Development of computational resources in ALICE

Latchezar Betev

24 February 2016

Development of high-performance computing to process the data of experiments of the Large Hadron Collider - at St. Petersburg State University in Russia and in the world

Where are we today?

Some figures

- **80** computing centers
  - **T0** (23%) + **8 T1s** (24%) + **71 T2s** (53%)
- Up to **100K** concurrently running jobs
- **600M** executed jobs
  - 1600 users
- **9 Tape** + **56 Disk** storage elements
  - 25 + 25 PB of data, **1B+ files**
- Up to **40 GB/s** read rates (10GB/s avg)
  - Writing at 1/10th the read rate

Distributed resources

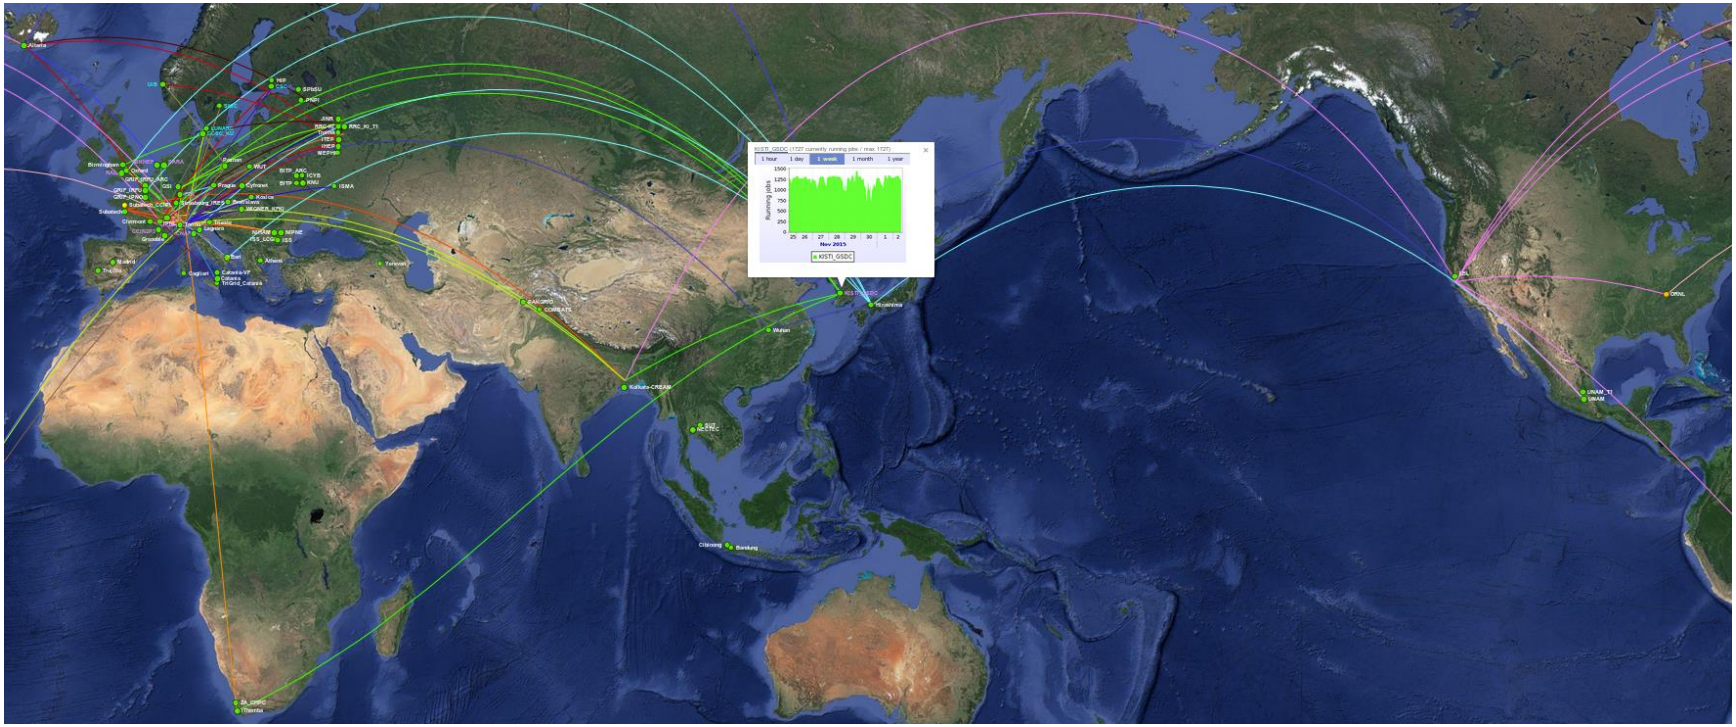

- Federated computing and storage resources
  - Users interact with the entire Grid through AliEn
- Tightly coupled central task queue and file catalogue
  - Tasks are typically sent to where a copy of the input data is but can also read from anywhere in the world

ALICE Grid evolution

10 years or history

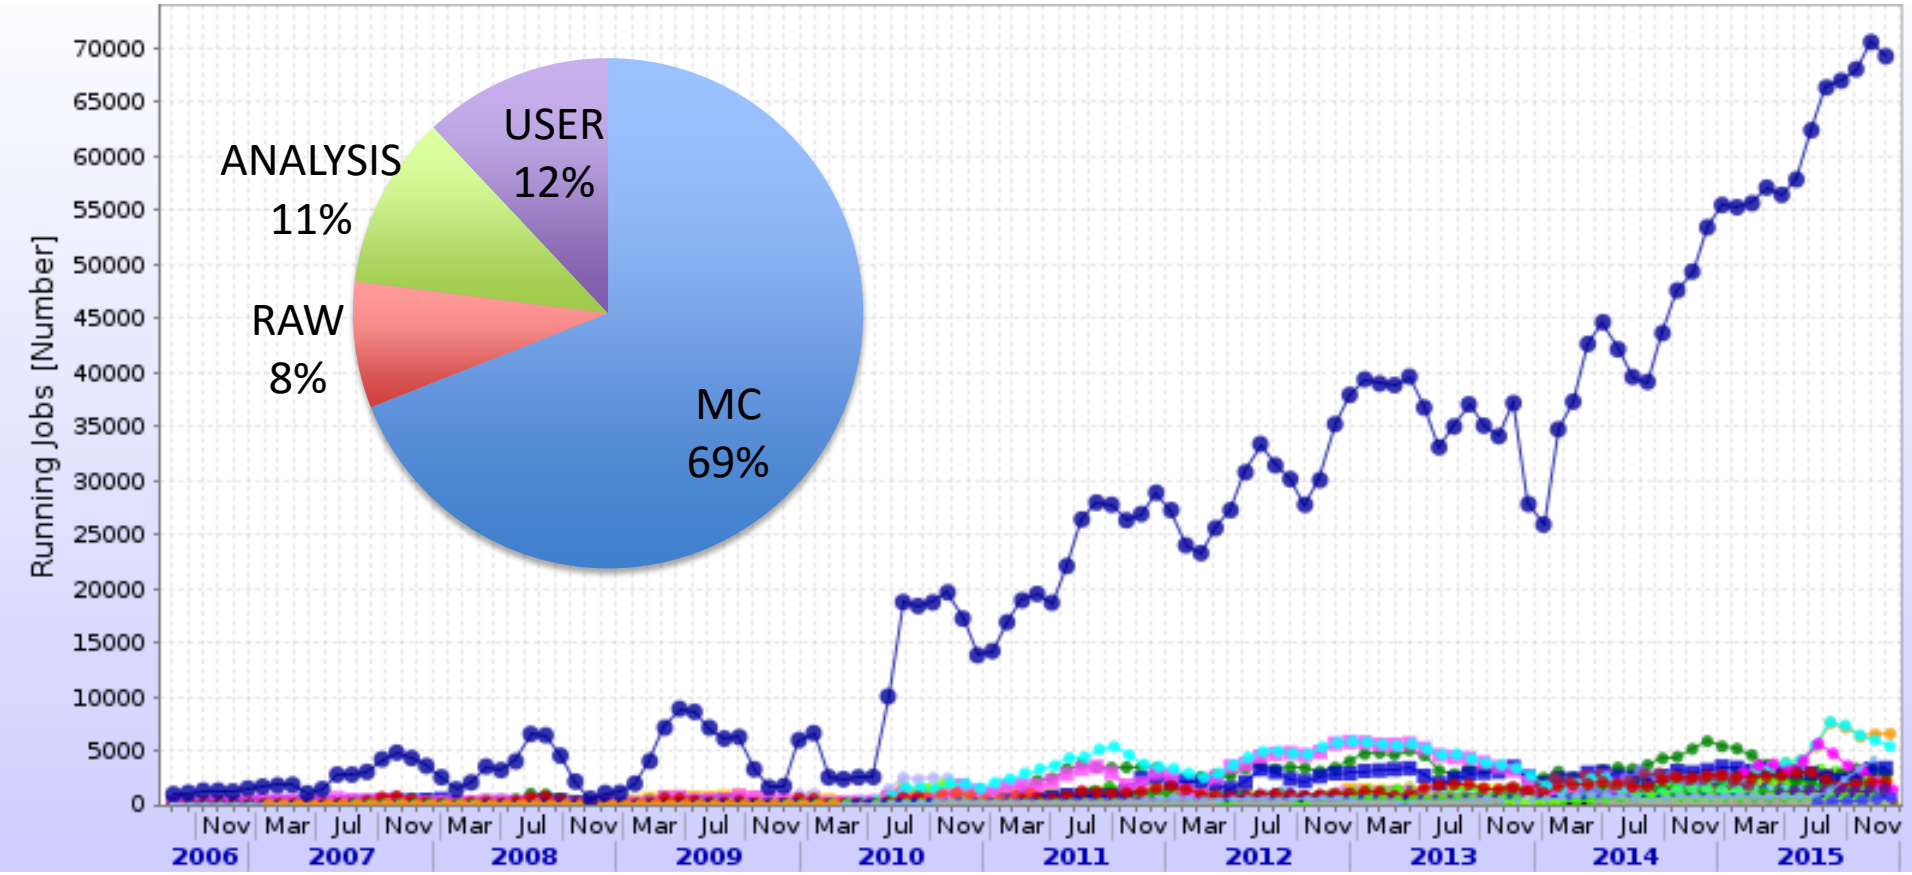

- Exponential growth of resources with ~flat budget

Start of Run 3 resource forecast

200K cores

- Expecting the 20%+ year-to-year growth to continue in the next years

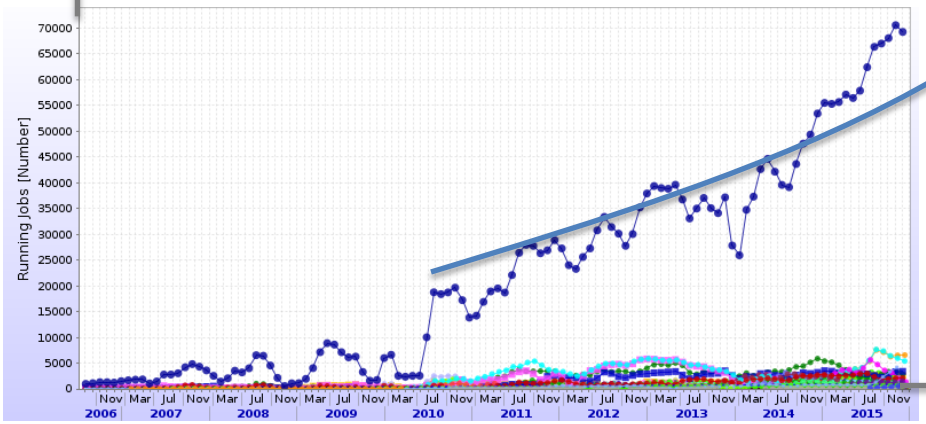

trade disk for CPU

+ flat budget
- old hardware

2016 2017 2018 2019 2020

Storage usage

+21PB since start of Run 2

- Sharp increase during Pb-Pb data taking
- 6.5PB of new raw data
  - 1.2PB of which was replicated to KISTI_GSDC::TAPE

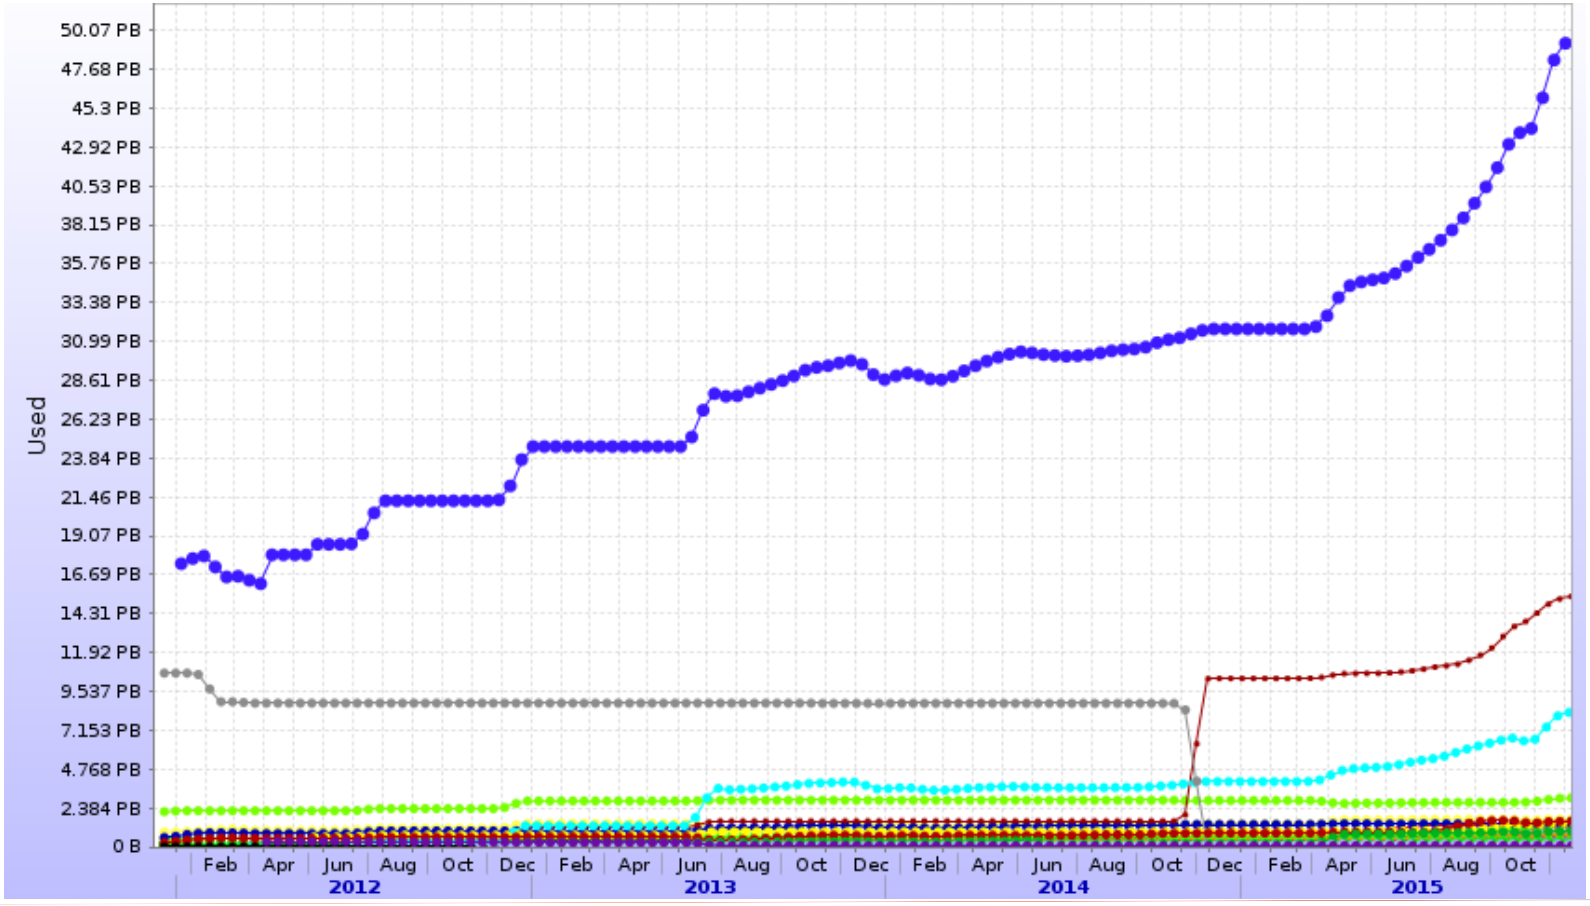

Storage forecast

Estimated disk requirements

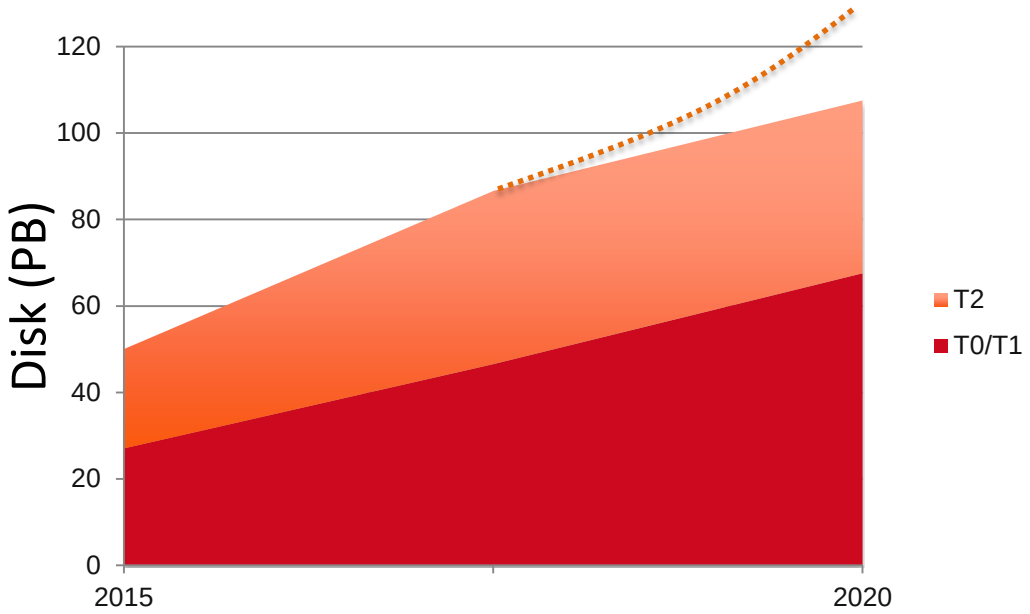

- Similar growth rate for the distributed storage space
  - Trade disks for CPUs at T2s during LS2 (one scenario)
  - Expected to be filled by the regular activity until the start of Run 3

Data management

- To scale to the volume, rates and number of files in Run 3 we need a uniform solution supported across the Grid
- Adopting EOS for both distributed data management and a global namespace
  - Instead of the current set of many independent storage elements
  - To be gradually deployed
    - 7 ALICE sites already run EOS
- EOS in production managing CERN data for all experiments
  - 117 PB of space, 40K disks

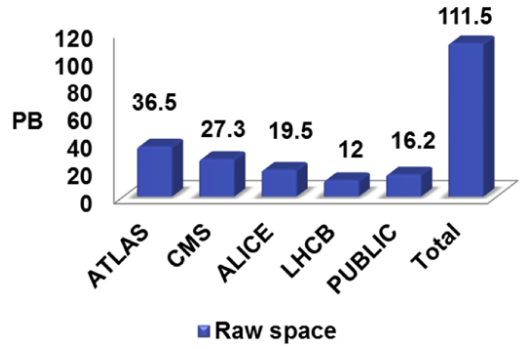

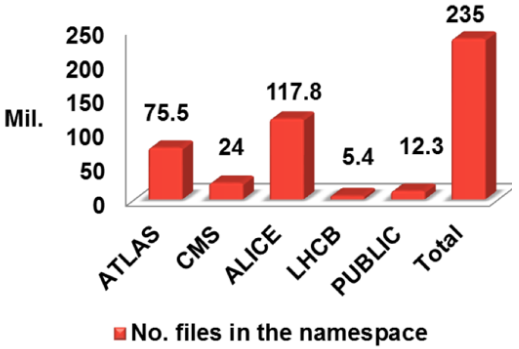

Complexity management

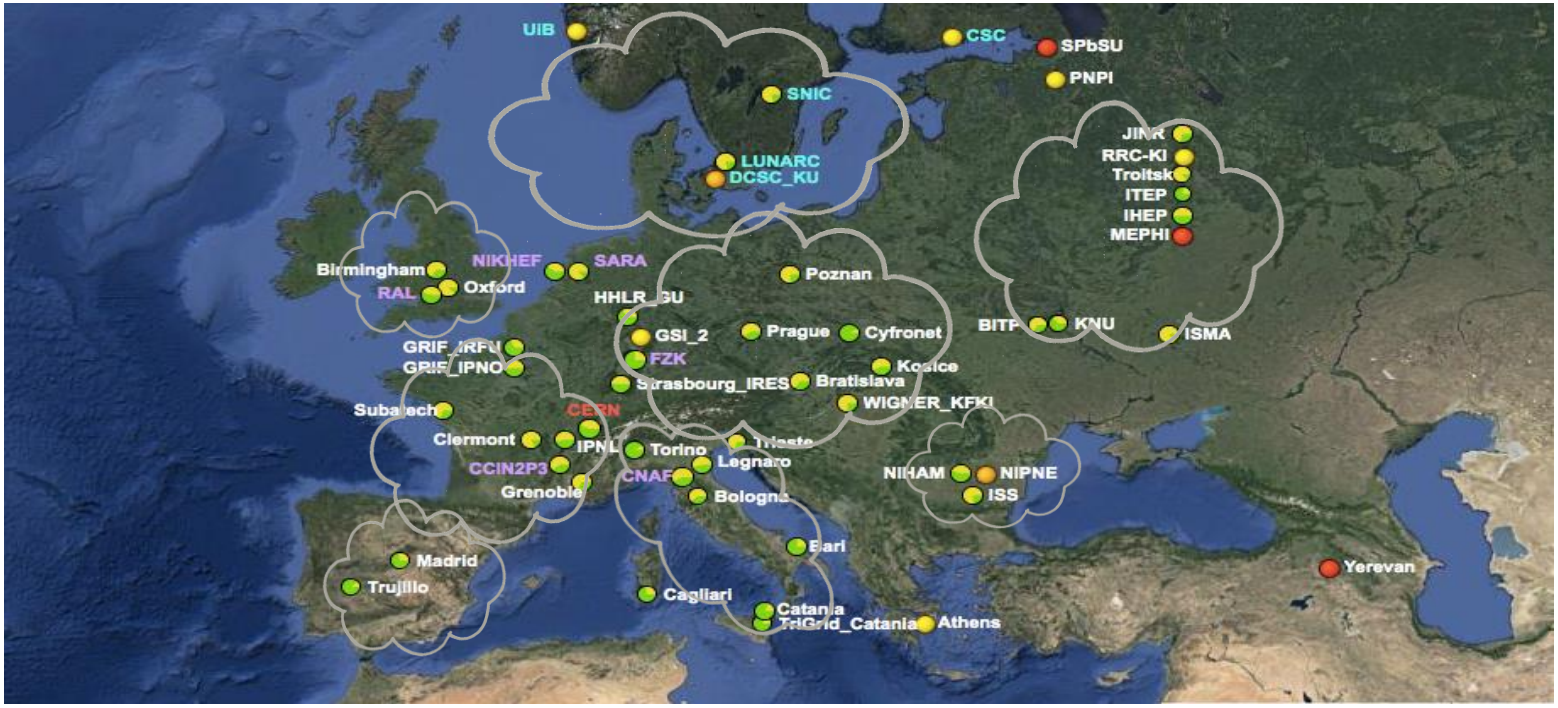

- We need to transform (logically) 100s of individual sites to 10s of clouds/regions
- Each cloud/region should provide reliable data management and sufficient processing capability to simplify scheduling and high level data management

New O2 facility

- 463 FPGAs
  - Detector readout and fast cluster finder
- 100K CPU cores
  - To compress 1.1 TB/s data stream 14x
- 5000 GPUs
  - Reconstruction speed-up
  - 3 CPU [1] + 1 GPU [2] == 28 CPU
- 60 PB disk space
  - Buffer space to allow for a more precise calibration
- The current Grid and more in a single computing center
  - Heterogeneous computing capacity
- Identical software should work in both Online and Offline environments

[1] Intel Sandy Bridge, 2GHz, 8 core, E5-2650
[2] AMD S9000

Computing shares in Run 3

Initial O2 facility as powerful as the entire Grid

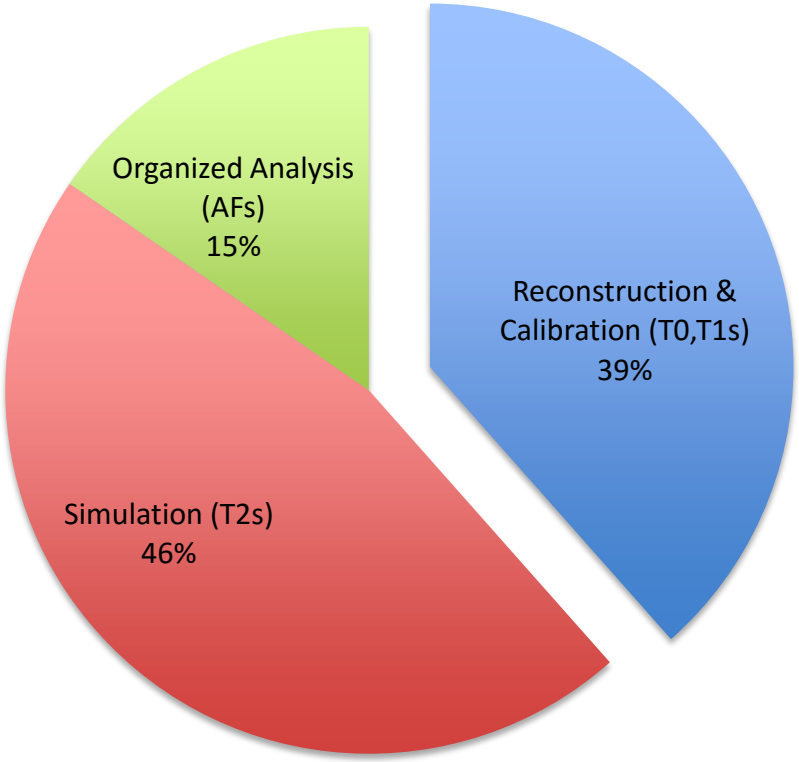

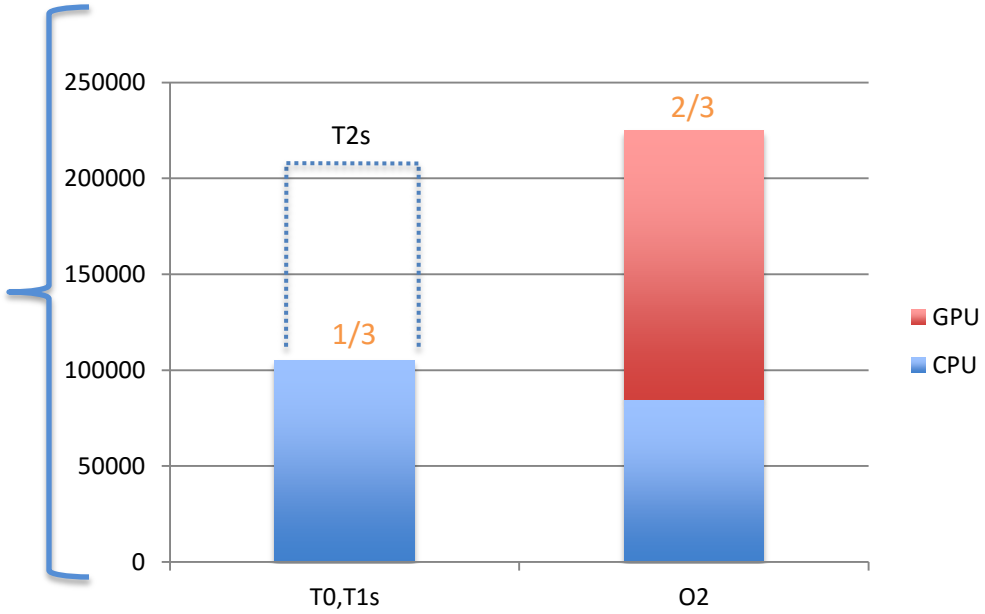

- Relative importance of O2 vs Grid will change in time if Grid resources continue to grow as expected

Analysis facilities

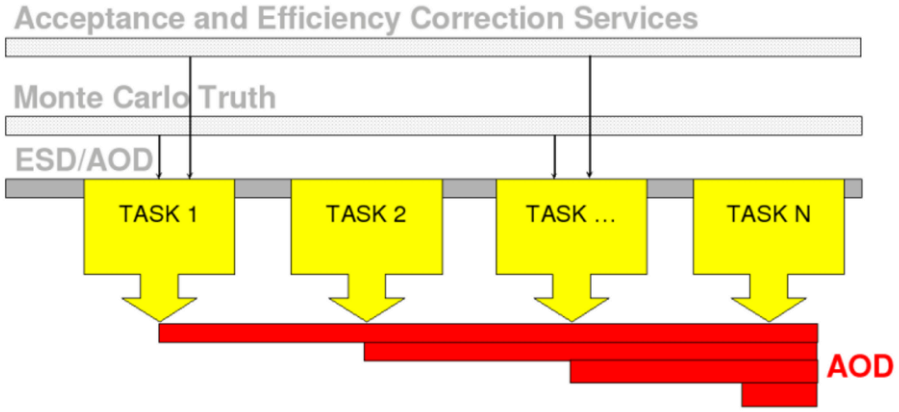

- Analysis is still an I/O bound operation, even after adopting the analysis trains
- Merging stages could be sped up on well connected, high memory machines
  - Leading to shorter turn-around time for entire trains
- Solution is to have dedicated analysis facility/facilities
  - Sites optimized for fast processing of large local datasets
  - Run organized analysis on local data, similar as today's Grid
  - Requires **20-30K CPUs** and **5-10 PB** of very well connected persistent storage space
  - Could be any of the T1s or T2s, but ideally this would be a purpose build facility optimized for such workflow

Analysis IO requirements

Preparing for the worst case scenario

- Legacy separation between CPUs and disks is not going to scale to the target AF size
- We need to bring the data closer to the CPU, eg. by combining disk and CPU in the same box
- And make the workload aware of this location so that jobs read not only from the site where they are running but also from the same rack or even the same machine
- Requires a paradigm change when building the computing centers and changes at several levels in the software stack

Replication policy

Can't aford to …

- Due to a substantial increase in data volume, in RUN 3, there will be only one instance of each raw data file (CTF) stored on disk with a backup on tape
  - EOS with CEPH backend with erasure encoding provides 20% overhead in available storage space with 100% availability in large installations such as CERN
  - In case of data loss, we will restore lost files from the tape
- The size of O2 disk buffer should be sufficient to accommodate one year of CTF data
  - As soon as it is available, the CTF data will be archived to the Tier 0 tape buffer or moved to the Tier 1s

=> Need 20 Gbps to Tier1s during data taking

Disk storage requirements

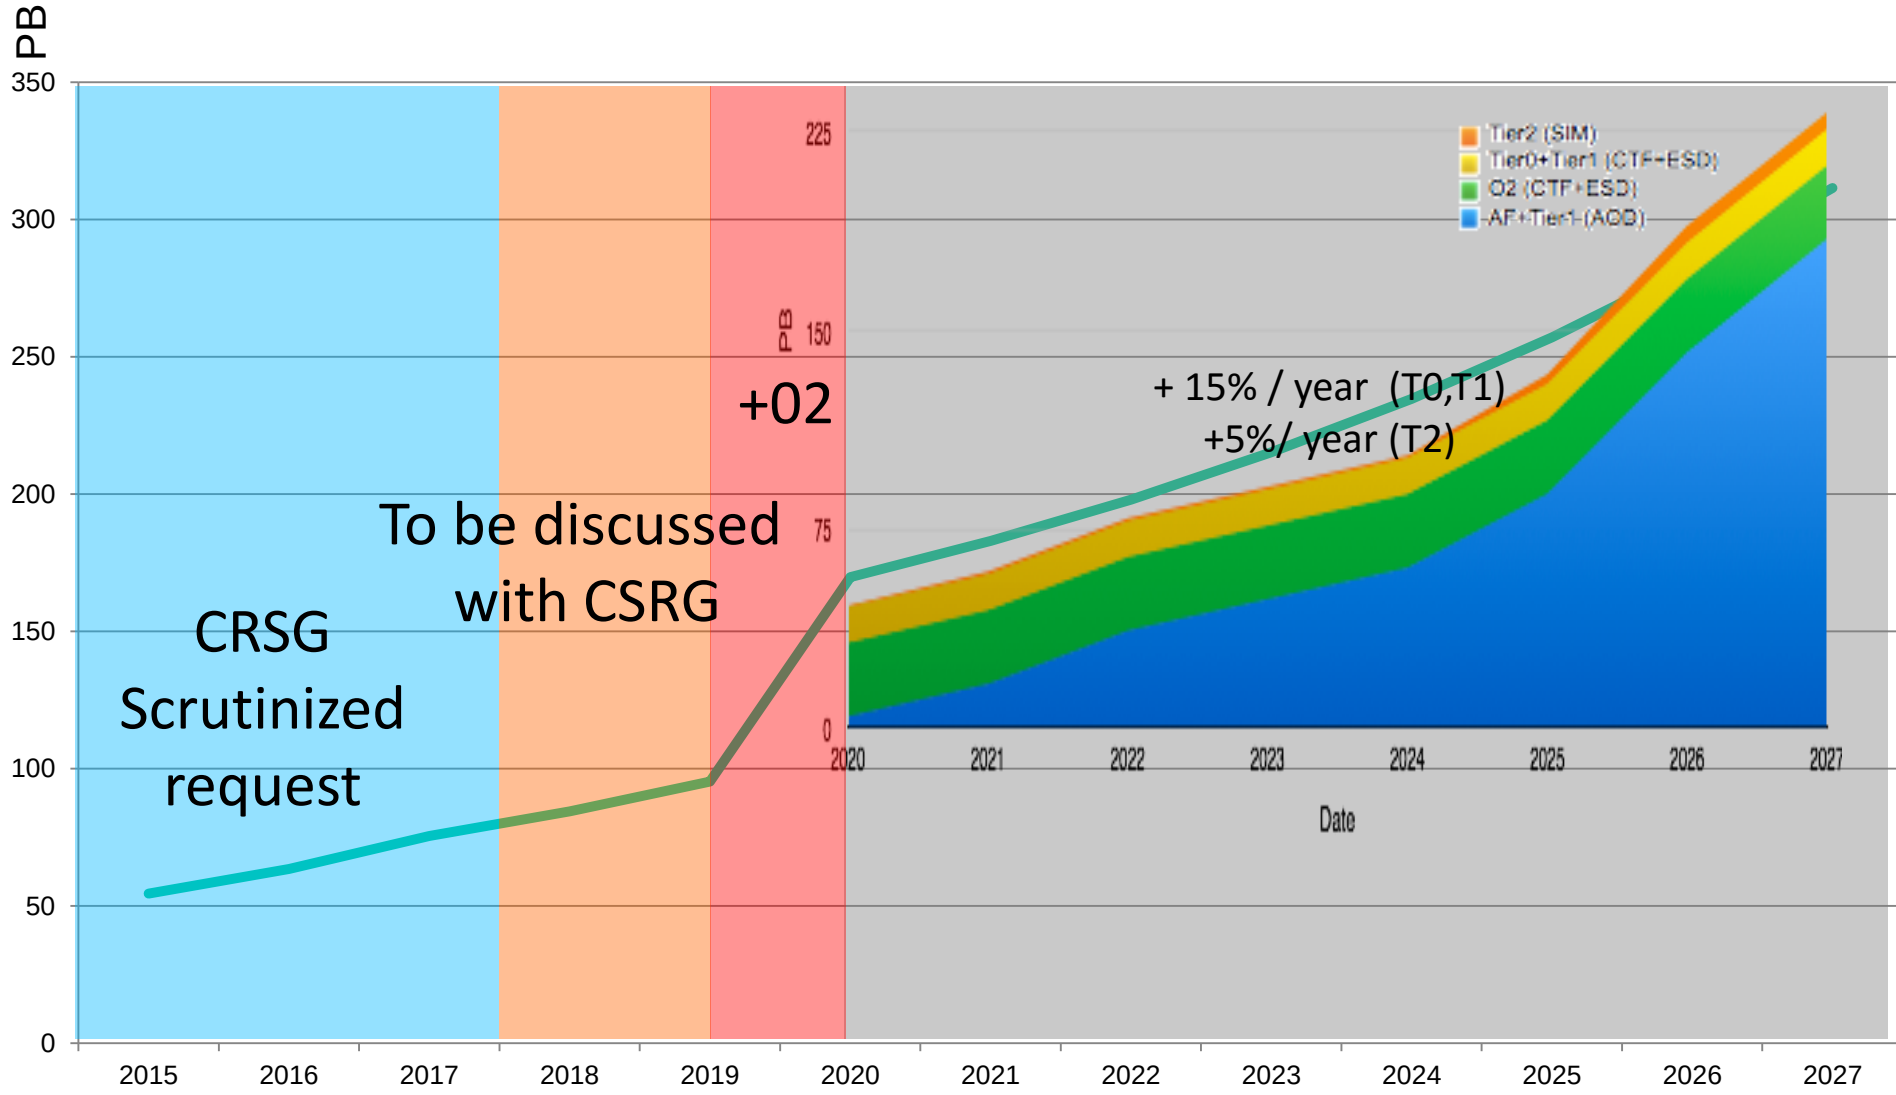

Summary

- We will continue to keep most of the data access local to the site
  - But target purpose-built resources with specific workloads
- Grid resources expected to increase with 20% per year
  - In 2020 we should be able to process 1/3 of raw data on the Grid and 2/3 in O2
  - Limited by the network and storage/archive capacity at T1s
- Analysis Facilities will improve the analysis efficiency and assure guaranteed turnaround time for processing of very large datasets