

INFN and the Future of Scientific Computing
4-may-2018

DPM-based distributed caching system for multi-site storage in ATLAS

Bernardino Spisso
INFN Napoli

About me

- Master degree in theoretical physics at Federico II di Napoli (2007).
- Ph.D. in mathematics at Westfälische Wilhelms-Universität of Münster (2012).
- I level University Master Course in “Technologies for high-performance scientific computing” at Federico II di Napoli (2014).

My previous occupation was in the ASTERICS-Km3Net collaboration, responsible for the projects:

- ROAst (ROOT extension with ASTrophysical).
- CORELib (COsmic Ray Event Library).

Since December 2017 I collaborate with Naples ATLAS computing Group:



- Alessandra Doria and Giampaolo Carlino.

And with the ATLAS Tiers-2 at LNF and ROMA1.

- Elisabetta Vilucchi (INFN-LNF) and Alessandro De Salvo (INFN-Roma1).

Motivations

- HL-LHC storage needs are above the expected technology evolution (15%/yr) and funding (flat).
- We need to optimize storage hardware usage and operational costs.



Plots from the last Joint WLCG-HSF Workshop 2018 in Naples.

Our scenario

- Explore distributed storage evolution to improve overall costs (storage and ops) taking in account:
 - Single common namespace and interoperability.
 - User analysis is often based on clusters hosted on medium sites (Tier2) and small sites (Tier3).
- **In order to reconcile these two trends, the target of my activity is to study a distributed storage system featuring a single access point to large permanent storage and capable to provide efficient and dynamic access to the data. In this view, medium sites like Tier2 and small sites like Tier3 will not necessarily require large storage systems, simplifying local management.**


This can be achieved by the adoption of a distributed storage and caching technologies.
- This activity takes place in the same context of the Data Lake project having very similar motivations.

Our implementation

- The Disk Pool Manager (DPM) is a data management solution widely used within ATLAS, in particular in three Italian Tier2.
- The latest versions of DPM are used in our implementation, that offer the possibility to manage volatiles pools to be used as caches.

By exploiting the fast connections between sites, we are deploying a first testbed among Naples, Frascati and Roma-1 using DPM. The aim is to study and develop a configuration in which a primary site represents a single entry point for the entire archiving system and each site can use its storage as permanent storage or as local cache.

Using a cache system the local site administrators can be dispensed from managing a complete storage system. The site became transparent for the central operations of the experiment.



Conclusions

- A first testbed using DPM among Naples, Frascati and Roma-1 is almost ready.
- Study of the best caching policy for the volatile pools.
- Evaluation of the performance of the developed prototype.
- System integration in the current ATLAS data management infrastructures.
- Synergies:
 - collaborations with the Naples BELLEII computing group (Silvio pardi (INFN-NA), Davide Michelino (GARR))
 - collaborations with the DPM development group.
- Create conditions for easy replication of the system on other sites or in other contexts.