A Large Ion Collider Experiment

ALICE Status Report

Predrag Buncic

RAW Data Volume (Run1)

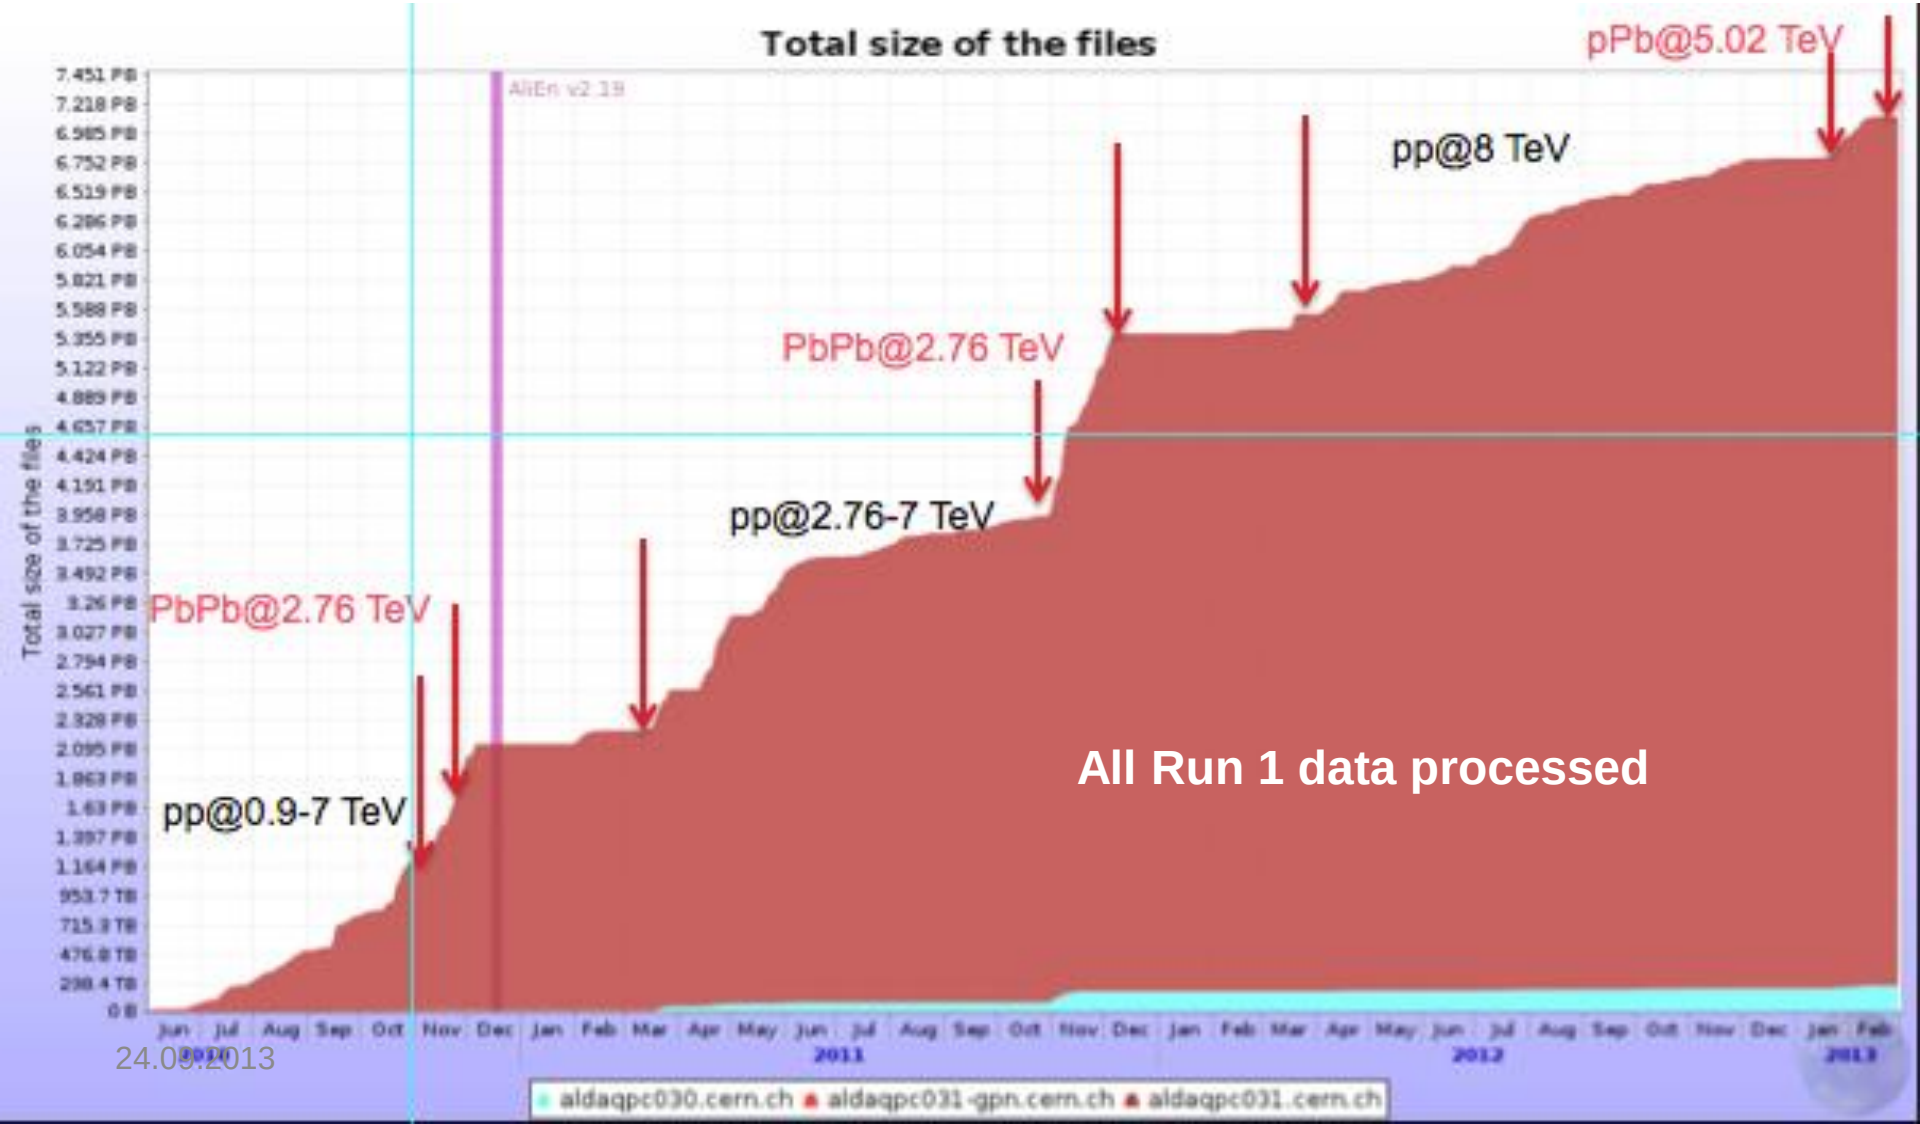

Disk Usage

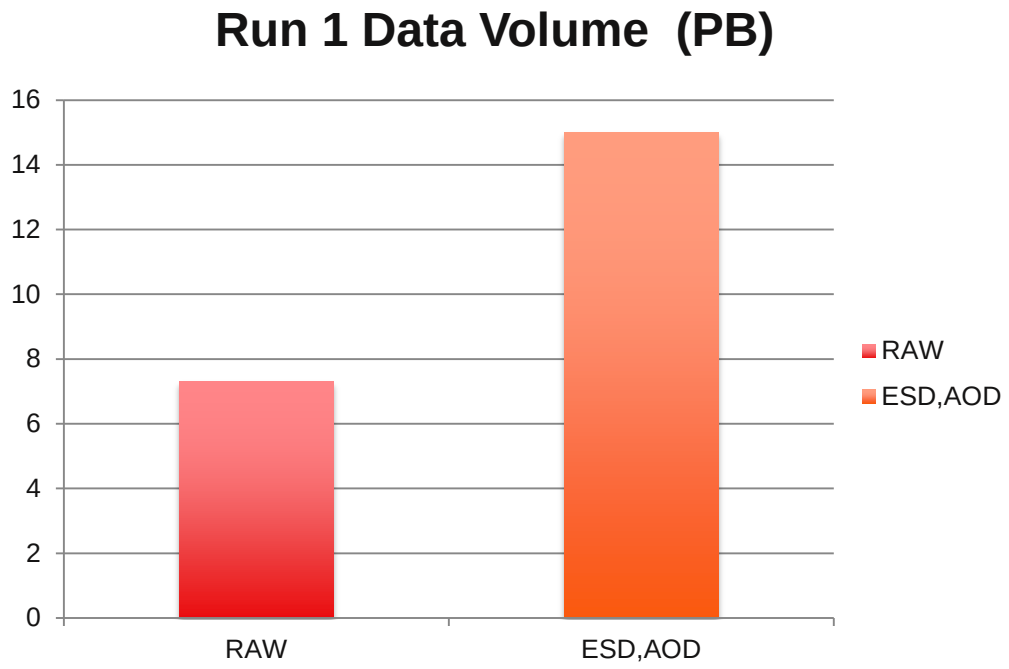

- 7.3 PB of raw data collected during RUN1
- 16 PB of derived data produced (MC, ESD, AOD)
- AOD are replicated 3x, ESD 2x

CPU Usage

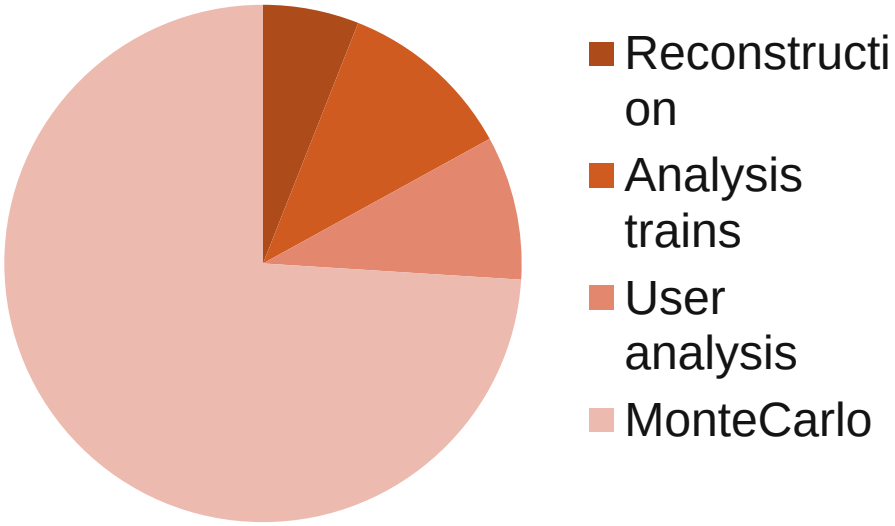

- Average (April-August) CPU usage
  - 6% raw data reconstruction
  - 11% centrally organized analysis (Analysis Train)
  - 9% end user analysis
  - 74 % Monte-Carlo production

Job Profiles

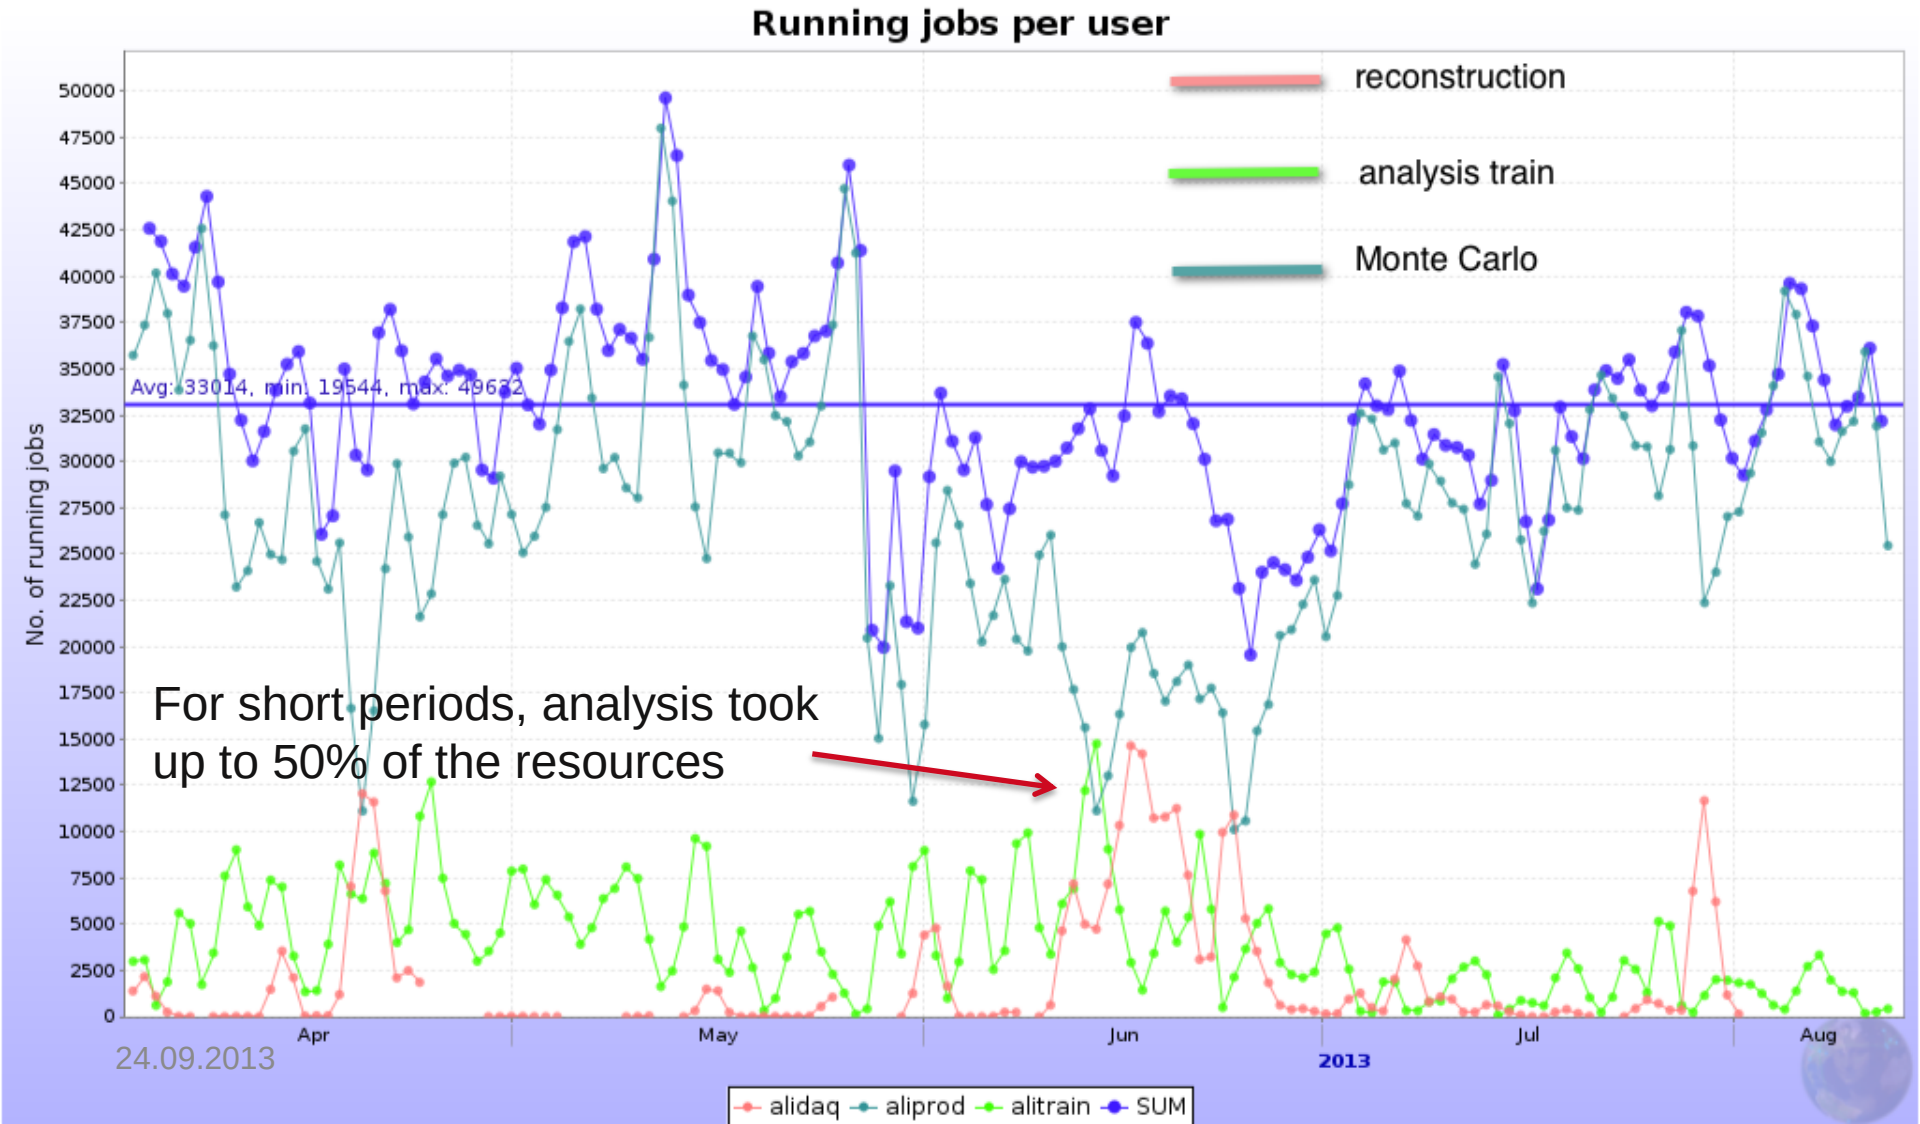

For short periods, analysis took up to 50% of the resources

CPU efficiency

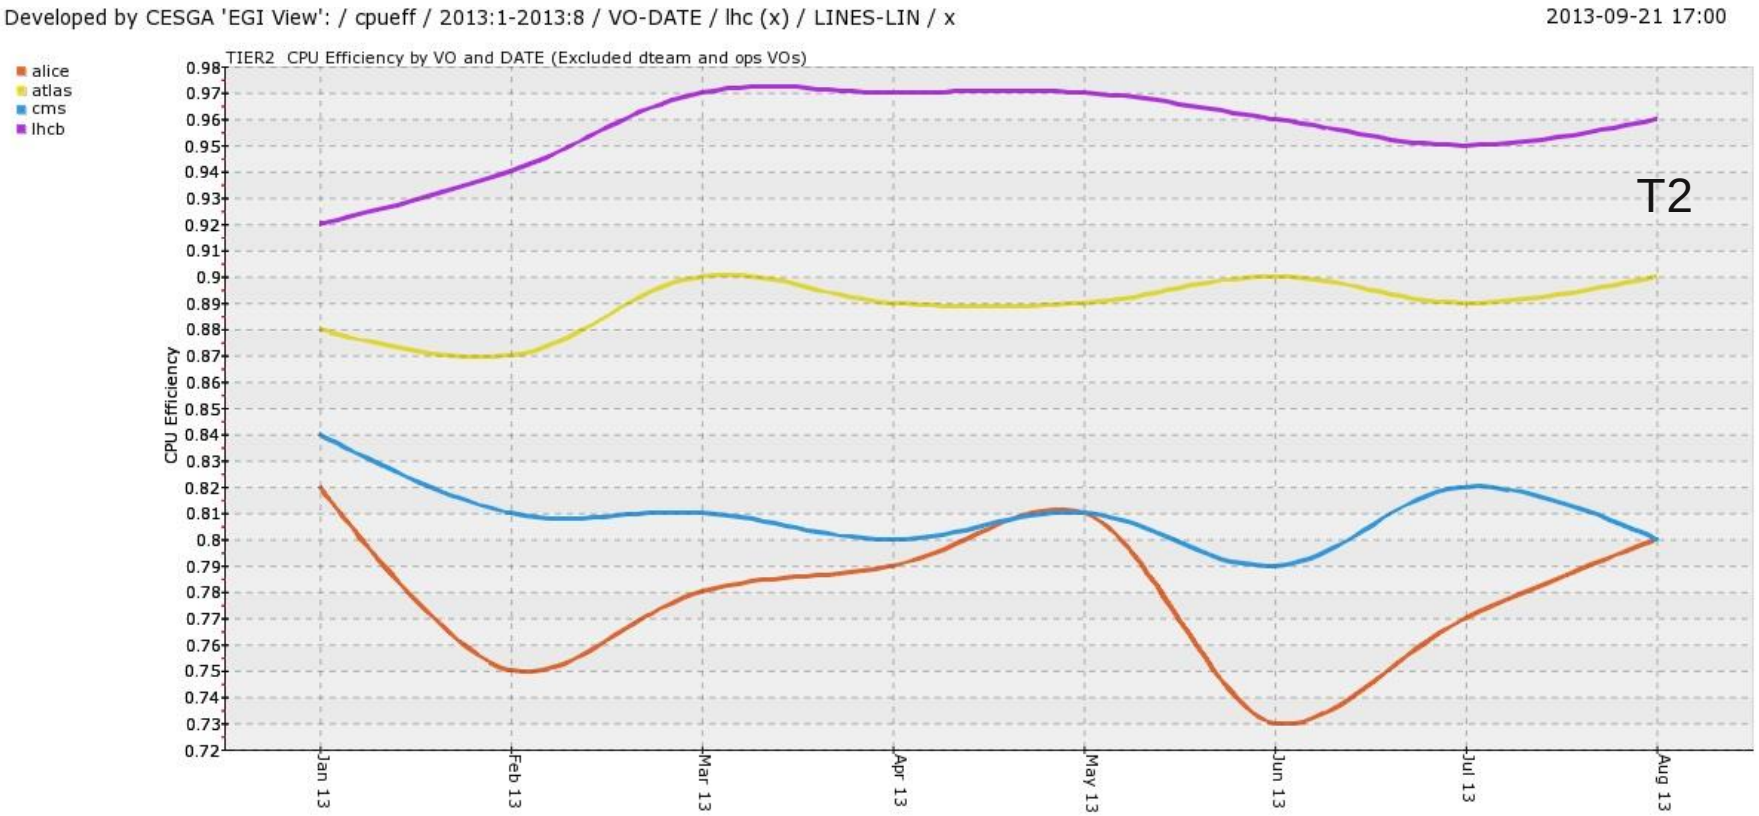

CPU efficiency is still not as good as LHCb and ATLAS, comparable to CMS (~80%)

Dataset popularity service

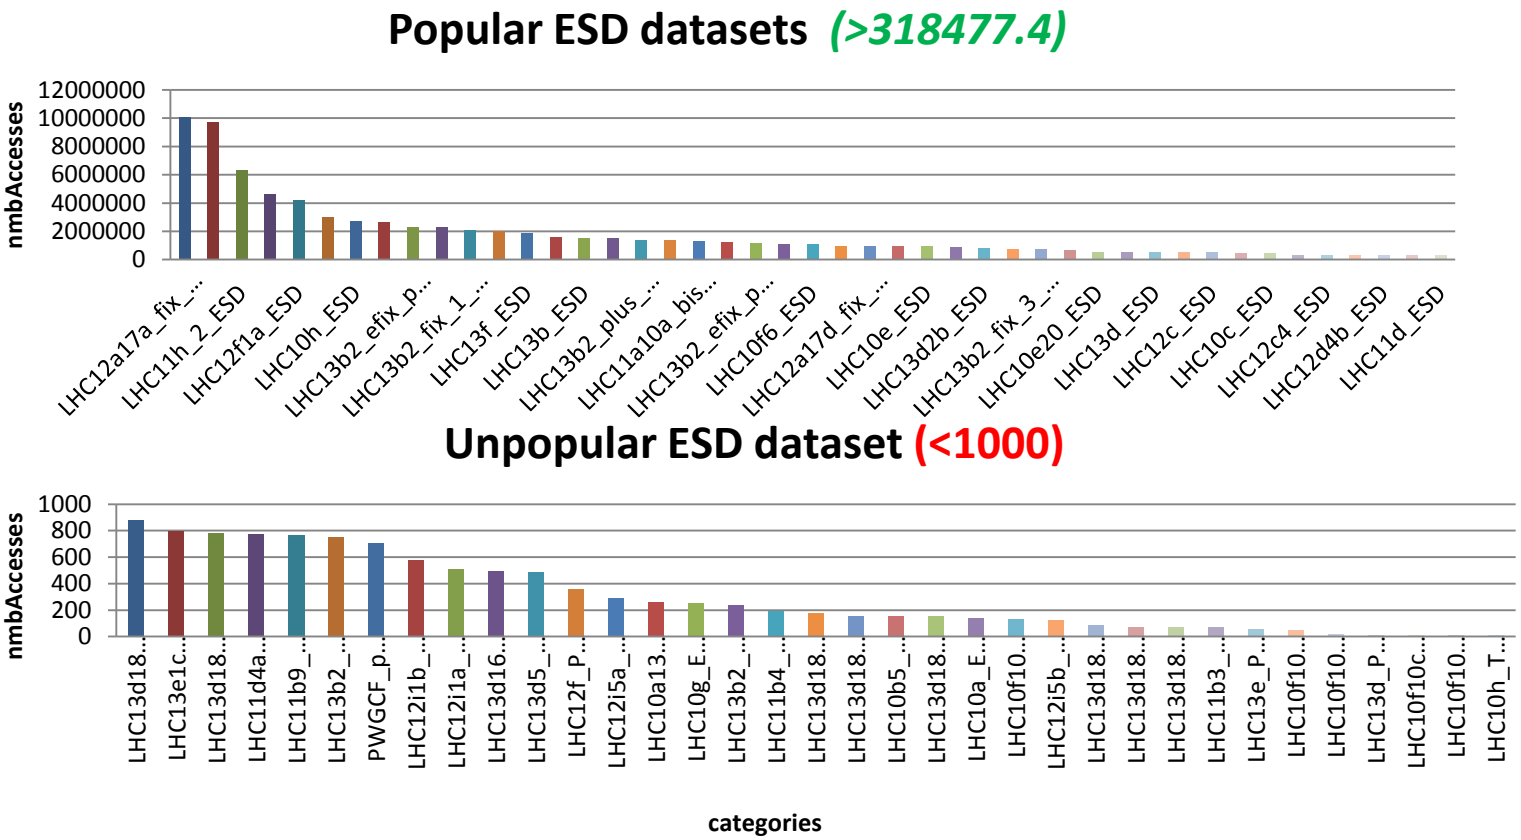

Knowing which datasets are (un)popular allows us to adjust number of file replicas and to eventually reclaim disk space or migrate entire dataset to tape

Improving analysis train performance

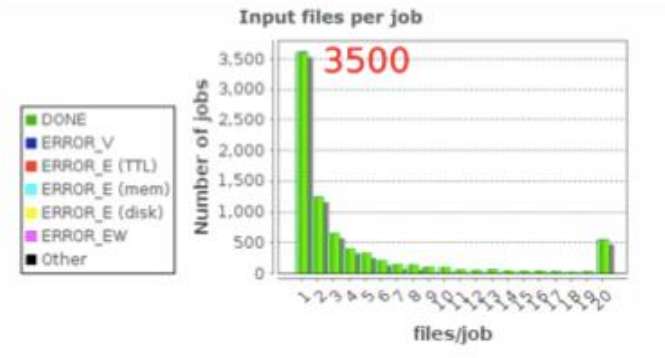

Figure 1: Splitting before clean up.

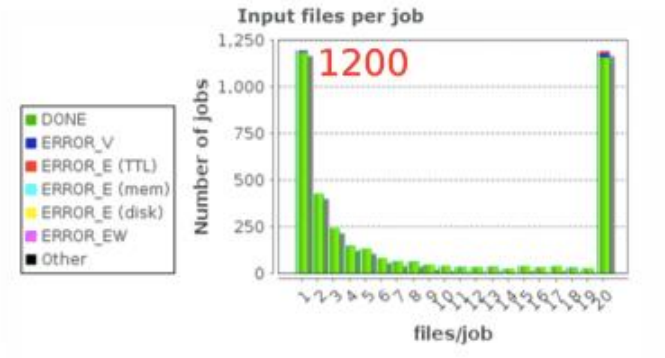

Figure 2: Splitting after clean up.

- LHC11a10a AOD090
- reduced number of jobs from 8000 to 3935
- saving time reduced from 9d 4:38 to 3d 18:33 (merging not included)

By spreading the popular datasets over fewer storage elements we can reduce number of individual jobs needed to process given dataset and reduce overheads of I/O and speed up the final merging step.

CVMFS deployment

- 40 sites installed CVMFS, 11 pending. 5 running in production
- 10700 out of 47777 jobs (22%) running this morning were running on sites where CVMFS was put in production

Upgrade activities

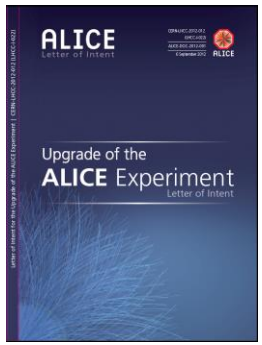

LoI

Panel Report

O^2 Technical Design Report

TDR

- Table of Contents ready
- Setting up Editorial Board

CWG Highlights

- CWG1 - Architecture
  - System Requirements Specifications document due for October

- CWG2 - Tools, guidelines and procedures
  - Reports and presentations templates created
  - Evaluation procedure completed and approved
  - Tools evaluated and recommended: git JIRA
  - Ongoing activities
    - C++ coding guidelines and standard
    - Tools evaluations:
      - Wiki and Web documentation
      - Code and API documentation
    - Licensing (Copyright and distribution of ALICE O^2 software)

CWG Highlights

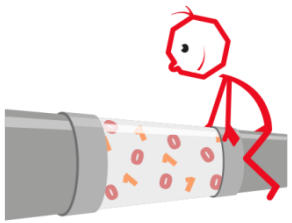

- CWG3 – Dataflow
  - Detector Read-Out
    - Different link protocols under investigation:
      DDL3 (custom, 10Gb/s), Ethernet (10 - 40 Gb/s)
      PCIe over cable (gen2, gen3; 16 - 128 Gb/s)
  - Data Processing
    - Framework prototype to evaluate local and remote flow mechanisms
    - Performance tests on-going
    - Using the open-source packages 0MQ and Apache Zookeeper
  - Simulation
    - OMNeT++ selected as a discrete-event simulation tool, Ptolemy selected for network simulation

CWG Highlights

- CWG4 – Data model
  - In Run 3 we will work with “time frames” (continuous read-out)
    - Collect and process data in well-delimited time intervals
  - Internal note initiated by CWG4, discussed with trigger and detector electronics coordination

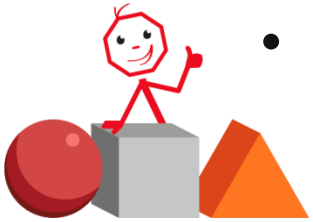

- CWG5 – Computing platforms
  - Defined benchmarks: TPC track finder (compute and memory latency) and TPC track fit (compute)
  - In progress: I/O, memory, IPC
  - Platforms: *Opteron / Xeon, Atom, ARM, AMD Fusion, AMD GPU. NVIDIA GPU, Intel Xeon Phi*
  - Programming models: OpenCL, OpenMP 4, C++11, Vc, (0MQ)

CWG Highlights

- CWG 7 - Calibration
  - Exploring different approaches
    - Synchronous, running on FLPs and EPNs
    - Asynchronous, running on EPN after data has been stored to local disk buffer

- CWG 7 - Reconstruction
  - Concept of calibration/reconstruction with continuous read-out and the expected space-charge distortions demonstrated with a toy model.
  - The tracking based of the CBM experiment's Cellular Automaton algorithms is under evaluation.

CWG Highlights

- CWG8 - Simulation
  - Transport Codes
    - Consolidation of transport code (ongoing)
      - Testing with Geant3, Geant4 and FLUKA (ongoing)
    - Profiling results available by October and used as input for
      - Fast simulation requirements (see below)
      - Tuning (= minimisation of computing time)
  - Fast simulation
    - Revival of barrel tracking parameterisation (ongoing)
  - Monte Carlo Generators
    - Integration of NLO Generators (ongoing)
    - New generators not included in AliROOT as it was the case till now

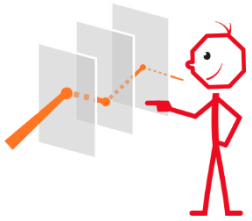

CWG Highlights

- CWG 9 – Visualisation, DQM
  - Reviewing existing ALICE systems (QA analysis, QA Yves, AMORE, event display, etc…)
  - Refine mandate and relationship between Run 2 and Run 3

- CWG 10 – Control, Configuration and Monitoring
  - Topics assigned to members of CWG10 for further conceptual development

- CWG 11 – Software Lifecycle
  - CMake proposed as build system

Summary

- All RUN1 data reconstructed
- The share of MC jobs continues to increase
- Less end user jobs compared to analysis trains
- We continue to work on improving job efficiency by doing proactive data management
- Ongoing Upgrade activities are beginning to deliver the first concrete output useful for TDR due in one year from now