

## The data processing in the ALICE experiment

---

**Dagmar Adamová<sup>1</sup>**

*Nuclear Physics Institute, Czech Academy of Sciences*

*Řež near Prague, CZ 25068, Czech Republic*

*E-mail: adamova@ujf.cas.cz*

**for the ALICE Collaboration**

We present a short overview of the RAW data processing chain in the ALICE experiment at the LHC. This represents transfer of RAW from Point2 to the CERN disk buffer and tapes, registration of data in the AliEn catalogue, replication to Tier-1s, reconstruction passes and the tools for analysis. Also, some specific information from the current data taking is included.

*XLVIII International Winter Meeting on Nuclear Physics - BORMIO2010*

*Bormio, Italy*

*January 25–29 2010*

---

<sup>1</sup> Speaker

## 1. Introduction

The ALICE experiment is a dedicated heavy-ion (HI) experiment at the CERN LHC [1] and will take data with a bandwidth of up to 1.25 GB/s (during the HI data taking). It also has its proton-proton (pp) physics program and has been and will be taking pp data, with a bandwidth of up to 500 MB/s. According to the LHC and experiments planning, for one Standard Data Taking Year (STDY) it is foreseen 7 months of pp running, 1 month of HI running and 4 months for a technical stop, for maintenance and upgrades. Altogether, the LHC experiments should be prepared to process data of the order of PetaBytes/SDTY.

To store, process and provide access to such enormous volumes of data, the LHC Computing Grid (LCG) [2] has been built. In addition, each of the LHC experiments including ALICE has developed its own distributed computing infrastructure, a Grid system, which in the case of ALICE is ALIce grid ENvironment (AliEn) [3].

In this paper, we will describe the system of data processing in the ALICE experiment. As an illustration of the performance of the described system, we will mention some results from the processing of data from the first pp collisions at the LHC as observed by ALICE.

## 2. Data Processing chain

ALICE consists of 18 subdetectors that interact with 5 online systems. During data taking, the data is read out by the Data Acquisition (DAQ) system as raw data streams produced by the subdetectors, and is moved and stored over several media. On this way, the raw data is formatted, the events are built, the data is objectified in the ROOT [4] format and then recorded on a local disk. The disk buffer at the ALICE experimental area (called Point2) is operated by DAQ and this storage is only intermediate: the final destination of the collected raw data is the CERN Advanced STORAge system (CASTOR) [5], the permanent data storage (PDS) at the CERN Computing center.

From the intermediate storage at Point2, the raw data is transferred to the disk buffer adjacent to CASTOR at CERN [see Fig.1]. The transfer rates are up to 500MB/s for the pp data taking period and up to 1.25GB/s for the HI data taking period. To overcome possible catastrophic situations like failures of the liaison between DAQ system and CASTOR or of the CASTOR itself, the ALICE on-site disk buffer size is sufficient to store an entire LHC spill of data [6]. After the migration of raw data to CERN, the data is registered in the AliEn catalogue. Not all the data transferred to CERN is copied to tapes of the CERN PDS, this happens automatically only for the data recorded during so called PHYSICS runs. There are over 90 types of runs [7] in ALICE, absolute majority of which are STANDALONE runs. The PHYSICS runs represent the global partition, when several subdetectors are taking data simultaneously. The data from other types of runs stay on the disk buffer for some time and can in principle be reconstructed and analyzed. Unfortunately, ALICE is lacking storage resources in general and tapes in the CERN PDS in particular and not all the data interesting for detectors can be stored on tapes. After the registration in AliEn, the raw data from PHYSICS runs is automatically queued for the Pass1 reconstruction which is performed at CERN (Tier-0 center in the ALICE Grid system [8]). The reconstruction is launched after the run has been processed by the Shuttle (see later in the text) and the migration of raw data to the CASTOR buffer has been completed. To optimize the processing of raw data, the data is split into parts of approximately same size called "data chunks" which are processed separately and in parallel. This makes the processing of data

distinctly faster and more efficient. The size of chunks depends on parameters of the processed run.

The raw data from other types of runs can be reconstructed on a request. The automatic reconstruction is typically completed within a couple of hours after the end of the run (EOR). The output files from the reconstruction are registered in AliEn and are available on the Grid for further processing.

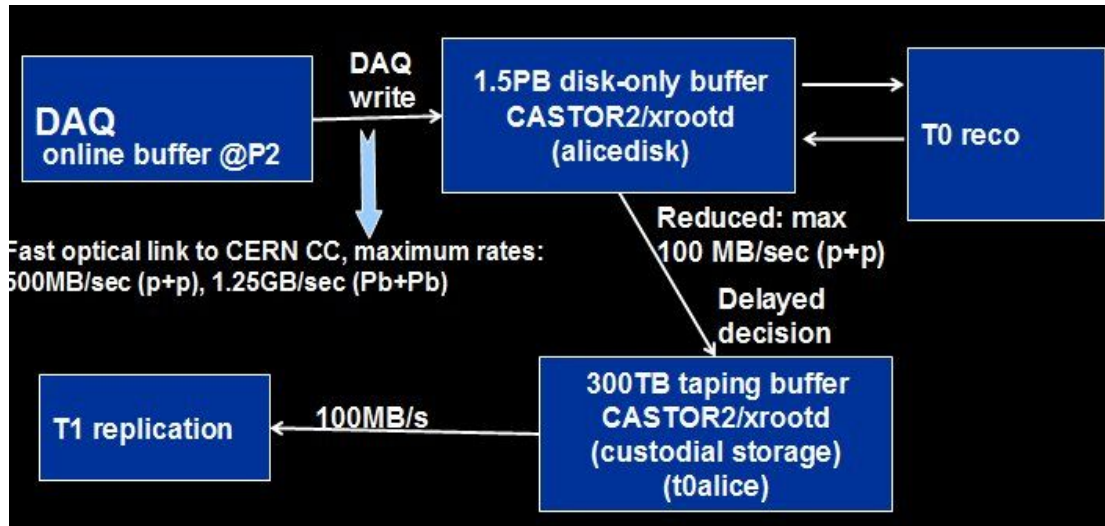


Figure 1. The data processing chain

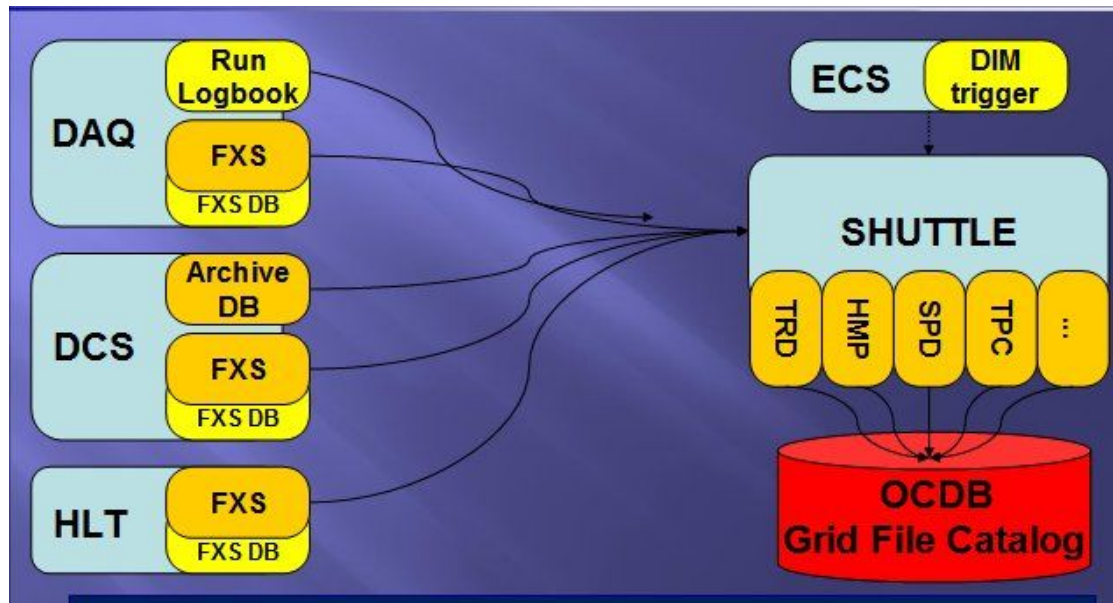


Figure 2. The schema of the SHUTTLE

All the raw data files stored in the CERN PDS are automatically replicated to remote Tier-1 [8] centers.

### 3. SHUTTLE

During data-taking, the subdetectors also produce conditions data which is information about the detector status and environmental variables. This data must be available from the beginning of the reconstruction. For a good performance of the conditions data gathering, in ALICE this data is extracted during data-taking using a software framework named Shuttle [9], see Figure 2. It collects data produced by online systems, feeds this data into detector specific procedures called pre-processors, which run on machines in the DAQ system and produce the conditions data objects, then converts the output into ROOT format and stores the output in the Grid Offline Conditions Data Base (OCDB). The Shuttle is launched automatically after the EOR and processes particular run types, mostly PHYSICS runs. Only the runs processed by the Shuttle can be reconstructed.

### 4. Monitoring and Information System

All the stages of the data processing chain described above are extensively monitored using a Java-based monitoring and control framework called MonALISA [10]. In general, the complexity of the ALICE Computing Grid infrastructure requires permanent control and management of an extensive set of services, which is performed using automated agents. An example of one of the ALICE MonALISA monitoring pages is shown on Figure 3. It shows a part of the table with information about the status of the processing of raw data taken by ALICE during the first pp collisions at the LHC in November/December 2009. You can find a very complete information here including:

Run number	Chunks		Events (reco)	AliEn job		QA	Software versions		Partition	Pass	Output dir	Comment
	OK/All	%		Job ID	Err		ROOT	ALIROOT				
104792	18 / 18	100%	74,378	39299982	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104792/ESDs/pass2/	900 GeV, - 0	
104618	3 / 3	100%	1,629	39299983	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104618/ESDs/pass2/	900 GeV, 0 0	
104439	3 / 3	100%	14,207	39261499	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104439/ESDs/pass2/	900 GeV, - 0	
104321	13 / 13	100%	88,208	39261500	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104321/ESDs/pass2/	900 GeV, - 0	
104320	3 / 3	100%	2,395	39261527	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104320/ESDs/pass2/	900 GeV, - 0	
104316	3 / 3	100%	7,231	39300007	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104316/ESDs/pass2/	900 GeV, - 0	
104315	3 / 3	100%	6,398	39261533	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104315/ESDs/pass2/	900 GeV, - 0	
104160	3 / 3	100%	9,161	39261536	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104160/ESDs/pass2/	900 GeV, ++	
104159	3 / 3	100%	1,854	39261537	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104159/ESDs/pass2/	900 GeV, ++	
104157	3 / 3	100%	11,990	39261547	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104157/ESDs/pass2/	900 GeV, ++	
104155	3 / 3	100%	5,637	39300026	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104155/ESDs/pass2/	900 GeV, ++	
104083	3 / 3	100%	749	39261549	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104083/ESDs/pass2/	900 GeV, ++	
104080	3 / 3	100%	616	39261560	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104080/ESDs/pass2/	900 GeV, ++	
104073	3 / 3	100%	911	39261559	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104073/ESDs/pass2/	900 GeV, ++	
104070	2 / 2	100%	1,921	39261561	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104070/ESDs/pass2/	900 GeV, ++	
104068	3 / 3	100%	1,615	39300027	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104068/ESDs/pass2/	900 GeV, ++	
104065	3 / 3	100%	8,387	39261572	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104065/ESDs/pass2/	900 GeV, ++	
104044	3 / 3	100%	5,676	39261579	✓	v5-25-04-1	v4-17-Rev-24	LHC09d	2	/alice/data/2009/LHC09d/000104044/ESDs/pass2/	900 GeV, ++	
<b>TOTAL</b>	<b>394/402</b>	<b>98%</b>	<b>1,024,058</b>	<b>43 jobs</b>	<b>9 err</b>							

Figure 3. Monitoring with MonALISA: Raw data reconstruction status

- number of each processed run
- number of processed data chunks
- percentage of successfully reconstructed data
- number of reconstructed events
- reconstruction job ID
- specification of software packages used for the reconstruction

- production partition
- path in the AliEn catalogue where the output files from this processing can be found

The additional features of this display include the link to the complete information about each run as recorded in the DAQ Electronic logbook (when clicking on the run number) or the link to the information about the progress of the reconstruction (when clicking on the job ID) and also the possibility to browse the output files inventory (when clicking on the AliEn path).

## 5. Concluding remarks

Starting on November 23rd 2009, the LHC has been delivering stable proton beams and pp collisions, in 2009 at the  $\sqrt{s}=900$  GeV and since March 30st 2010 at  $\sqrt{s}=7$  TeV. The ALICE data processing chain has been performing excellently during both mentioned periods, delivering the reconstructed data for the further analysis immediately after the runs were finished. This contributed to a remarkable activity of the ALICE experiment:

After the very first collisions in ALICE on November 23rd 2009, the first 284 events of pp collisions at  $\sqrt{s}=900$  GeV were immediately successfully reconstructed and meticulously analyzed with such enormous effort that in 5 days after the data was recorded, the paper with the Physics results was written and another 3 days later it was accepted for publication in the European Journal of Physics C [11]. It was the first ever published paper with Physics results from collisions at the LHC.

This fast progress in delivering the Physics results would of course be unthinkable without a steady work on the ALICE software framework for reconstruction, simulation and analysis of the data, AliRoot [12]. For understanding and improvement of the AliRoot performance, as well as for understanding the behavior of the ALICE detectors, the fast feedback given by the offline reconstruction is essential.

Currently (second week of April 2010), a draft of the ALICE paper on the charged particle multiplicity measurement in pp collisions at  $\sqrt{s}=7$  TeV is to appear in a couple of days. It represents processing and analysis of 12.5 mil. of events with the total size of raw data of about 6TB within 2 weeks. This fast progress, similar to the processing of the data from the 900 GeV pp collisions, illustrates the excellent functioning of the ALICE Offline framework and infrastructure, in these first months of the LHC operations.

### Acknowledgements

I highly appreciate the critical reading and help with the manuscript by Latchezar Betev (CERN). The work was supported by the MSMT CR contracts No. 1P04LA211 and LC 07048.

### References

- [1] ALICE Collaboration: *ALICE: Physics Performance Report, Volume I, J. Phys. G30* (2004) 1517; *Volume II, J. Phys. G32* (2006) 1295.
- [2] *LHC Computing Grid: Technical Design Report*, <http://lcg.web.cern.ch/LCG/tdr/>
- [3] P. Saiz et al., *AliEn - ALICE environment on the GRID, Nucl. Instrum. Meth. A502* (2003) 437; <http://alien.cern.ch/twiki/bin/view/AliEn/Home>
- [4] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework, Nucl. Instrum. Meth. A389* (1997) 81; <http://root.cern.ch>

- [5] G. Cancio et al., *Data Management Evolution and Strategy at CERN*, to appear in proceedings of the 17th Int. Conf. CHEP 2009, Prague, March 21-27, 2009.
- [6] E. Divia et al., *The ALICE online data storage system*, to appear in proceedings of the 17th Int. Conf. CHEP 2009, Prague, March 21-27, 2009.
- [7] *ALICE Experiment Control System - ECS: Run types*, <http://alice-ecs.web.cern.ch/alice-ecs/>
- [8] *ALICE Experiment Computing TDR*, <http://aliceinfo.cern.ch/Collaboration/Documents/TDR/Computing.html>.
- [9] J.F. Grosse-Oetringhaus et al., *The ALICE Online-Offline Framework for the Extraction of Conditions Data*, to appear in proceedings of the 17th Int. Conf. CHEP 2009, Prague, March 21-27, 2009.
- [10] C. Grigoras et al., *Automated agents for management and control of the ALICE Computing Grid*, to appear in proceedings of the 17th Int. Conf. CHEP 2009, Prague, March 21-27, 2009.
- [11] K. Aamodt et al. (ALICE Collaboration), *First proton--proton collisions at the LHC as observed with the ALICE detector: measurement of the charged particle pseudorapidity density at  $\sqrt{s} = 900$  GeV*, *Eur. Phys. J.* **C65** (2010) 111.
- [12] <http://aliceinfo.cern.ch/Offline>