R&D session: Software and Computing Requirements and possible architecture

Simone Campana – CERN

CERN ATLAS EXPERIMENT

The data rate, volume and complexity challenge

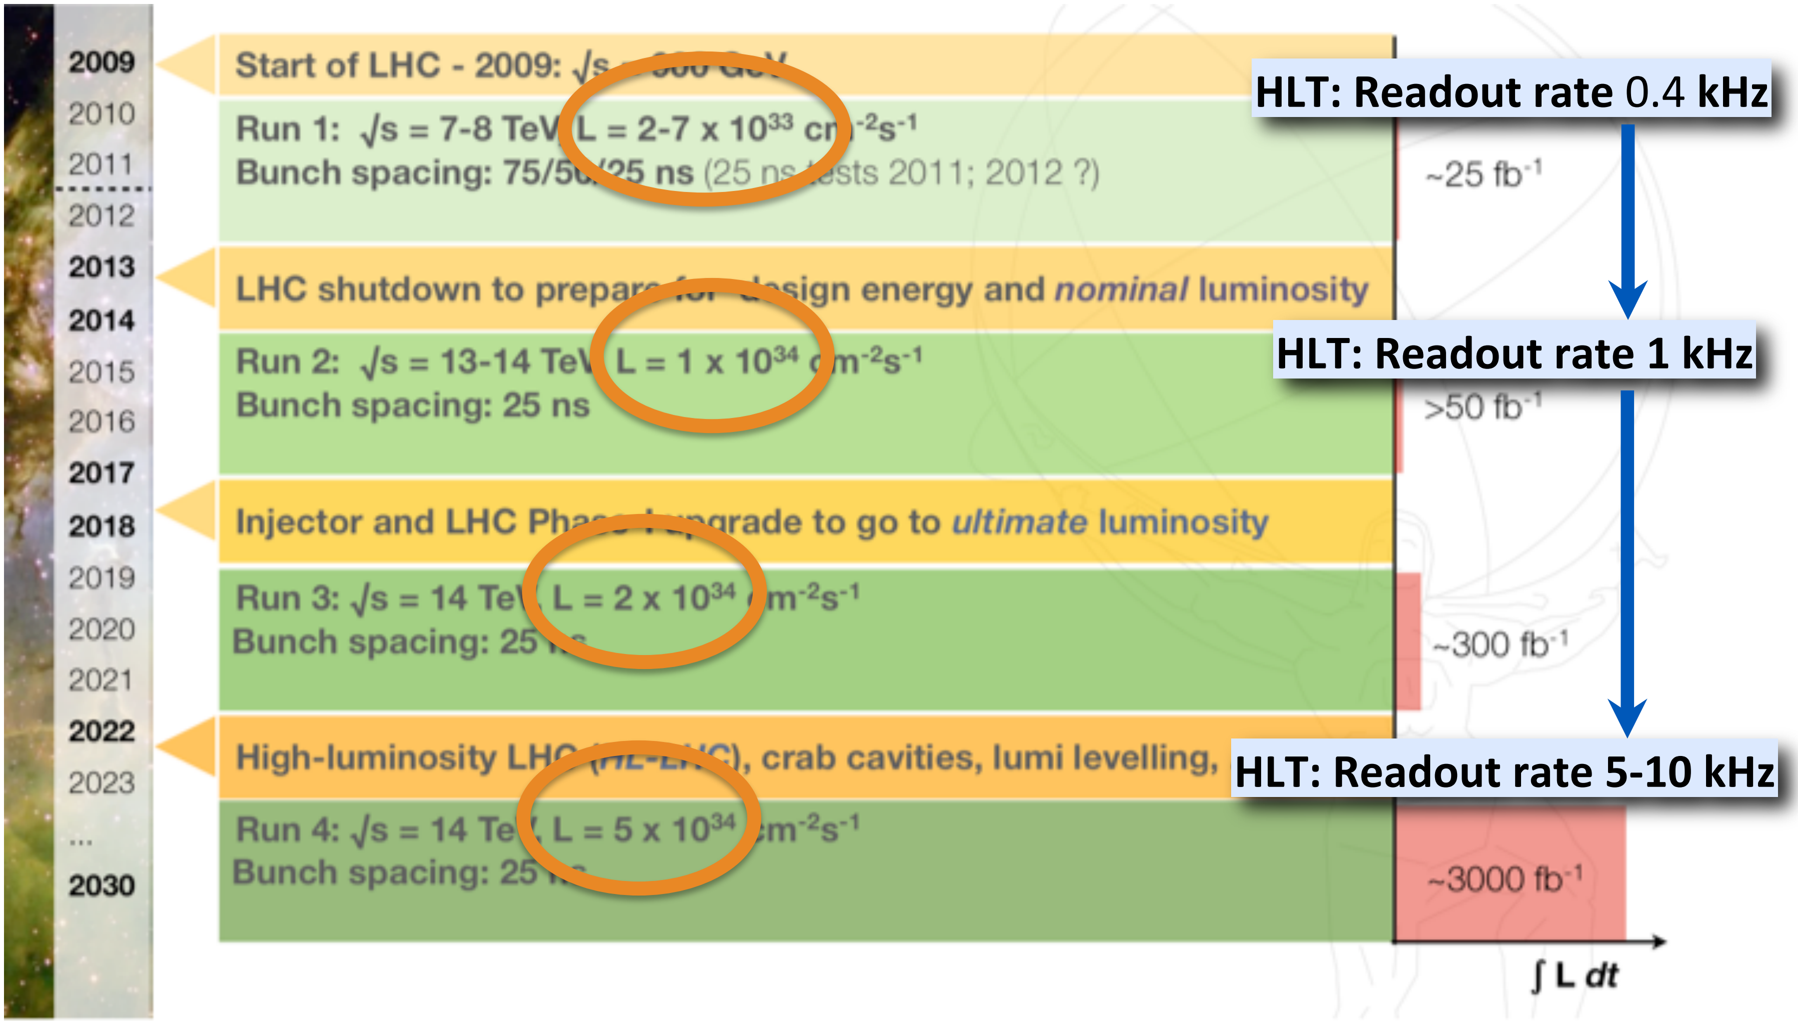

Effect of pile-up increase

The average pile-up:

<mu>=14 in 2015
<mu>=23 in 2016
<mu> ≈ 35 in 2017
…
<mu> up to 200 in HL-LHC (10 years)

Higher pileup means:

<u>Linear</u> increase of digitization time
<u>Exponential</u> increase of Reco time
Larger events
Lots of more memory

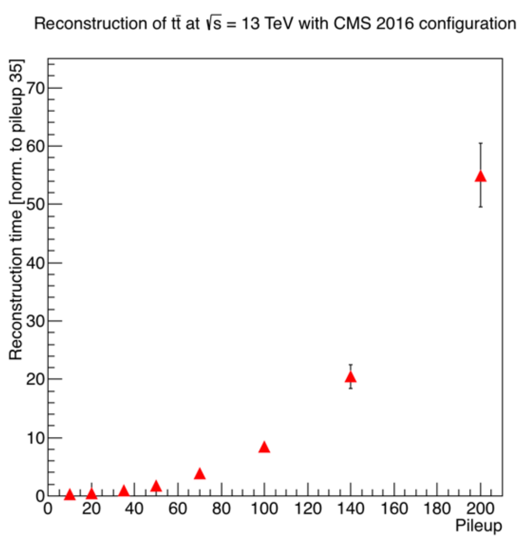

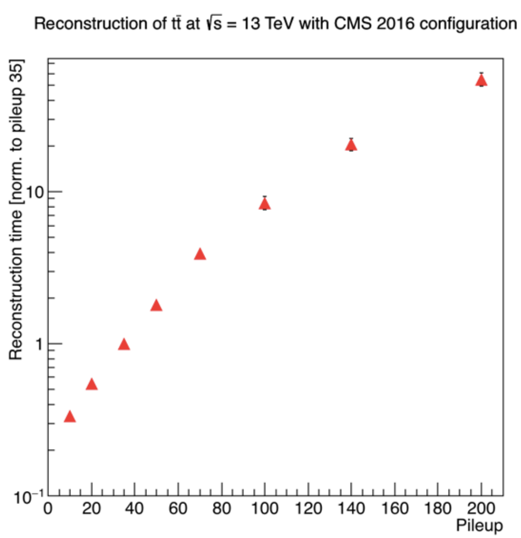

The exponential increase in reconstruction time saturates beyond Run-3 conditions (mu=80)

Indicate a loss of tracking efficiency of the current detector layouts at HL-LHC

Estimates of resource needs for HL-LHC

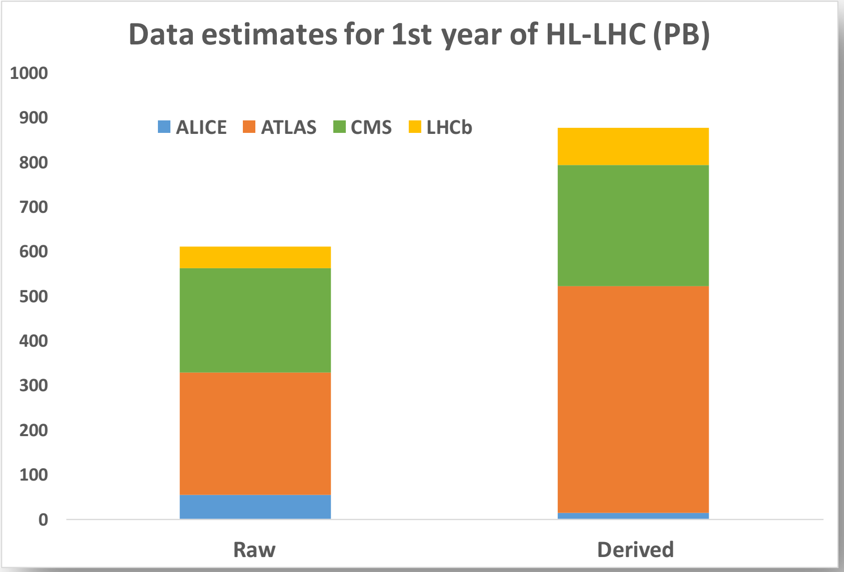

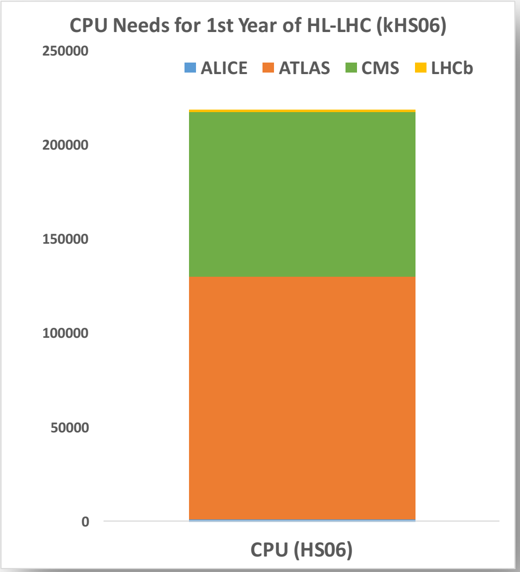

Storage

Raw 2016: 50 PB → 2027: 600 PB

Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU

x60 from 2016

Technology at ~20%/year will bring **x6-10** in 10-11 years

=> x10 above what is realistic to expect from technology with constant cost

In this presentation…

- The resources needed for HL-LHC will be driven by ATLAS and CMS
- Alice and LHCb will face a challenge in LHC Run-3 and already evolved their computing model
- … I will focus on ATLAS and CMS computing at HL-LHC
- I am more familiar with the ATLAS computing model and the tools to project it to the future.
- Many plots will be based on those tools and the ATLAS computing model, but the conclusions apply to both ATLAS and CMS

Input parameters, assumptions, disclaimers

Simple model based on today's computing models, but with expected HL-LHC operating parameters

ATLAS Input Parameters at HL-LHC (LOI = the ATLAS Letter of Intent for Upgrade Phase-2)

Output HLT rate: 10kHz (5 to 10 kHZ in LOI)
Reco and Simul Time/Evt: from LOI
Nr Events MC / Nr Events Data = 2
Fast Simulation: 50% of MC events
LHC live seconds /year: 5.5M

CMS Input Parameters at HL-LHC

Output HLT rate 7.5 kHz
LHC live seconds /year: 6.0M
Dataset overlap factor: 1.2
Reco and Simul Time at mu=200
Nr Events MC / Nr Events Data = 1.3
Analysis estimated as +60% of all other CPU usage

Simplified Computing Model with respect to 2016/2017 resource requests:

Legacy from previous years not taken into account
=> Little difference at the beginning of the Run-4 but huge difference for Run-2 and Run-3

HL-LHC baseline resource needs

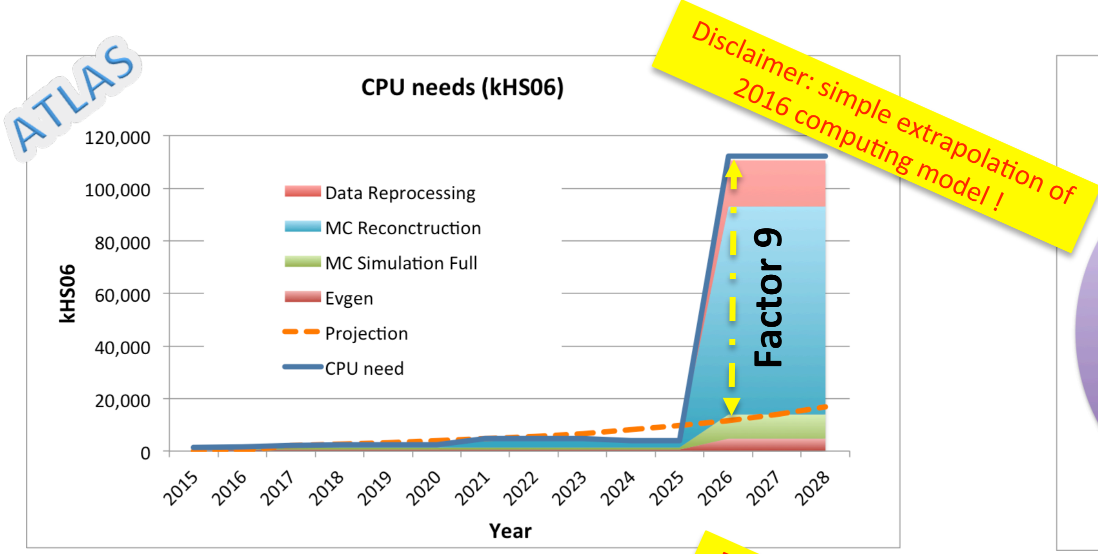

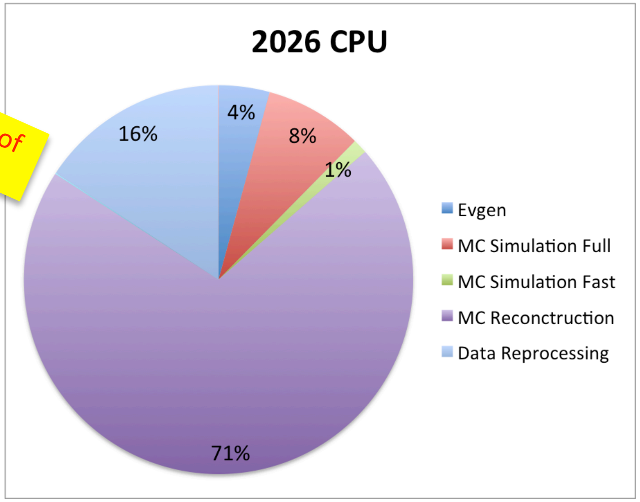

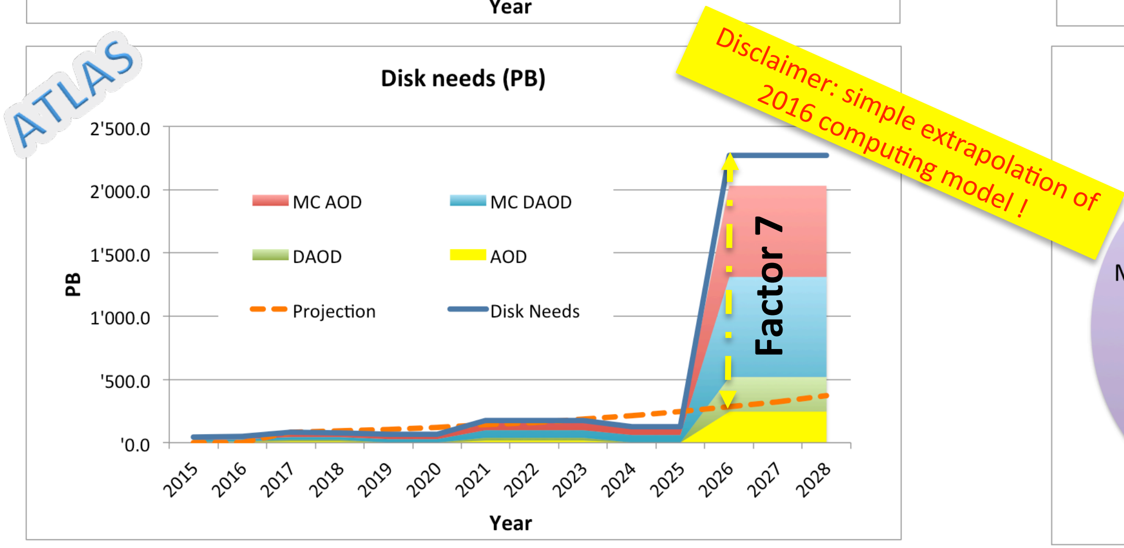

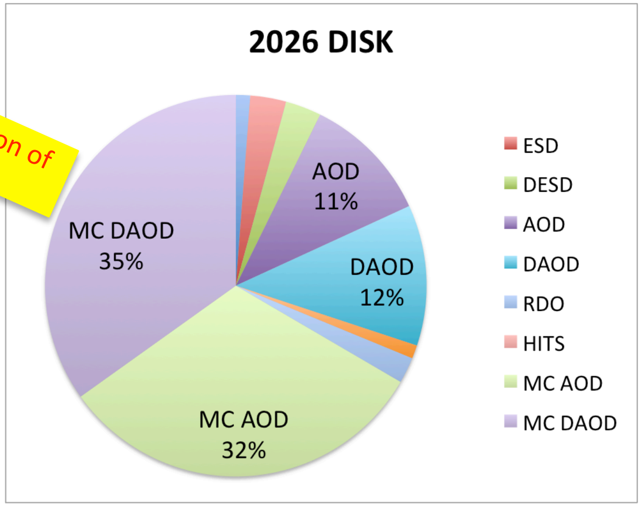

events: HLT output rate and MC needs

The output trigger rate does not determine only the amount of data per year but also the amount of Monte Carlo to be produced.

We foresees a value between 5 kHz and 10kHz. ATLAS baseline is 10kHz, CMS is 7.5kHz

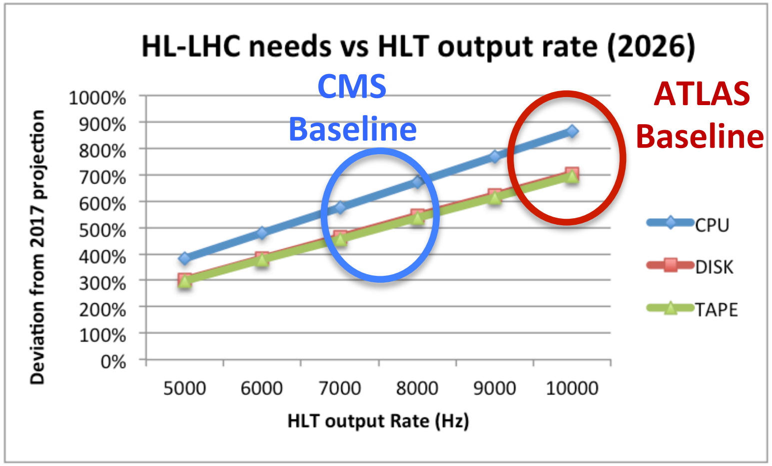

The physics case for HL-LHC will evolve in the next years.

One might assume a lower need of MC with respect to data, but generators might become more expensive seeking precision

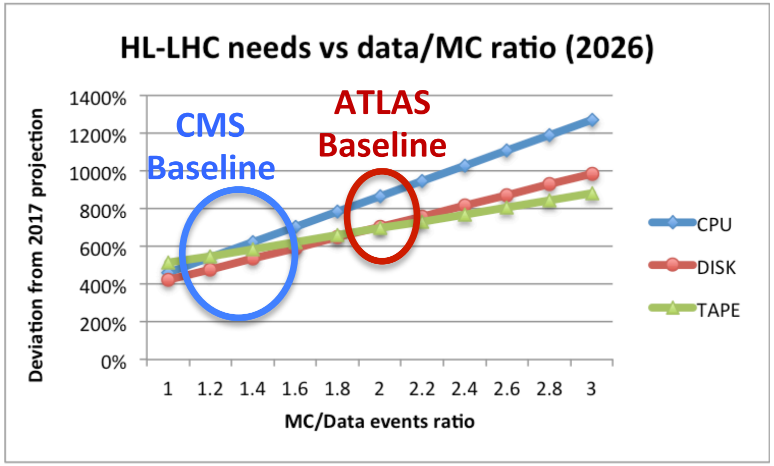

Fast Simulation and Fast Chain

G4 Fast Simulation will moderately help in HL-LHC. CPU is driven by reconstruction

Both ATLAS and CMS invested in a Fast Chain. x10 (++) faster than standard simulation

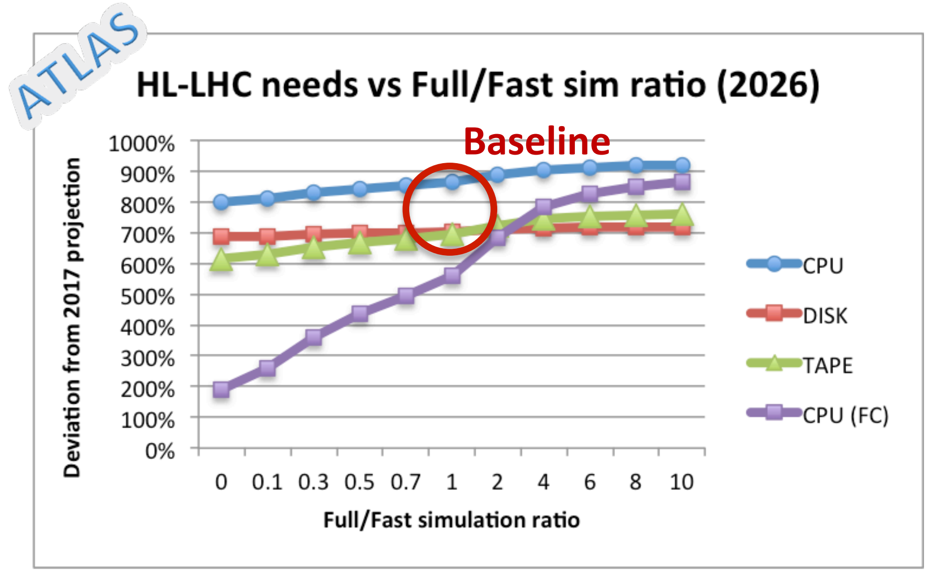

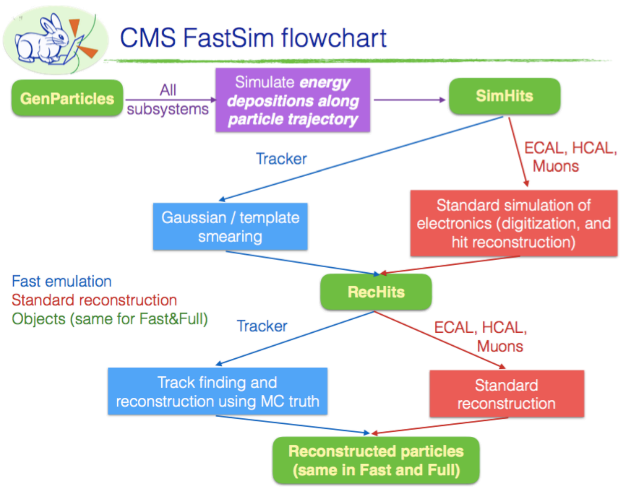

ATLAS Fast Chain

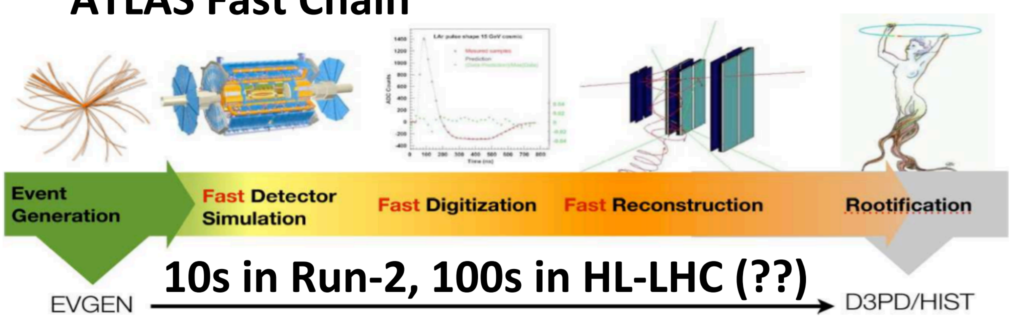

Layouts and Reconstruction

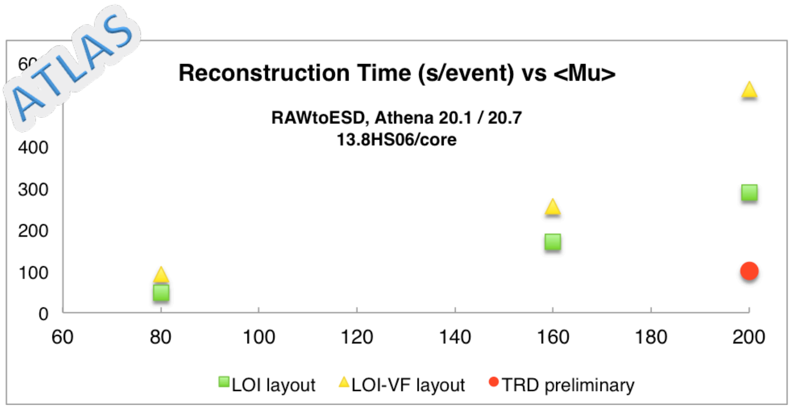

Reconstruction time dominates the CPU consumption in HL-LHC

The detector layout will play an important role, together with the optimization/tuning of algorithms. Tracking will be the main consumer

It is important to consider computing performance in designing the HL-HLC detectors. Good that this is happening

LOI Layout

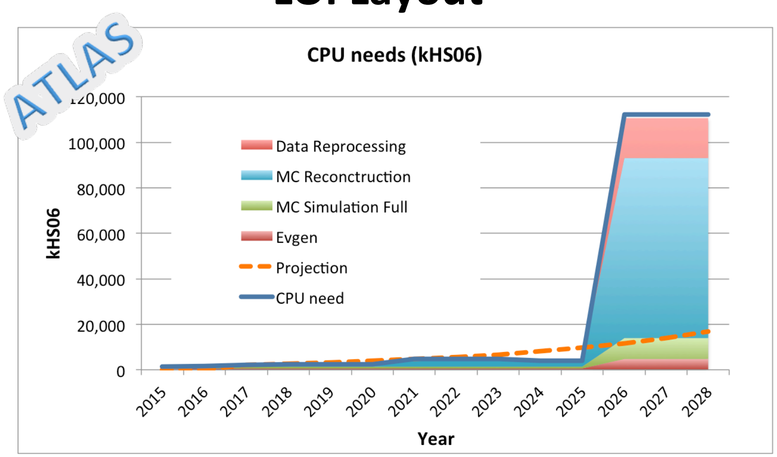

Possible TDR Layout

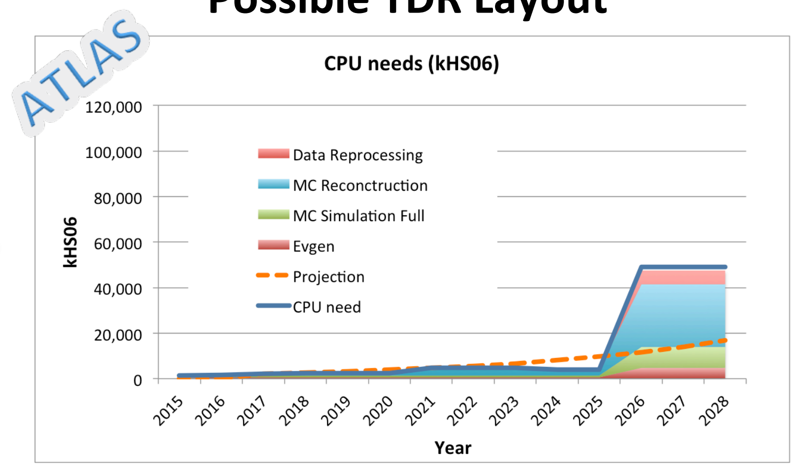

Preliminary conclusion

- The CPU needs for HL-LHC could exceed x10 the projection of today's resources in 2026 in a pessimistic scenario
- In reality, large gains are foreseeable and we are on the right path
- Hardware trends will play a crucial role and our software will need to adapt to them
- So please listen carefully to the next two presentations

What about Storage ?

No AOD on disk

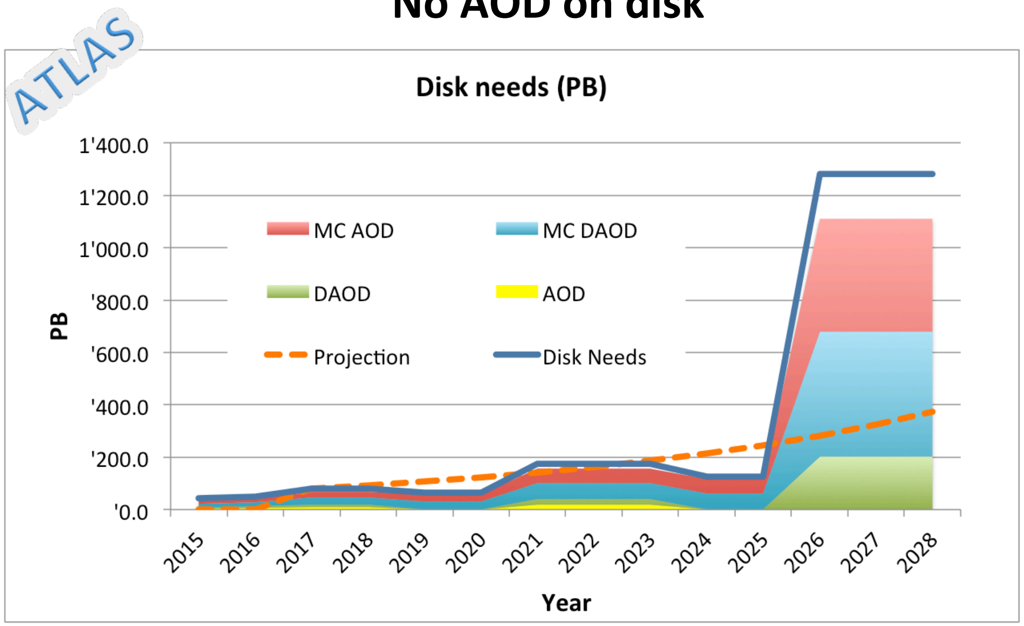

Storage is really the hard part. Even in an optimistic scenario, we are still far from solving the problem

AODs and their derived formats are the main consumers.

With no AOD on disk you get x4 above the resource projection (left plot)

The remaining gain must come from re-thinking of distributed data management, distributed storage and data access. A network driven data model allows to reduce the amount of storage, particularly for disk. Tape today costs at least 4 times less than disk.

Computing infrastructure in HL-LHC

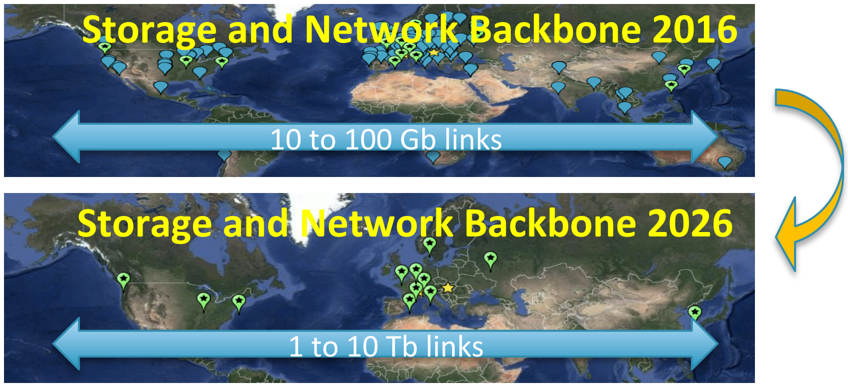

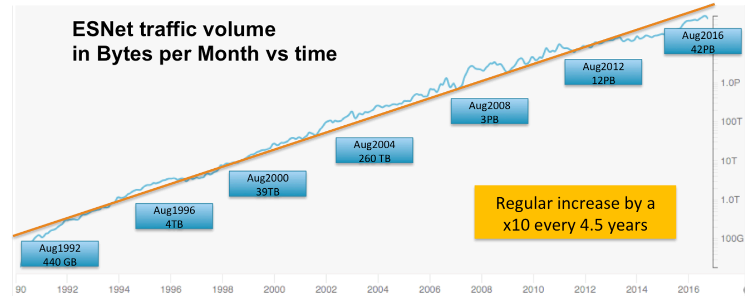

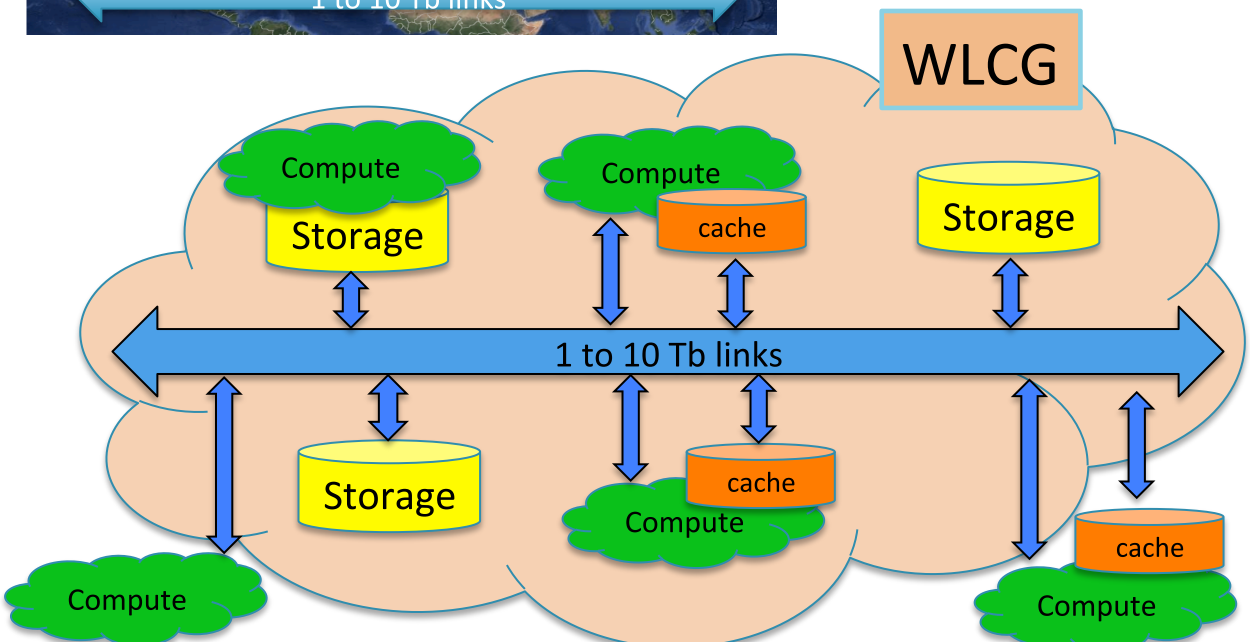

A data cloud for science

Storage and Compute loosely coupled but connected through a fast network

Heterogeneous Computing facilities (Grid/Cloud/HPC/ …) both in and outside the cloud

Different centers with different capabilities, fo different use cases

Data Management: Challenges and Opportunities

- "Funny how tape never seems like the cheap option when you have to pay for it". One could say the same about network
- A fast WAN does not imply fast data access. The infrastructure and the I/O layers need to be optimized from end to end
- Multilevel caching should be built **IN** the infrastructure rather than **ON** top of it
- A unique opportunity to define and implement a common data management and data access layer
- Today WLCG is a data Grid. Tomorrow we will have a data cloud The challenge is always the data

Conclusions

- We identified a concrete set of steps in preparation for computing at HL-LHC
- To keep cost of computing under control in 2026 we need to invest effort from now. Data will be the challenge.
- The effort spans many areas: online, offline software, distributed computing, physics, infrastructure and facilities. The detector layout will play a crucial role
- It is important to consider cost of computing when choices are made
- We are on schedule to define a computing model for HL-LHC in the next three years