Meeting of the Architecture Working Group with F.Touchard
27th September

Summary by Saul Gonzalez

Main topics discussed or raised during the meeting:

1. EF dependence in offline and PESA
2. SFI to sub-farm mapping
3. Sub-farm to SFO mapping
4. SFO to mass storage: Event classification
5. SFO to mass storage: Who is responsible?
6. Confusion regarding EF-related parameters (rates, etc.)
7. Online services

Summary:

The EF will use offline selection software and algorithms. This implies that the EF will have a very strong dependence in offline software. Although much progress has been achieved in the past months, the boundaries between the EF/PESA/offline are still unclear.

The EF interface to DataCollection is the SFI, while the SFO pulls events from the EF sub-farms. The EF 'dialog' with the DAQ is very simple because the EF deals with full events: Events are put in memory, analyzed, and augmented (with EFresult); the Processing Task notifies the SFO that the event is ready, and the SFO pulls the event from the sub-farm.

In the present prototypes, there is a correspondence between the SFI and an EF sub-farm. It is possible to introduce a network between all SFIs and all sub-farms to allow load balancing, however, a supervisor would be needed to handle sub-farm assignment. Event monitoring may be easier in a LAN-based SFI; presently monitoring is handled by a proxy process pulling events from the sub-farm, the SFI, or the SFO. It was concluded that there is a need for 'use cases' of filtered or unfiltered events at the EF level.

There was a discussion of what constitutes a 'sub-farm' from an online point of view. It is presently not clear if it is a single control entity; Discussions are ongoing with the online.

It is foreseen that multiple EF sub-farms can talk to a single SFO. Possible use cases are, e.g., a stream of monitoring events or a discovery stream going to a single SFO for priority dispatching. It is planned to store events as bytes-stream to disk arrays attached to the SFO. These disk arrays must be able to handle anywhere from 20 minutes to one day's worth of data. The responsibility for these disks (TDAQ or offline) is still to be defined. It was pointed out that detector groups might want to stream events for hardware debugging, particularly at start-up. It was agreed that use cases needed to be collected for detector activities during early running.

During the discussion, there was some confusion with the system parameters being used by different people (e.g., # of SFIs, total number of processors, event building rate). The confusion led to the conclusion that the architecture needs to be scaleable. It was also agreed that an AWG session would be scheduled with a PESA representative to clarify the status of the PESA-dependent input parameters.

Presently, processing tasks do not use online services directly. It still needs to be clarified how databases will be handled in the context of Athena. Histogram use cases also need to be developed to understand the load on the control network. In addition, the control and supervision tasks need to be unified across the HLT.