Chapter 9

Computing and Resources

This section describes the dataflow of the LHCb computing model for all stages in the processing of the real and simulated LHCb events [22]. The roles of the various Tier centres are discussed and the distribution of the processing load and storage are outlined.

There are several phases in the processing of event data. The various stages normally follow each other in a sequential manner, but some stages may be repeated a number of times. The workflow presented here reflects the present understanding of how to process the data. A schematic of the logical dataflow is shown in figure 9.1 and is described in more detail below.

The raw data from the detector is produced via the Event Filter farm of the online system. The first step is to collect data, triggering on events of interest. The raw data are transferred to the CERN Tier 0 centre for further processing and archiving. The raw data, whether real or simulated, must then be reconstructed to form physical quantities such as the calorimeter clusters needed to provide the energy of electromagnetic and hadronic showers, tracker hits to be associated to tracks whose position and momentum are to be determined. Information about particle identification (electron, photon, π^0, charged hadrons, muon) is also reconstructed from the appropriate sub-systems. The

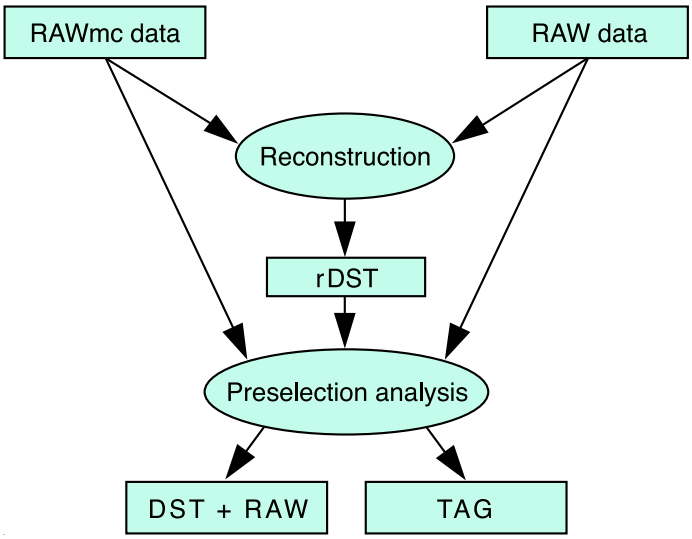

Figure 9.1: The LHCb computing logical dataflow model.

event reconstruction results in the generation of new data, the *Data Summary Tape* (DST). Only enough data will be stored in the DST to allow the physics pre-selection algorithms to be run at the next stage: this is known as a reduced DST (rDST). The first pass of the reconstruction will happen in quasi-real time. It is planned to reprocess the data of a given year once, after the end of data taking for that year, and then periodically as required. This is to accommodate improvements in the reconstruction algorithms and to make use of improved determinations of the calibration and alignment of the detector in order to generate new improved rDST information.

The rDST is analysed in a production-type mode in order to produce streams of selected events for further individual analysis. This activity is known as *stripping*. The rDST information is used to determine the four-momentum vectors corresponding to the measured particles, to locate primary and secondary vertices and reconstruct composite particles such as B candidates. A pre-selection algorithm will be provided for each channel of interest. Since these algorithms use tools that are common to many different physics analyses they are run in production-mode as a first step in the analysis process. The events that pass the selection criteria will be fully re-reconstructed, recreating the full information associated with each event. The output of the stripping stage will be referred to as the (full) DST and contains more information than the rDST. Before being stored, the events that pass the selection criteria will have their raw data added in order to have as detailed event information as needed for the analysis. An event tag collection will also be created for faster reference to selected events. The tag contains a brief summary of each event's characteristics as well as the results of the pre-selection algorithms and a reference to the actual DST record. The event tags are stored in files independent of the actual DST files. It is planned to run this production-analysis phase 4 times per year: once with the original data reconstruction; once with the re-processing of the raw data, and twice more, as the selection cuts and analysis algorithms evolve.

The baseline LHCb computing model is based on a distributed multi-tier regional centre model. It attempts to build in flexibility that will allow effective analysis of the data whether the Grid middleware meets expectations or not. A schematic of the LHCb computing model is given in figure 9.2.

CERN is the central production centre and will be responsible for distributing the raw data in quasi-real time to the Tier-1 centres. CERN will also take on the role of a Tier-1 centre. Six additional Tier-1 centres have been identified: CNAF (Italy), FZK (Germany), IN2P3 (France), NIKHEF (The Netherlands), PIC (Spain) and RAL (United Kingdom): there is also a number of Tier-2 computing centres. CERN and the Tier-1 centres will be responsible for all the production-processing phases associated with the real data. The raw data will be stored in its entirety at CERN, with another copy distributed across the other 6 Tier-1 centres. The second pass of the full reconstruction of the raw data will also use the resources of the LHCb online farm. As the production of the stripped DSTs will occur at these computing centres, it is envisaged that the majority of the distributed analyses will be performed at CERN and at the Tier-1 centres. The current year's stripped DST will be distributed to all centres to ensure load balancing. It should be noted that although no user analysis is envisaged at the Tier-2 centres in the baseline model presented, it should not be proscribed, particularly for the larger Tier-2 centres.

It is expected that the reconstruction and the first stripping of data at CERN and at the Tier-1 centres will follow the production in quasi real-time, with a maximum delay of a few days. The DST output of the stripping will remain on disk for analysis and be distributed to all other Tier-1 centres and CERN, whilst the raw and rDST will be migrated to the mass storage system, MSS.

2008 JINST 3 S08005

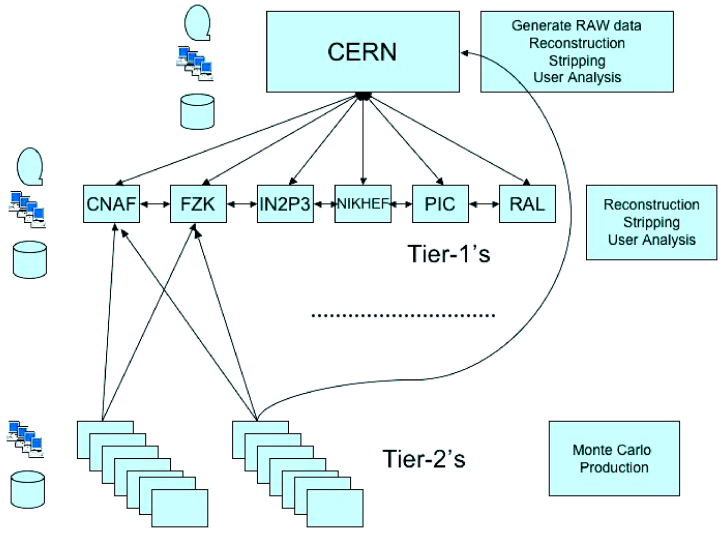

Figure 9.2: Schematic of the LHCb Computing Model.

Table 9.1: Projected resource usage in 2008 at CERN, the Tier-1 and Tier-2 centres. A 1 GHz PIII processor is equivalent to 400 KSI2k.

Site	CPU (MSI2k.years)	Disk (TB)	Tape (TB)
CERN	0.36	350	631
Tier-1	1.77	1025	860
Tier-2	4.55	-	-

The re-processing of the data will occur over a two-month period. During this process the raw data will need to be accessed from the MSS both at CERN and the Tier-1 centres. The CPU resources available at the pit allow a significant fraction of the total re-processing and perhaps the subsequent stripping to be performed there. Hence at CERN there is an additional complication that the raw data will also have to be transferred to the pit; similarly, the produced rDST will have to be transferred back to the CERN computing centre. To enable later stripping, it is necessary to distribute a fraction of the rDST produced at CERN during this re-processing to the Tier-1 centres; this is a consequence of the large contribution from the online farm.

The (two) stripping productions outside of the reconstruction or of the re-processing of the data will be performed over a one-month period. Both the raw and the rDST will need to be accessed from the MSS to perform this production. The produced stripped DSTs will be distributed to all production centres.

The Tier-2 centres will be primarily MonteCarlo production centres, with both CERN and the Tier-1 centres acting as the central repositories for the simulated data. The MonteCarlo production is expected to be an ongoing activity throughout the year and is the mainstay of the Tier-2 centres. The whole of the current year's MonteCarlo production DST will be available on disk at CERN and another 3 copies, on disk, distributed amongst the other 6 Tier-1 centres.

The 2008 resource requirements needed for the LHCb computing model at CERN and integrated across the Tier-1 centres and the Tier-2 centres are given in table 9.1.

2008 JINST 3 S08005