

Abstract

A wide range of detector commissioning, calibration and data analysis tasks is carried out by members of the Compact Muon Solenoid (CMS) collaboration using dedicated storage resources available at the CMS CERN Tier-2 centre.

Relying on the functionalities of the EOS disk-only storage technology, the optimal exploitation of the CMS user/group resources has required the introduction of policies for data access management, data protection, cleanup campaigns based on access pattern, and long term tape archival.

The resource management has been organised around the definition of working groups and the delegation to an identified responsible of each group composition.

In this poster contribution we illustrate the user/group storage management, and the development and operational experience at the CMS CERN Tier-2 centre in the 2012-2015 period.

Introduction

We administer the eos-based [1] storage resources CMS provides to the collaborators resident at CERN:

- **user space:** 1 TB of storage for personal analysis or development tasks; it's made available to the CMS users who are resident at CERN. **370 single user areas, 128 TB in total.**
- **group space:** most of the non centrally managed (PhEDEx [3]) CMS storage is used for analysis activities shared by a team of users, e.g. detector groups (e.g. ECAL), physics object (muon) or analyses teams (exotica). The typical use case is the production of datasets (e.g. ntuples) by one or few users, used by the whole group. **330 users, 33 working teams, 1073 TB**
- **working teams:** defined according to the existing detector or physics groups of CMS, are each assigned an amount of group space, in proportion to the T2 pledge of their group; the detector or physics groups conveners are in charge of defining which users can use the eos resources and of cleanup and archival campaigns

Group Space Administration


The group space was initially accessible by all CERN-resident users; a **single quota** was set on the disk space of all the groups combined: over-use by one group filled everyone else's quota, bringing operations to a stall; such structure provided no incentive to cleanup old files and make room for new ones



drwxr-sr+ aguy	zh	Feb-7-2015	Apr-2-2015	/eos/cms/ecal/f1.txt
drwxr-sr+ otherguy	zh	Jan-6-2014	Feb-9-2014	/eos/cms/ecal/tupl_a.root
drwxr-sr+ otherguy	zh	Oct-4-2012	Feb-9-2015	/eos/cms/ecal/tupl_b.root
drwxr-sr+ trdguy	zh	May-3-2013	Mar-5-2014	/eos/cms/muon/eff1.root
drwxr-sr+ fthguy	zh	Jun-1-2012	Jun-2-2014	/eos/cms/exo/plot.eps
drwxr-sr+ fthguy	zh	Jun-1-2012	Mar-5-2014	/eos/cms/exo/plot.root

Introduced **working teams with a separate quota:**


- CMS detector or physics groups conveners manage the team members using e-group's [2], the CERN-wide official interface to manage groups at CERN
- A sub-set ('cleaners') of the team members granted permissions to delete files from other users of the team to implement cleanup campaigns, via e-group



drwxr-sr+ aguy	zh	Feb-7-2015	Apr-2-2015	/eos/cms/ecal/f1.txt
drwxr-sr+ otherguy	zh	Jan-6-2014	Feb-9-2014	/eos/cms/ecal/tupl_a.root
drwxr-sr+ otherguy	zh	Oct-4-2012	Feb-9-2015	/eos/cms/ecal/tupl_b.root
drwxr-sr+ trdguy	zh	May-3-2013	Mar-5-2014	/eos/cms/muon/eff1.root
drwxr-sr+ fthguy	zh	Jun-1-2012	Jun-2-2014	/eos/cms/exo/plot.eps
drwxr-sr+ fthguy	zh	Jun-1-2012	Mar-5-2014	/eos/cms/exo/plot.root

Data popularity to cleanup stale data:


- Single file most-recent-time of access was provided by CERN IT, based on the monitoring of eos access via XRootD; together with the file creation time, popularity was aggregated per working team and distributed to detector or physics groups conveners and 'cleaners'
- bi-monthly file removal campaigns (candidates for deletion: created > 4 mo, accessed > 3 mo), essential in the phase before the single team quotas were deployed



drwxr-sr+ aguy	zh	Feb-7-2015	Apr-2-2015	/eos/cms/ecal/f1.txt
drwxr-sr+ otherguy	zh	Jan-6-2014	Feb-9-2014	/eos/cms/ecal/tupl_a.root
drwxr-sr+ otherguy	zh	Oct-4-2012	Feb-9-2015	/eos/cms/ecal/tupl_b.root
drwxr-sr+ trdguy	zh	May-3-2013	Jun-2-2014	/eos/cms/muon/eff1.root
drwxr-sr+ fthguy	zh	Jun-1-2012	Mar-5-2014	/eos/cms/exo/plot.eps
drwxr-sr+ fthguy	zh	Jun-1-2012	Mar-5-2014	/eos/cms/exo/plot.root

Archival to tape to free space for new data

- stale data possibly needed in the future can be cleaned up by archiving it out of eos onto tape storage (castor [5]), thanks to a recent eos feature
- 'cleaners' command the archival; restoring back to disk is expected to happen sparsely



drwxr-sr+ aguy	zh	Feb-7-2015	Apr-2-2015	/eos/cms/ecal/f1.txt
drwxr-sr+ otherguy	zh	Oct-4-2012	Feb-9-2015	/eos/cms/ecal/archive/tupl_b.arch
drwxr-sr+ trdguy	zh	May-3-2013	Mar-5-2014	/eos/cms/muon/archive/eff1.arch
drwxr-sr+ fthguy	zh	Jun-1-2012	Jun-2-2014	/eos/cms/exo/plot.eps

Status and Prospects

Following the success with the CERN Tier-2, also the **CMS CERN T3** and the **CMS CERN Analysis Facility (CAF)** are to be organised with the same team-driven design and eos+egroups technology choices

The **CERN Analysis Facility (CAF):**

- Reserved resources for high priority and short turn-around time workflows (prompt-feedback, fast calibrations and analyzes). In runII relied on castor storage and a pool of dedicated CPU's
- Refurbishment for runIII, ongoing: cleanup of old/stale users, switch to eos and management of disk and cpu based on working teams
- CERN batch queues, managed through dedicated LSF [4] groups. We'll use the new capability of mapping e-groups into LSF groups, to have a single entry point for the self-management of the CAF and CMS CERN T3

Results

The CMS CERN T2 user and group storage are re-organized to delegate most of the administration to the working teams (conveners and 'cleaners'), remove interference across different teams and minimize the administration operations.

References

- [1] EOS: <http://eos.cern.ch/>
- [2] Cern egroups: <https://espace.cern.ch/e-groups-help/default.aspx>
- [3] CMS PhEDEx: <https://cmsweb.cern.ch/phedex/>
- [4] LSF: <http://www-03.ibm.com/systems/platformcomputing/products/lsf/>
- [5] CERN castor: <http://castor.web.cern.ch/>



Download this poster: