



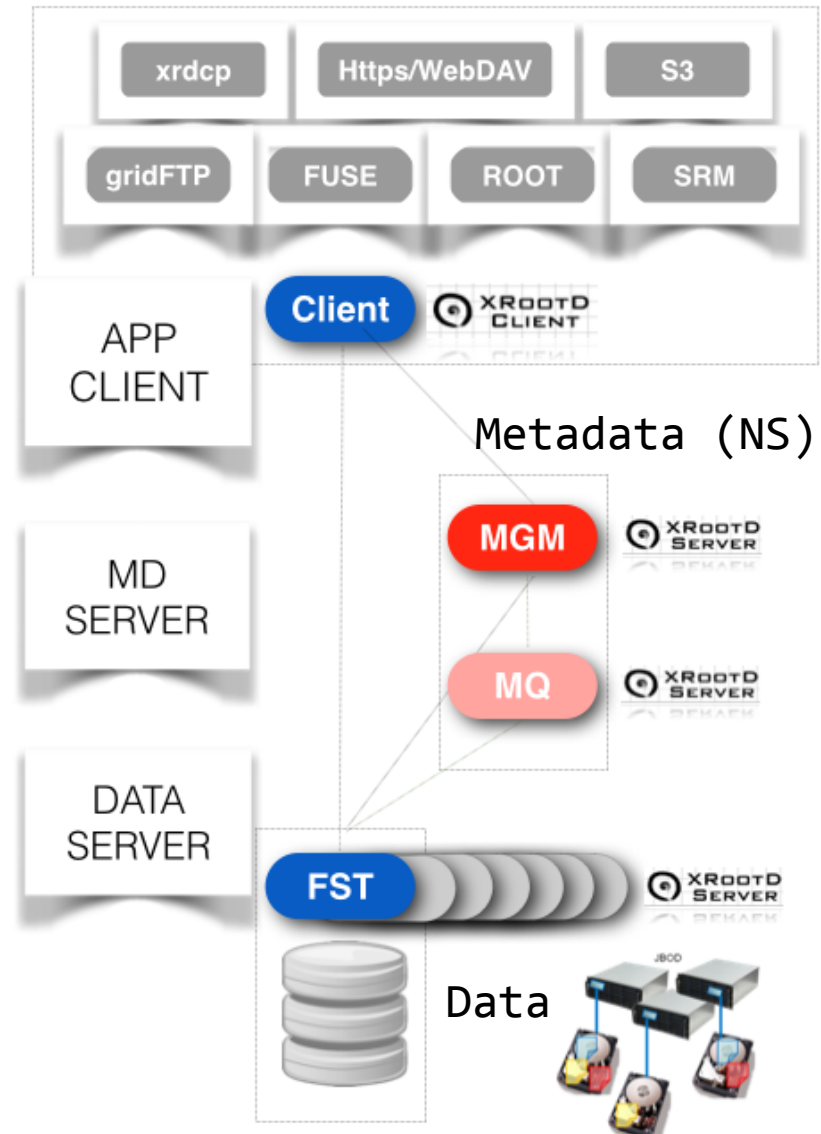
EOS and CERNBox Update

**Luca Mascetti
CERN - IT Storage**


Description and Architecture

EOS


- Project started in 2010
- Licence free
- Simple and scalable solution
- Easy to operate
- In-memory namespace
- Secure access (krb5, gsi)
- Quotas (user/group)
- Network RAID (RAIN)
- Tuneable QoS
- Dev&Ops in CERN/IT-ST



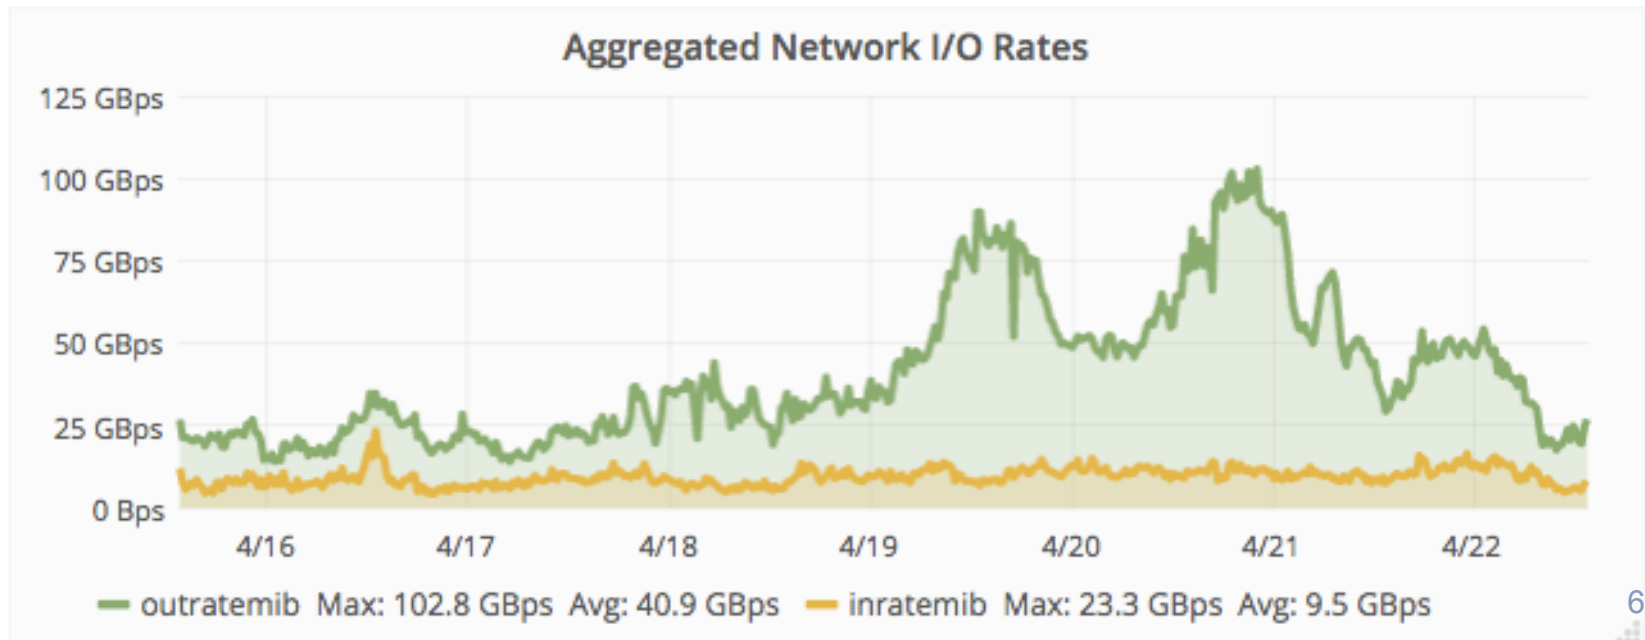



EOS Service @ CERN



- Disk Storage for all LHC and physics data
 - and for CERNBox
- Deployment across two computer centres
 - CERN and Wigner RCP
 - Third link recently added



EOS Service @ CERN



EOS Service @ CERN



Public
AMS
COMPASS
NA61/NA62
Others

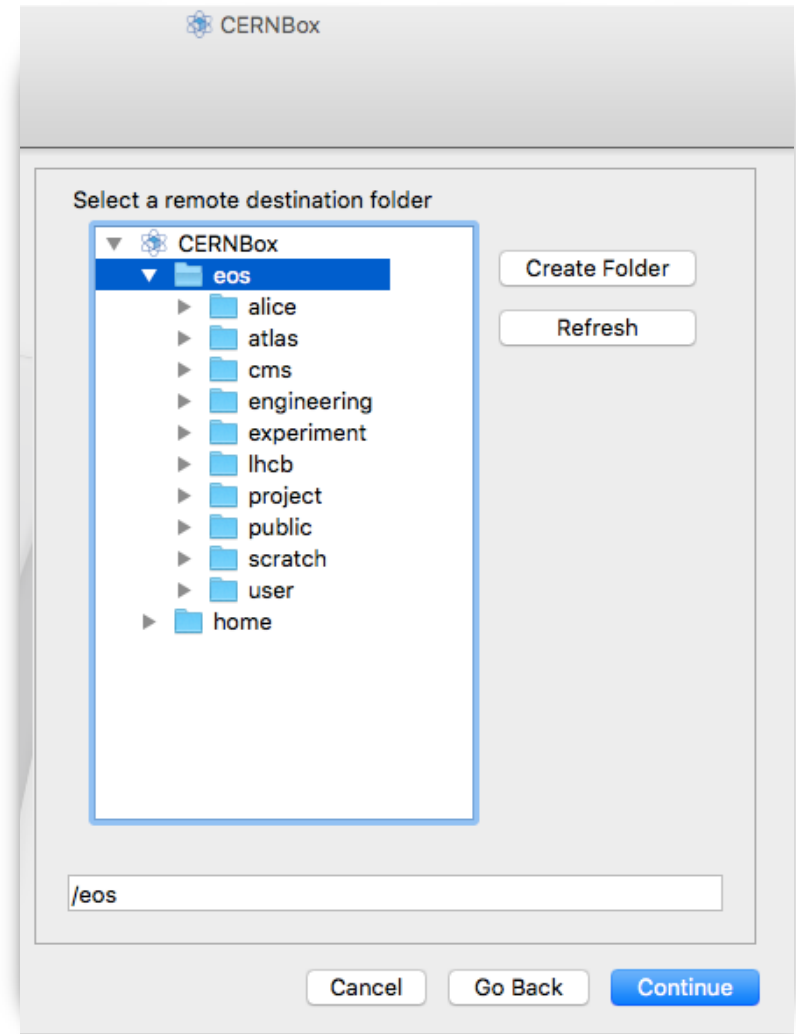

CERNBox




One instance per LHC Experiment

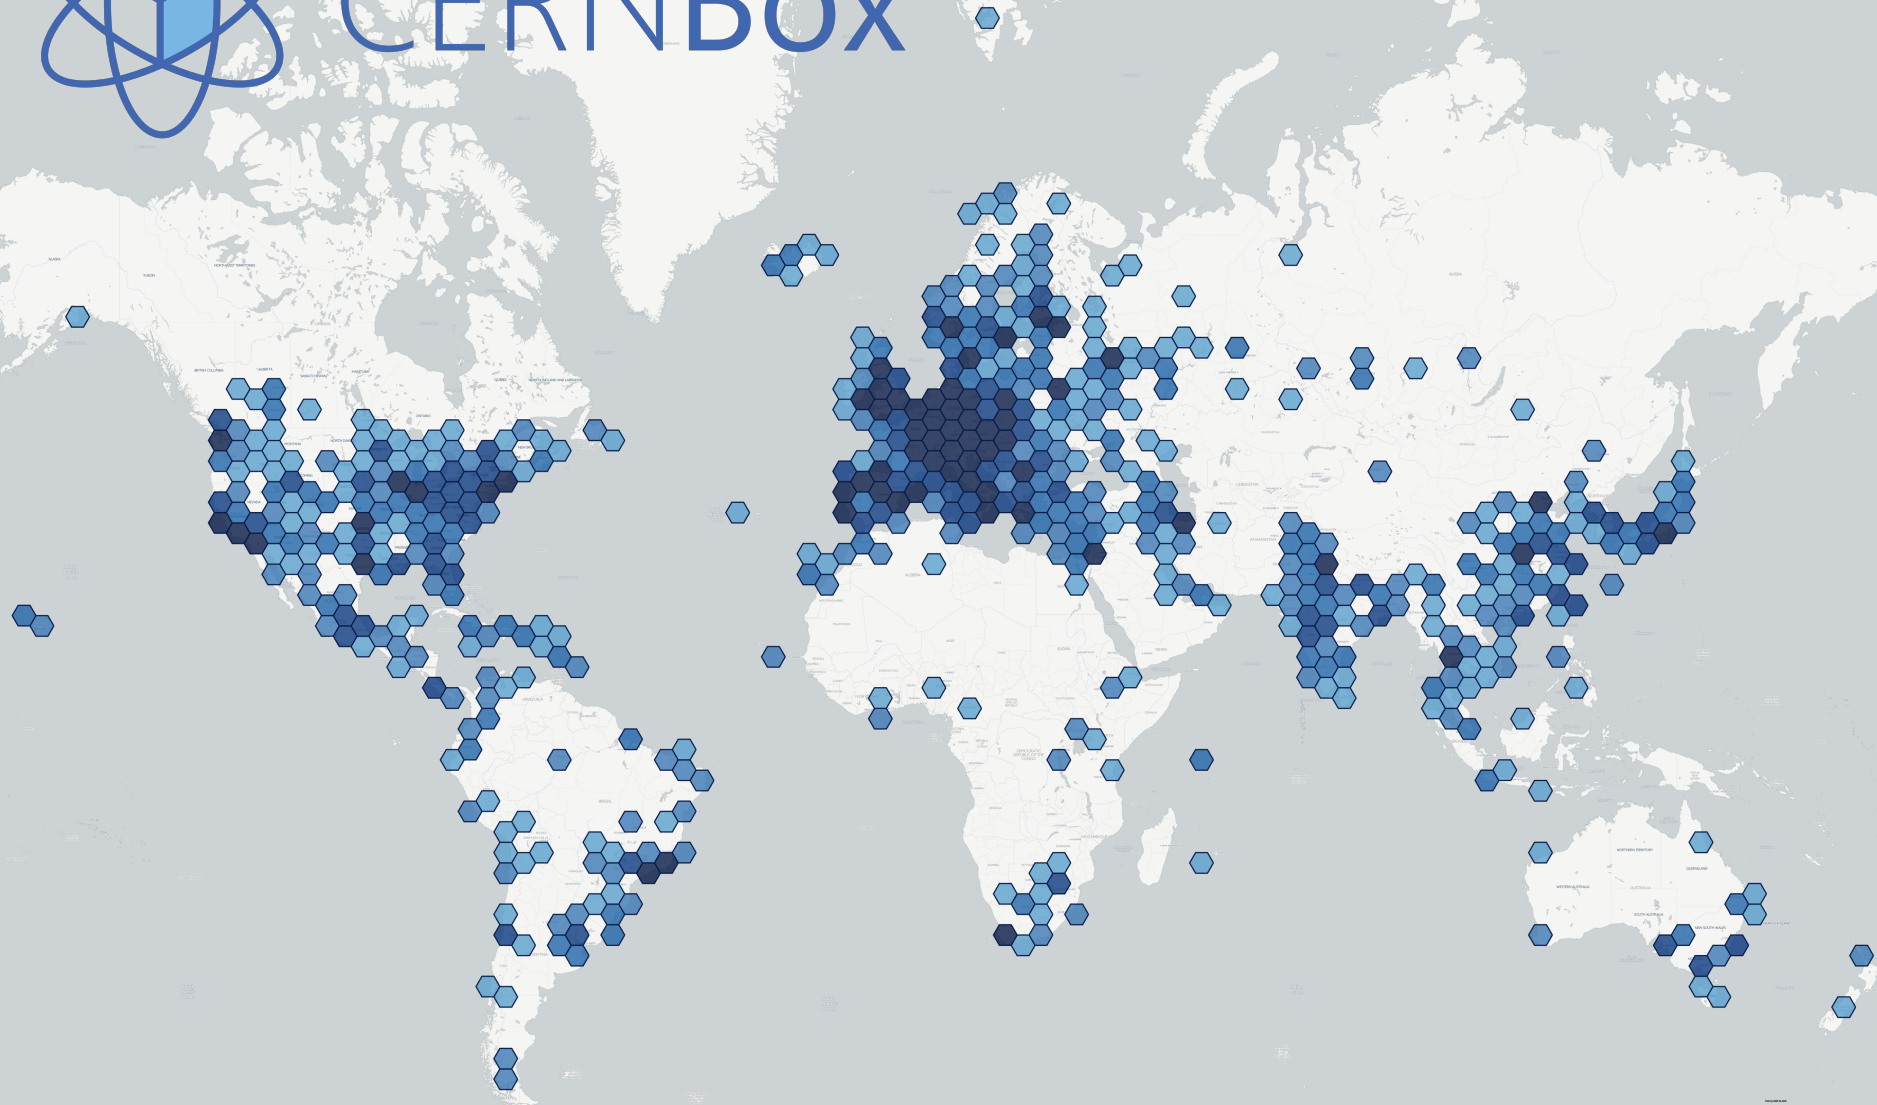



Bring data closer to our users: CERNBox





CERNBox






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**




CERNBox Service Numbers


	Jan 2016	Jan 2017	Apr 2017
Users	4074	8411	9501
# files	55 Mil	176 Mil	218 Mil
# dirs	7.2 Mil	19 Mil	20.6 Mil
Used Raw	208 TB	806 TB	1110 TB
Deployed Raw	1.3 PB	3.2 PB	3.3 PB




Number of files and directories

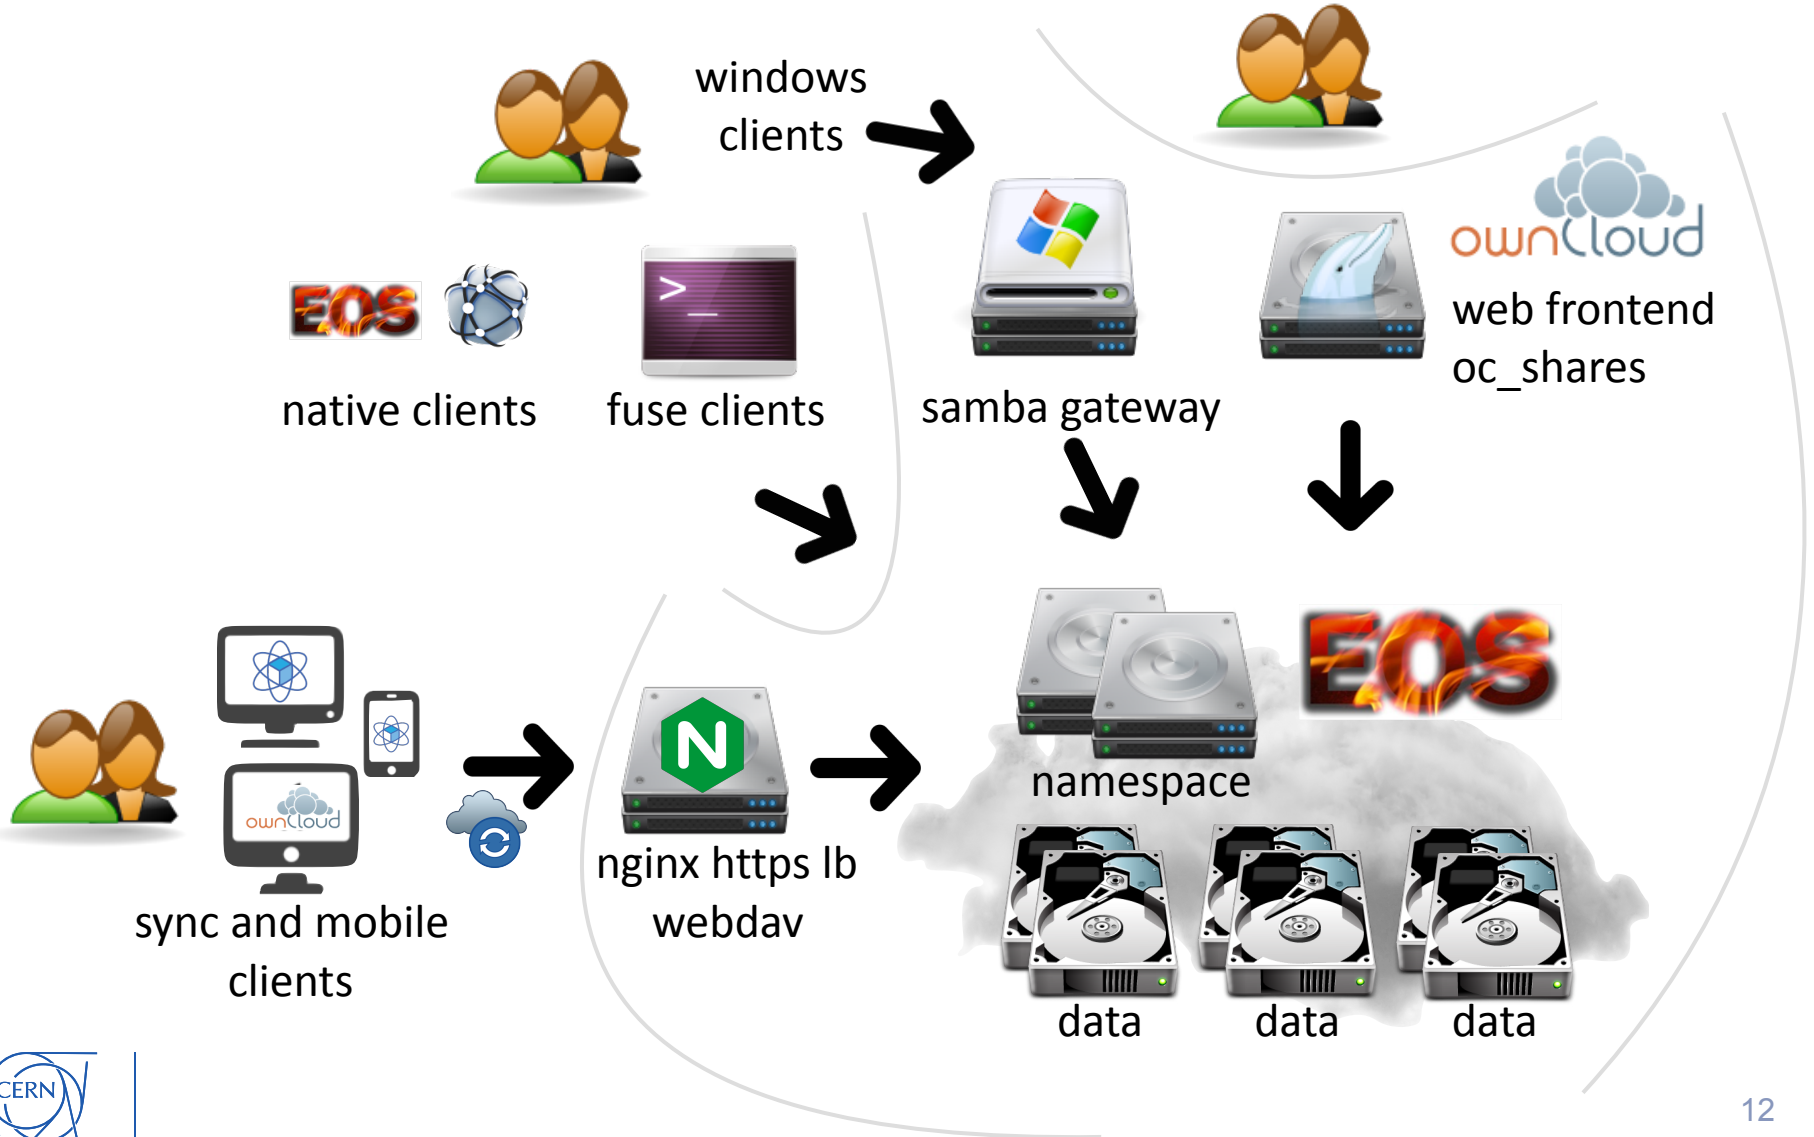


CERNBox users

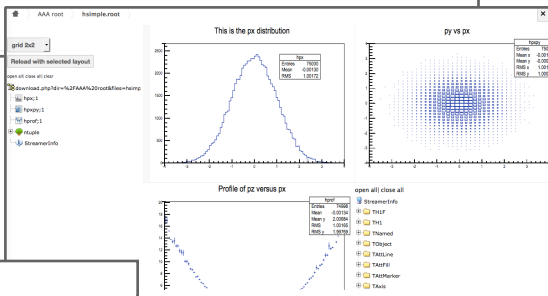
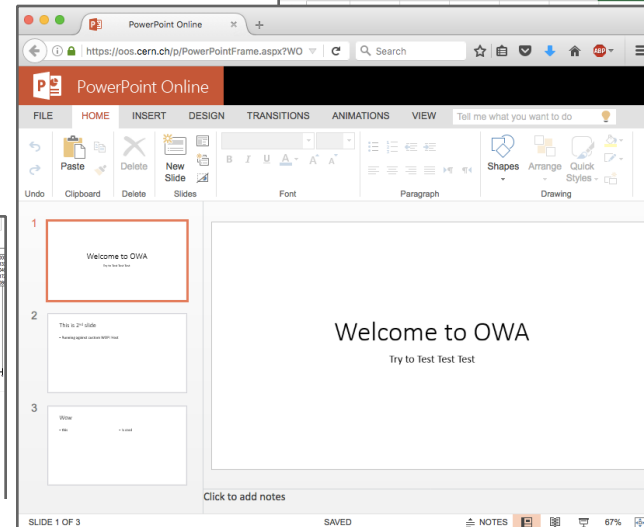
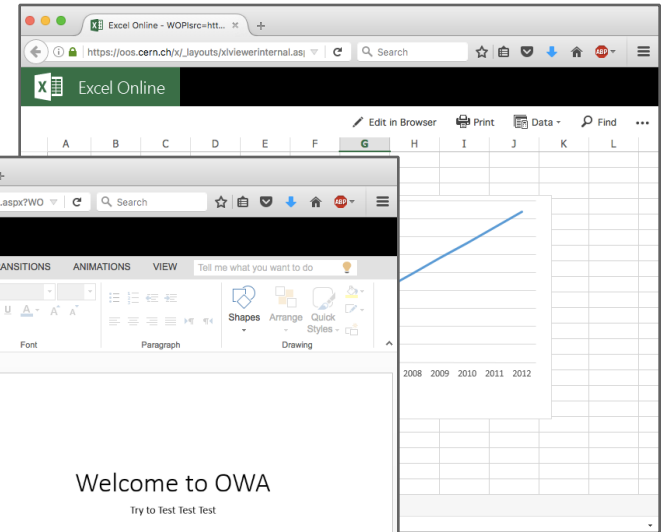
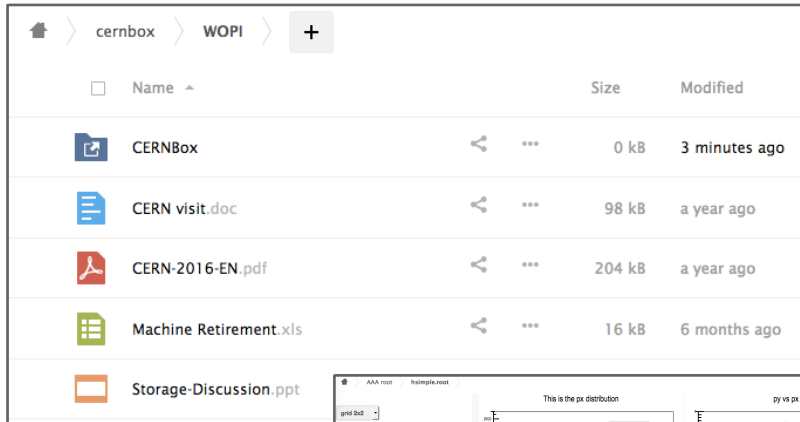




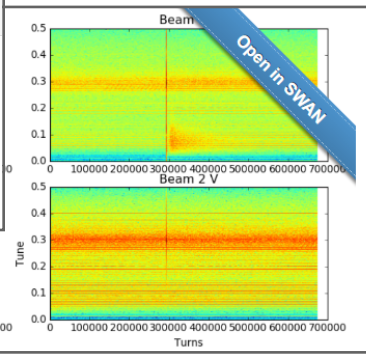
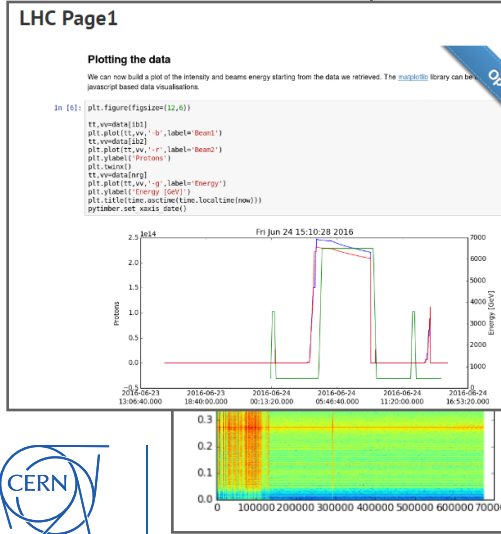
CERNBox Architecture



CERNBox as main entry-point



Open in SWAN




- Integration with multiple services
- ROOTJS
- Jupiter notebooks (SWAN)
- Microsoft Office 365
- and others to come...



Hardware Evolution

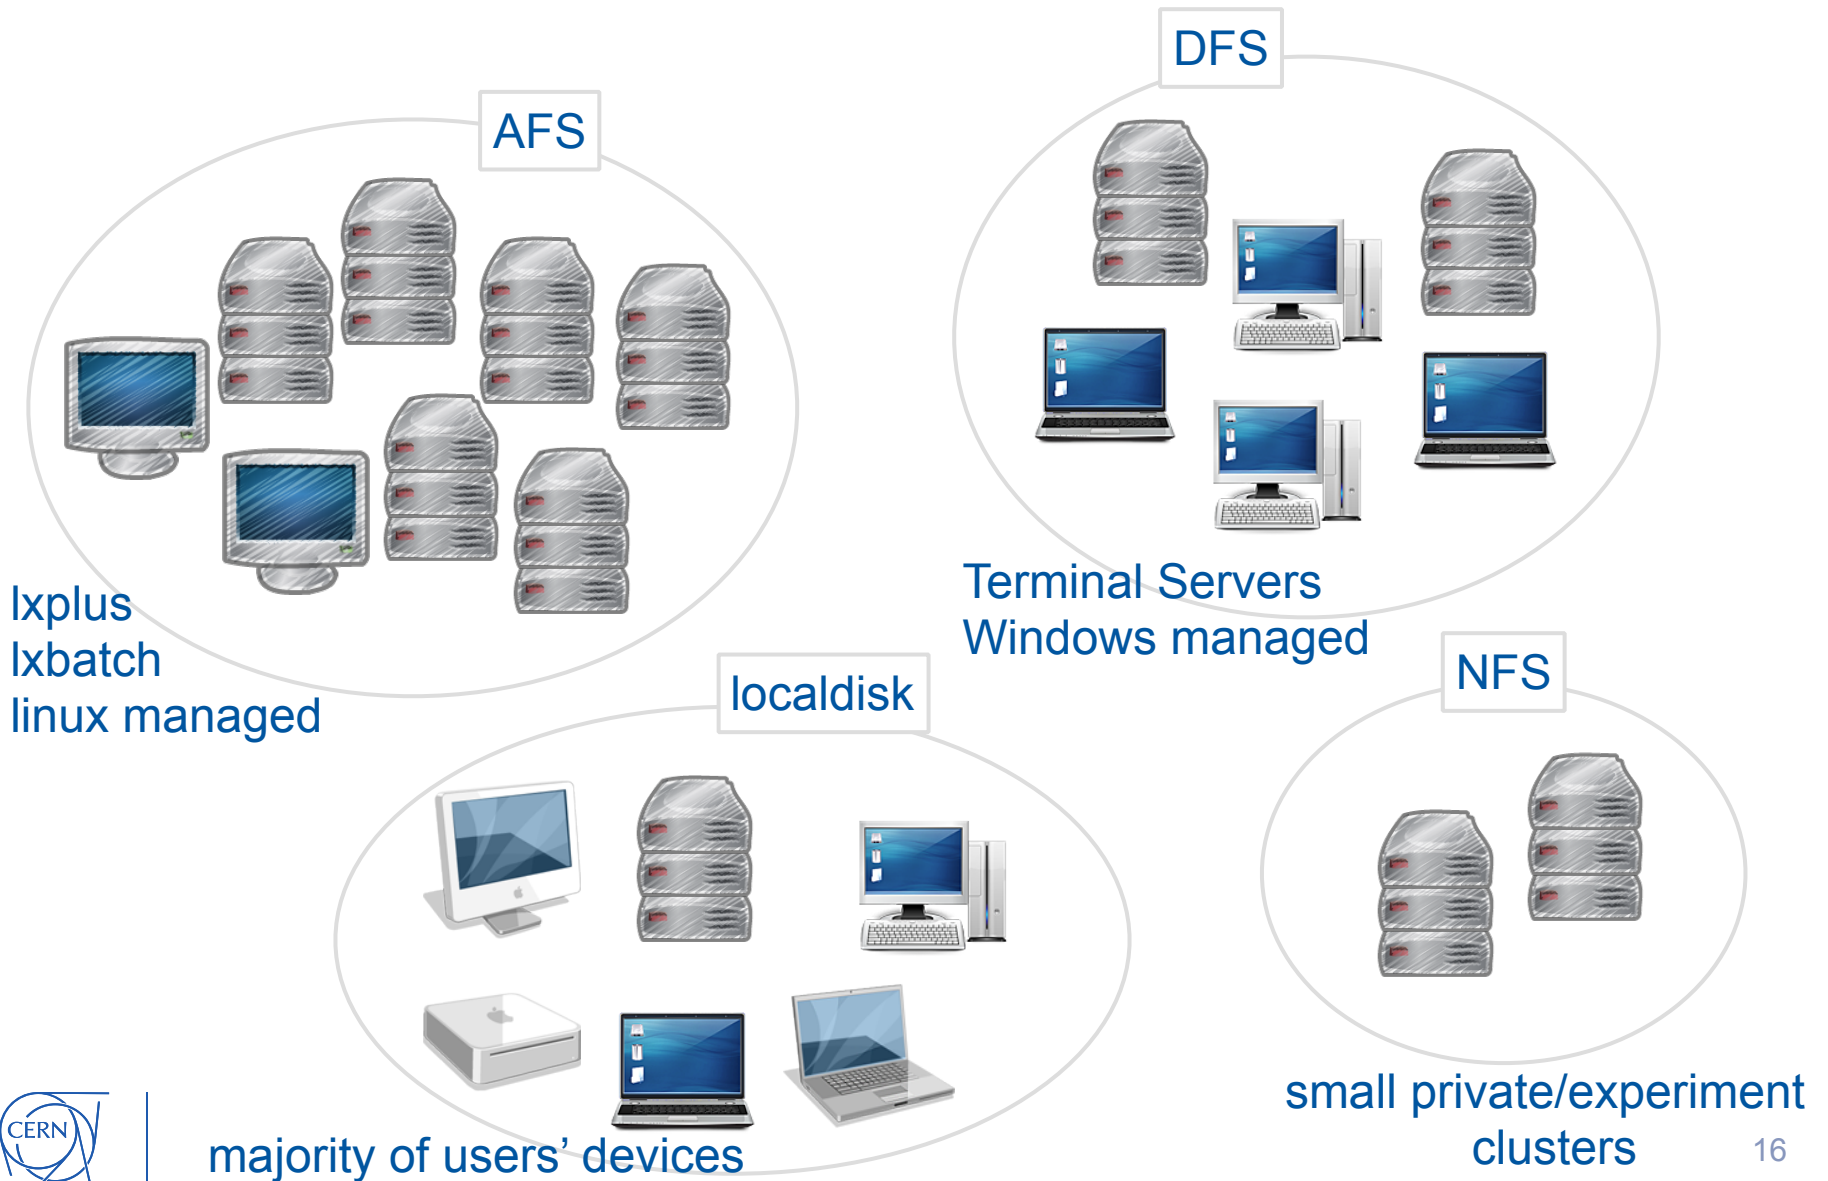
- Uniform storage for all our main services
- Profiting from economy of scale
 - minimise price per GB
- System Unit:
 - 8 physical cores (16 virtual) 64/128GB of RAM
 - 2x disk-tray of 24x 6TB HDDs (~**290 TB** brick)




- Prepared to receive 100PB on EOS (end of April/May)
- Planned test for the next generation
 - 4 trays per system unit + cascading (~600 TB brick)

Generic HOME Directory Service

\$HOME Current Status



\$HOME Future Vision



EOS\CERNBox



Ixplus
Ixbatch
linux managed



Terminal Servers
Windows managed



majority of users' devices

small private/experiment
clusters




Trigger

Current AFS Status

- Slow decline of AFS
- Concern for the long-term viability of OpenAFS
 - Future highly visible issues (IPv6, DES)
 - Increasing issues for client support
 - Data corruption or low performance on latest Linux versions
- 2015 IT Dept. level decision phaseout AFS
 - EOS FUSE as replacement strategy

CERN \$HOME




- Strategic choice to use EOS\CERNBox
 - flexible and powerful storage infrastructure
 - master the full storage development
 - supported by existing user base (> 9k users)
- Same HOME for multiple operating systems
 - enable synchronisation and disconnected access
 - improve user experience (no split-brain home)
- Serve with the same technology multiple community
 - Avoid to run separate/isolate storage cluster
 - better system interoperability and profit from future synergy (DFS)

```
[jdoe@lxplus066 ~]$  
[jdoe@lxplus066 ~]$ pwd  
/eos/user/j/jdoe
```

\$HOME will be set to `/eos/user/<u>/<username>/`

Desktop/Laptop



Each directory is a Sync Folder for CERNBox

```
Desktop -> /eos/user/j/jdoe/Desktop  
Documents -> /eos/user/j/jdoe/Documents  
Pictures -> /eos/user/j/jdoe/Pictures  
Music -> /eos/user/j/jdoe/Music  
Videos -> /eos/user/j/jdoe/Videos
```

Clusters

Ixplus, Ixbatch and terminals



\$HOME is set to /eos/user/j/jdoe/
prototype Ixplus-eos.cern.ch available soon

Cluster Desktops




If needed cluster desktop could be configured
with \$HOME in /eos/user/j/jdoe/ like Ixplus

Windows Terminal Server



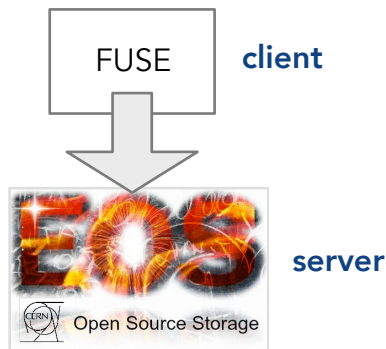
File access is via SAMBA



EOS FUSE Current Status


- /eos already mounted on lxplus and lxbatch
 - adding features + bugfix during the last months
- Most of the issues will be solved with the new eos fuse client
 - decision to start with a clean rewrite (EOSFUSEng)
- Behaviour will be similar to a local filesystem

V2 implementation



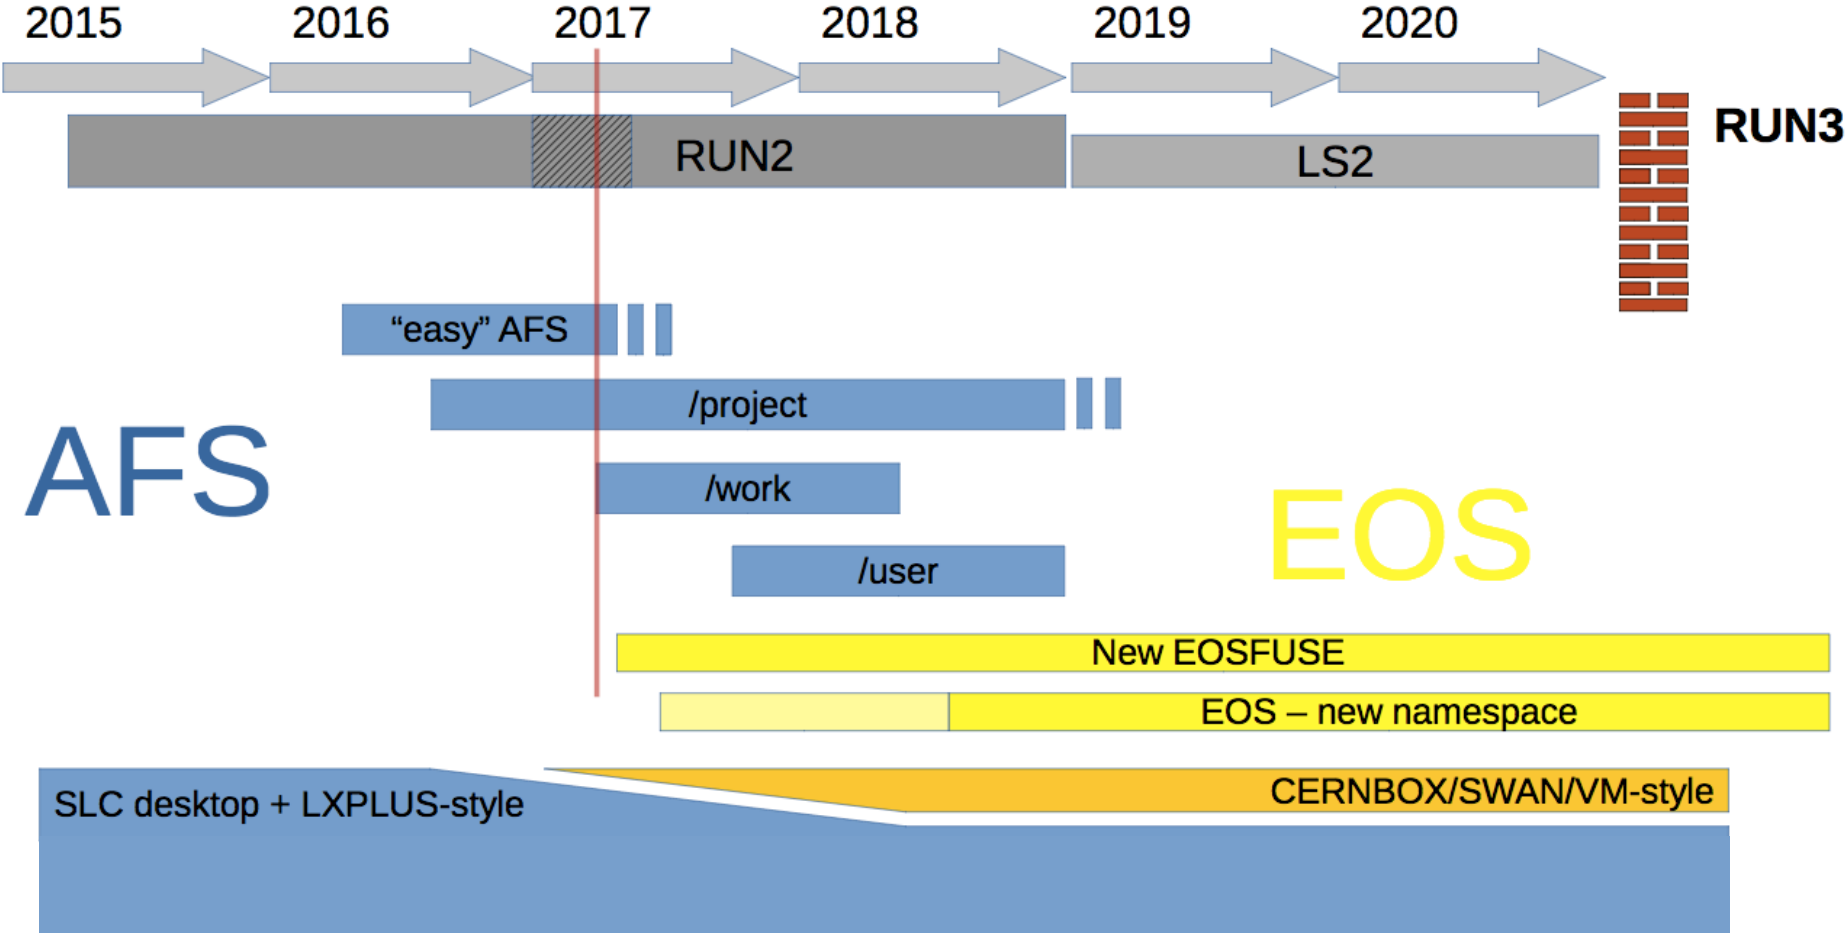
FUSE filesystem implemented as **pure client side** application without dedicated server side support.

V3 implementation



Dedicated server-side support providing a fully asynchronous server->client communication, leases, locks, file inlining, local meta-data and data caching

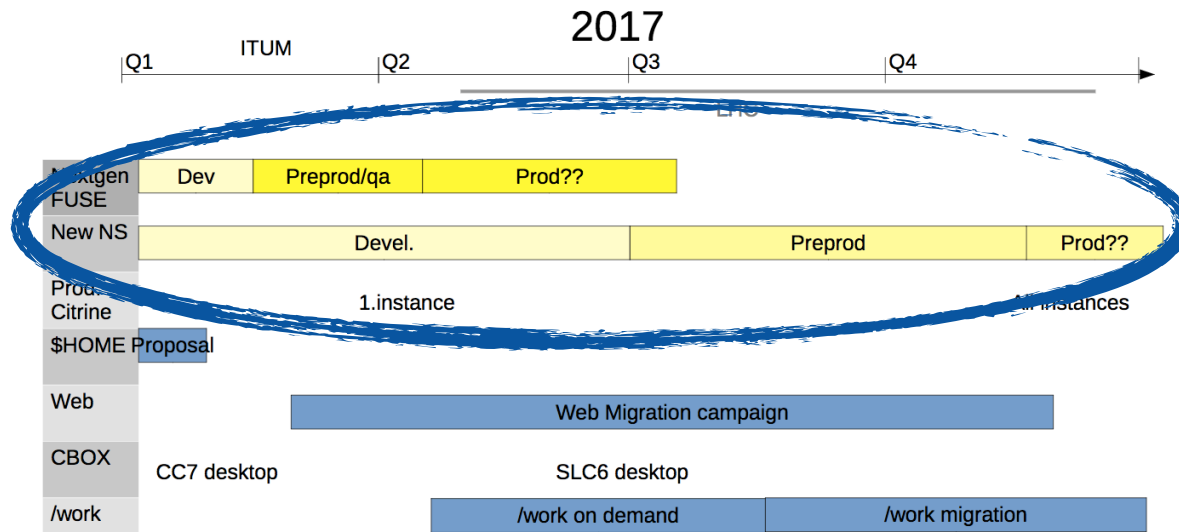
AFS Migration Plan



Major Milestones - Roadmap



- FUSE performance and localFS behaviour
 - development and deployment of **EOSFUSEng**
 - ixplus-eos.cern.ch** testbed
 - Namespace scalability (> 1 Billion Files)
 - Boost booting performance (x4 achieved)
 - New EOS server major upgrade
- Deployed in few instances
- Scheduled on May 15th for eoslhcb



Summary and Outlook

Summary and Outlook

EOS provides a **very flexible** platform for large communities

- integrated in Tier-0 workflow by ATLAS & CMS
- more than 9k users storing data today via **CERNBox**

Demonstrated unprecedented scalability

- largest **low-cost** HEP storage installation site today

CERNBox as an extension of the Desktop

- Bring data closer to our users
- New ways to interact with the data
- Ongoing integration with multiple services

Strategic direction for CERN based disk storage

- for physics data (user/group/grid)
- as generic home directory via **EOS fuse** and **CERNBox**



www.cern.ch

Again... This is really really amazing!!!! I am figuring out what I can do now!
Freedom!

All the best,

