

A comparison of different database technologies for the CMS AsyncStageOut transfer database

Monday, 10 October 2016 15:30 (15 minutes)

AsyncStageOut (ASO) is the component of the CMS distributed data analysis system (CRAB3) that manages users' transfers in a centrally controlled way using the File Transfer System (FTS3) at CERN. It addresses a major weakness of the previous, decentralized model, namely that the transfer of the user's output data to a single remote site was part of the job execution, resulting in inefficient use of job slots and an unacceptable failure rate.

Currently ASO manages up to 600k files of various sizes per day from more than 500 users per month, spread over more than 100 sites and uses a NoSQL database (CouchDB) as internal bookkeeping and as a way to communicate with other CRAB3 components. Since ASO/CRAB3 were put in production in 2014, the number of transfers constantly increased up to a point where the pressure on the central CouchDB instance became critical, creating new challenges for the system scalability, performance, and monitoring. This forced a re-engineering of the ASO application to increase its scalability and lower its operational effort.

In this contribution we present a comparison of the performance of the current NoSQL implementation and a new SQL implementation, and of how their different strengths and features influenced the design choices and operational experience. We also discuss other architectural changes introduced in the system to handle the increasing load and latency in delivering the output to the user.

Primary Keyword (Mandatory)

Databases

Secondary Keyword (Optional)

Tertiary Keyword (Optional)

Primary authors: BALCAS, Justas (California Institute of Technology (US)); MASCHERONI, Marco (Fermi National Accelerator Lab. (US))

Co-authors: CIANGOTTINI, Diego (Universita e INFN, Perugia (IT)); RUPEIKA, Emilis Antanas (Vilnius University (LT)); VAANDERING, Eric (Fermi National Accelerator Lab. (US)); RIAHI, Hassen (CERN); MARRA DA SILVA, Jadir (UNESP - Universidade Estadual Paulista (BR)); HERNANDEZ, Jose (CIEMAT); BELFORTE, Stefano (Universita e INFN, Trieste (IT))

Presenter: VAANDERING, Eric (Fermi National Accelerator Lab. (US))

Session Classification: Track 2: Offline Computing

Track Classification: Track 2: Offline Computing