
INFN-Box

— Giacinto Donvito, Stefano Stalio —
Istituto Nazionale di Fisica Nucleare

Experiences with ownCloud

- Many INFN sites have their own ownCloud installation
- Used mainly for personal data synchronisation and group sharing
- **Different technologies** chosen as for storage backend, authentication, db engine, virtualization platform, load balancing, high availability, ...
- **Different experiences** are precious
- Time to **wrap up and rationalize**

Experiences with other sync & share tools



- Presently INFN offers its staff and associates a **centrally managed cloud storage** system based on Pydio
- Used mainly for software and document distribution
- Relying on the INFN AAI for user authentication/authorization
- Deployed on a single site
- The INFN expertise relies heavily on technological projects carried on in these past years

The INFN Corporate Cloud



- **INFN-CC** is a project for a **multi-regional OpenStack installation** where some services are common to all regions while other services are local and associated to a single region
- **INFNbox** and **INFN-CC** will probably be deployed on the same INFN sites and will share part of the infrastructure: posix storage, object storage, network, authentication. INFN-box might exploit the INFN-CC IaaS capabilities.
- **INFN-CC** and **INFN-Box** are not strictly connected: both could exist by its own. But **INFN-CC** will be a good base for implementing **INFN-Box**.

INFN-CC highlights

- One of the highlights of INFN-CC is the **distributed object storage infrastructure** that allows for transparent geographic data replication. Has many useful applications
- **Geographically distributed SQL server** with Percona XtraDB

Both technologies have been thoroughly tested and are ready to be put in production. Both very useful for the deployment of a **geographically distributed cloud storage** service.

- **INFN-AAI** provides a single AuthN and AuthZ infrastructure for the whole INFN
- A huge work, mainly for the integration of Identity Management of 4 National LABs and 30 units into a single one: GODiVA (Gestione Ospiti, Dipendenti, Visitatori ed Associati)
- **INFN-AAI** services are accessible via LDAP and SAML2 protocols, through a network of distributed 389-DS servers and some SAML2 IdPs.
- All INFN Kerberos5 REALMs involved into AuthN phase, via a private plug-in of 389-DS
- AuthZ data (entitlement and group) provided via LDAP and SAML2, comes directly from the INFN's IDM

INFN DNS-HA



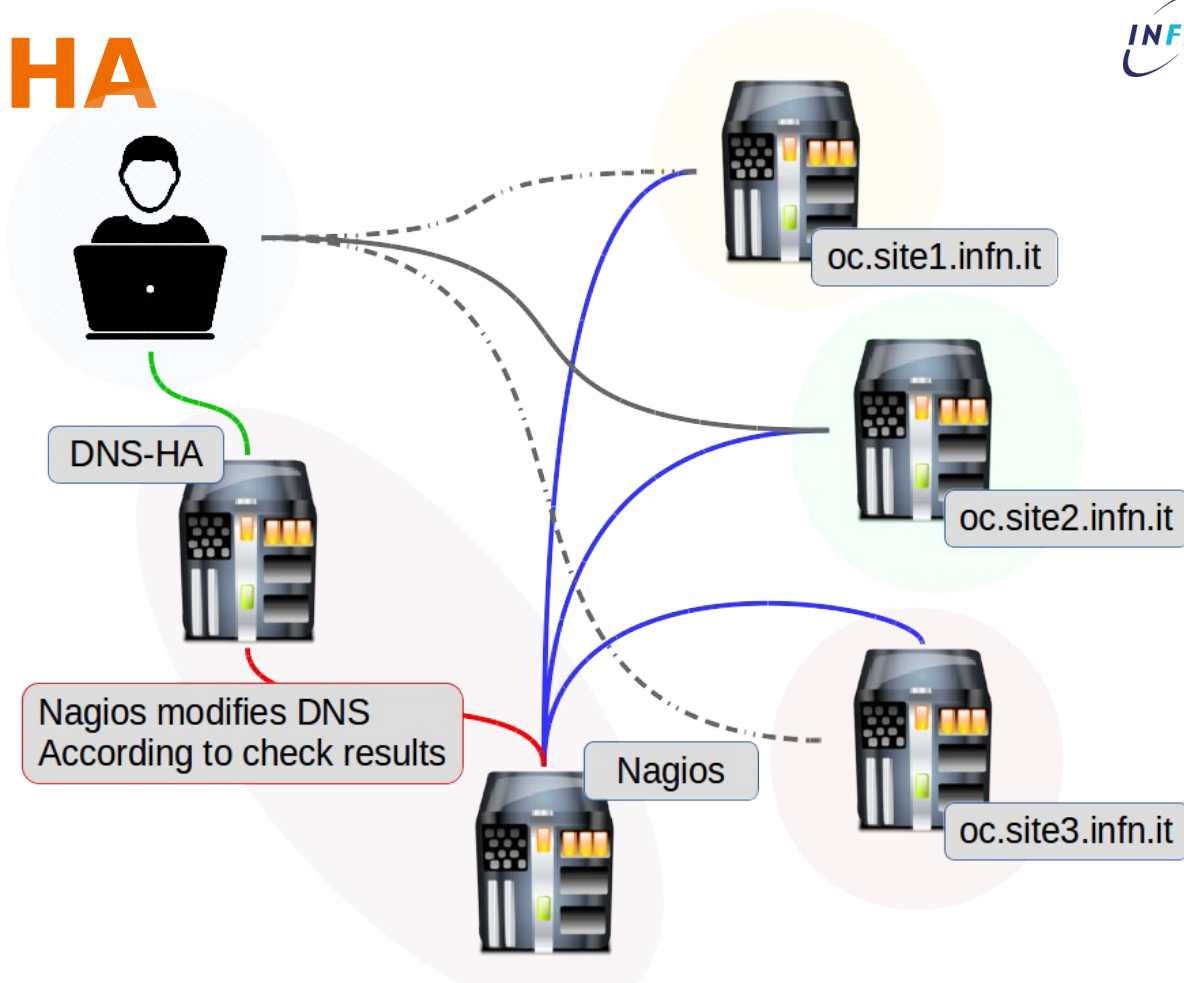
- **A distributed, resilient DNS infrastructure** able to maintain full functionality in case of “disaster” in any of the INFN Computing centers hosting it;
- Nagios (or similar) can dynamically update the ip address(es) pointed to by a hostname providing a certain network service;
- **INFN DNS-HA acts as HA system and/or load balancer** for geographically redundant services.

[https://agenda.infn.it/getFile.py/access?
contribId=20&sessionId=2&resId=1&materialId=slides&confId=7443](https://agenda.infn.it/getFile.py/access?contribId=20&sessionId=2&resId=1&materialId=slides&confId=7443)

INFN DNS-HA

The nagios service for DNS-HA is deployed in a separate INFN site and has an *external view* of the system status

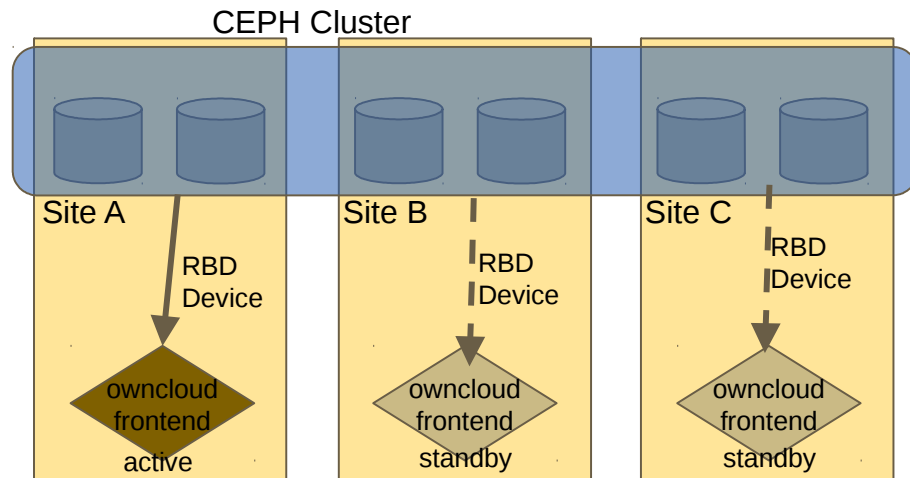
DNS-HA is distributed itself



- CEPH can provide very big block devices that the ownCloud frontend can use as local file-system
- In addition to the standard posix storage, ownCloud users or admins could add external CEPH (object) storage
- Using CEPH block storage is possible in order to create different devices for different scientific groups
- Limits of a geographically distributed CEPH cluster, although with large bandwidth and relatively low latency, still to be fully understood

CEPH

The CEPH Cluster will be split across different sites in order to provide the redundancy needed for the high-availability of data



Advantages of an organization-wide s&s system

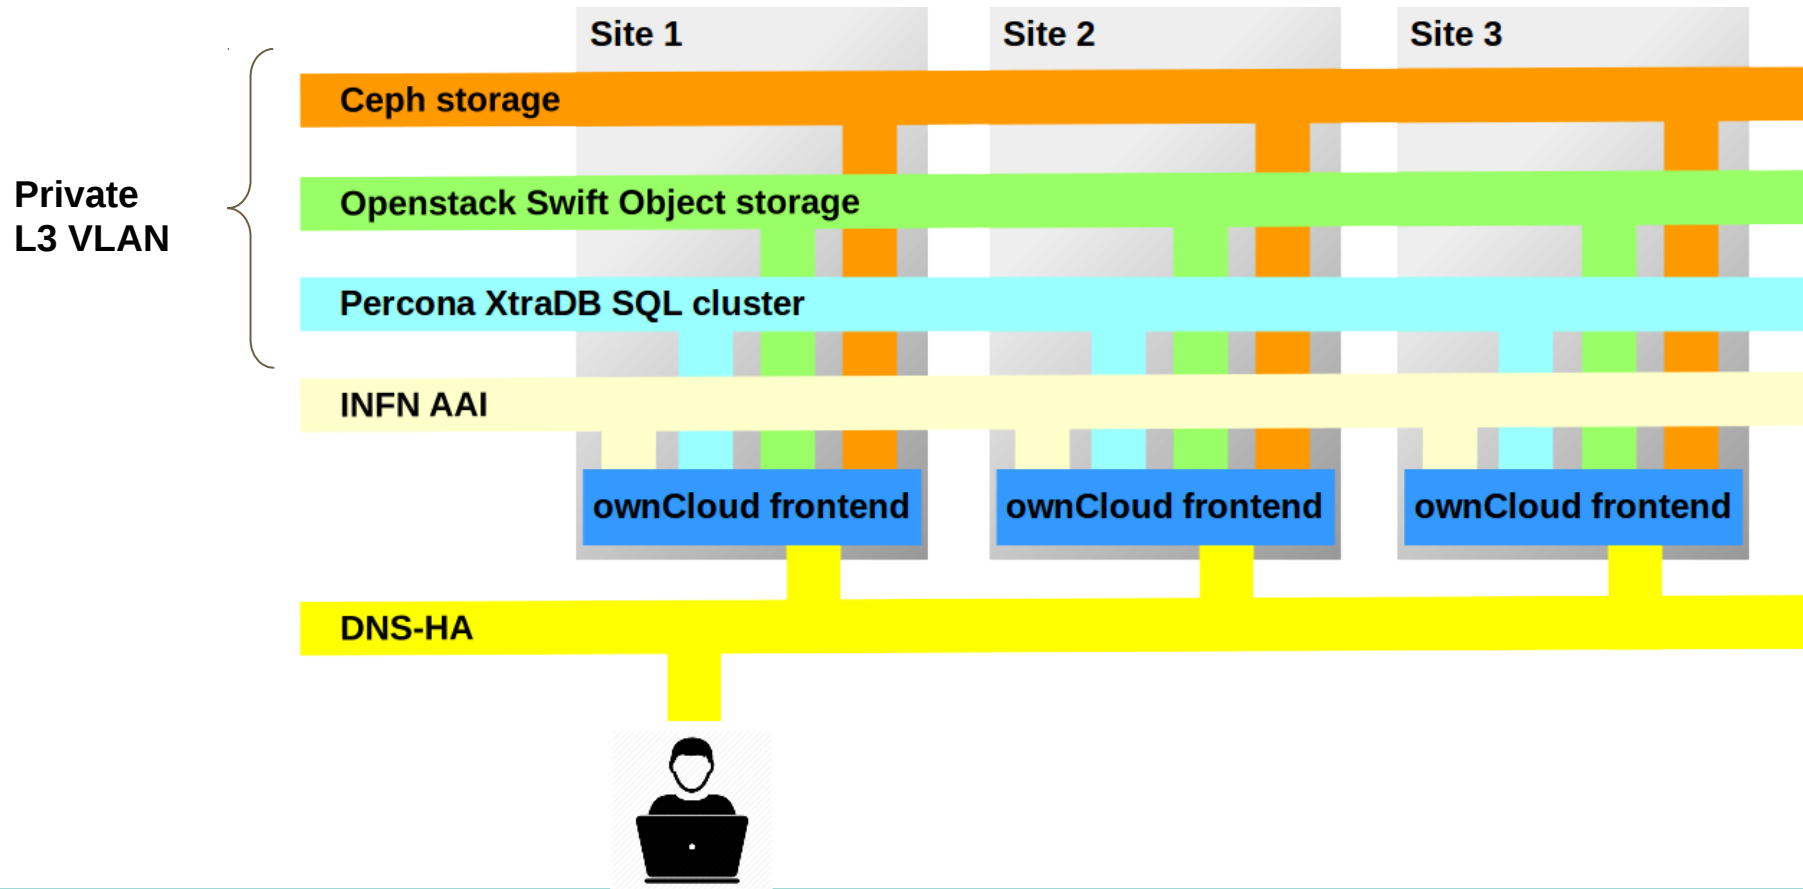


- The GARR network connection will provide a **highly performant and reliable backbone** that will also be used to provide disaster recovery & geographically distributed HA solutions
- **Exploiting skills and expertise** coming from persons based in different INFN departments
- End-users will use the service for **sharing documents and small experiments datasets** with colleagues all over the agency
- The service could be easily exploited to **share files with colleagues from other organizations**

Putting it all together

- Single sync & share infrastructure
- Distributed over ~~two~~ or three INFN sites
- Full geographic data and access redundancy
- Exploiting INFN/GARR network infrastructure
- Exploiting INFN-CC object storage infrastructure
- CEPH
- Percona Xtradb Cluster
- INFN DNS-HA for geographic load balancing and HA
- Integrated with INFN-AAI

General layout



What is happening around us?




What is CERNBox ?



CERNBox provides a cloud synchronisation service

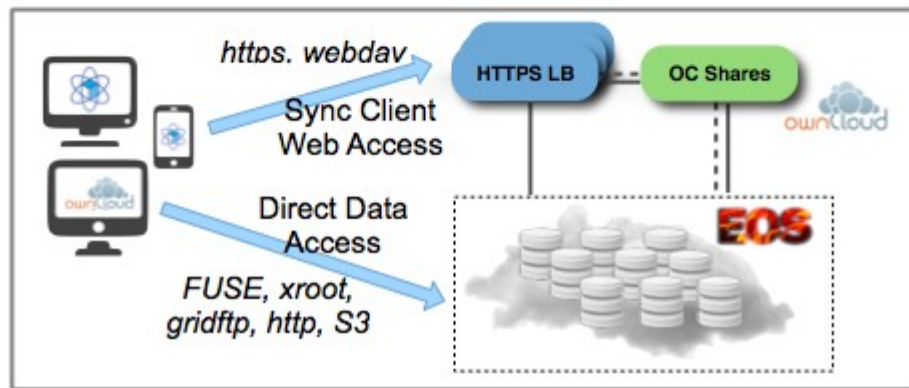
- Available for all CERN users (1TB/user)
- Synchronise files (data at CERN) and offline data access
- Easy way to share with other users
- All major platforms supported
- Based on **ownCloud** integrated with **EOS**



CERN BOX


$$\text{CERNBox} = \text{EOS} + \text{owncloud}$$

- EOS offers “virtually unlimited” cloud storage for end-users
- Full solution compatible with **ownCloud clients**
- **Remove dependency** on the ownCloud Database
 - EOS has a very fast **in-memory namespace**
- Integrate **web-access** and **sharing, versions, trash bin**
- Significant improvement of the external storage concept
 - **Great scalability**, as good as the underlying backend
- **Great performance**, efficient hardware usage



Access Methods: FUSE



```
[lmascott@lplus2015 ~]$
[lmascott@lplus2015 ~]$ df -H -t fuse
Filesystem      Size  Used Avail Use% Mounted on
eosuser         500T   70T  437T   14% /eos/user
eosatlas        36P   17P   20P   45% /eos/atlas
eosalice        20P   11P   8.5P   57% /eos/alice
eoscms          20P   14P   15P   49% /eos/cms
eoslhcb         13P   7.6P  4.6P   63% /eos/lhcb
eospublic       16P   5.0P  11P   36% /eos/public
[lmascott@lplus2015 ~]$
[lmascott@lplus2015 ~]$ ls -lc /eos/user/L/lmascott/
total 6644
drwx----- 1 lmascott c3      5 Dec 10 15:58 CERN
drwx----- 1 lmascott c3      0 Jan 26 18:18 debug
drwx----- 1 lmascott c3      0 Dec 11 09:43 download
drwx----- 1 lmascott c3      0 Oct 31 18:24 pdf
drwx----- 1 lmascott c3      1 Dec 11 09:44 personal
drwx----- 1 lmascott c3      8 Dec 10 12:11 pictures
```



CERNBox

fs

webdav

xroot

sync

share

mobile

web

ACLs


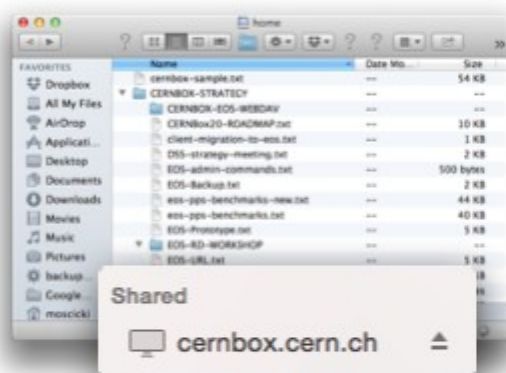



Namespace

Physical Storage


010110
110011
101000
0001

Access Methods: WebDAV







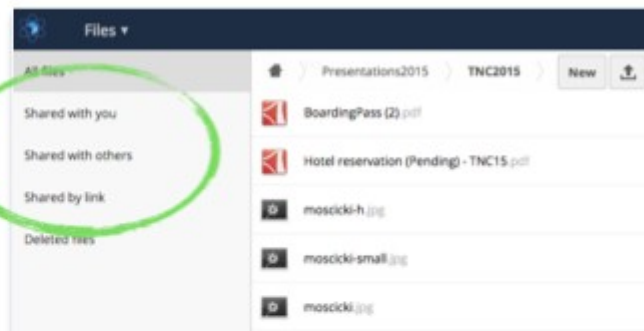
Embedded ROOT viewer in CERNBox browser



Access Methods: Sync



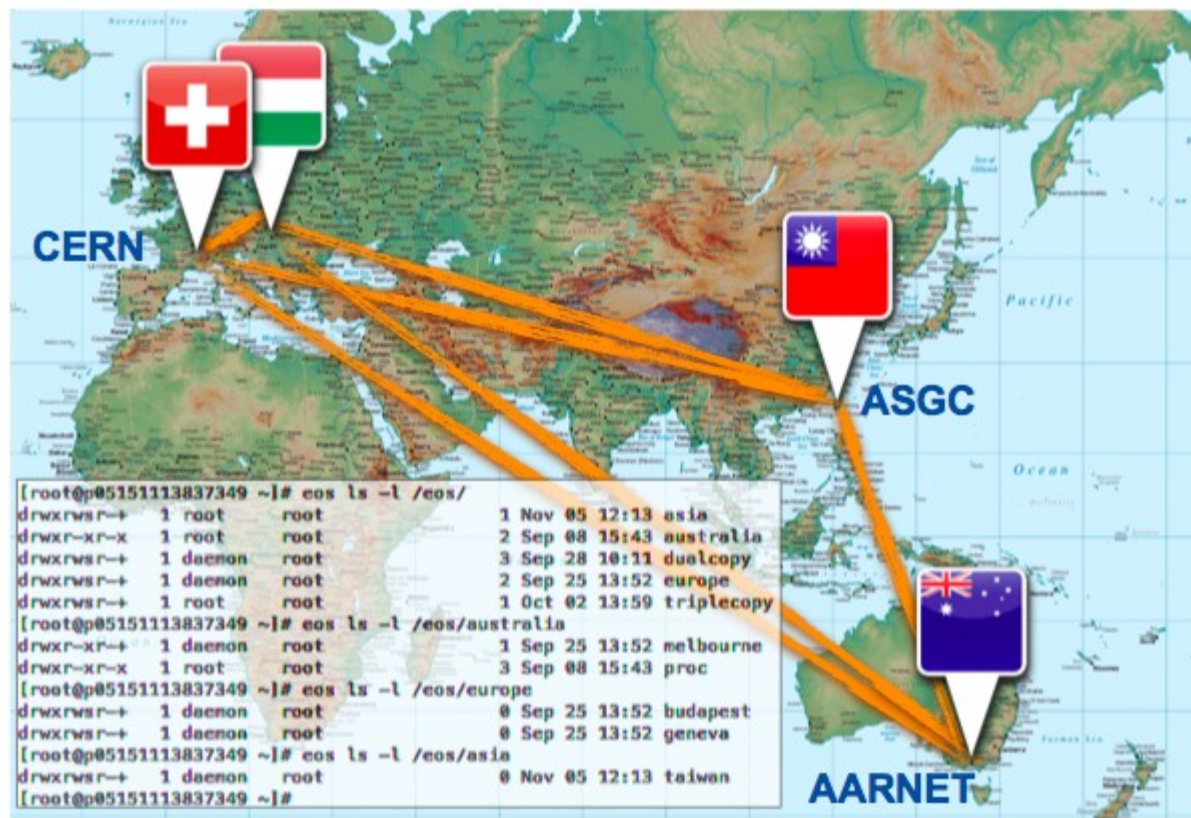
Access Methods: Sharing



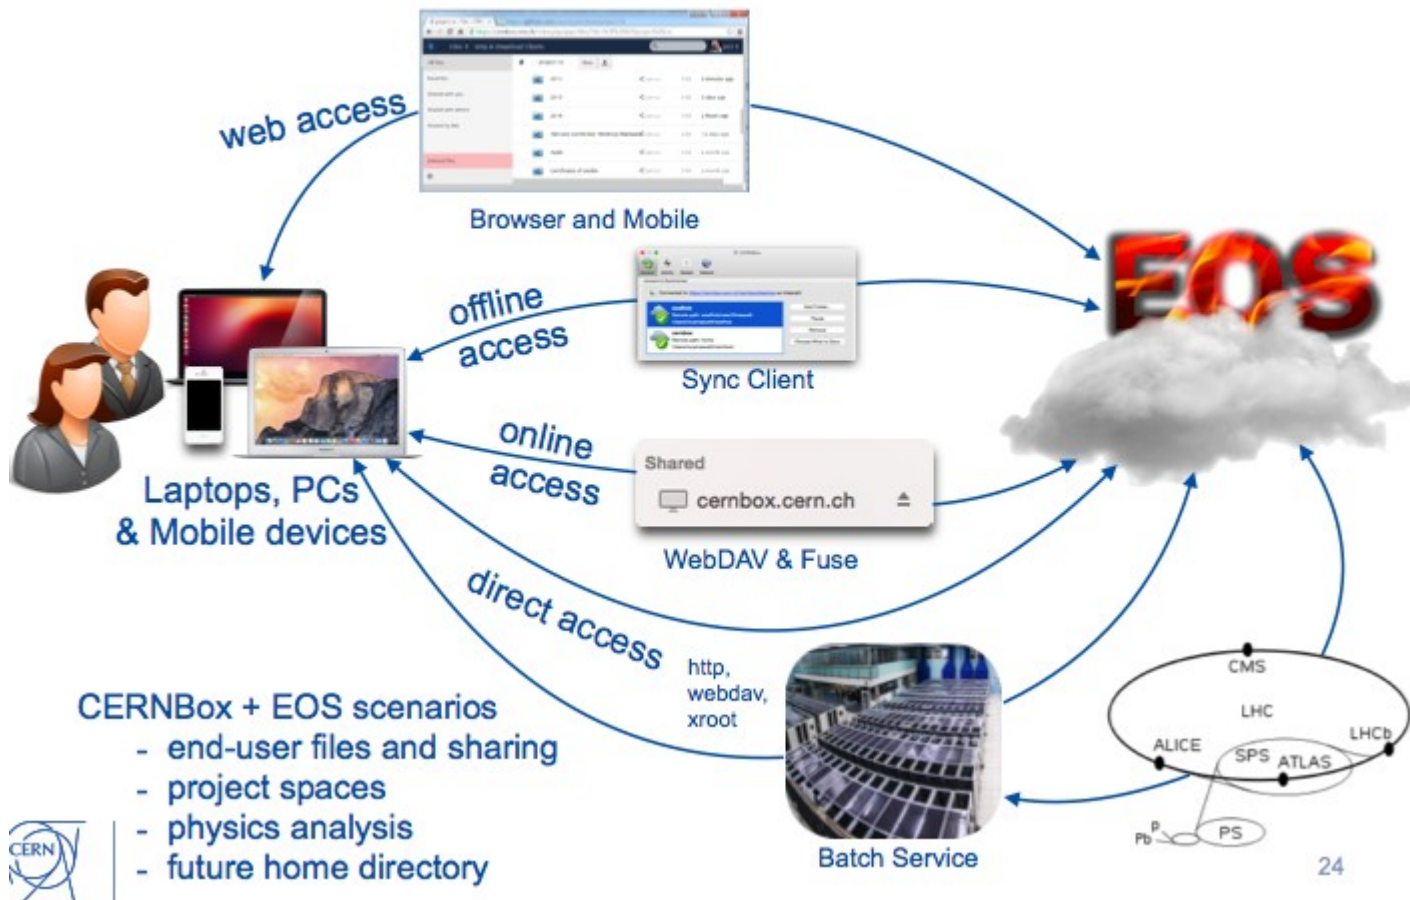
Access Methods: Mobile & Web



R&D - EOS World-Wide Deployment



Future Home Directory (\$HOME)



OpenCloudMesh

OpenCloudMesh is a joint international initiative under the umbrella of the GÉANT Association that is built on ownCloud's open Federated Cloud sharing application programming interface (API) taking Universal File Access beyond the borders of individual Clouds and into a globally interconnected mesh of research clouds - without sacrificing any of the advantages in privacy, control and security an on-premises cloud provides. OpenCloudMesh provides a common file access layer across an organization and across globally interconnected organizations, whether the data resides on internal servers, on object storage, in applications like SharePoint or Jive, other ownClouds, or even external cloud systems such as Dropbox and Google (syncing them to desktops or mobile apps, making them available offline).

Key stakeholders

Name	Organization	Interest / Involvement / Role	RACI	Stakeholder Comments
Peter Szegedi	GÉANT	Project management (about 200h/year)	A, R	Committed
Jakub Moscicki Massimo Lamanna	CERN	Project management	A, R	Committed
Frank Karlistschek Christian Schmitz	ownCloud Inc.	Project management	A, R	Committed
Charles du Jeu David Gillard	Pydio	Contribute to the specifications and development	R, C	Committed
Rogier Spoor	SURFnet	Contribute to the specifications and development	R, C	
Ron Trompert	SURFsara	Contribute to the specifications and development	R, C	Committed
Christoph Herzog Simon Leinen	SWITCH	Contribute to the specifications and development	R, C	Committed
Guido Aben David Jericho	AARNet	Contribute to the specifications and development	R, C	Committed
Holger Angenent	Sciebo / Uni Münster	Contribute to the specifications and development	R, C	Committed
David Antoš	CESNET	Contribute to the specifications and development	R, C	
Frederik Orellana	DeIC	Contribute to the specifications and development	R, C	
Kurt Bauer	ACOnet	Contribute to the specifications and development	R, C	
Christian Kracher	University of Vienna	Contribute to the specifications and development	R, C	Committed
Jari Miettinen	CSC/Funet / EUDAT	Contribute to the specifications and development	R, C	
Andreas Eckey	Technische Universität Berlin	Contribute to the specifications and development	R, C	
Woojin Seok	KISTI	Interest from South Korea	I	
Benedikt Wegmann Ralph Krimmel	GWDG	Interest	I	

The new ownCloud 9.0



Comment on your files

Tag your files

Application Works

Both Calendar and Contacts have been updated for ownCloud 9 while the Mail app remains a beta feature. Collaborations with Spreed and Collabora Cloudsuite are bringing web conferencing and real time online document editing to ownCloud.

New Storage API

ownCloud is designed to store all file metadata in the ownCloud database. This includes information like sharing information, file IDs and more. As the database can become a bottleneck in extremely large systems, offloading these file handling activities and metadata to an alternative location can improve performance, allowing the same hardware to scale to support even more users.

The new ownCloud 9.0



- Updated Sharing
- Collaborate Across Servers
 - with 9.0 introducing auto-complete of names across servers.
- Securely Connect ownCloud Servers
- Improved Developers APIs
- Activities and Notifications

Status, roadmap and conclusions



- CERN offered us to collaborate using EOS under INFNBox
- ownCloud offered us to talk with RedHat to test CEPH Object Storage as backend
- We are now finalizing the general architecture as well as that of each component of the infrastructure
- In a couple of months we will submit our project to INFN National Computing Commission (CCR) for the formal approval
- At least 3 different sites have already provided the availability to work on this project (Bari, CNAF, LNGS)
- INFNbox service would like to be part of **OpenCloudMesh** initiative
 - To facilitate the interactions and sharing among the major research institution

References

<https://opensource.com/business/16/3/cern-and-owncloud>

<http://cs3.ethz.ch/program.html>

<https://wiki.geant.org/display/OCM/Open+Cloud+Mesh>

Contacts

giacinto.donvito@ba.infn.it

stefano.stalio@lngs.infn.it