

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 29

Type: **oral presentation**

The ATLAS Data Flow system for the second LHC run

Tuesday, April 14, 2015 2:30 PM (15 minutes)

After its first shutdown, LHC will provide pp collisions with increased luminosity and energy. In the ATLAS experiment the Trigger and Data Acquisition (TDAQ) system has been upgraded to deal with the increased event rates. The Data Flow (DF) element of the TDAQ is a distributed hardware and software system responsible for buffering and transporting event data from the Readout system to the High Level Trigger (HLT) and to the event storage. The DF has been reshaped in order to profit from the technological progress and to maximize the flexibility and efficiency of the data selection process.

The updated DF is radically different from the previous implementation both in terms of architecture and expected performance. The pre-existing two level software filtering, known as L2 and the Event Filter, and the Event Building are now merged into a single process, performing incremental data collection and analysis. This design has many advantages, among which are: the radical simplification of the architecture, the flexible and automatically balanced distribution of the computing resources, the sharing of code and services on nodes. In addition, logical farm slicing, with each slice managed by a dedicated supervisor, has been dropped in favour of global management by a single farm master operating at 100 kHz.

The Data Collection network, that connects the HLT processing nodes to the Readout and the storage systems has evolved to provide network connectivity as required by the new Data Flow architecture. The old Data Collection and Back-End networks have been merged into a single Ethernet network and the Readout PCs have been directly connected to the network cores. The aggregate throughput and port density have been increased by an order of magnitude and the introduction of Multi Chassis Trunking significantly enhanced fault tolerance and redundancy.

We will discuss the design choices, the strategies employed to minimize the data-collection latency, the results of scaling tests done during the commissioning phase and the operational performance after the first months of data taking.

Primary author: HAUSER, Reiner (Michigan State University (US))

Presenter: HAUSER, Reiner (Michigan State University (US))

Session Classification: Track 1 Session

Track Classification: Track1: Online computing