

Experience with dynamic resource provisioning of the CMS online cluster using a cloud overlay

Tuesday, 10 July 2018 11:45 (15 minutes)

The primary goal of the online cluster of the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) is to build event data from the detector and to select interesting collisions in the High Level Trigger (HLT) farm for offline storage. With more than 1100 nodes and a capacity of about 600 kHEPSpec06, the HLT machines represent up to 40% of the combined Tier0/Tier-1 capacity on the Worldwide LHC Computing Grid (WLCG). Moreover, it is currently connected to the CERN IT datacenter via a dedicated 160 Gbps network connection and hence can access the remote EOS based storage with a high bandwidth. In the last few years, a cloud overlay based on Openstack has been commissioned to use these resources for the WLCG when they are not needed for data taking. This online cloud facility was designed for parasitic use of the HLT, which must never interfere with its primary function as part of the DAQ system. It also allows to abstract from the different types of machines and their underlying segmented networks. For the monitoring, an infrastructure based on Graphite for metric storage and aggregation, and Grafana for the visualisation has been deployed. During the LHC technical stop periods, the HLT cloud is set to its static mode of operation where it acts like other grid facilities. The online cloud was also extended to make dynamic use of resources during periods between LHC fills. These periods are a-priori unscheduled and of undetermined length, typically of several hours, once or more a day. For that, it dynamically follows LHC beam states and hibernates Virtual Machines (VM) accordingly. Finally, this work presents the design and implementation of a mechanism to dynamically ramp up VMs when the DAQ load on the HLT reduces towards the end of the fill.

Primary authors: HOLZNER, Andre Georg (Univ. of California San Diego (US)); PETRUCCI, Andrea (Rice University (US)); RACZ, Attila (CERN); MECIONIS, Audrius (Vilnius University (LT)); DELDICQUE, Christian (CERN); WERNET, Christian (University of Applied Sciences (DE)); PAUS, Christoph (Massachusetts Inst. of Technology (US)); SCHWICK, Christoph (CERN); VAZQUEZ VELEZ, Cristina (CERN); SIMELEVICIUS, Dainius (Vilnius University (LT)); DA SILVA GOMES, Diego (CERN); RABADY, Dinyar (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); PAPAKRIVOPOULOS, Ioannis (National Technical Univ. of Athens (GR)); BRANSON, James Gordon (Univ. of California San Diego (US)); ANDRE, Jean-Marc Olivier (Fermi National Accelerator Lab. (US)); HEGEMAN, Jeroen (CERN); FULCHER, Jonathan (CERN); ORSINI, Luciano (CERN); GLADKI, Maciej Szymon (Ministere des affaires etrangeres et europeennes (FR)); STANKEVICIUS, Mantas (Fermi National Accelerator Lab. (US)); DOBSON, Marc (CERN); PIERI, Marco (Univ. of California San Diego (US)); LETTRICH, Michael (Technische Universität Muenchen (DE)); DOUALOT, Nicolas (Fermi National Accelerator Lab. (US)); CHAZE, Olivier (CERN); ZEJDL, Petr (Fermi National Accelerator Lab. (US)); MOMMSEN, Remi (Fermi National Accelerator Lab. (US)); ERHAN, Samim (University of California Los Angeles (US)); CITTOLIN, Sergio (Univ. of California San Diego (US)); MOROVIC, Srecko (Fermi National Accelerator Lab. (US)); REIS, Thomas (CERN); BEHRENS, Ulf (Deutsches Elektronen-Synchrotron (DE)); RAPSEVICIUS, Valdas (Fermi National Accelerator Lab. (US)); O'DELL, Vivian; DEMIRAGLI, Zeynep (Massachusetts Inst. of Technology (US))

Presenter: DA SILVA GOMES, Diego (CERN)

Session Classification: T7 - Clouds, virtualization and containers

Track Classification: Track 7 –Clouds, virtualization and containers