

TECH WEEK  
STORAGE 24

Open Source  
Storage Technologies  
for the Exabyte Era

---



# Collaborative Cloud Storage

*Giuseppe Lo Presti*

On behalf of the CERNBox team

*CERN Storage Day, 14 March 2024, CERN, Geneva*

# General Home, Media & Project spaces

CERNBox  
33K User Accounts

ROOT Data Analysis Framework  
Microsoft 365  
diagrams.net  
CodiMD

Sync/Share  
WebApps

EOS File Store

indico  
zenodo  
CDS  
AV  
restic

30K clients  
lxplus  
lxbatch  
Windows

Samba  
DFS

fusex

# Data Analysis & Physics Infra

WLCG  
~26PB Physics Data worldwide  
150K+ clients worldwide

DPM

indico  
restic

CVM FS

S3

CEPH Object Store

# General Infrastructure & Services

EP-ESE  
OPENSIFT  
puppet

Backups  
Preveessin  
Zone

AFS  
FILER

RBD

openstack.

CERN Tape Archive

GitLab

HPC

Ceph FS



# Contents

- The Project
- The Technology
- The Service
- Future Prospects





# What is 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

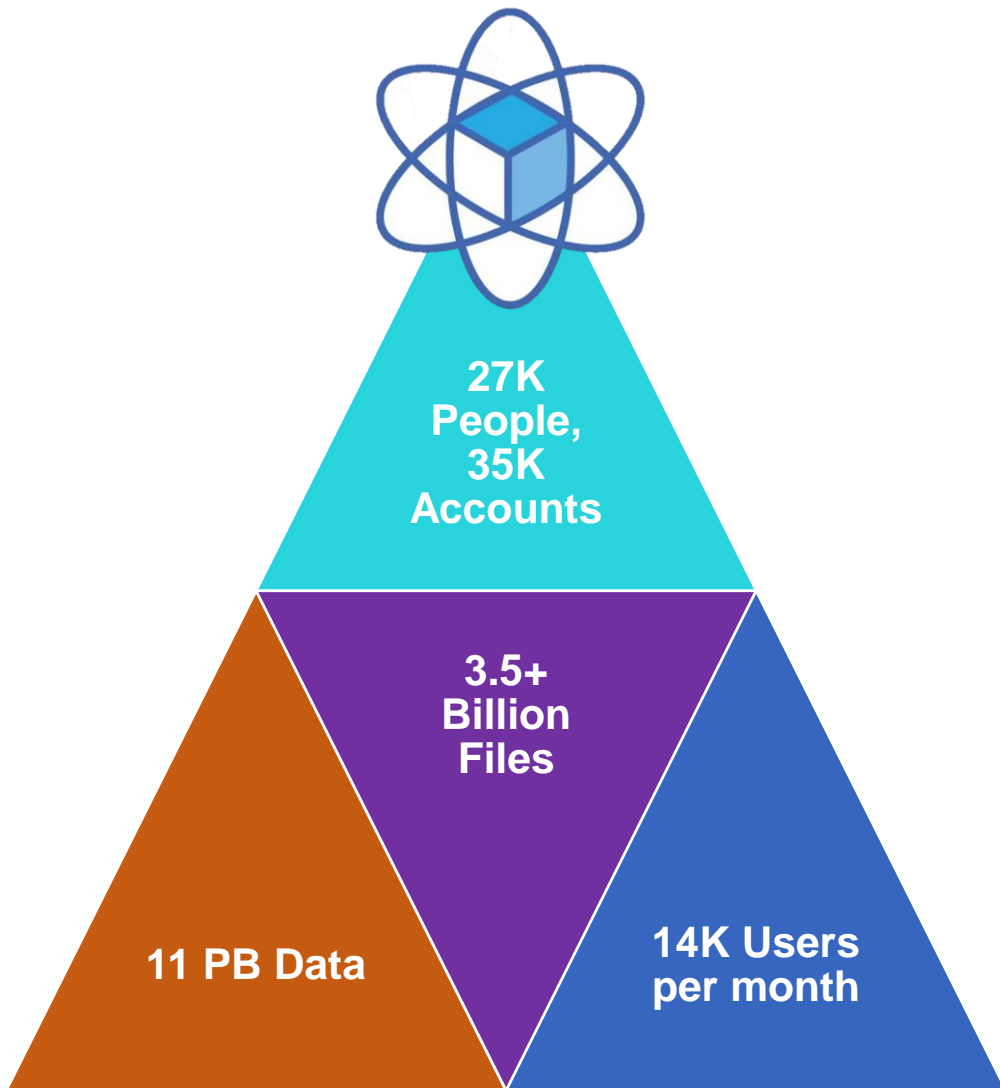


- Data safely stored on CERN premises
- Security policies to keep data safe and confidential





# 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)



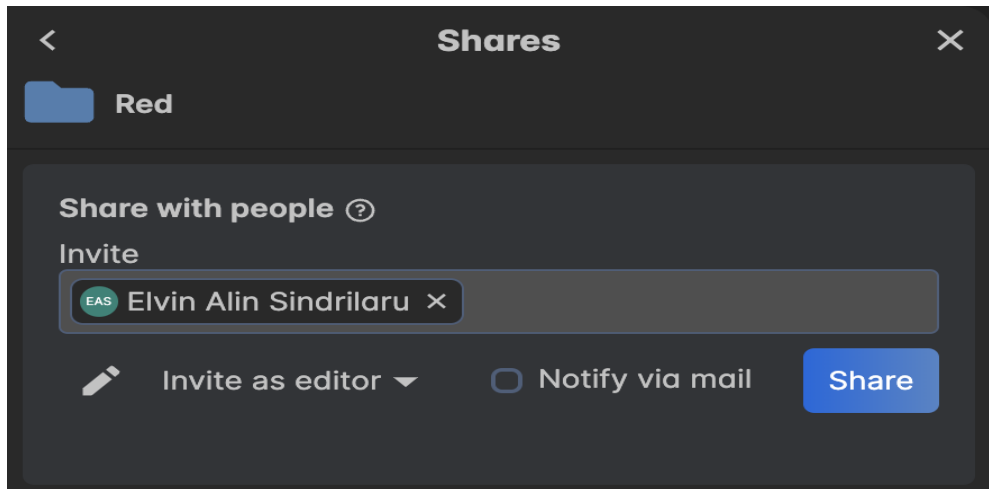
Scientific Computing: integration tier for multiple scientific platforms (Jupyter Notebooks, computing farms, HPC clusters, ...)





# What makes CERNBox different?

- Namespace metadata lives in the storage
  - Setting metadata on files from Web or mobile/desktop applications reflects it on the storage
- Storage is directly accessible and mountable
  - This allows convergence of access via multiple access methods

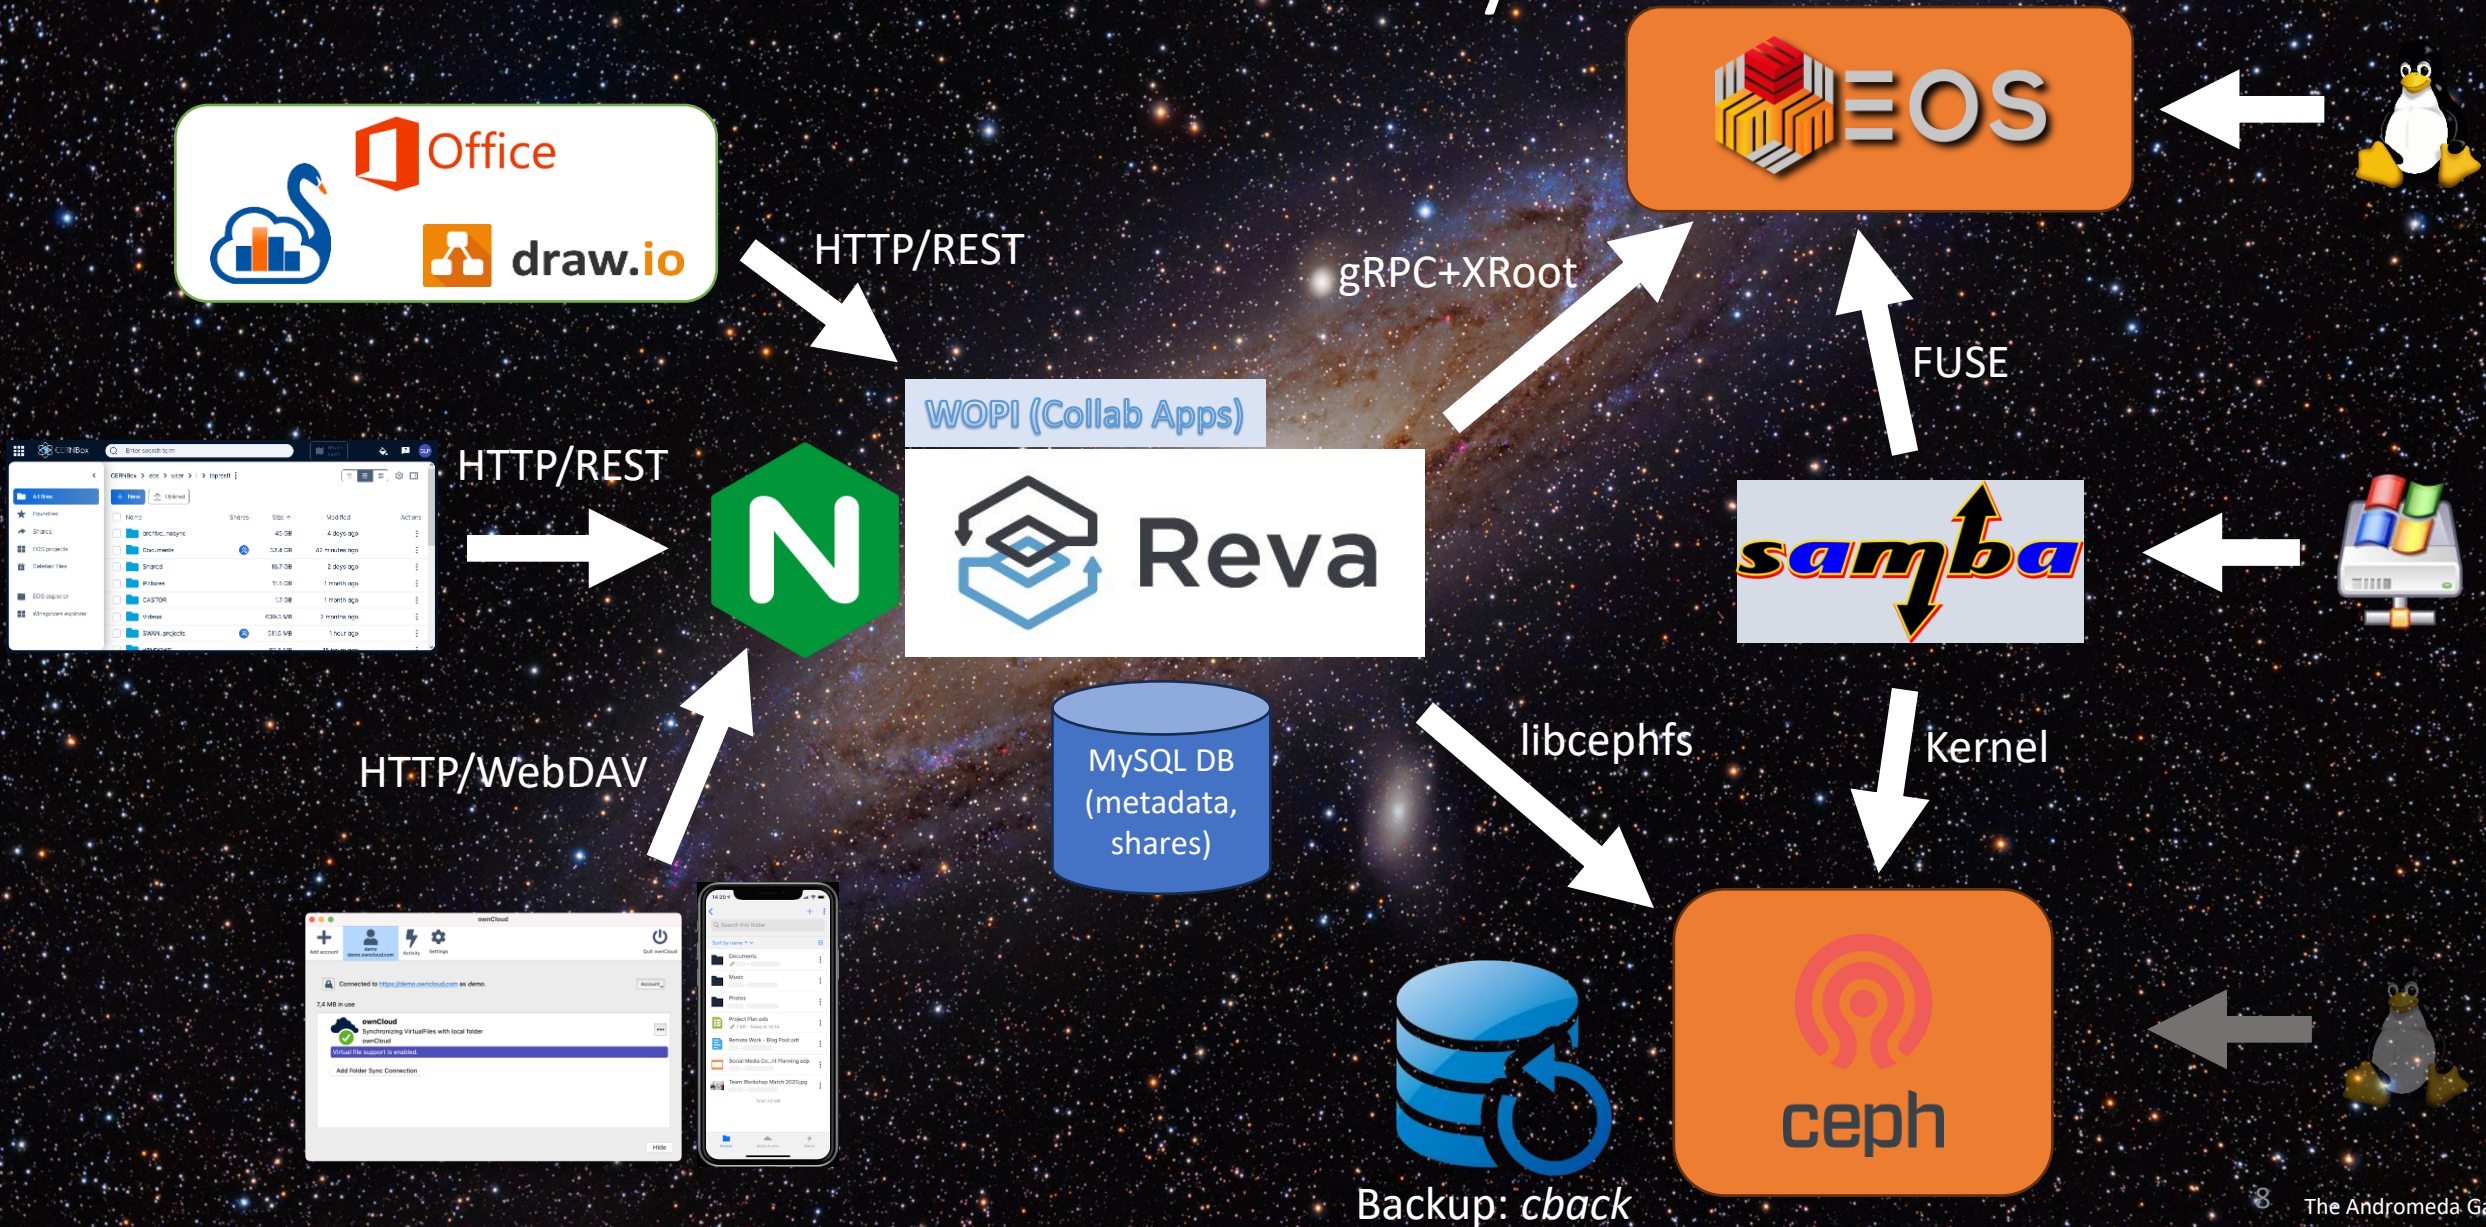


```
gonzalhu — 1:1 - ssh - Hugos-MacBook-Pro-401.local — tmux — 93x48
[root@eosblu (mgm:master mq:master) ~]$ eos attr ls /eos/user/g/gonzalhu/Red
sys.acl="u:elvin:rwxm+dq"
```





# Architecture: The CERNBox Galaxy








# 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: [Golang](#) and [Python](#), [Protobuf](#)







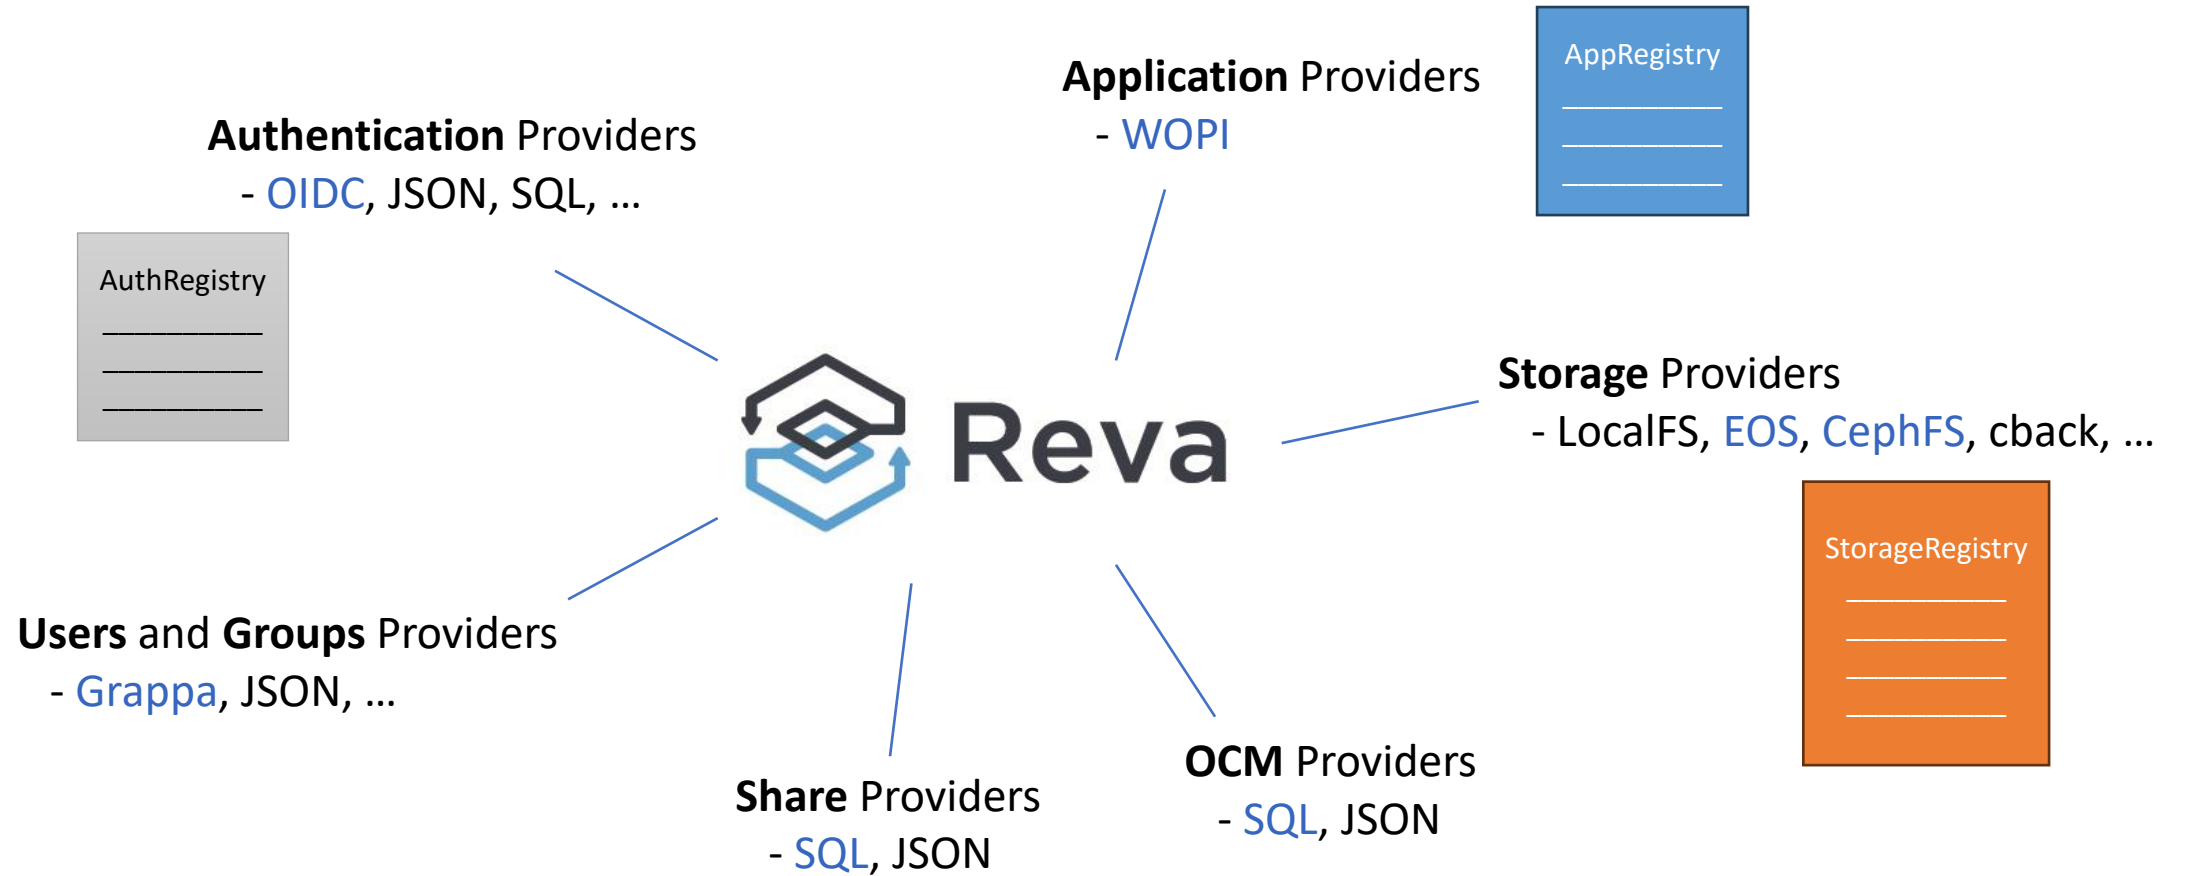
# The Technology: protocols and APIs


- Core backend written in Go
  - Public API available at <https://github.com/cs3org/cs3apis>, implemented as well by ownCloud Infinite Scale (oCIS)
  - Exposes **HTTP** and **gRPC** services
  - Exposes **OCM** HTTP endpoints (<https://github.com/cs3org/OCM-API>) for federating storages
- 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
- Sync clients provided by ownCloud, branded with the CERNBox logo
  - 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**





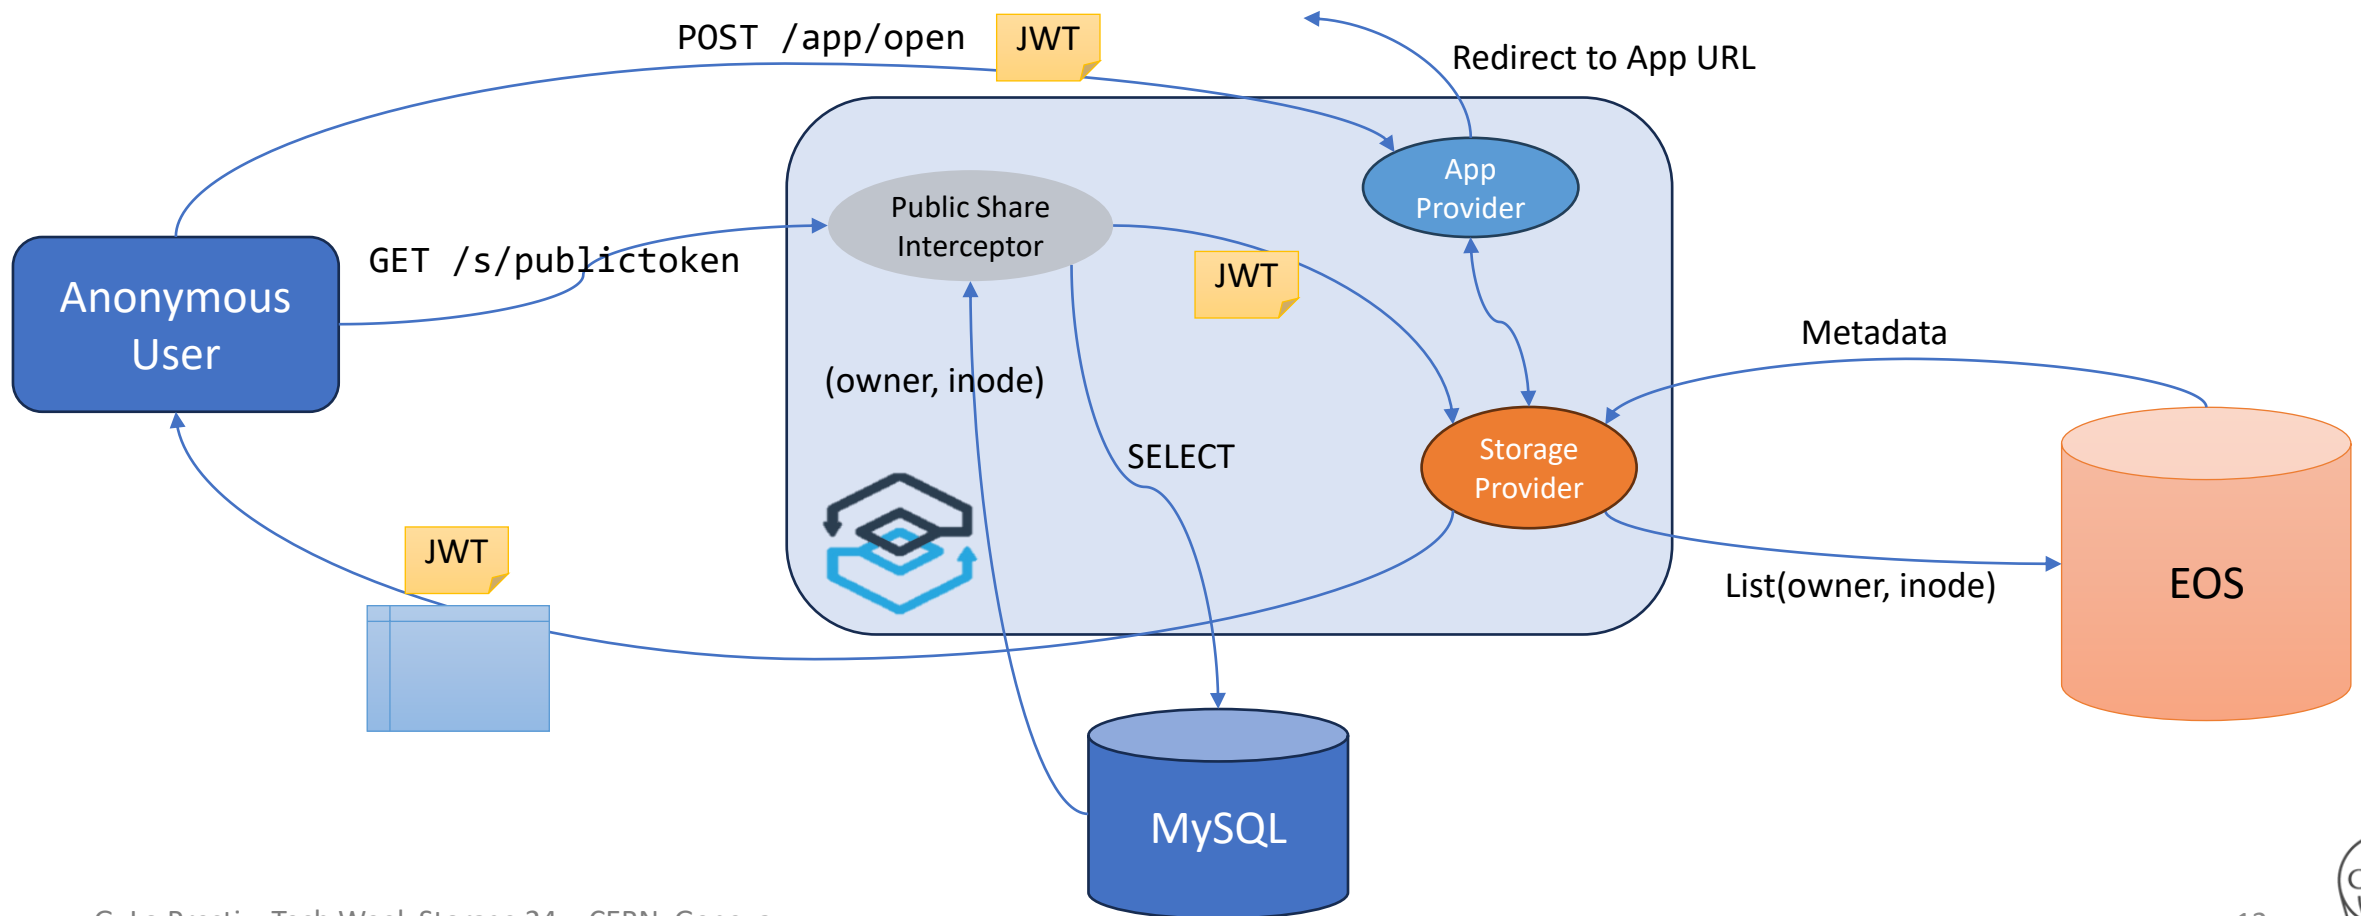
# Zooming into Reva





# Journey of a web request

A mini-dive into the web system: a user clicks on a “public link” and opens an app







# 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






# WOPi for Collaborative Apps

- A journey started in 2017: integrate Microsoft Office Online in CERNBox
  - Deployed on-premises back then
- Web-application Open Platform Interface: a HTTP REST protocol, specified by Microsoft
  - Adopted by a bunch of other vendors, notably Collabora Online and Onlyoffice
- Currently, CERNBox supports the full standard
  - Works with **Microsoft Office 365** (on-cloud), Collabora Online, and Onlyoffice
  - Integrates **CodiMD** with custom extensions
  - Could integrate further apps such as **Diagrams.net** in collaborative mode





# The Storage: EOS and CephFS

- The storage plays a crucial role in supporting all cloud storage use-cases
- **EOS is the reference implementation**, as it includes all desired capabilities
  - Rich ACLs management
  - Atomic uploads for the sync clients
  - Automatic file **versioning** upon overwrites
  - **Recycle bin**
  - **File locking**
  - Mounted filesystem (FUSE) access, “near-POSIX” interface
- CephFS has been recently integrated
  - Not all capabilities have been implemented yet: no atomic uploads, no versioning, no recycle
  - Those may come as a community effort








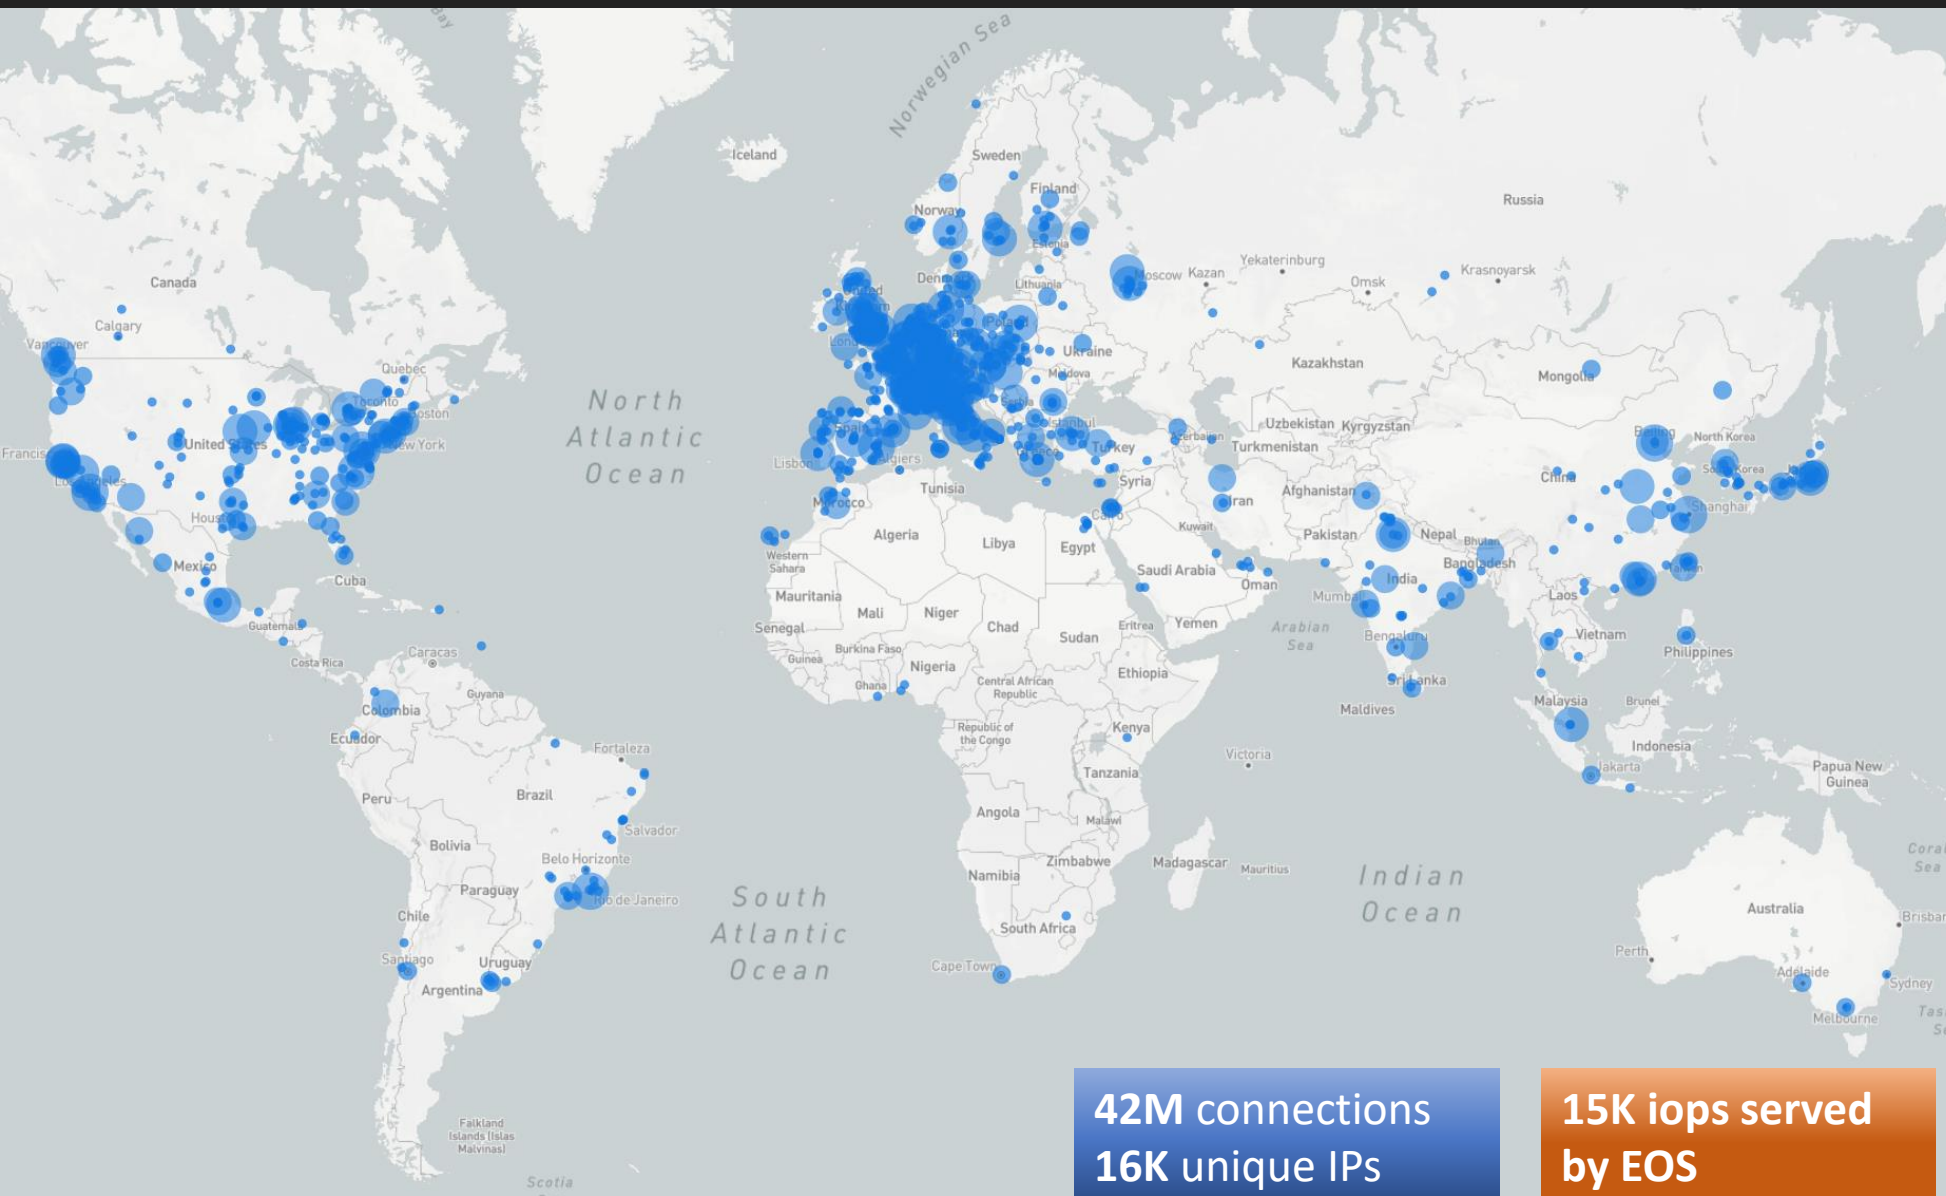
# The Service

- **14K** monthly users, **6K+** daily users
- **420K** public links, **160K** shares, **1.3K** project areas
- Spanning the whole Organization
  - Physics, Engineering, Administration
- A great diversity of use-cases and workflows
  - Generating tens of thousands of different file types
  - Most represented: ROOT, followed by source code
- *More statistics will be presented during the EOS workshop*






# The Service: usage on a single day (Tue Feb 27<sup>th</sup>)



**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



# Future Prospects

- Plenty of nice ideas about extensions and additional features
  - Thanks to the rich feedback from our users
- However, the limited development effort available right now imposes to consolidate and streamline
  - We are a relatively small team!
- Focus for this year:
  - Commission CephFS to onboard Windows use-cases (**+1B files**, more later today)
  - Bring in the new Spaces concept, to benefit from upstream ownCloud web interface
  - Improve and automate operations procedures
- Happy to dive in more details in the afternoon session!







# CERNBox Docs and Resources

- CERNBox repository: <https://github.com/cernbox>
- Reva server: <https://github.com/cs3org/reva>
- ownCloud Infinite Scale: <https://doc.owncloud.com/ocis/next>
- CERNBox publications:  
[https://scholar.google.com/scholar?hl=en&as\\_sdt=0%2C5&q=cernbox&btnG=](https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&q=cernbox&btnG=)





# Thanks for your attention!

## Credits to the excellent team that makes CERNBox a reality!

