

Community Distribution for Analysis Facilities

Dr. Giuseppe Lo Presti, CERN/IT

On behalf of the CERNBox team





Contents

- Cloud Sync&Share Storages at CERN and beyond
- The Technology
- Community and Outlook:
Storage and Sync&Share for Analysis Facilities





Sync&Share Storage at CERN: CERNBox

CERN's cloud collaboration storage platform

- **1 TB** for every user, up to **10 TB** on demand
- Built purely with open-source components
 - ownCloud for “Dropbox” use-case
 - EOS and Ceph for storage

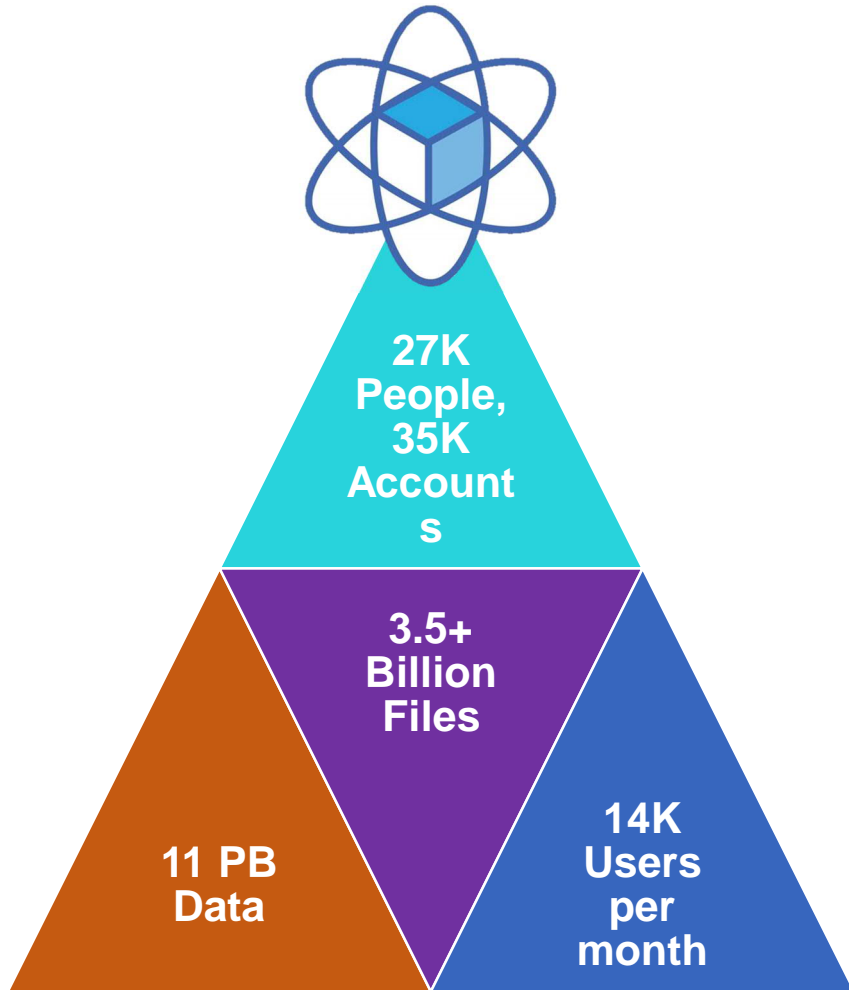


- Data safely stored on CERN premises
- Security policies to keep data safe and confidential





Usage and Features in a nutshell



Provides storage and sync for users



Users have control over access permissions, sharing links, and can collaborate in real-time



Versioning and File History



Security and Authentication (KRB, OIDC, tokens)



Robust backup in two computer centres



Integration with CERN baseline services (auth, groups, mail)




Scientific Computing: integration tier for multiple scientific platforms (Jupyter Notebooks, computing farms, HPC clusters, ...)





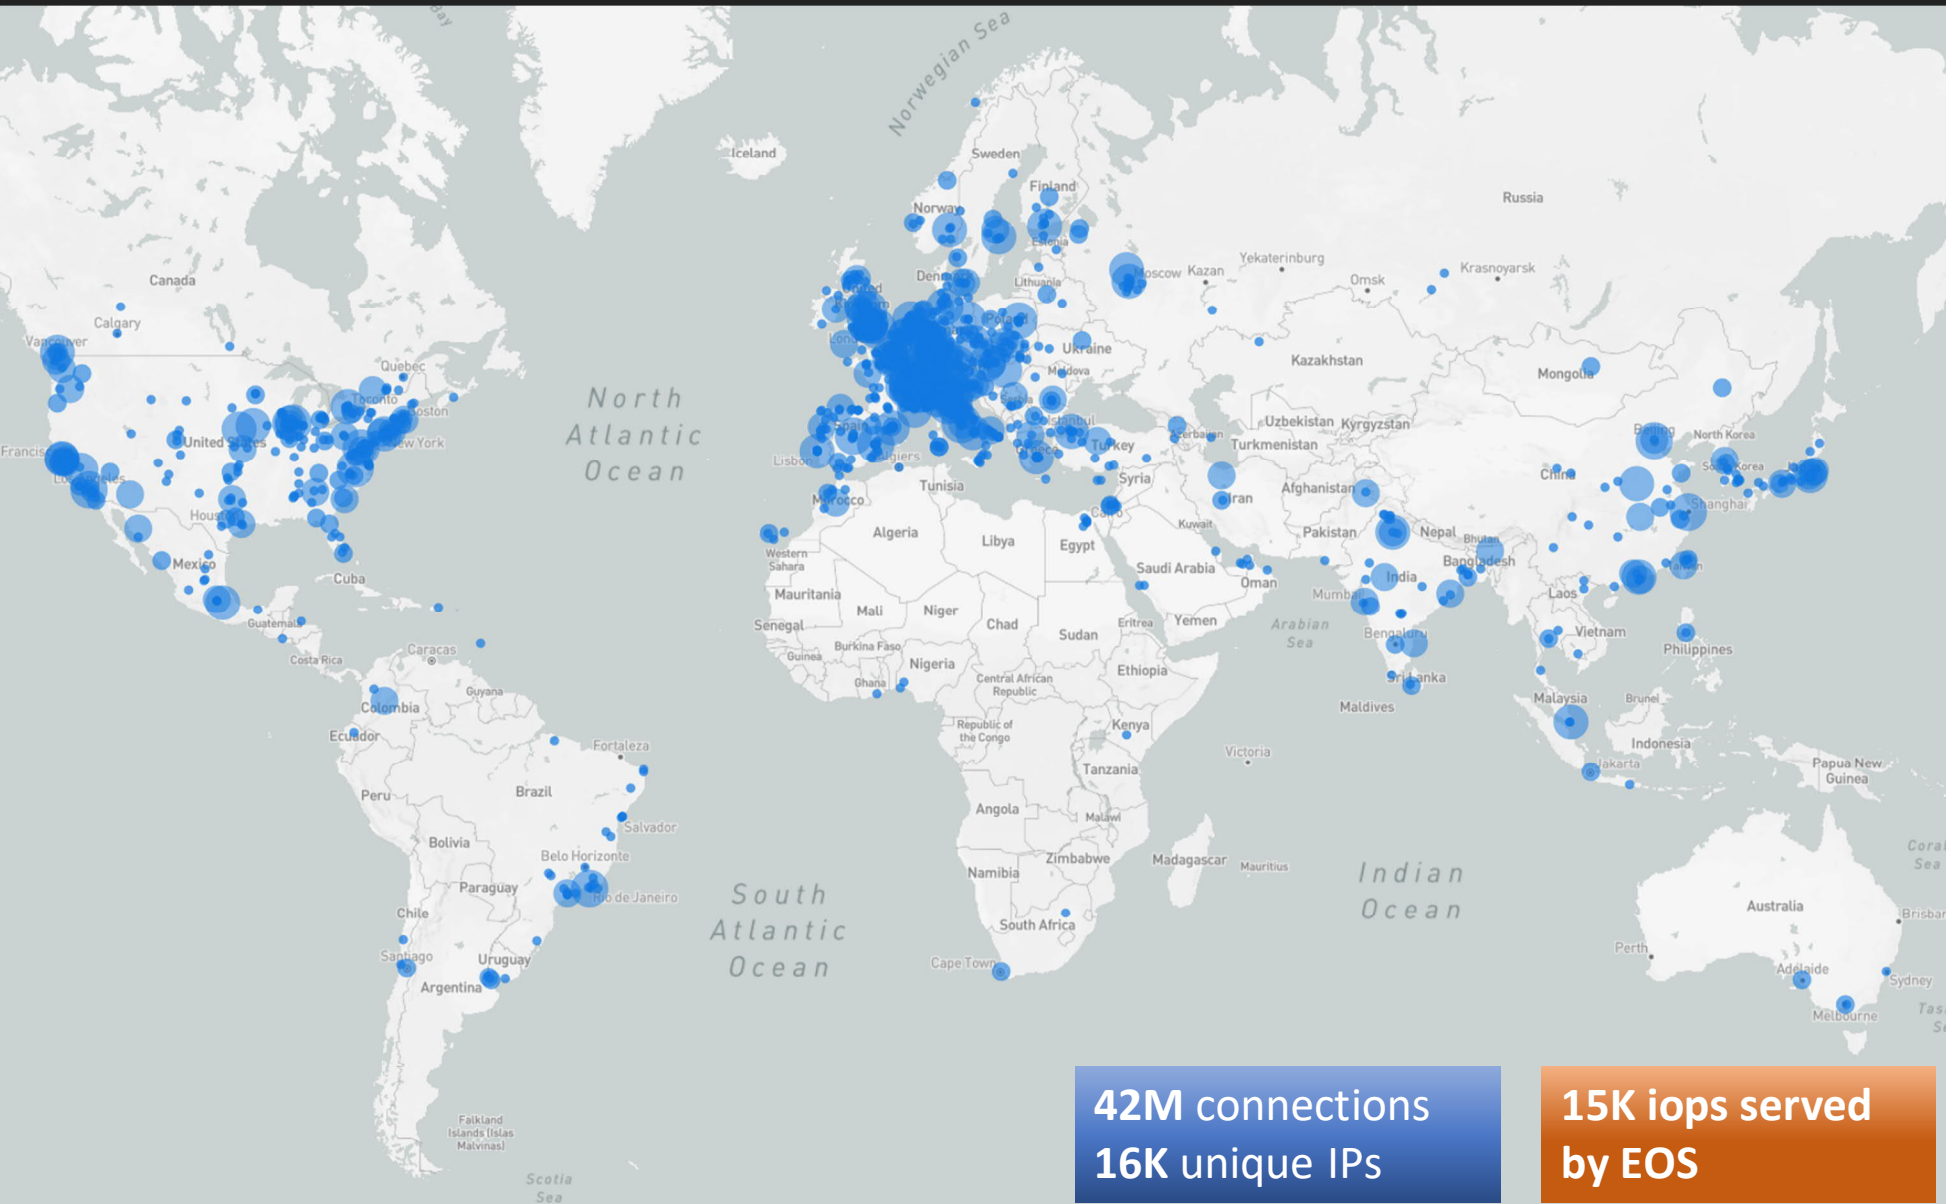
A Cloud Storage for Scientific Computing

- Usage spanning the whole Organization
 - Physics, Engineering, Administration
 - **420K** public links, **160K** shares, **1.3K** project areas
- A great diversity of use-cases and workflows
 - Generating tens of thousands of different file *types*
 - **Most represented: sources** (all together), **images**, **ROOT** files
 - A large fraction of files are machine-generated
- Combined **FUSE + cloud** access very popular
 - About **1.9K** unique users access the storage via FUSE daily
 - **40-45%** of them also use cloud (offline sync or web) access **at the same time**
 - Storage throughput dominated by **FUSE traffic from Batch farm (70%** of total I/O)





World-wide usage on a single day (Tue Feb 27th)



89 COUNTRIES **1467 CITIES**

[Copy URL](#) [New Map](#)

| | |
|----------------|----------|
| Switzerland | 7540 IPs |
| France | 1819 IPs |
| United States | 1277 IPs |
| Italy | 917 IPs |
| United Kingdom | 763 IPs |
| Germany | 760 IPs |
| Spain | 316 IPs |

42M connections
16K unique IPs

15K iops served
by EOS

The Cloud Sync&Share Storage (CS3) Community



- An active community, **significant intersection with HEP sites**
- Sites survey updated in March 2024
 - **32 sites** running an EFSS
 - **500K+** users in Education & Research
- Meet up at the CS3 Workshop, this year hosted at CERN as part of [TechWeekStorage24](https://techweekstorage24.org)



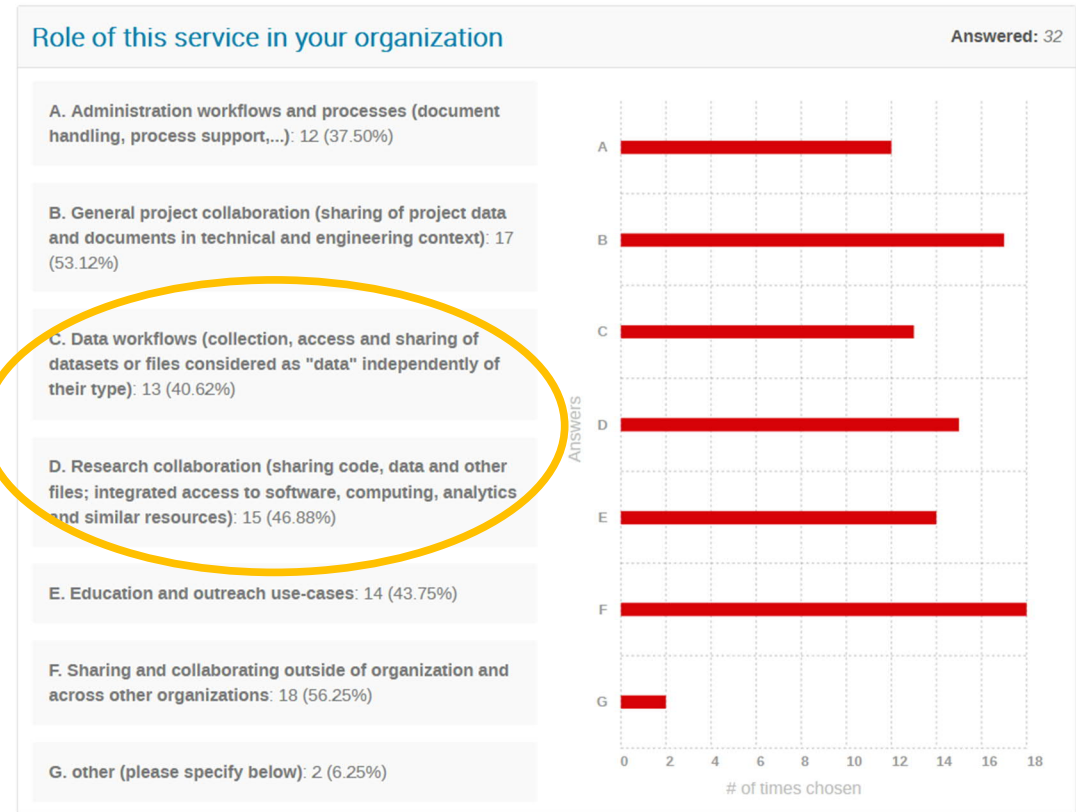
From <https://cs3map.ethz.ch>





Usage of CS3 Services in the Community

- Survey run at the CS3 Workshop 2024
- **Significant interest in Data workflows and Scientific Computing at large**
 - More than Administration workflows, corresponding to pure Dropbox-like use-cases
 - **Collaboration across other organizations** even more important
- Integration of Sync&Share with platforms such as JupyterHub increasingly popular




The Technology




The Technology

- Architecture paradigm: **microservices**
- Multiple components and daemons, each responsible for (a simple) part of the service, easing horizontal scalability
 - Most parts are distributed, metadata is kept on the storage and partly cached on a MySQL database
 - Technologies of choice: [Vue.js](#), [Golang](#) and [Python](#), [Protobuf](#)





The Technology: protocols and APIs

-  **Reva**: the backend, written in Go
 - Public API available at <https://github.com/cs3org/cs3apis>, in collaboration with ownCloud
 - Provides **HTTP** and **gRPC** services, including federation ([OpenCloudMesh](#)), WLCG-based TPC, ...
- Web frontend co-developed with ownCloud
 - Part of their new product oCIS
 - Currently on a fork, aim is to converge in the coming weeks
 - Mix of **WebDAV** and HTTP REST APIs
- (Branded) Sync clients provided by ownCloud
 - Fully based on **WebDAV**
- Storage: EOS offers **XRootD, HTTP, gRPC, FUSE**; CephFS offers native library
 - Reva interacts with EOS via **XRootD**, soon via **gRPC+HTTP**
- Satellite components, probes, etc. written in **Python and Go**





Same APIs, different implementations



/eos/project/m/ml -for-hep/my_great_gpt/dataset. h5

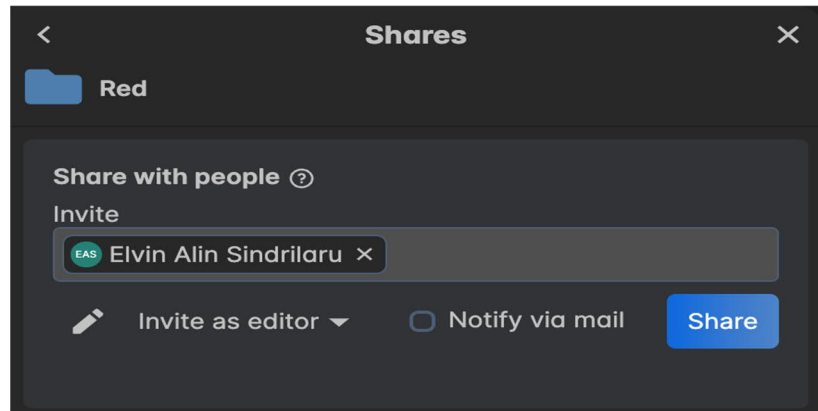
0d4d04eb-d081-41a9-92cd-8b49c540d08c! df62cd7f-538d-4b15-8fb3-b046ee6bbe77\$630dea7f-f2b1-4971-82fa-63a8eb005dc1





What makes CERNBox different?

- Storage is exposed transparently & it's directly accessible and mountable
 - This allows convergence of access via multiple access methods



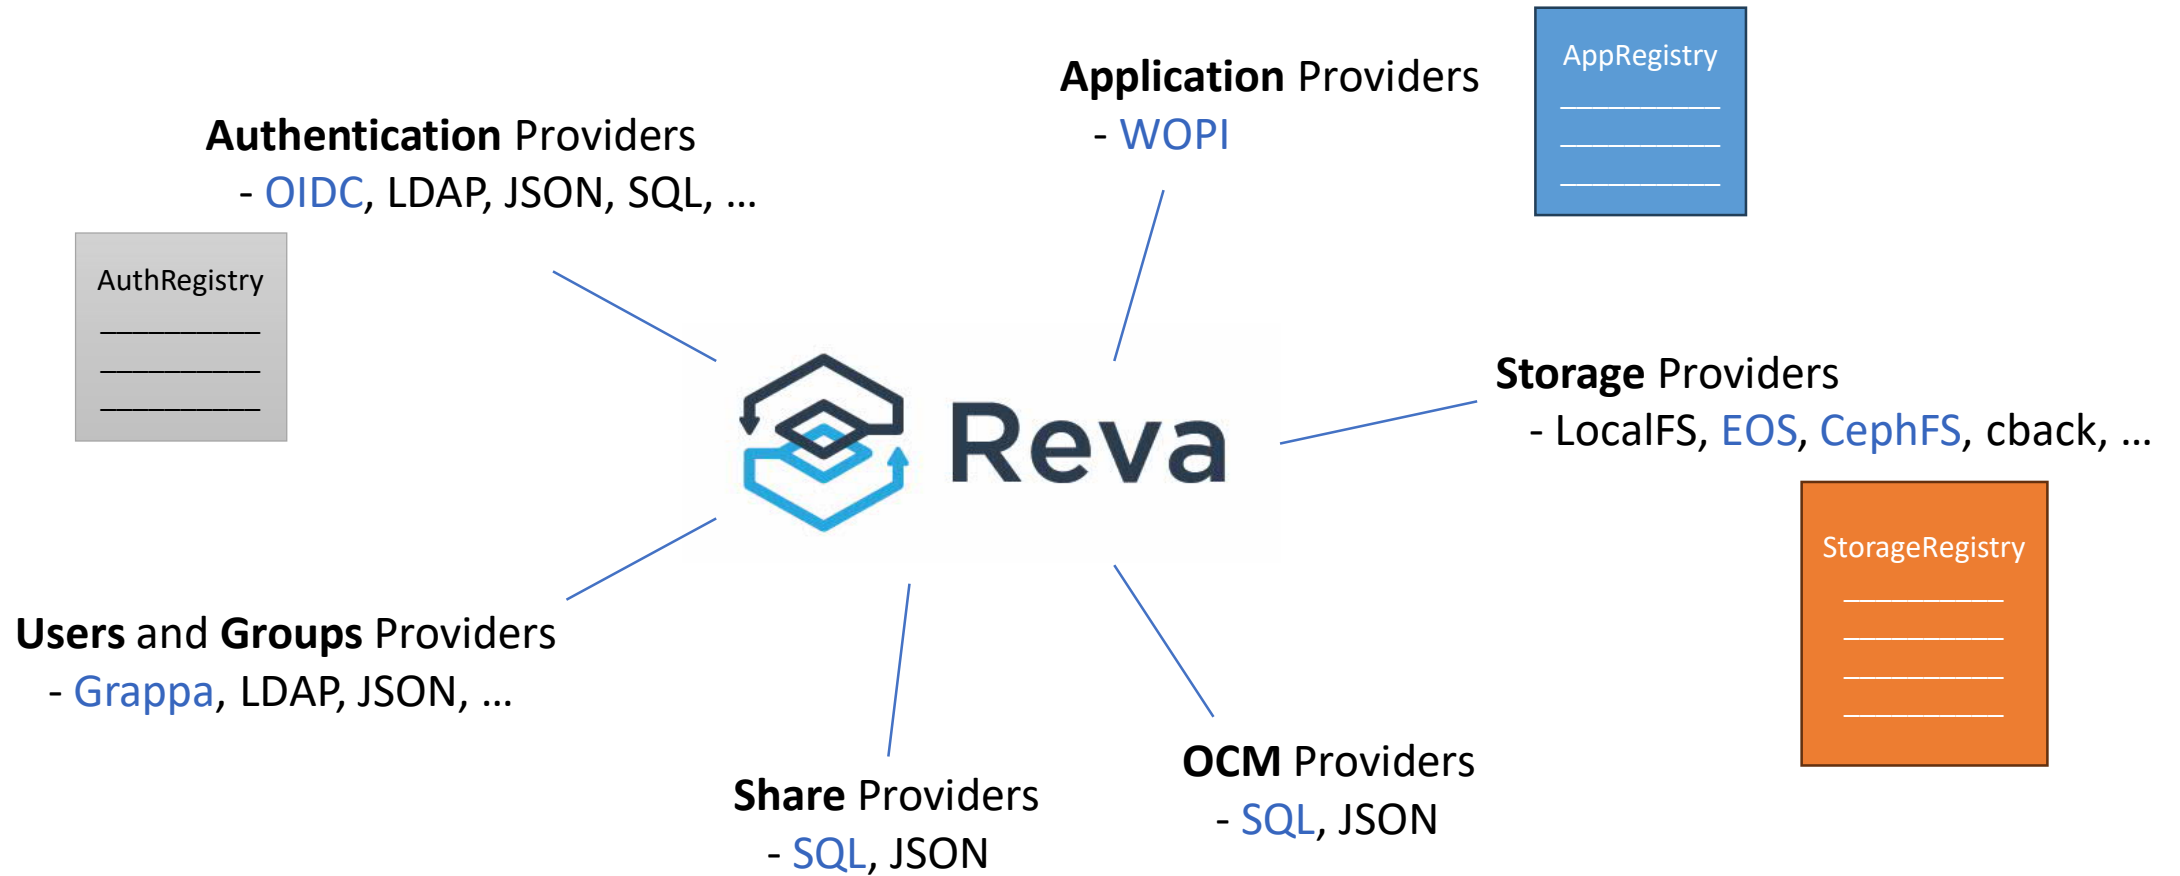
```
gonzalhu — 1:1 - ssh - Hugos-MacBook-Pro-401.local — tmux — 93x48
[root@eosblublu (mgm:master mq:master) ~]$ eos attr ls /eos/user/g/gonzalhu/Red
sys.acl="u:elvin:rwxm+dq"
```

- As a corollary, namespace metadata lives in the storage
 - Setting metadata on files from Web or mobile/desktop applications is reflected on the storage







Zooming into Reva





The underlying Storage: EOS and CephFS

- The storage plays a crucial role in supporting all cloud storage use-cases

| Feature | EOS | CephFS |
|---|-----|---|
| Rich ACLs for users&groups, authenticated shares | ✓ | ✓ |
| Atomic uploads for sync clients | ✓ | ✓ |
| Automatic file versioning | ✓ | ✗ |
| Recycle bin | ✓ |  |
| File locking (support for web apps) | ✓ |  |
| Mounted filesystem access, “near-POSIX” interface | ✓ | ✓✓ |

Reva provides the rest:

- ✓ Desktop/Mobile sync clients
- ✓ Non-authenticated shares
- ✓ Hooks for web apps (from text editors to Jupyter and external services)



Community and Outlook



A growing interest

- Building on the CS3 Community, and on the HEP Community at large
 - [BoF event at CHEP 2023](#) gathered significant interest
- First in-person Meet the Team Event @ CERN TechWeekStorage24
 - With KIT, IN2P3, SURF, LBL, PIC, RAL-STFC representatives, among others
- What can we offer?
A CERN-agnostic CERNBox setup + consultancy
- What do we aim for?
A self-supporting community (cf. EOS), to address together common issues, and contribute to the development of drivers (e.g. CephFS) or specific capabilities
 - A plugin system is already in place for Reva





Joining the Community

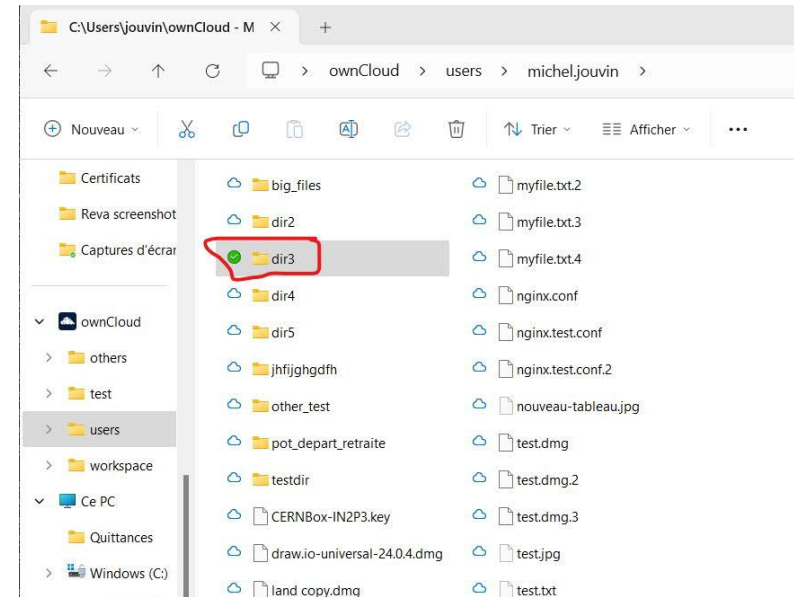
- **Please get in touch!**
- Send a mail to reva-community@cern.ch to join the mailing list
- Reva server repository: <https://github.com/cs3org/reva>
 - Discussion page <https://github.com/cs3org/reva/discussions> recently launched
- CERNBox repository: <https://github.com/cernbox>





An Example

- IN2P3 prototype with **CephFS** (courtesy Michel Jouvin):
upstream **ownCloud** client over Reva + direct access to underlying filesystem



```
root@ijc-rgw1 ~]# ls /cephfs/users/michel.jouvin/
big_files      dir3      myfile      nginx.conf      pot_depart_retraite  test.jpg
CERNBox-IN2P3.key  draw.io-universal-24.0.4.dmg  myfile.txt  nginx.test.conf  testdir              test.txt
dir2           jhfjghgdfh  myfile.txt.2  nginx.test.conf.2  test.dmg             text_rbytes
dir3           'land copy.dmg'  myfile.txt.3  nouveau-tableau.jpg  test.dmg.2          test.dmg.3
dir4           land.dmg      myfile.txt.4  other_test         test.dmg.3
```





Outlook: which Storage for Analysis Facilities

- High-performant storage is key for the high-throughput computing needed at HEP Analysis Facilities
 - See also [talk by A. Sciabà on Friday](#)
 - Mountable storage for interactive use-cases
- High Performance I/O for HEP @ Compute & Accelerator Forum ([last week!](#)): ROOT's RNTuple being promoted to replace TTree for [optimized storage and I/O](#)
 - => The heavy lifting remains with the underlying storage
 - EOS is being evolved, in coordination with ROOT and xrootd developers
- Ongoing discussions about Analysis Facilities in HEP...
CERNBox and its ecosystem, including **SWAN**, is a proven implementation for a **Cloud Storage underlying an Analysis Facility**



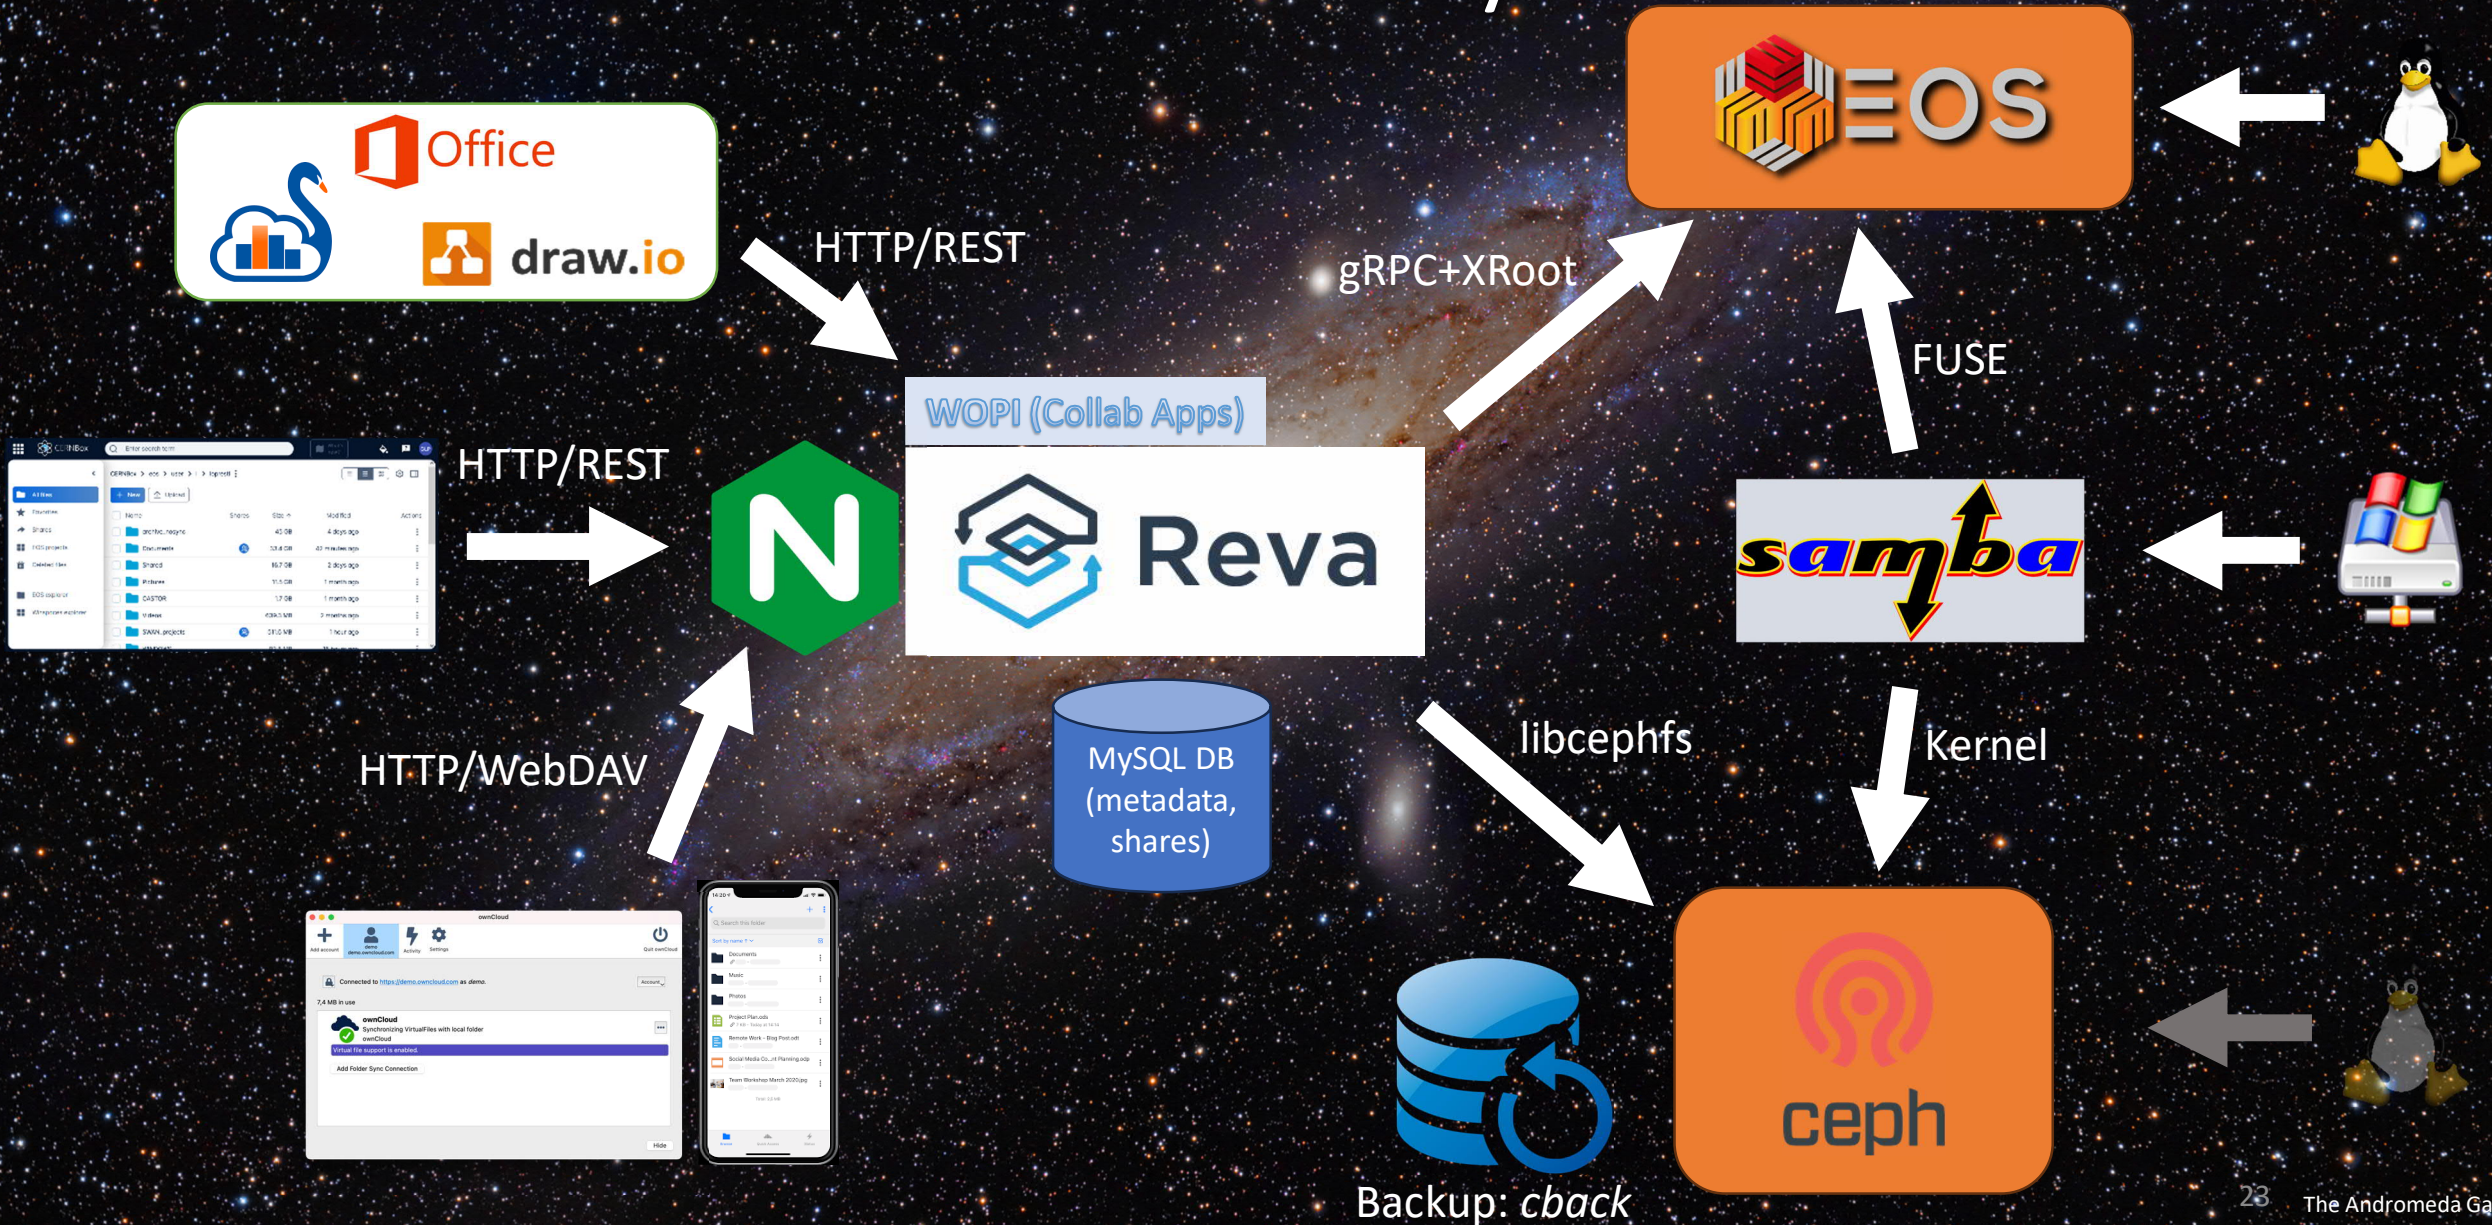
BACKUP SLIDES



Reva CephFS resources

- Reva CephFS community deployment forum:
<https://github.com/cs3org/reva/discussions/4610>
- Reva CephFS deployment documentation (online soon):
<https://reva.link/docs/tutorials/cephfs-tutorial>

Architecture: The CERNBox Galaxy





CERNBox Applications

- Long track of apps integrated in CERNBox
- For Research / Data Science
 - SWAN, Apache Spark (via FUSE mount!)
 - ROOT viewer
- For productivity
 - Microsoft Office 365 Online
 - Collabora, CodiMD, Onlyoffice
 - Draw.io
- For Engineering
 - CAD viewer

