



# News from **XRootD** &

# EOS

Disk Storage @ CERN

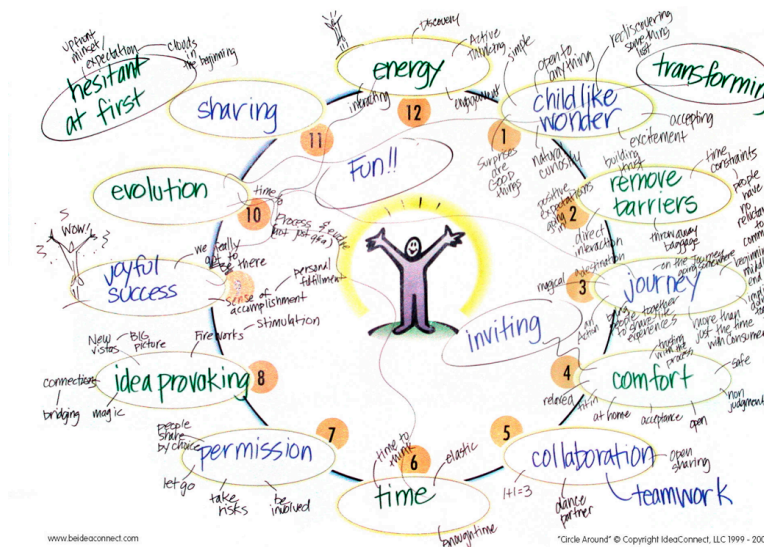
Andreas-Joachim Peters IT-DSS  
ALICE Tier 1/2 Workshop Turin  
23.-25.2.2015

**XRootD**

HTTPS SRM  
WebDAV  
HTTP OwnCloud  
FUSE gridFTP

EOS SRM  
HTTP OwnCloud

- News & Plans **XRootD** (4.0 / 4.1 / 4.2)
- Status of **EOS** Development & Operation
- **CERNBOX**







- **XRootD** Workshop held end of January 2015 in San Diego
- **XRootD 4.0** see <https://indico.cern.ch/event/330212>
  - major refactoring exercise - IPV6 dual-stack support
  - client replaced - XrdCl (multiprotocol support - client-side plugins)
  - caching proxy
  - remote debugging facility
- **XRootD 4.1**
  - EPEL packaging ( XRootD 3+4 compatibility packaging )
  - OSG packing missing
  - Proxy support for out-going traffic server-side
- **XRootD 4.2** ( March? )
  - CEPH/Rados plug-in files-only ( IT-DSS Sebastian Ponce ) [CASTOR]
    - full storage listing supported, no readV support (yet)
  - Throttling Plug-in ( Nebraska )
  - Proxy support for out-going traffic client-side






- Lukasz Janyst leaving IT-DSS in April  
(client developer)
  - maintenance taken over by Elvin <[Elvin.Sindrilaru@cern.ch](mailto:Elvin.Sindrilaru@cern.ch)>
- Summer student project proposal at CERN
  - work on HTTP 2.0 interface based on Facebook Proxygen HTTP library







- Current production version  
EOS Beryl 0.3.97 Aquamarine 
- @CERN
  - do many 'small' updates with adiabatic changes for continuous service stability and feature improvements (all instances 2014 > 99% available)
- focus in Aquamarine release
  - Archiving tools
    - freeze part of the EOS namespace and im- export to any other XRootD enabled storage ( e.g. CASTOR )
    - archive not accessed group data to free online space
    - incremental backup support for CERNBOX
  - OwnCloud integration  
(WebDAV, OC-DAV)



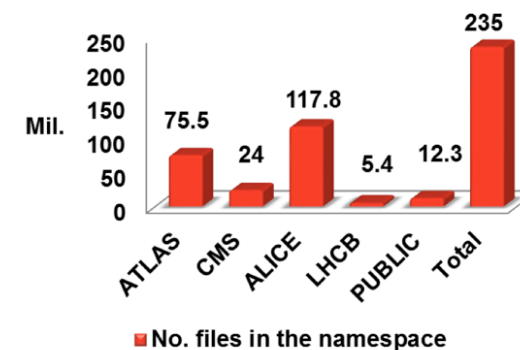
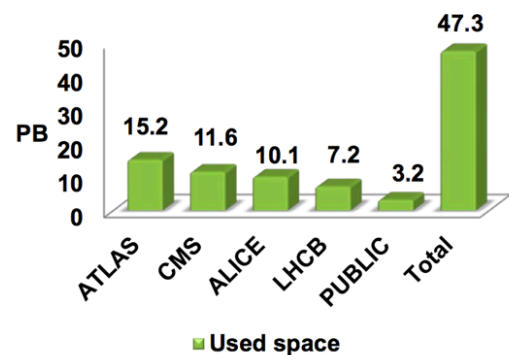
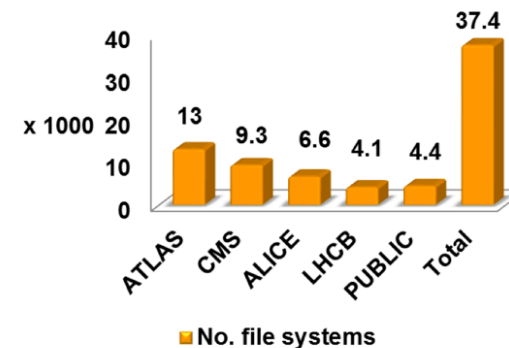
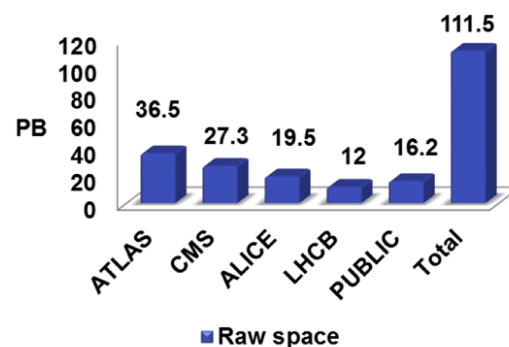


- **240 Nodes, 6.500 disks, 19 PB** raw volume
- **>0.5 TB/s** (local) disk bandwidth
- **190 GB/s** theor. network connectivity
- **267 Mio.** replicas
- no obvious core bug in production release - connectivity glitches & updates responsible for downtimes
- transition of the monitoring system (SLS to Kibana) ... currently no nice plots to show for longer history





- deployed **111.5 PB** raw space at CERN
- **37.400** disks




a lot of new small scale experiment customers served in public instance ...

*News on XRootD & EOS*





- **XRootD4** based version  
EOS Citrine 0.4.x **Citrine** 
- postponed due to XRootD release schedule
- main add-on: multi-site capabilities
  - geographic aware access & placement
    - host, rack, building, site awareness
  - scattered (max distance), hybrid and gathered (minimum distance) placement policies
  - CERN target: CERN-Wigner CC
- readV support in RAIN files

EOS native **primary** protocol is

## 1 XRootD

- used for File IO, internal communication/replication, CLI
- currently all EOS disk servers world-wide visible

LAN/WAN opt.

**Secondary** protocols

## 2,3 gridFTP-SRM

- provided via gateway machines using VDT implementation + DSI plug-in
- provided via BestMan, FUSE mount & quota scripts

WAN opt.

## 4 FUSE \*not a protocol

- client: implemented over XRootD protocol
- no locks, no symlinks, 95% compatible

POSIX flavor

## 5 HTTP

- currently provided via embedded HTTP server and overlay network

## 6 HTTPS

- provided via NGINX gateway, proxy module & add. Auth headers

## 7 WebDAV

- implemented in embedded HTTP server
- incomplete implementation

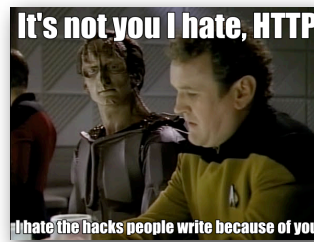
5 http flavors

## 8 S3

- implemented in embedded HTTP server
- incomplete implementation

## 9 OwnCloud HTTP

- implemented in embedded HTTP server
- compatible with OC http server



**9** more or less **supported protocols**

- used by experiments  
**XRootD > gridFTP >> FUSE, HTTP(S)**
- XRootD most reliable and performant implementation in EOS
- no new protocols foreseen  
deployment change on the horizon  
all external traffic in- & out via XRootD gateways
- planned to provide XrdCl plug-ins to enhance features and performance  
loaded by client according to plug-in policies

[News on XRootD & EOS](#)



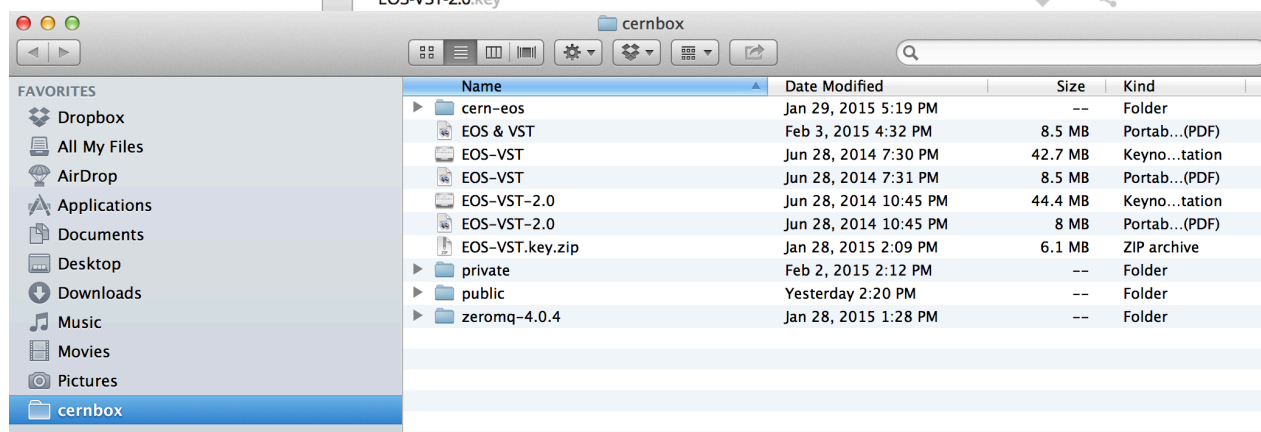
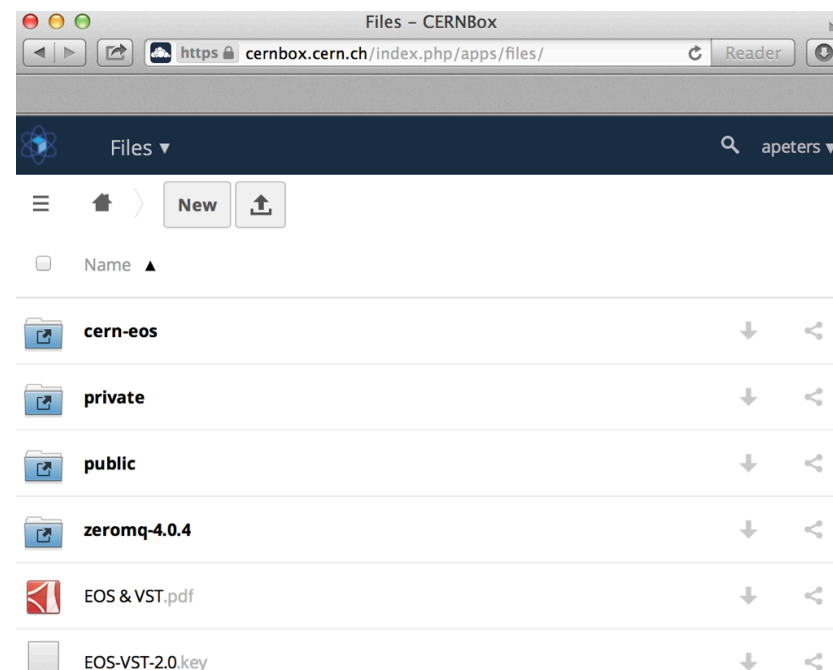
- add CEPH based object-store based namespace to in-memory namespace
  - remove scalability limit of namespace
  - reduce service start-up time ( e.g. 10 minutes for ALICE )
  - development of namespace (RadosFS) complete, interfacing to EOS now
- be prepared: **CEPH/XRootD** might be soon proposed as site **SE**
  - + CEPH at CERN 100% availability
  - count on 80% usable storage space
  - + EC encoding





# DSS

# CERNBOX



CERN IT Department  
CH-1211 Geneva 23  
Switzerland  
[www.cern.ch/it](http://www.cern.ch/it)

News on XRootD & EOS





- **dropbox-like beta** service based on OwnCloud with EOS back-end [ new EOSUSER instance ]
  - file versioning
  - garbage bin
- **1 TB storage per user** ( use it now: [cernbox.cern.ch](http://cernbox.cern.ch) )
- **accessible via**
  - mobile clients (Android, iPhone)
  - Web Portal ( access & sharing )
  - desktop sync clients ( Linux/Windows/Mac )
  - **soon**: as a remote drive
- **file/directory sharing by**
  - anonymous links
  - by CERN ID
  - **soon**: by EGROUP
- **soon**: SSO



- **what do you do with that?**

- **share/edit documents**, n-tuples **with collaborators**
  - world-wide accesible
- **share** contents **across your devices** using a trustable CERN service
- **store** private documents **using an encrypted file system** add-on
- develop **new workflows**
  - code and job output originates and ends up on YOUR laptop with offline availability
- get some extra free space accessible to you via XRootD & WebDAV protocol world-wide





DSS

# Get CERNBOX yourself...

CERN IT  
Department

- **start at** <http://cernbox.cern.ch>



RESEARCH

TECHNOLOGY

OPERATIONS

DEVELOPMENT

TESTING

DEPLOYMENT

MAINTENANCE

MONITORING

SECURITY

COMPLIANCE

ARCHITECTURE

INFRASTRUCTURE

OPERATIONAL

RESEARCH


*News on XRootD & EOS*



CERN IT Department  
CH-1211 Geneva 23  
Switzerland  
[www.cern.ch/it](http://www.cern.ch/it)

Tuesday, February 24, 15

- **grab** a virtual machine **openstack.cern.ch**
- **download** <http://eos.cern.ch/rpms/eos-deploy.v2>
  - run it and follow instructions
  - include 'alice' in the instance name to get a ready ALICE SE installed
- **read** EOS documentation  
<http://eos.readthedocs.org>
- pass by at CERN bld. 31 to install your site with an expert in IT-DSS in less than 30 minutes or via remote sessions
- YUM repository production release  
<http://eos.cern.ch/rpms/eos-aquamarine/>

A black and white photograph of a typewriter keyboard. A sheet of paper is placed over the keys, and the word "Questions?" is typed in a classic typewriter font. The paper is slightly wrinkled and the lighting creates soft shadows across the keys and the paper.

Questions?

Thank You