live with high failure rates

more (1)

Improve your **workflows / algorithms**

- often x2 per year so far, it is not nuts.. low-hanging fruits are taken already. ~DONE.
- more in a HSF context as a community joined effort? Perhaps.. BEING DONE.

Improve the **experiment software** you work with

- CMSSW multicore/thread-safe frameworks migration DONE and more BEING DONE
  - for completeness: AthenaMP (ATLAS), CMSSW (CMS), AliROOT (ALICE), Gaudi (LHCb)

Increase diversity in the **architecture(s)** you work on (GPUs, ARM, ..)

more (2)

- more power per money unit is good, but a complex ecosystem, and not painless transitions

Otherwise…

- accept to impact physics throughput (at least in terms of "collect now, process much later")
- "rethink" something radically

more (3, 4)

CMS on AmazonWS

more (I)

Exploiting the Amazon Web Services (AWS) grant (Jun 2015 - Mar 2016)

- ✦ Goal: demonstrate ability of executing any of the CMS centrally-organised workflows
- ✦ start-up at FNAL was intense and successful, **stably running on >50k cores**
  - official production requests, close to 1G evts, delivery in time for Moriond 2016
- ✦ contributed ~25% to the global CMS concurrent capacity
- ✦ lessons learned to be fully documented as a deliverable of the project

AmazonWS communication:
https://aws.amazon.com/blogs/aws/experiment-that-discovered-the-higgs-boson-uses-aws-to-probe-nature/

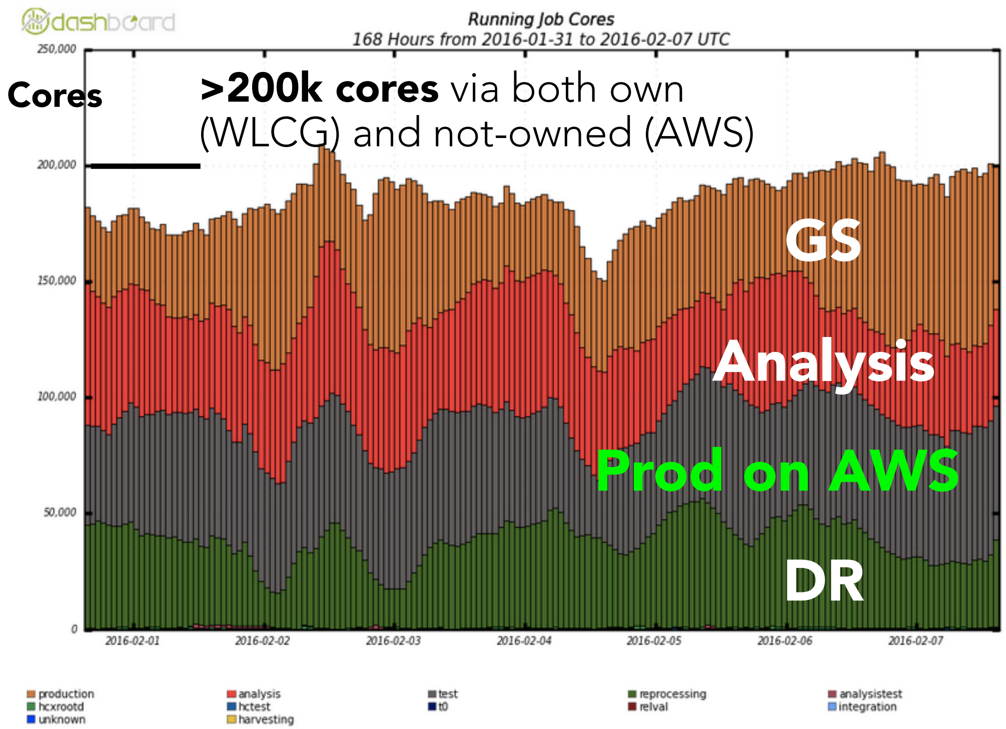

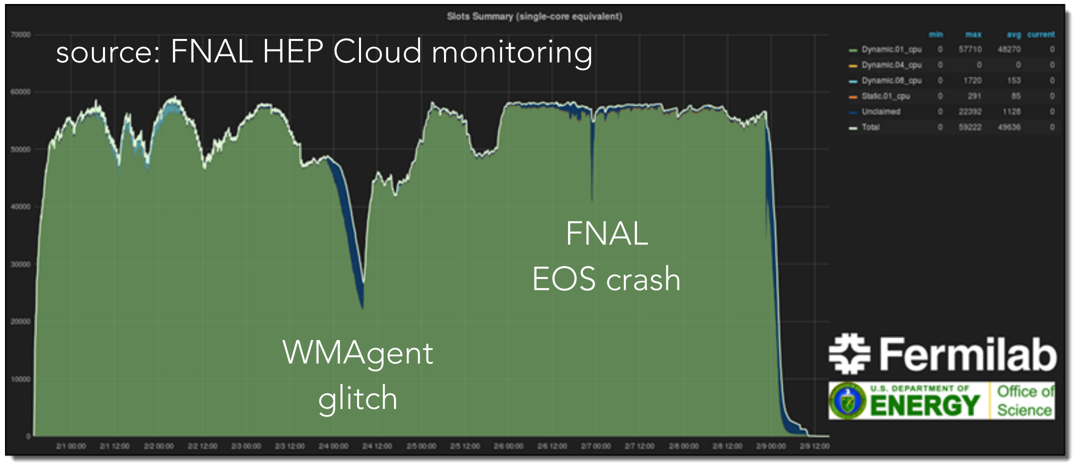

CMS official note being written…

“New” architectures

more (2)

R&D on software environment for:

- **low-power** computing architectures - such as Intel Atom or ARM processors
- **co-processors** - such as Intel Xeon Phi. Or GPUs

CMS approach in R&D on this

- not in production in the short-term, more crucial to gain experience in using them on the longer time scale perspective of Phase-II
- short-term focuses on making needed modifications to CMS SIM/RECO to run on such platforms, plus small-scale demonstrators
- particularly interesting for the low-power platform studies is evaluating and working to improve actual performance per Watt as well as memory usage of CMS software
- novel hardware for development and testing accessed via direct collaborations with vendors as well as through the CERN open lab framework

Multi-direction tests at small scale, but looks to the future:

- how developments could be used in a full-scale CMS SIM/RECO application
- benefits and constraints coming from a more heterogenous computing ecosystem

CMS "active evolving" dataset

more (3)

Typically the triggering system is limited by # evts it must reasonably reject

- and not the # evts it can accept

All CMS triggered evts are "equal"

- triggering decision ~100 ms, each evt treated as having same physics potential as any other evt
- this from collection to curation (i.e. evts are processed/archived/reprocessed/analysed)
- _nearly all_ evts are bkg, most are later rejected by analysis selections; some amount of skimming is done for analysis, but _all_ evts are re-processed and re-skimmed when needed

<u>Collecting and archiving</u> evts is **<u>not</u> expensive**

- archive of 1 yr worth of Phase-II data at 5 kHz is 400 kCHF [*]

Cost comes from resources to <u>process, reprocess, serve, analyse</u> the evts

- every evt collected in CMS is processed on average 2.5 times, hosted in AODs 1-3 times on disk
- # evts accessed in analysis hard to compute (e.g. Run-1 data in 2012 was read dozens of times)

Investigate alternative **data reduction and selection techniques**

- understand better the evts, and reduce the actively processed sample
- tag evts as candidate bkg; fill them into a distribution but stop to continuously reprocess/access them; the more you do it, the better you understand evts, the better choices over time you do
- challenging R&D, but it may decrease the scale of computing needs by significant factors

[*] WLCG evolution doc

Networks

more (4)

Estimated grow rate of global data centre traffic is +30% /year.

HEP/global traffic is ~10^{-3}

- within data centres: ~2000 EB/yr (HEP traffic is ~2 EB/yr)
- between data centres: ~200 EB/yr (HEP traffic is ~0.3 EB/yr)
- One may compare the 0.3 EB/yr above with a ~1k EB/yr expected for general global IP traffic in 2015-16

Rapid growth of enterprise networking market

- e.g. +22% of 1/10/40/100 Gbps ports deployed in 2013 wrt previous year

Costs decrease, and 100G costs becoming affordable

- 10G WAN connectivity cost decreased by a factor of ~3 in 3 years
- several academic backbone networks are already deploying 100G connections

It is not unlikely that HEP traffic growth will not exceed the technical capabilities we will have. We may need to expect "soon":

- 100G connectivity to T1s and very large T2s; 10G connectivity to large T2s; ≲1G connectivity for the rest, with HEP community and network providers needing to work closely together to ensure adequate provisioning in all badly connected regions

Basically, kind of problem is similar as before Run-1 started, despite the scale is at least ~1 order of magnitude higher

Network evolution(s)

more (4)

aka: evolution in the way we think of / use network(s)

Needless to say, "**la rete non è un problema**"©

Networks just OK for HEP/LHC Run-1 (and Run-2 so far)

- NRENs brilliancy, LHC{OPN,ONE} evolutions, monitoring and troubleshooting strenght, digital divide reduction, ..

We frequently face issues which networks may be part of the solution for

- "owned" resources facing input data availability latencies (SDN?)
- beyond data locality, heavier reliance on WAN access, benefit in workflows execution (flexibility) at a low cost (CPU eff)
- specialised systems? must be efficiently fed with large data volumes
- cloud storage? categorised as size and I/O capabilities: if bandwidth is sufficient, storage-CPU may not be close

CMS AAA/DDM as a first step towards a more advanced system (e.g. CDN)?

- Dynamic replication and clean up, predictive placement may be also based on network awareness
- could CDN be the primary source for the bulk, non-specialised computing?
- Moving to a CDN for data management with no expectation from the application for data locality simplifies the use of cloud resources, opportunistic resources, local user controlled systems.
  - the deterministic pre-placement may be preserved for specialised systems where the hardware is expected to be too fast to be served over WAN, or so specialised that the loss of data access could not be tolerated.

Summary

Run-2 core challenges are being addressed via workflows execution improvements, operations tactics, efficiency gains

- remarkable work in LS1 so far to get Run-2 started

Once this is done, though, *no easy factors* are left for after Run-2

- no space space for adiabatic evolutions of the model in all its parts
- we need to do a lot better with the same amount of CHF/€/$

Rethink something radically will be a must

- drive the CMS-internal changes of paradigms (as much as we realistically can)
- track and be able to quickly exploit promising technology (or breakthrough)
- equip CMS with a <u>LHC overall cost model</u> and <u>computing system(s) simulation</u>

Being able to focus on <u>both</u> Run operations and the CMS R&D program is a must.