Track 1 - Summary

ACAT 2019. Saas-Fee. March, 11-15

Conveners: Patricia Mendez Lorenzo & Gordon Watts

Track 1 : Highlights

- 39 oral -> 780 min of Computing Technology for Physics Research
- 40 posters
- >50 persons per session

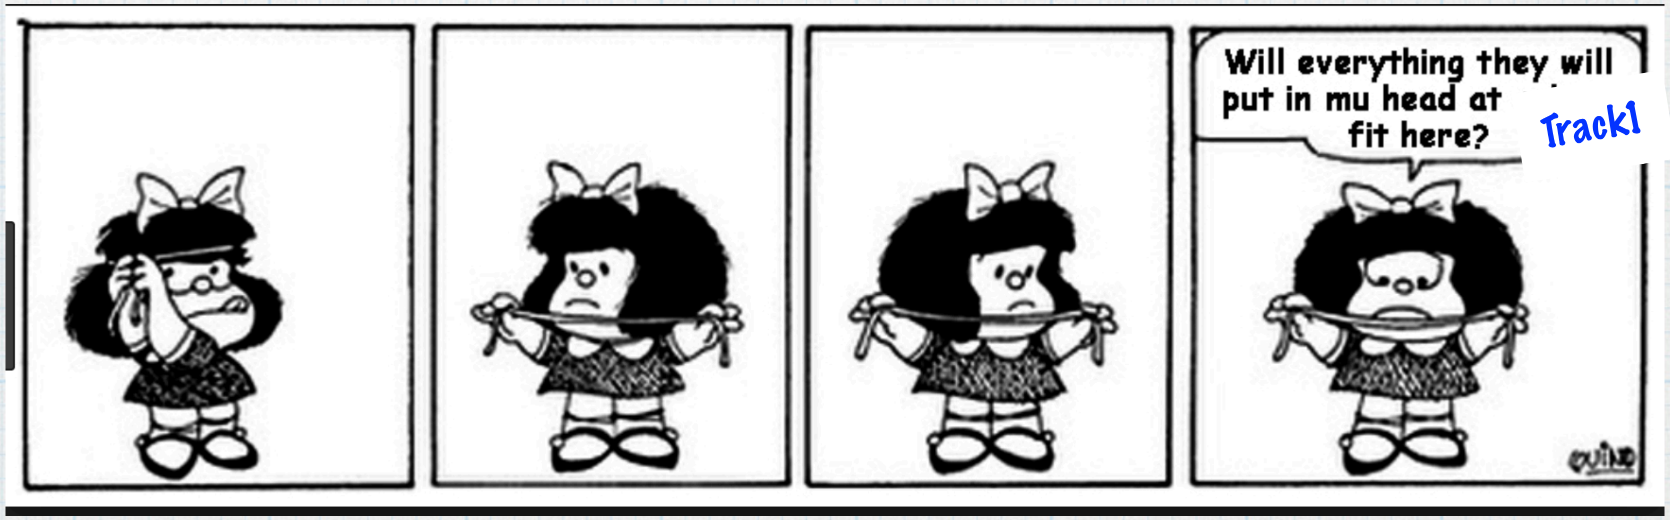

Parallel evolution

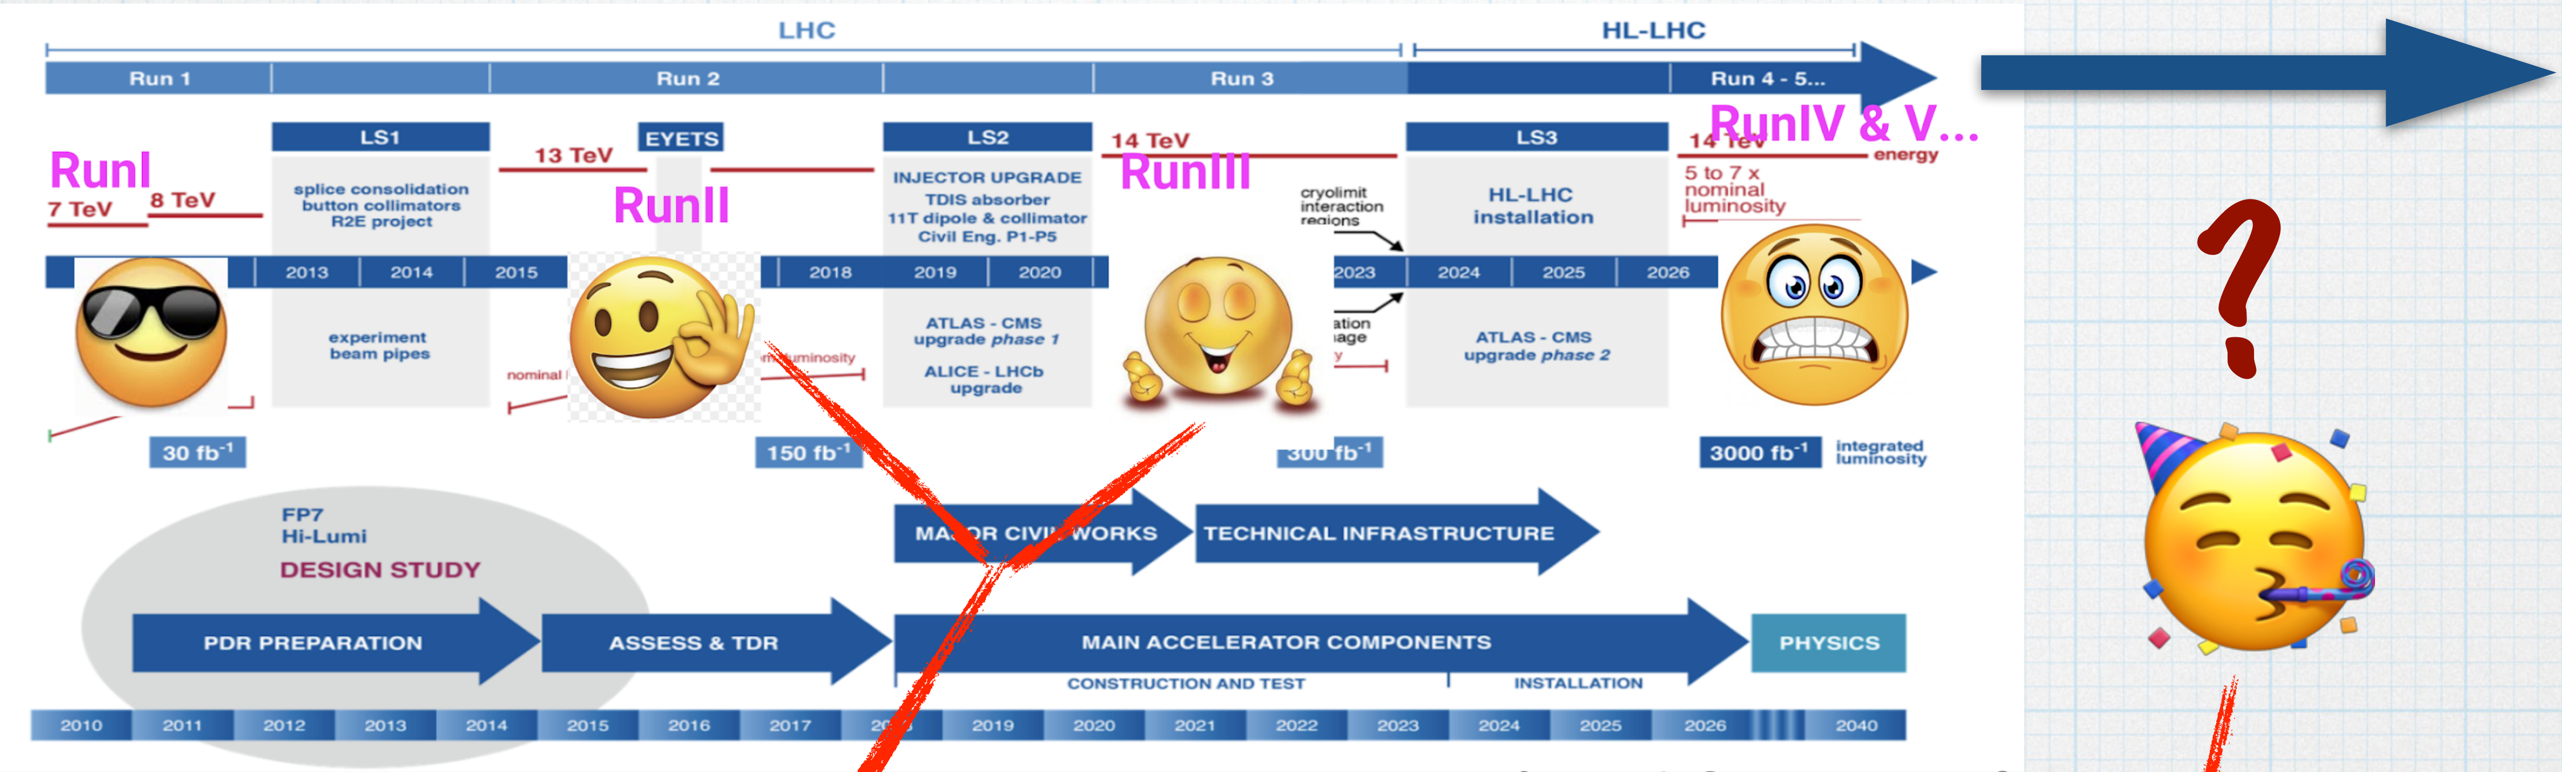

(CMS slides)

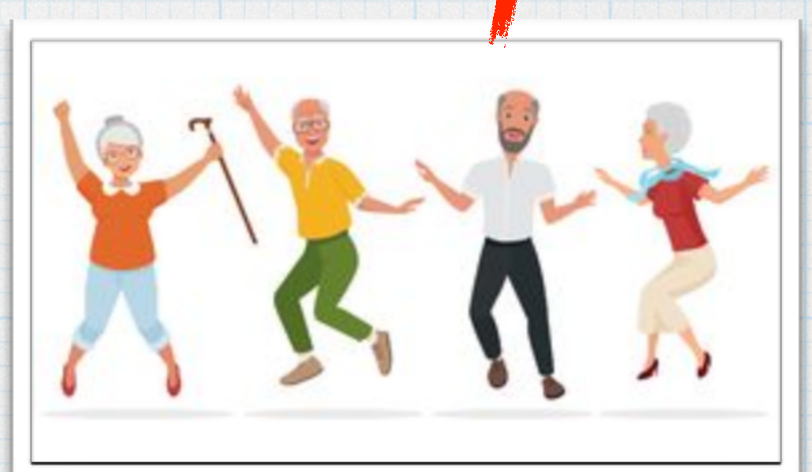

Central messages

* Run3 contained , HL-LHC explotes →

Instantaneous luminosity X 5-7
particles per collision x 5
more data x 15
more granular detector x 10 readout channels

* Funding agencies will not "assist" in this explosion
  * HL-LHC data paradigms unsustainable with flat budget
  * No accelerator HW —> no allocations
* BUT: computing workflows have made little use of heterogeneous resources

Let's be more concrete

Good bye LHCb
Let's welcome LHCb

New data formats: NanoAOD(CMS), Turbo(LHCb)

New detector interfaces: FELIX (ATLAS)

C++20 (or C++2.0) New ROOT implm. GeantV

Implementation of ML and heterogeneous HW at all levels

Data paradigm in the Grid: Data-Lake
Workflow can always improve: Containers

Highlights from experiments (1)

* Good-bye LHCb, hello LHCb

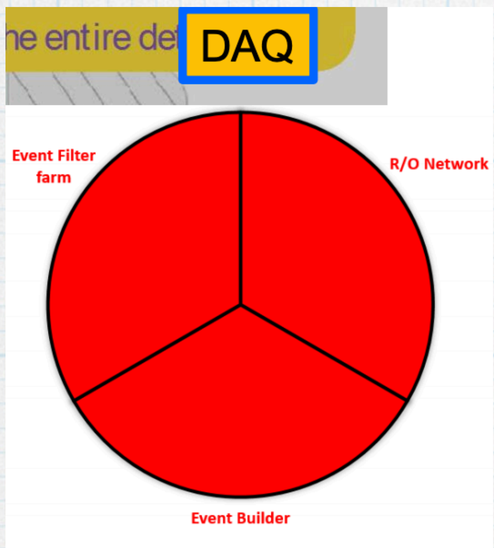

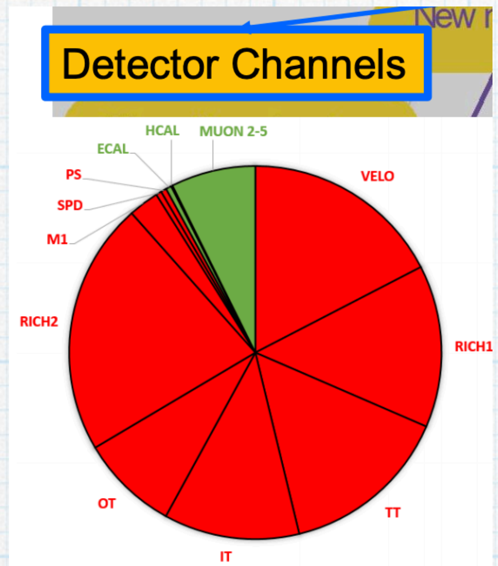

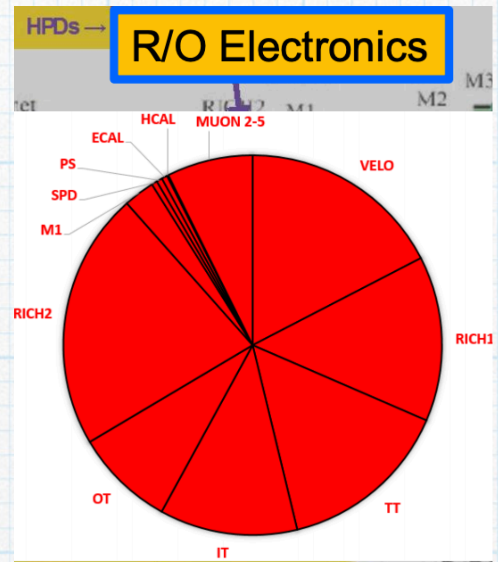

* 5-fold increase instantaneous lumi
    * 4x10^32 —> 2x10^33 (cm^-2s^01)
* Full SW trigger at 30MHz inelastic collision rate

* New SW based HLT (sync. HTL1 + async. HLT2)
    * Replacement of a HW based HTL by a new SW based HTL
        * Allen FW for HTL1 for GPUs usage with visible impacts in the memory handling, scalability, data access and cost reductions
    * Physics Analysis based in the information provided at trigger level with no raw data access
* Turbo paradigm for reduction of event data format
    * Keep of each event as much as needed (similar approach in CMS)
        * Reduction of the output data rate in a factor 2 (Run2 using 30% of events in Turbo to 70% expected in Run3) —> Only high-level reconstructed information is saved to offline (no raw)
        * Only 26% of data will still be kept in full stream mode (no raw)
        * 6% raw kept for calibration purposes

Highlights from experiments (2)

* Changes in the core Sw structure: GAUDI
    * Re-engeniering needed to cope with previous quantities
        * Multi-threading and multi/many cores structures supported
        * New scheduler and a re-design of the data structures towards optimisation of memory resources and speed-up of data access

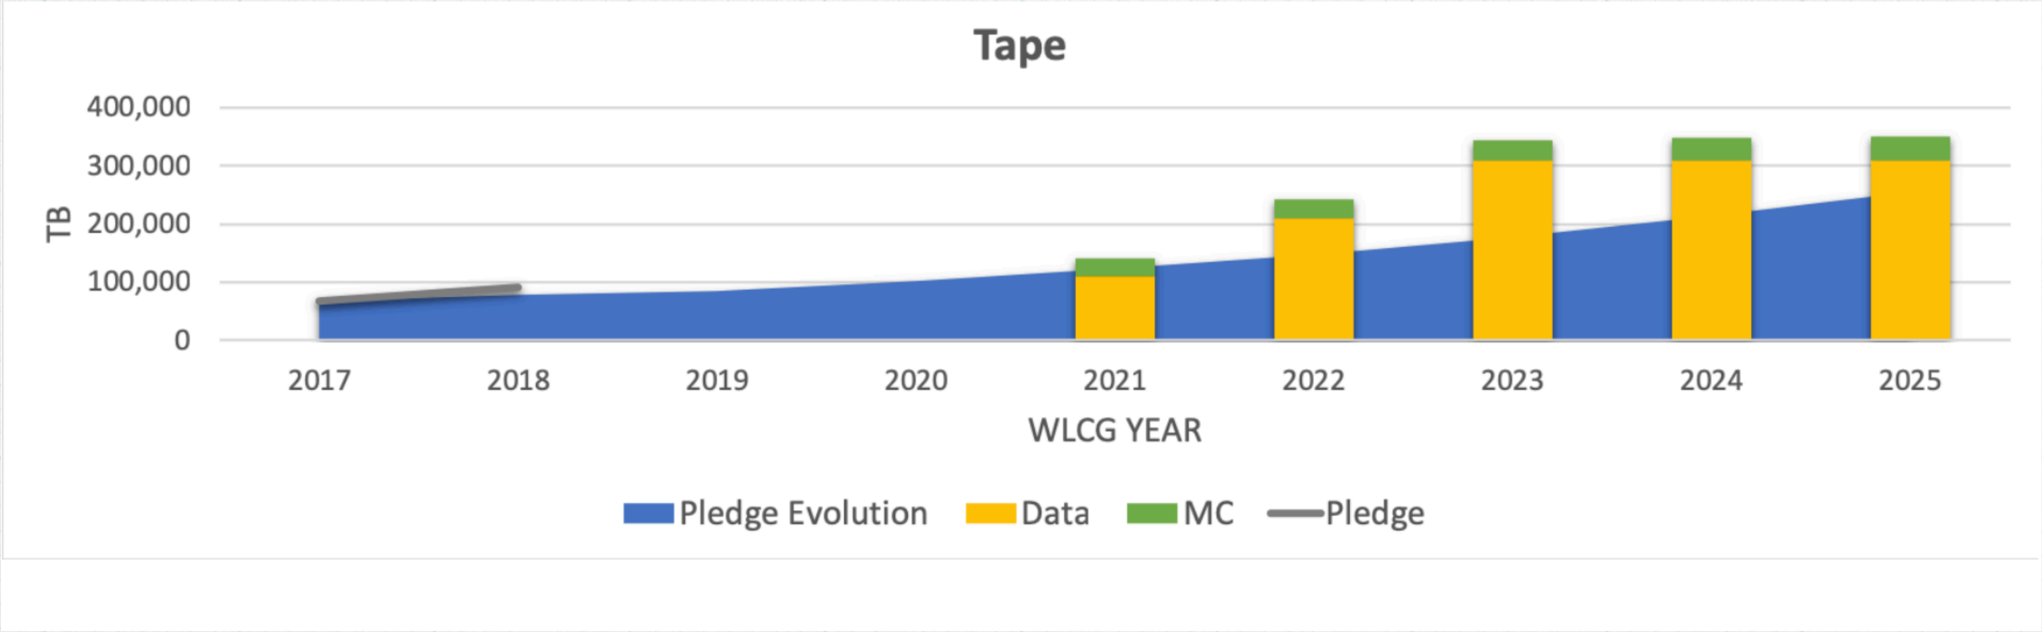

Highlights from experiments (2)

* From MiniAOD (Run2) to NanoAOD(Run3): New Data Format for CMS
    * More portable and scalable approach
        * Maximal flexibility to define contains
    * ntuple-like format including relevant information for most of the generic analysis
        * 50% of analysis expected with NanoAODs
    * Full MiniAODs under transformation to NanoAODs

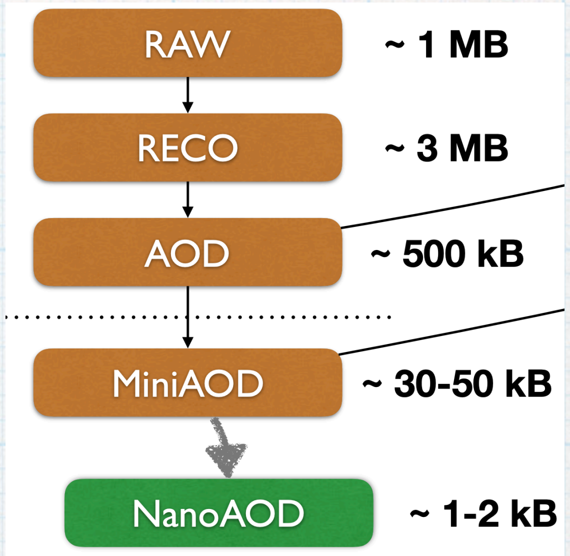

Highlights from experiments (3)

* HCAL/ECAL Calorimeters: from >85K to >91k channels in Run3
  * Current situation: calorimeters take 20-25% of reco time
  * Towards implementation of GPUs for parallelisation purposes —> very promising results
  * Testing FPGAs —> work still ongoing
* Global effort in CMS to exploit power of heterogeneous computing —> Modelling needed for improving the accuracy of CMS computing resource needs for HL-LHC
  * Initial group setup for the evaluation of data usage and justification of changes in the data model

Highlights from experiments (5)

* ATLAS: parallel evolution of the systems to cope with the new experiment upgrades:
  * phase I (Run3) and phase II (Run 4)
* Evolution of the ATLAS SW system
  * Software system: From PANDA to ProdSys2
    * Support of HPCs and Clouds
  * DAQ level: ML techniques
    * Anomalies detection
    * Connection instabilities
  * FELIX system: the interface between the data acquisition; detector control and TTC (Timing, Trigger and Control) systems; and new or updated trigger and detector front-end electronics
    * Phase I requirements already fulfilled

Pushing the brain: New ideas

* **Automated and Intelligent Data Migration Strategy for HEP storage**
    * Model based the prediction of future file access using deep learning algorithms and labelling data based on their access (hot/warm/cold)
* **FPGA and accelerated machine learning inference as a service for HEP computing**

ML power for physics processes | Parallelization approches | Computing challenges

    * Recast physics problems as ML problems taking advantage of the acceleration of such processes in specific HW: GPU, FPGAs, ASICs
    * Promising results from a proof-of-concept on cloud FPGA (Brainwave)
        * large computing tasks: >100 benefict over CPU only
        * latency-limited tasks (HTL): suitable for closer clouds and edge solutions

We cannot forget

* C++20 in 2020 (or c++2.0): Programming will never be the same
* PyROOT (ROOT Python bindings) modernisation:
    * New implementation of PYROOT based on Copyy
    * Interoperability with the Python data science ecosystem tools (Numpy and pandas)
    * coexistence and co-support of python2 and 3 per each build
* Simulation: GeantV: performance tag (beta) of GeantV demonstrator coming soon

There is physics beyond LHC

* Nuclear physics (Jefferson Lab)
    * Description of JANA2 as multi-threaded event reconstruction framework (c++1 1)
* STAR Data Production Workflow on HPC

Grid evolution: Data-Lake

* **Main challenge in HL-LHC is Storage**
  * Data-Lake is the evolution of the current WLCG infrastructure in terms of data handling towards the optimisation of storage costs
    * EUlake prototype: revise concepts of redundancy, caching, interoperability and reproducibility

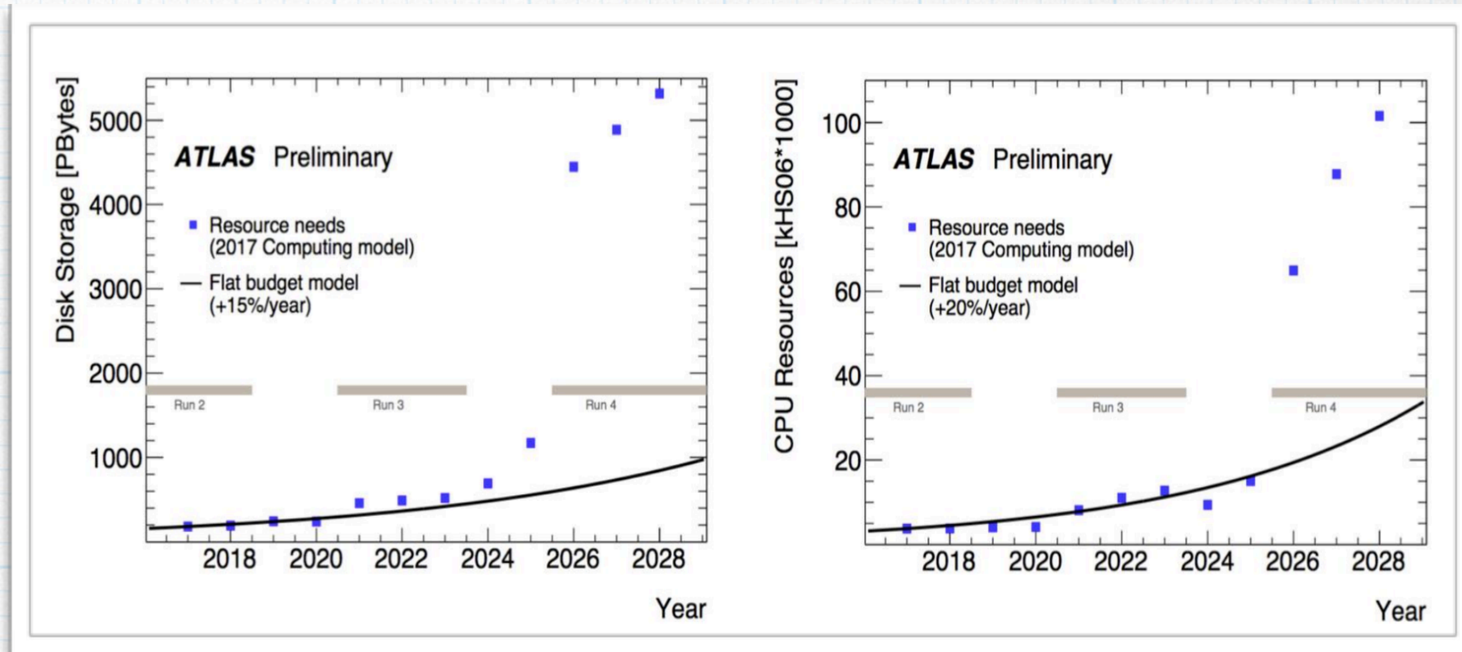

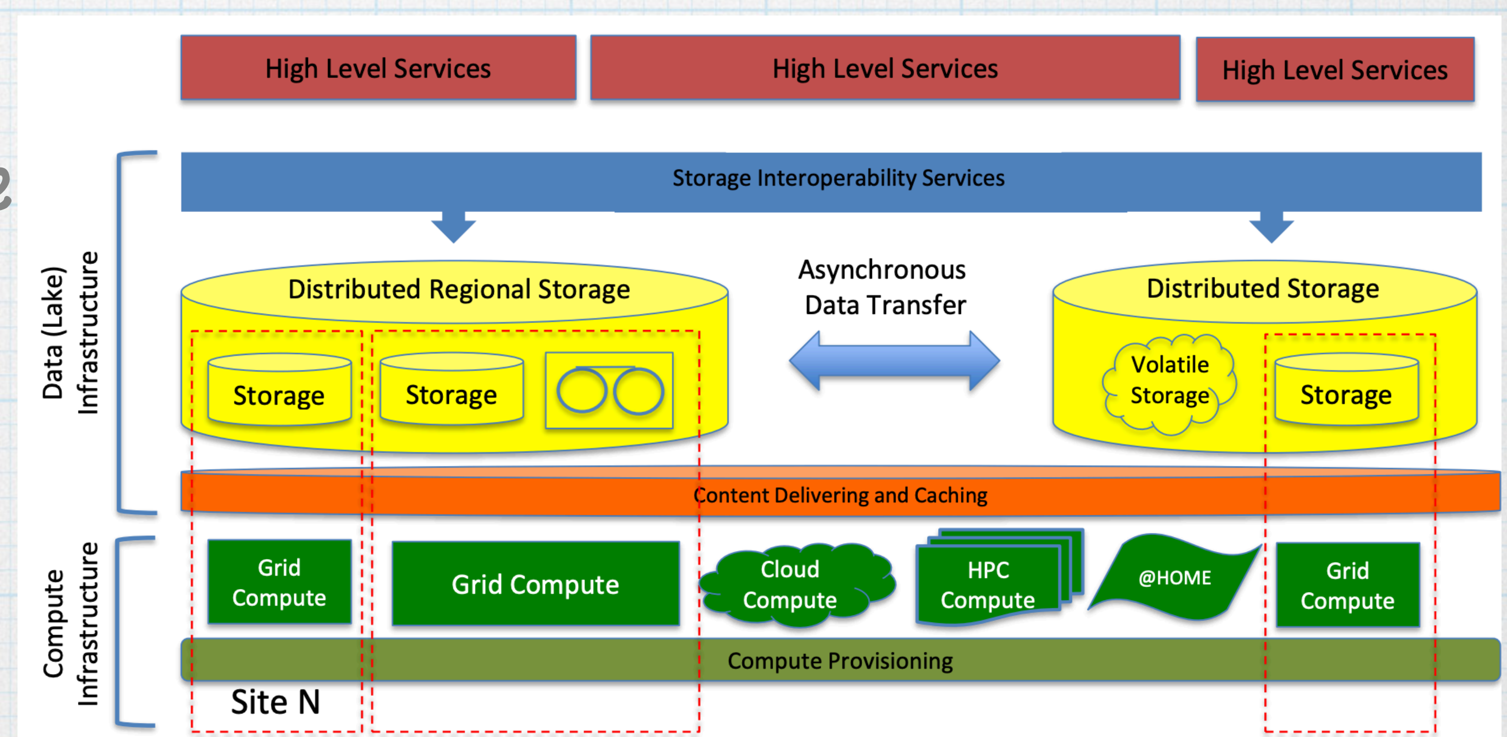

Presentations at the 3 levels:
* Cashing prototype: DODAS
* Workflow: Containers solution for ATLAS