



# ATLAS Data Storage

Dario Barberis

CERN & Genoa University



# Computing Model: central operations

- Tier-0:
  - Copy RAW data to Castor tape for archival
  - Copy RAW data to Tier-1s for storage and reprocessing
  - Run first-pass calibration/alignment (within 24 hrs)
  - Run first-pass reconstruction (within 48 hrs)
  - Distribute reconstruction output (ESDs, AODs & TAGS) to Tier-1s
- Tier-1s:
  - Store and take care of a fraction of RAW data
  - Run "slow" calibration/alignment procedures
  - Rerun reconstruction with better calib/align and/or algorithms
  - Distribute reconstruction output to Tier-2s
  - Keep current versions of ESDs and AODs on disk for analysis
- Tier-2s:
  - Run simulation
  - Run calibration/alignment procedures
  - Keep current versions of AODs on disk for analysis
  - Run user analysis jobs



# Computing Model and Storage

- General principle: each computing centre is responsible for the permanent archival of the data it produces
- Exceptions:
  - RAW data are archived on tape at CERN and at Tier-1s
    - Only case for having 2 archival copies of the same data
  - Simulated data produced by Tier-2s are archived on tape at Tier-1s
- Therefore:
  - RAW: T1D0 at CERN and at Tier-1s. Samples to TOD1 at Tier-1s, Tier-2s and CAF.
  - ESD: T1D0 at production site (CERN for 1st pass, Tier-1s for reprocessing). In addition, TOD1 at production site and paired Tier-1, and samples to Tier-2s and CAF.
  - AOD: T1D0 at production site (CERN for 1st pass, Tier-1s for reprocessing). In addition, TOD1 at all Tier-1s, Tier-2s and CAF
  - Simulated RAW: T1D0 at archival Tier-1.
  - Simulated ESD and AOD: same as for real data.



# Transient Considerations

- Before the beginning of data-taking, there are (almost) only simulated data
- In order to save storage space (as the pledges in the MoU seem to have little to do with reality) we are:
  - Saving only HITS for events without pile-up (initial running!)
  - Saving only RDOs (Raw Data Objects) for events with pile-up
  - Producing ESD and AOD as per Computing Model
- HITS and RDOs go to TOD1 as they are produced, are moved to T1D0, and removed from TOD1 as soon as reconstruction is run
  - It could be 3-6 months!
- Staging of RAW and SIM files from tape for re-processing:
  - We understand that SRM 2.2 will support staging.
  - Files will have to stay in the disk cache until processed (usually days).