20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013)

Contribution ID: **98** Type: **Oral presentation to parallel session**

The CMS openstack, opportunistic, overlay, online-cluster Cloud (CMSooooCloud)

Thursday 17 October 2013 12:06 (22 minutes)

The CMS online cluster consists of more than 3000 computers. It has been exclusively used for the Data Acquisition of the CMS experiment at CERN, archiving around 20Tbytes of data per day.
An openstack cloud layer has been deployed on part of the cluster (totalling more than 13000 cores) as a minimal overlay so as to leave the primary role of the computers untouched while allowing an opportunistic usage of the cluster. This allows running offline computing jobs on the online infrastructure while it is not (fully) used.
We will present the architectural choices made to deploy an unusual, as opposed to dedicated, "overlaid cloud infrastructure". These architectural choices ensured a minimal impact on the running cluster configuration while giving a maximal segregation of the overlaid virtual computer infrastructure. Openvswitch was chosen during the proof of concept phase in order to avoid changes on the network infrastructure. Its use will be illustrated as well as the final networking configuration used. The design and performance of the openstack cloud controlling layer will be also presented together with new developments and experience from the first year of usage.

Author: Dr COARASA PEREZ, Jose Antonio (CERN)

Co-authors: HOLZNER, Andre Georg; PETRUCCI, Andrea (CERN); SPATARU, Andrei Cristian (CERN); Dr RACZ, Attila (CERN); DUPONT, Aymeric Arnaud (CERN); NUNEZ BARRANCO FERNANDEZ, Carlos (CERN); DELDICQUE, Christian (CERN); HARTL, Christian (CERN); PAUS, Christoph (Massachusetts Inst. of Technology (US)); SCHWICK, Christoph (CERN); WAKEFIELD, Christopher Colin (Staffordshire University (GB)); GIGI, Dominique (CERN); MESCHI, Emilio (CERN); STOECKLI, Fabian (Massachusetts Inst. of Technology (US)); GLEGE, Frank (CERN); MEIJERS, Frans (CERN); BAUER, Gerry (Massachusetts Inst. of Technology (US)); Dr POLESE, Giovanni (University of Wisconsin (US)); SAKULIN, Hannes (CERN); BRANSON, James Gordon (Univ. of California San Diego (US)); SUMOROK, Konstanty (Massachusetts Inst. of Technology (US)); MASETTI, Lorenzo (CERN); ORSINI, Luciano (CERN); Dr DOBSON, Marc (CERN); PIERI, Marco (Univ. of California San Diego (US)); SANI, Matteo (Univ. of California San Diego (US)); CHAZE, Olivier (CERN); RAGINEL, Olivier (Massachusetts Inst. of Technology (US)); ZEJDL, Petr (CERN); Dr MOMMSEN, Remi (Fermi National Accelerator Lab. (US)); GOMEZ-REINO GARRIDO, Robert (CERN); ERHAN, Samim (Univ. of California Los Angeles (US)); CITTOLIN, Sergio (Univ. of California San Diego (US)); MOROVIC, Srecko (Institute Rudjer Boskovic (HR)); BEHRENS, Ulf (Deutsches Elektronen-Synchrotron (DE)); O'DELL, Vivian (Fermi National Accelerator Laboratory (FNAL)); OZGA, Wojciech Andrzej (AGH University of Science and Technology (PL))

Presenter: Dr COARASA PEREZ, Jose Antonio (CERN)

Session Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtualization

Track Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtualization