

The CMS Data Acquisition - Architectures for the Phase-2 Upgrade

Tuesday, 11 October 2016 11:00 (15 minutes)

In Long Shutdown 3 the CMS Detector will undergo a major upgrade to prepare for the second phase of the LHC physics program, starting around 2026. The HL-LHC upgrade will bring instantaneous luminosity up to $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (levelled), at the price of extreme pileup of 200 interactions per crossing. A new silicon tracker with trigger capabilities and extended coverage, and new high granularity forward calorimetry will enhance the CMS acceptance and selection power. This will enable precision measurements of the Higgs boson properties, as well as extend the discovery reach for physics beyond the standard model, while coping with conditions dictated by the HL-LHC parameters.

Following the tradition, the CMS Data Acquisition System will continue to feature two trigger levels.

The detector will be read out at an unprecedented data rate of up to 50 Tb/s read at a Level-1 rate of 750 kHz from some 50k high-speed optical detector links, for an average expected event size of 5MB. Complete events will be analysed by a software trigger (HLT) running on standard processing nodes, and selected events will be stored permanently at a rate of up to 10 kHz for offline processing and analysis.

In this paper we discuss the baseline design of the DAQ and HLT systems for the Run 4, taking into account the projected evolution of high speed network fabrics for event building and distribution, and the anticipated performance of many-core CPU and their memory and I/O architectures. Assuming a modest improvement of the processing power of 12.5% per year for the standard Intel architecture CPU and the affordability, by 2026, of 100-200 Gb/s links, and scaling the current HLT CPU needs for increased event size, pileup, and rate, the CMS DAQ will require about:

- 800 100 Gb/s links connecting 400 servers for event building
- 4500 processing servers, or a total of about 11 MHS for HLT

Implications on hardware and infrastructure requirements for the DAQ “data center” are analysed. Emerging technologies for data reduction, in particular of CPU-FPGA hybrid systems, but also alternative CPU architectures, are considered. These technologies may in the future help containing the TCO of the system, while improving the energy performance and reducing the cooling requirements.

Novel possible approaches to event building and online processing are also examined, which are inspired by trending developments in other areas of computing dealing with large masses of data.

We conclude by discussing the opportunities offered by reading out and processing parts of the detector, whenever the front-end electronics allows, at the machine clock rate (40 MHz). While the full detector is being read out and processed at the Level-1 rate, a second, parallel DAQ system would run as an “opportunistic experiment” processing tracker trigger and calorimeter data at 40 MHz. This idea presents interesting challenges and its physics potential should be studied.

Tertiary Keyword (Optional)

High performance computing

Secondary Keyword (Optional)

Trigger

Primary Keyword (Mandatory)

DAQ

Primary author: MESCHI, Emilio (CERN)

Co-authors: HOLZNER, Andre Georg (Univ. of California San Diego (US)); RACZ, Attila (CERN); CRAIGS, Ben (CERN, Geneva, Switzerland); DELDICQUE, Christian (CERN); PAUS, Christoph (Massachusetts Inst. of Technology (US)); SCHWICK, Christoph (CERN); CONTESCU, Cristian (Fermi National Accelerator Lab. (US)); SIMELEVICIUS, Dainius (CERN); GIGI, Dominique (CERN); GLEGE, Frank (CERN); MEIJERS, Frans (CERN); DARLEA, Georgiana Lavinia (Massachusetts Inst. of Technology (US)); GOMEZ CEBALLOS RETUERTO, Guillermo (Massachusetts Inst. of Technology (US)); SAKULIN, Hannes (CERN); BRANSON, James Gordon (Univ. of California San Diego (US)); ANDRE, Jean-Marc Olivier (Fermi National Accelerator Lab. (US)); HEGEMAN, Jeroen (CERN); FULCHER, Jonathan (CERN); MASETTI, Lorenzo (CERN); ORSINI, Luciano (CERN); DOBSON, Marc (CERN); PIERI, Marco (Univ. of California San Diego (US)); DOUALOT, Nicolas (Fermi National Accelerator Lab. (US)); CHAZE, Olivier (CERN); ZEJDL, Petr (Fermi National Accelerator Lab. (US)); BRUMMER, Philipp Maximilian (CERN, Geneva, Switzerland); JIMENEZ ESTUPINAN, Raul (CERN); MOMMSEN, Remi (Fermi National Accelerator Lab. (US)); ERHAN, Samim (Univ. of California Los Angeles (US)); CITTOLIN, Sergio (Univ. of California San Diego (US)); MOROVIC, Srecko (CERN); REIS, Thomas (CERN); BEHRENS, Ulf (Deutsches Elektronen-Synchrotron (DE)); O'DELL, Vivian (Fermi National Accelerator Laboratory (FNAL)); DEMIRAGLI, Zeynep (Massachusetts Inst. of Technology (US))

Presenter: MESCHI, Emilio (CERN)

Session Classification: Track 1: Online Computing

Track Classification: Track 1: Online Computing