Future Generation Computer Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Integration of end-user Cloud storage for CMS analysis

Hassen Riahi [a,*], Alberto Aimar [a], Alejandro Álvarez Ayllón [a], Justas Balcas [b], Diego Ciangottini [c], José M. Hernández [d], Oliver Keeble [a], Nicolò Magini [e], Andrea Manzi [a], Luca Mascetti [a], Marco Mascheroni [e], Andres Jorge Tanasijczuk [f], Eric Wayne Vaandering [e]

[a] European Organisation for Nuclear Research, CH-1211 Geneva 23, Switzerland
[b] California Institute of Technology, Pasadena, CA 91125, USA
[c] Universitá and INFN Perugia, Via Alessandro Pascoli, 06123 Perugia, Italy
[d] CIEMAT, Madrid 28040, Spain
[e] Fermi National Laboratory, Batavia, IL 60510, USA
[f] University of California, San Diego, La Jolla, CA 92093-0354, USA

HIGHLIGHTS

- A model for the integration of end-user Cloud storage in a Grid infrastructure for the distributed data analysis is proposed.
- The model relies on changes in the Grid infrastructure and developments in the data movement middleware.
- The model shows good performance when moving data between Grid and Cloud at scale.

ARTICLE INFO

Article history:
Received 1 April 2016
Received in revised form 19 January 2017
Accepted 11 April 2017
Available online xxxx

Keywords:
Data transfer protocols
Cloud storage
Distributed data management
Distributed data analysis

ABSTRACT

End-user Cloud storage is increasing rapidly in popularity in research communities thanks to the collaboration capabilities it offers, namely synchronisation and sharing. CERN IT has implemented a model of such storage named, CERNBox, integrated with the CERN AuthN and AuthZ services. To exploit the use of the end-user Cloud storage for the distributed data analysis activity, the CMS experiment has started the integration of CERNBox as a Grid resource. This will allow CMS users to make use of their own storage in the Cloud for their analysis activities as well as to benefit from synchronisation and sharing capabilities to achieve results faster and more effectively. It will provide an integration model of Cloud storages in the Grid, which is implemented and commissioned over the world's largest computing Grid infrastructure, Worldwide LHC Computing Grid (WLCG).

In this paper, we present the integration strategy and infrastructure changes needed in order to transparently integrate end-user Cloud storage with the CMS distributed computing model. We describe the new challenges faced in data management between Grid and Cloud and how they were addressed, along with details of the support for Cloud storage recently introduced into the WLCG data movement middleware, FTS3. The commissioning experience of CERNBox for the distributed data analysis activity is also presented.

1. Introduction

The Compact Muon Solenoid (CMS) is one of the two multi-purpose experiments that collect and analyse data from proton–proton collisions at the Large Hadron Collider (LHC). The analysis model of CMS foresees activities driven by data location to optimise the usage of the network and profit from LAN data access: data are distributed over many computing centres according to CMS data placement policies and the user's processing takes place at the sites where data are located. The user's results are stored and made available to the collaboration for later access in the Grid storage element of the user-designated site. Fig. 1 shows schematically the distributed data analysis model in CMS.

The distributed analysis represents a complex task in a distributed environment with low latency requirements. It involves more than 60 computing centres, geographically distributed all around the world used by more than 600 individual users per month. The average load reached in Run 1 was 20 k concurrently running jobs and 200 k completed jobs per day. Each job produces typically one output and one log file. With the increase of LHC energy and luminosity, the CMS model is expected to increment

* Corresponding author.
E-mail address: hassen.riahi@cern.ch (H. Riahi).

http://dx.doi.org/10.1016/j.future.2017.04.021
0167-739X/

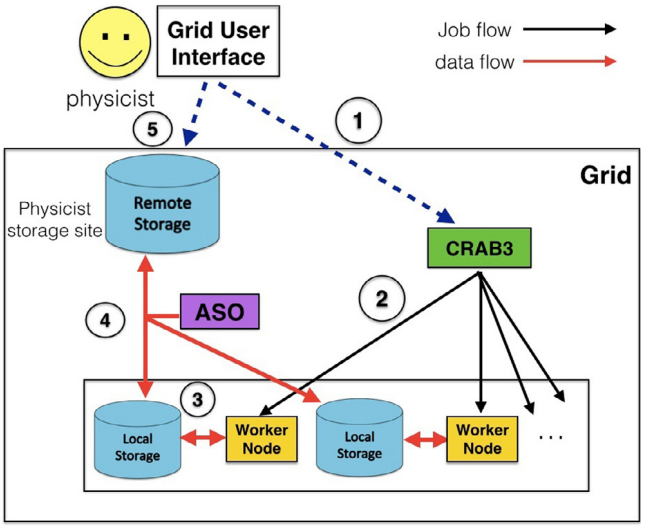

Fig. 1. Distributed data analysis model in CMS.

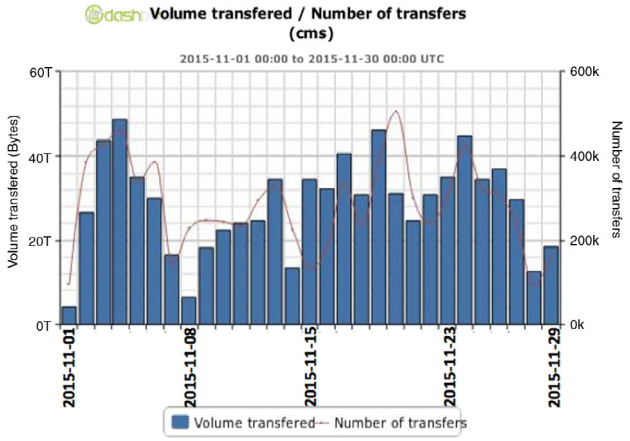

Fig. 2. Daily volume and number of files transferred by ASO during November 2015.

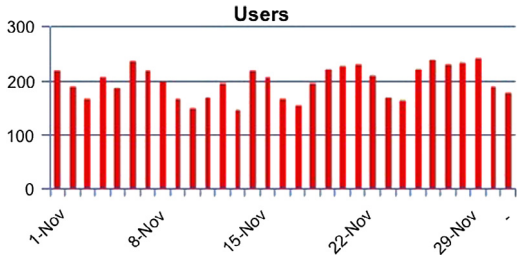

Fig. 3. Daily number of ASO users during November 2015.

its computing needs and expand into Cloud resources. This creates in particular challenging requirements for the integration of Cloud storage resources into the CMS computing model as well as for data management between Grid and Cloud.

CERN IT has implemented a model of end-user Cloud storage, named CERNBox. This will allow CMS users to make use of their own storage in the Cloud for their analysis activities as well as to benefit from synchronisation and sharing capabilities to achieve physics results faster and more effectively.

Sections 2 and 3 give an overview on CERNBox and the CMS user data management models. Section 4 explains the motivation for the integration of end-user Cloud storage in CMS computing model. In Section 5 we present the integration architecture, challenges and developments. Finally Sections 4 and 5 describe the commissioning phase and results.

2. Overview of the CERNBox model

CERNBox is one of the latest services provided by the CERN IT Storage group. It is a Cloud storage service for file synchronisation and sharing based on the ownCloud open-source software stack integrated with the multi-petabyte CERN disk-only system, EOS [1]. CERNBox provides a service layer on top of EOS and allows to maintain a coherent view of users' data, enabling them to handle data locally on their devices and also remotely via the Cloud storage.

At present, users have a default 1 TB quota and they have several ways of accessing their files: disconnected with the Desktop Synchronisation client, with mobile client app (iOS and Android), directly via command line tools or through more sophisticated Grid protocols (xrootd [2], GridFTP and SRM). For the following integration only the xrootd protocol was taken into consideration.

Users are able to selectively access or synchronise data with their devices while the entire data storage is accessible online by the batch/Grid jobs for massive processing or for the creation of big data archives.

CERNBox is an innovative service and is evolving to become the main entry point for the entire CERN data store.

3. User data management for CMS analysis

3.1. A syncStageOut

AsyncStageOut (ASO) [3] is the data management service of the CMS distributed data analysis system, CRAB [4]. It is designed to handle the transfer and bookkeeping of users' outputs to the final storage. It started production in June 2014. It relies on the WLCG [5] data movement middleware, File Transfer Service (FTS) [6].

Typically, ASO transfers 20 TB–40 TB of data in 200 k–400 k files per day for 100–200 users. Fig. 2 shows the daily data volume transferred and the number of transfers during November 2015, while Fig. 3 shows the daily number of ASO users during the same period.

3.2. FTS

FTS is the transfer engine at the heart of the global LHC data distribution system. CMS and other experiments use FTS as a low level data mover, integrating it into their higher level frameworks. The FTS infrastructure has been able to reliably move many petabytes of LHC data per month.

FTS is designed to optimally exploit the available storage and network resources in order to maximise reliable data throughput, employing adaptive optimisation and smart scheduling to achieve this. Its modular design was conceived to allow easy integration of new types of resources such as those described in this contribution.

4. Motivation of the integration

CMS counts more than 1000 users performing physics analysis activities. For each user, the CMS computing model foresees 1 TB of storage at his home institute. However, several CMS institutes do not have computing facilities, thus hampering analysis activities of the users belonging to those institutes. The integration of an end-user Cloud storage, such as CERNBox, overcomes this limitation by allowing CMS users to make use of their storages in the private and public Clouds.

Collaboration and low latency are crucial for the physics analysis in CMS. The analysis activity is divided into various physics

Please cite this article in press as: H. Riahi, et al., Integration of end-user Cloud storage for CMS analysis, Future Generation Computer Systems (2017), http://dx.doi.org/10.1016/j.future.2017.04.021.