



Contribution ID: 108

Type: oral presentation

The LHCb Data Acquisition and High Level Trigger Processing Architecture

Thursday, April 16, 2015 9:30 AM (15 minutes)

The LHCb experiment at the LHC accelerator at CERN collects collisions of particle bunches at 40 MHz. After a first level of hardware trigger with output of 1 MHz, the physically interesting collisions are selected by running dedicated trigger algorithms in the High Level Trigger (HLT) computing farm. This farm consists of up to roughly 25000 CPU cores in roughly 1600 physical nodes each equipped with 2 TB of local storage space.

This work describes the LHCb online system with an emphasis on the developments implemented during the current long shutdown (LS1). We will elaborate the architecture to treble the available CPU power of the HLT farm and the technicalities to determine and verify precise calibration and alignment constants which are fed to the HLT event selection procedure. Precise calibration and alignment constants are determined and verified in a separate data acquisition activity as soon as data from particle collisions are delivered by the LHC collider. We will describe how the constants are fed into a two stage HLT event selection facility using extensively the local disk buffering capabilities on the worker nodes. With the installed disk buffers, the installed CPU can be used during periods of up to ten days without beams. These periods in the past accounted to more than 70 % of the total time.

Primary author: FRANK, Markus (CERN)

Co-authors: JOST, Beat (CERN); GASPAR, Clara (CERN); NEUFELD, Niko (CERN)

Presenter: FRANK, Markus (CERN)

Session Classification: Track 1 Session

Track Classification: Track1: Online computing