

Evolution, design, management and support for the CMS Online computing cluster

Tuesday, 11 October 2016 16:30 (15 minutes)

After two years of maintenance and upgrade, the Large Hadron Collider (LHC) has started its second four year run. In the mean time, the CMS experiment at the LHC has also undergone two years of maintenance and upgrade, especially in the field of the Data Acquisition and online computing cluster, where the system was largely redesigned and replaced. Various aspects of the supporting computing system will be addressed here.

The increasing processing power and the use of high end networking technologies (10/40Gb/s Ethernet and 56Gb/s Infiniband) has reduced the number of DAQ event building nodes, since the performance of the individual nodes has increased by an order of magnitude since the start of LHC. The pressure on using the systems in an optimal way has increased accordingly, thereby also increasing the importance of proper configuration and careful monitoring to catch any deviation from standard behaviour. The upgraded monitoring system based on Ganglia and Icinga2 will be presented with the different mechanisms used to monitor and troubleshoot the crucial elements of the system.

The evolution of the various sub-detector applications, the data acquisition and high level trigger, following their upgraded hardware and designs over the upgrade and running periods, require a performant and flexible management and configuration infrastructure. The puppet based configuration and management system put in place for this phase, will be presented, showing it's flexibility to support a large heterogeneous system, as well as, it's ability to do bulk installations from scratch or rapid installations of CMS software cluster wide. A number of custom tools have been developed to support the update of rpm based installations by the end users, a feature not typically supported in a datacenter environment. The performance of the system will also be presented with insights into its scaling with the increasing farm size over this data taking run.

Such a large and complex system requires redundant, flexible core infrastructure services to support them. Details will be given on how a flexible and highly available infrastructure has been put in place, leveraging various high availability technologies, from network redundancy, through virtualisation, to high availability services with Pacemaker/Corosync.

To conclude, a roundup of the different tools and solutions used in the CMS cluster administration will be given, pulling all the above into a coherent, performant and scalable system.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

DAQ

Primary Keyword (Mandatory)

Computing facilities

Primary author: DOBSON, Marc (CERN)

Co-authors: HOLZNER, Andre Georg (Univ. of California San Diego (US)); RACZ, Attila (CERN); CRAIGS, Benjamin Gordon (CERN); 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); MESCHI, Emilio (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); 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); JIMENEZ ESTUPINAN, Raul (CERN); MOMMSEN, Remigius K. (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))

Session Classification: Posters A / Break

Track Classification: Track 1: Online Computing