

Integration of the Italian cache federation within CMS computing model

Friday, 5 April 2019 09:00 (20 minutes)

The next decades at HL-LHC will be characterized by a huge increase of both storage and computing requirements. A factor 20 is expected for the storage, while on computing the estimation is about a 30x CPUs. Moreover, we foresee a shift on resources provisioning towards the exploitation of dynamic (on private or public cloud and HPC facilities) solutions. In this scenario, the computing model of the CMS experiment is pushed towards an evolution for the optimization of the amount of space that is managed centrally and the CPU efficiency of the jobs that run on “storage-less” resources. In particular, the computing resources of the “Tier2” sites layer, for the most part, can be instrumented to read data from a geographically distributed cache storage based on unmanaged resources, reducing, in this way, the operational efforts by a large fraction and generating additional flexibility.

One of the benefits behind a distributed cache space is the possibility to leverage the national high bandwidth network from NRENs to reduce the number of file replica around the WLCG sites. The cache system will appear as a distributed and shared file system populated with the most requested data; in case of missing information data access will fallback to the remote access.

Moreover in a possible future scenario based on the data-lake model, it is reasonable to imagine that many satellite computing centers might appear and disappear dynamically. In this sense, a protection layer against a central managed storage might be a key factor along with the control of the data access latency. A Content Delivery Network has many affinities with the cache architecture desired and, in each region, duplicate files within its federated servers are not foreseen. The cache storages will be by definition “non-custodial”, thus reducing the overall operational costs.

The objective of this contribution is to present the first implementation of an INFN federation of cache servers, developed also in collaboration with the eXtreme-DataCloud EU project. The CNAF Tier-1 plus Bari and Legnaro Tier-2s provide unmanaged storages which have been organized under a common namespace. This distributed cache federation has been seamlessly integrated with the CMS computing infrastructure.

The technical implementation of this solution is based on XRootD, largely adopted in the CMS computing model under the “Anydata, Anytime, Anywhere project” (AAA). The technology is compliant with the most common storage protocol present in the WLCG and several activities already demonstrated the effective management of cold and warm cache scenarios. Moreover, the possibility to plug custom caching algorithms will provide the ability of playing with historical and online data access metrics for smart data placement decision. The results in terms of CMS workflows performances will be shown. In addition, a complete simulation of the effects of the described model under several scenarios, including dynamic hybrid cloud resource provisioning, will be discussed. Finally, a plan for the upgrade of such a prototype towards a stable INFN setup seamlessly integrated with the production CMS computing infrastructure will be discussed.

Primary authors: Dr FALABELLA, Antonio (INFN); CESINI, Daniele (INFN-CNAF); CIANGOTTINI, Diego (INFN Perugia); MAZZONI, Enrico (INFN Pisa); Mr DONVITO, Giacinto (INFN); BAGLIESI, Giuseppe (INFN Pisa); BIASOTTO, Massimo (INFN Legnaro); TRACOLLI, Mirco (INFN Perugia); Dr SPIGA, daniele (INFN-PG); Dr BOCCALI, tommaso (INFN)

Presenter: CIANGOTTINI, Diego (INFN Perugia)

Session Classification: Physics & Engineering Application

Track Classification: Physics (including HEP) and Engineering Applications