



Contribution ID: 217

Type: oral presentation

A scalable monitoring for the CMS Filter Farm based on elasticsearch

Tuesday, April 14, 2015 3:45 PM (15 minutes)

A flexible monitoring system has been designed for the CMS File-based Filter Farm making use of modern data mining and analytics components. All the metadata and monitoring information concerning data flow and execution of the HLT are generated locally in the form of small “documents” using the JSON encoding. These documents are indexed into a hierarchy of elasticsearch (es) clusters along with process and system log information. Elasticsearch is a search server based on Apache Lucene. It provides a distributed, multitenant-capable search and aggregation engine. Since es is schema-free, any new information can be added seamlessly and the unstructured information can be queried in non-predetermined ways. The leaf es clusters consist of the very same nodes that form the Filter Farm thus providing “natural” horizontal scaling. A separate “central” es cluster is used to collect and index aggregated information. The fine-grained information, all the way to individual processes, remains available in the leaf clusters. The central es cluster provides quasi-real-time high-level monitoring information to any kind of client. Historical data can be retrieved to analyse past problems or correlate them with external information. We discuss the design and performance of this system in the context of the CMS DAQ commissioning for LHC Run 2.

Primary authors: MESCHI, Emilio (CERN); MOROVIC, Srecko (CERN)

Co-authors: ANDRONIDIS, Anastasios (Aristotle Univ. of Thessaloniki (GR)); HOLZNER, Andre Georg (Univ. of California San Diego (US)); PETRUCCI, Andrea (CERN); Dr RACZ, Attila (CERN); DUPONT, Aymeric Arnaud (CERN); STIEGER, Benjamin (CERN); NUNEZ BARRANCO FERNANDEZ, Carlos (CERN); DELDICQUE, Christian (CERN); PAUS, Christoph (Massachusetts Inst. of Technology (US)); SCHWICK, Christoph (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)); Mr VEVERKA, Jan (Massachusetts Inst. of Technology (US)); ANDRE, Jean-Marc Olivier (Fermi National Accelerator Lab. (US)); Dr HEGEMAN, Jeroen (CERN); 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)); CHAZE, Olivier (CERN); ROBERTS, Penelope Amelia (Nottingham Trent University (GB)); ZEJDL, Petr (CERN); MOMMSEN, Remi (Fermi National Accelerator Lab. (US)); ZAZA, Salvatore (Sezione di Pisa (IT)); ERHAN, Samim (Univ. of California Los Angeles (US)); CITTOLIN, Sergio (Univ. of California San Diego (US)); BAWEJ, Tomasz Adrian (University of Wisconsin (US)); BEHRENS, Ulf (Deutsches Elektronen-Synchrotron (DE)); O'DELL, Vivian (Fermi National Accelerator Laboratory (FNAL))

Presenter: MOROVIC, Srecko (CERN)

Session Classification: Track 1 Session

Track Classification: Track1: Online computing