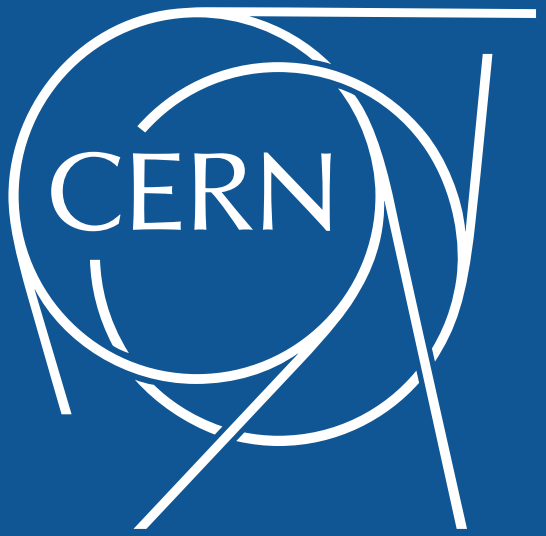
CERN

Experience in integrating end-user cloud storage for CMS Analysis

Hassen Riahi
CERN IT

CS3 Workshop
Zürich, 18th January 2016

Outline

- Overview
- Goals
- Integration architecture
- Challenges & Developments
- Commissioning experience
- Conclusions

CERNBox

- CERNBox is a production CERN service providing a cloud synchronization service as commercial cloud storages for end-users (Dropbox, Amazon Cloud Drive, Google Drive, etc.)
  - Data is located at CERN
  - Offline access on all devices/all platforms
  - Synchronize files
    - **10% of CMS users download data to their desktop**
  - Easy sharing with other users
- Available for all CERN users (1TB/user)
- It has been integrated with:
  - CERN batch
  - E-Groups
  - Root Viewer

See related work: L. Mascetti *et al.* "CERNBox: Cloud Storage for Science", CS3 workshop, Oral presentation, 1/19/16.

Distributed Data Analysis in CMS

- +1000 individual user/month
- +50 sites
- +20k concurrent jobs
- Typically 1 output/job
  - Files vary in size
- +200k completed jobs/day
- Minimal latencies
- Chaotic environment

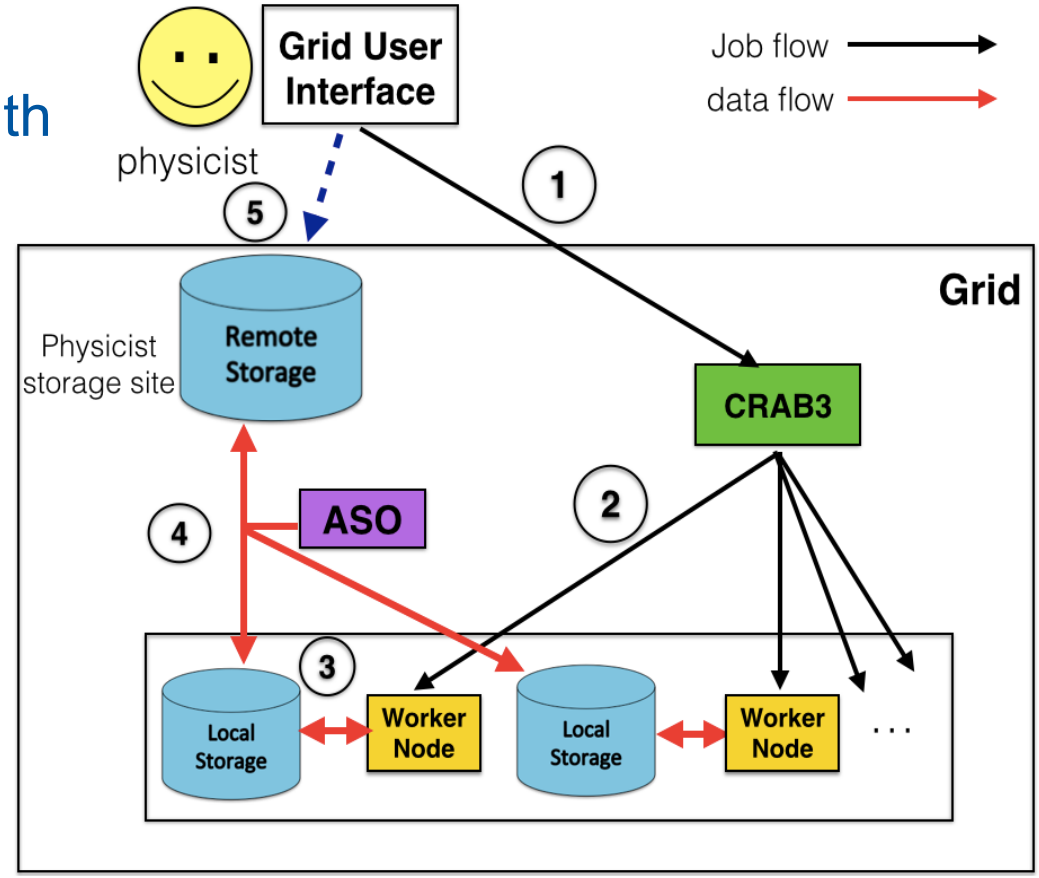

AsyncStageOut (ASO)

- ASO is the distributed user data management system for CMS Analysis
- It has started production in June 2014
- It relies on File Transfer Service (FTS3)
- ~ 800 TB/8 M of files transferred during the last month
  - ~ 200k transfers/day
    - Peaks of ~ 500k transfers/day

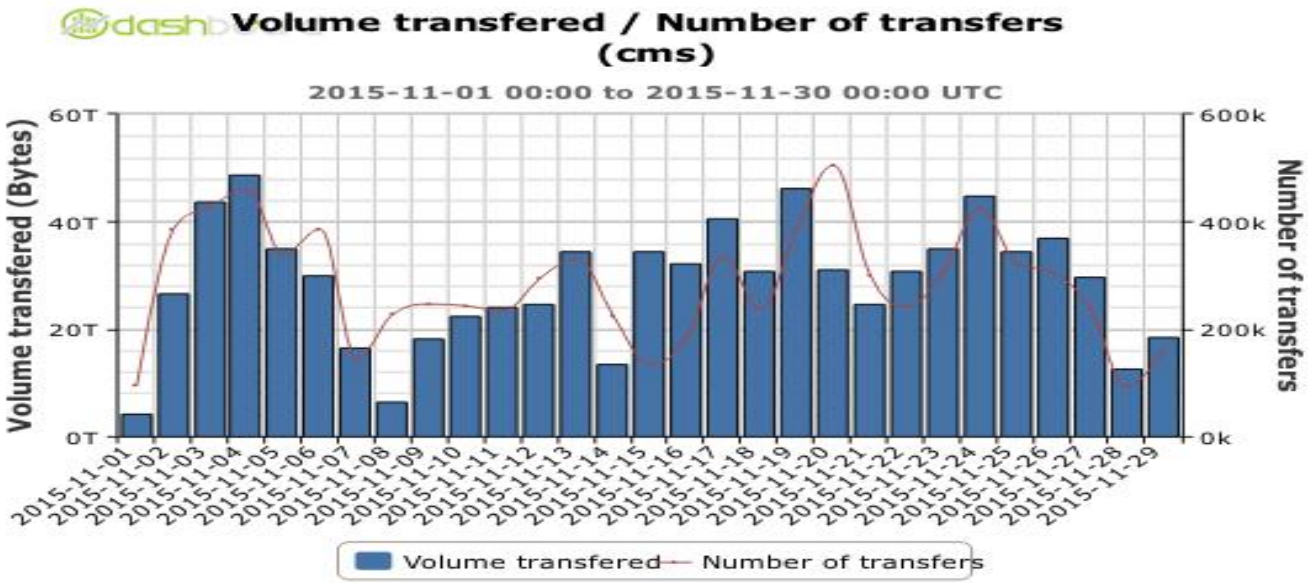

FTS3

- FTS3 is a low level data movement service
- It is the WLCG middleware for data transfers
  - Started officially production at August 1st, 2014
- Transfers more than 20PB/month (record of 2PB/day)
- Capable of auto-adapt the concurrency based on multiple parameters

Goals

- Provide cloud storage to CMS users lacking of storage area in their home institute to perform physics analysis
- Enable CMS users to extend their storage in the grid by using their own storage in private/ public clouds
- Exploit Sync. and Share capabilities in CMS Analysis for:
  - Prompt and offline access from the laptop to physics results produced by grid jobs
  - Easy and fast sharing of results among the physics group

Integration Architecture

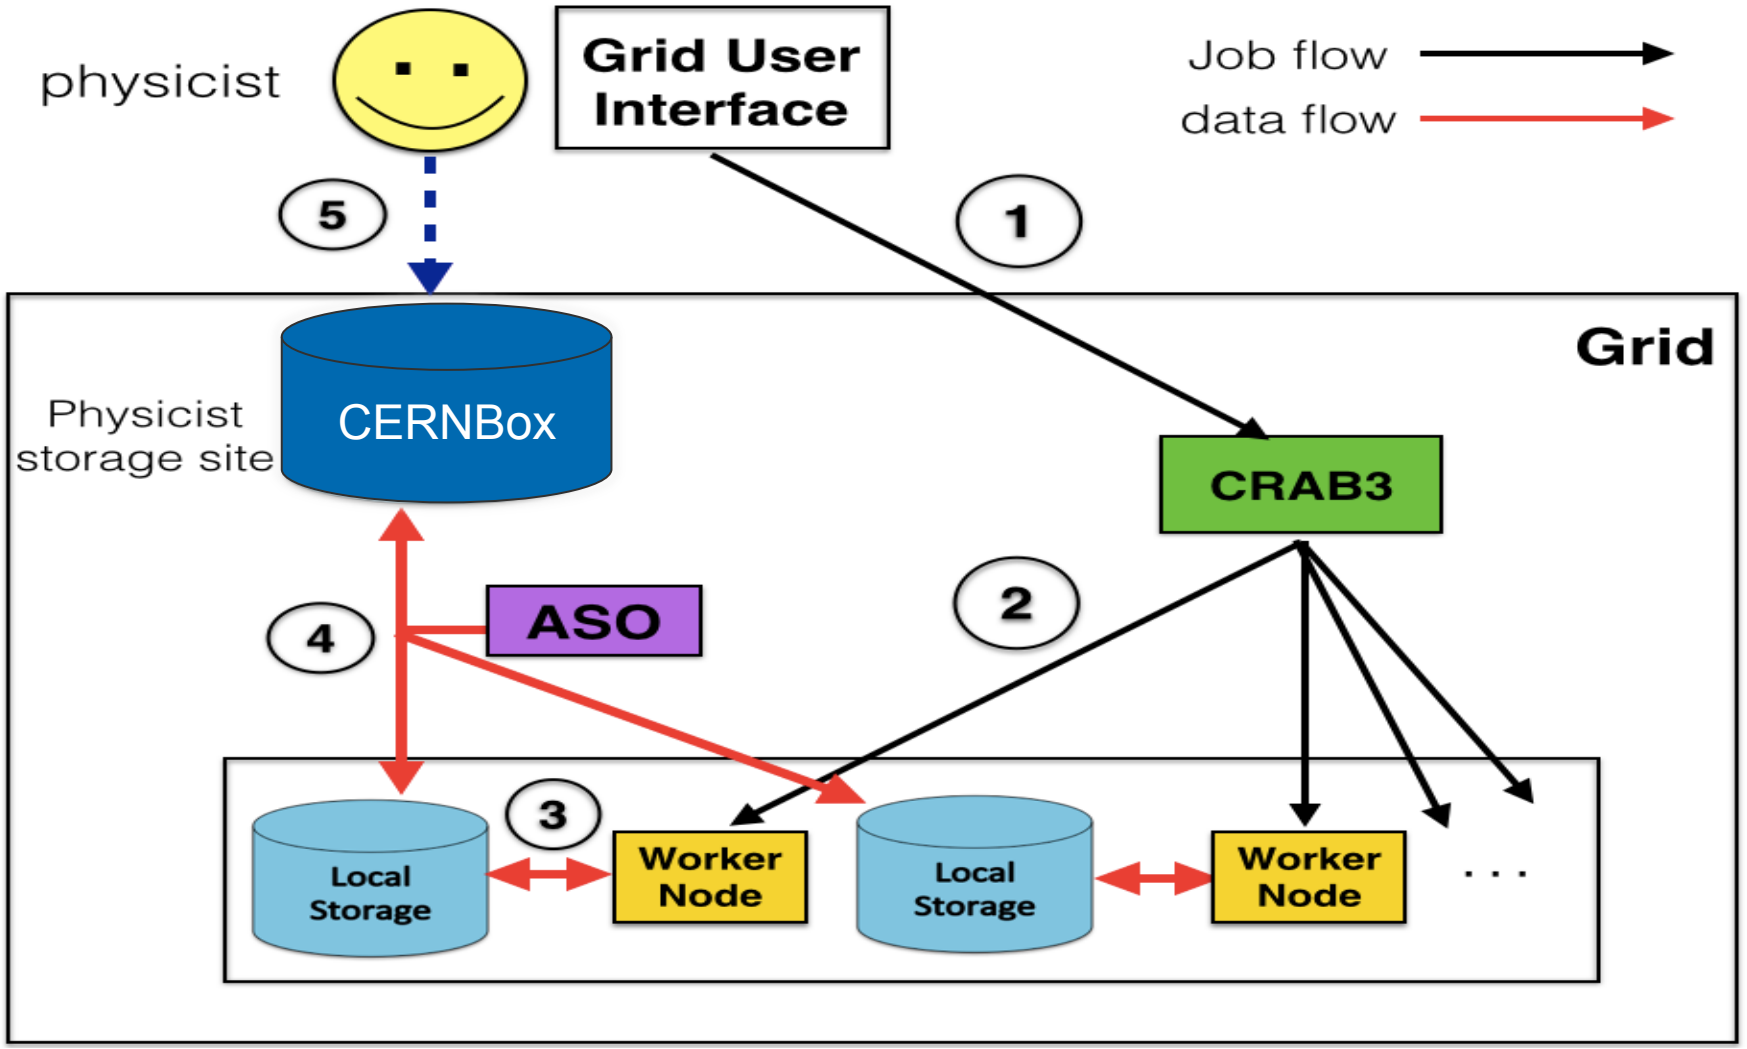

Infrastructure: Challenges & Developments

- CMS supports only the transfer between SRM/ GridFTP storage endpoints for users' data movement in the grid while cloud sites could expose other transfer protocols
  - The support of data movement from/to any transfer protocol exposed by the storage has been included in ASO
- CMS sites are required to follow CMS naming convention (/store/user/username) when configuring the storage paths for their users
  - The support of ad-hoc users' paths in the storage has been included

DM Challenges

- CMS storages are exposing SRM/GridFTP interface for users' data transfer while it is not necessarily the case for cloud storages
  - CERNBox relies so far on XROOT/HTTP protocols
  - Dropbox exposes its own protocol
  - AWS S3
  - …
- 3rd party copy between 2 storages could be possible only if source and destination are exposing the same transfer protocol

FTS3 Developments For Cloud

- FTS3 supports multiple protocols
  - Normally used with 3rd party enabled protocols
  - Supports acting as a proxy between incompatible protocols
- S3, Dropbox and plain DAV/HTTP are supported
  - Third party copies can be done if the other endpoint supports it
    - Dropbox/S3 ⇔ DPM/dCache WebDAV
  - Pipelining is done otherwise
    - Dropbox/S3 ⇔ FTS3 ⇔ Storage
  - S3 credentials need to be configured on the FTS3
  - Dropbox access require OAuth negotiation via WebFTS or manual configuration
  - 3.5 will include support for short-lived S3 credentials
- For further details: fts-support@cern.ch

Commissioning Exercise

- The exercise ran for 4 days
- Jobs were configured to stage-out the outputs via ASO from the grid to CERNBox
- ~ 10 expert users were submitting continuous load
- FTS3 @pilot has been used to not interfere with the production activity
- ~ 8 TB in ~ 1M of files have been moved from the grid to CERNBox during this exercise

Results

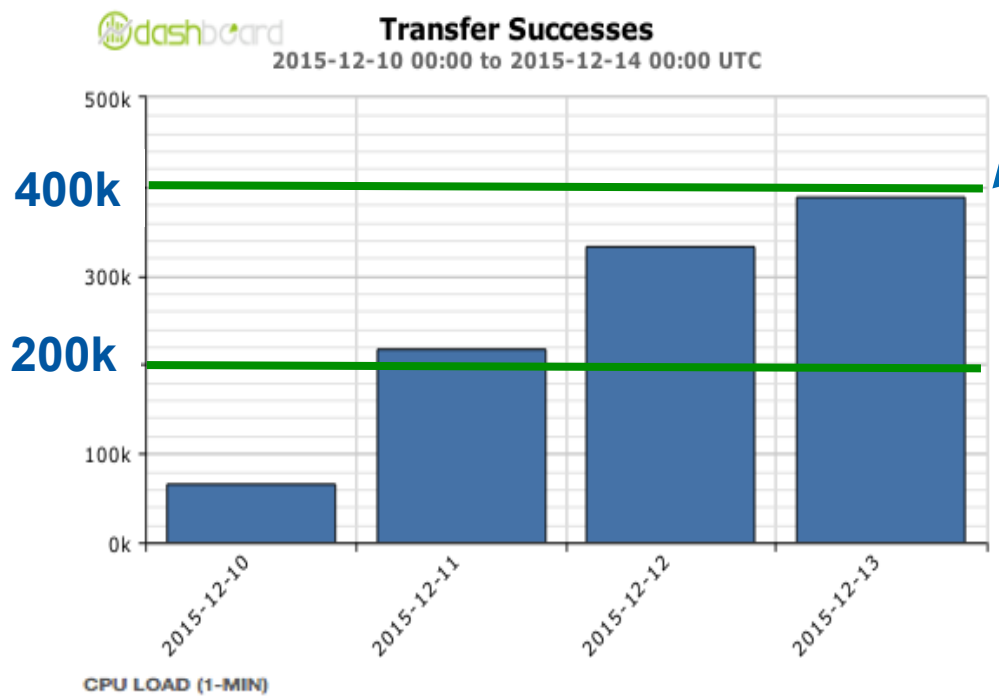

~ 2 times the total number of transferred files per day for the overall analysis activity in CMS

~ 10 times the number of users' files transferred per day to the CMS site @CERN

- CPU usage in FTS production instances can reach up to 85% in high load conditions
  - This traffic could interfere with other transfers

Conclusions

- Integrating the usage of the end-user cloud storages for CMS analysis activity allows to users to:
  - make use of their own storage in private/public clouds
  - benefit of “Sync and Share” capabilities
- CERNBox has been integrated with success as a grid resource for the distributed data analysis
  - Gained experience for transparent integration of end-users cloud storage in CMS computing infrastructure
- The data movement from grid to cloud, CERNBox, using FTS3 has shown good performance during the commissioning exercise
- Next:
  - Expose and ramp-up the service in production for real physics analysis
  - Explore the actual level of interest of CMS physics groups to “Sync. and Share” capabilities

Thank you!

hassen.riahi@cern.ch