

# CERNBox: Storage gateway for CERN and beyond

*Hugo Gonzalez Labrador*<sup>1,\*</sup>, *Gianmaria Del Monte*<sup>1</sup>, *Javier Ferrer*<sup>1</sup>, and *Oliver Keeble*<sup>1</sup>; on behalf of the CERNBox team

<sup>1</sup>CERN

**Abstract.** The CERN IT Storage group ensures the coherent design, development, operation and evolution of storage and data management services at CERN for all aspects of physics, user and project data and general needs of the Laboratory. CERNBox is one of the services that actively contributes to this objective.

CERNBox is a cloud collaboration platform providing storage space (18PB) to users (37K accounts), projects (1.1K) and to LHC experiments (ATLAS, CMS, LHCb, ALICE) and other smaller experiments (33). We present the evolution of the platform from its humble origins 10 years ago and how the HEP landscape has influenced the expansion of the service to continue satisfying the ever growing demand of our users. We describe the challenges and opportunities presented by integrating a sync and share solution (ownCloud) with CERN's custodial storage solution (EOS) and how this stack boosted a coherent integration with analysis farms (LXBATCH), Linux (LXPLUS) and Windows (TS) interactive clusters and Web analysis facilities (SWAN). We then turn to the part CERNBox has played in consolidation efforts in the IT department (DFS, AFS) and in providing end-users with a portfolio of applications (ROOT, CAD, MS365) while guaranteeing a maximum level of privacy.

We conclude by presenting the future of the service with its heterogeneous storage federation capabilities (EOSHOMEs, EOSPROJECTs, EOSATLAS, CEPHFS) and WLCG transfer mechanisms (WebDAV TPC).

## 1 Origins of CERNBox

A decade ago, CERNBox was born as a sync and share platform, offering an alternative to commercial solutions like Dropbox. Its initial mission was to provide CERN's scientific community with a secure data-sharing platform. Over time, it evolved into a sophisticated data management solution tailored to the unique needs of the High Energy Physics community. It began as a basic file synchronization and sharing platform for individual users, emphasizing security and integration with CERN's ecosystem, including Single Sign-On (SSO) and e-groups. Later, recognizing the demand for enhanced collaboration, CERNBox expanded its capabilities to support project-based collaboration and integrated with Web Analysis platforms like SWAN[9], enhancing its versatility. Continuing its evolution, CERNBox became the designated storage solution for many engineering and physics workflows, solidifying its role in scientific data management.

---

\*e-mail: hugo.gonzalez.labrador@cern.ch

## 1.1 Decade-long Evolution

As shown in Figure 1, the service has reached a plateau in number of user accounts. CERN's user population is constant and the service has managed to reach out to the entire user community. Data volume continues to grown at a manageable pace.

Figure 2 also shows the number of unique active users in the system over the last 6 months. The same figure also shows the user distribution across different access protocols.

Over the last 6 months, almost 20,000 unique users have used the system, accounting for 1,300,000 daily accesses.




Figure 1: Evolution of CERNBox

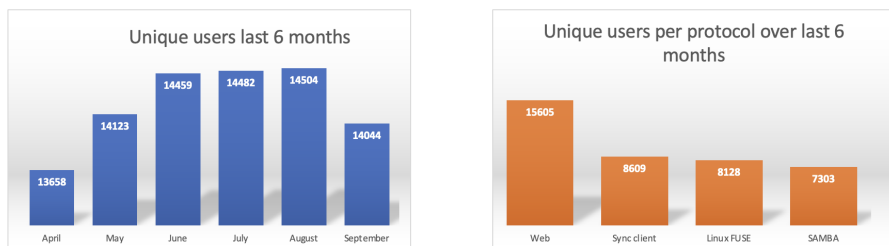


Figure 2: Unique users of CERNBox in the last 6 months

## 2 Integration with Scientific Environments

The fusion of synchronized and shared storage with scientific analysis and processing environments came before similar capabilities were developed within the industry. While commercial solutions such as Google Workspace & Google Collab offer comparable functionality at the individual user level, **CERNBox presents a comprehensive integration of synchronized and shared storage, providing both online and offline access to the entire scientific storage repository at the institutional level, accessible to all users.** CERNBox is actively shaping a continuum for storage and data analysis, bridging the gap between data center services and personal devices, thereby unlocking entirely new possibilities for data collection, access, and utilization. Notable examples include the automatic uploading of images from mobile devices within the accelerator tunnel to a central processing repository for quality control checks of the LHC machine, effortless sharing and access to notebooks across laptops, central compute services, and Grid processing, as well as the direct extraction of data

from laboratory devices like oscilloscopes into a centrally available repository. Here are a few key integrations:

- Coherent integration with analysis farms (LXBATCH): LXBATCH mounts the CERNBox storage back-end (EOS) as a user-space file system, making it available for user or grid submitted jobs. Results from these jobs are automatically stored on CERNBox and propagated to end-user devices. This approach marks a paradigm shift from the traditional model where users must periodically check and manually copy results.
- Integration with Linux (LXPLUS) and Windows (TS) interactive clusters: CERN offers two platforms for interactive computing, one for Linux users (LXPLUS) and another for Windows users (Terminal Servers). The user can connect to these systems for their computing needs. In LXPLUS, EOS is mounted in the same way as LXBATCH, while in Windows TS, via a SAMBA service[12].
- Ideal connectivity with Web data analysis facilities (SWAN): it provides users with an interactive environment accessible through a standard web browser, eliminating the need for complex setups or specific software installations. EOS is mounted in SWAN, and CERN-Box REST APIs are accessible through specific SWAN plugins, enabling users to effortlessly share their notebooks and collaborate.

## 2.1 Integration Tier for Analysis Facilities

As presented during the CERNBox BOF session[2], CERNBox’s capabilities extend beyond data storage and management. It can serve as an integration tier for analysis facilities within CERN. This integration facilitates the transfer of data between CERNBox and various analysis environments, enabling researchers to conduct in-depth data analysis with ease. Figure 3 highlights the majority of functionalities available in the platform that wire with the different science fabric utilities available at CERN. For example: offline access, collaborative environment, GRID Storage Element, sharing of data, cross-institutional collaboration, automatic websites, physics data gathering, integration with batch computing and web analysis (SWAN) and interactive clusters (LXPLUS, Windows TS).

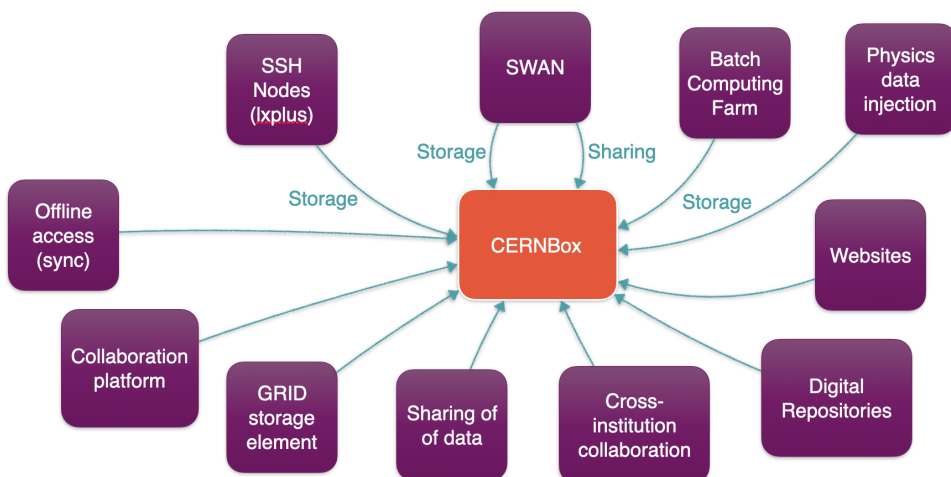


Figure 3: A major subset of CERNBox functionalities that makes it attractive as an integration tier for analysis facilities

### 3 Sync and Share with Custodial Storage

CERNBox's integration journey extends to the successful fusion of a sync and share solution, ownCloud[3], with CERN's custodial storage solution, EOS. This integration has created a powerful stack that supports both data sharing and preservation. Additionally, it enables CERNBox to operate as a bridge between user-driven data sharing and the preservation of valuable research data. A key aspect of the synergy between the front-end layers of the service and the storage back-end is the possibility to have an homogeneous access to the data. Figure 4 illustrates this behaviour by sharing arbitrary data with a colleague. The access control list (ACL) is honoured across multiple access protocols, providing access to shared data from the most suitable environment (mounted file-system in Linux, network drive in Windows, or locally synchronized copy).

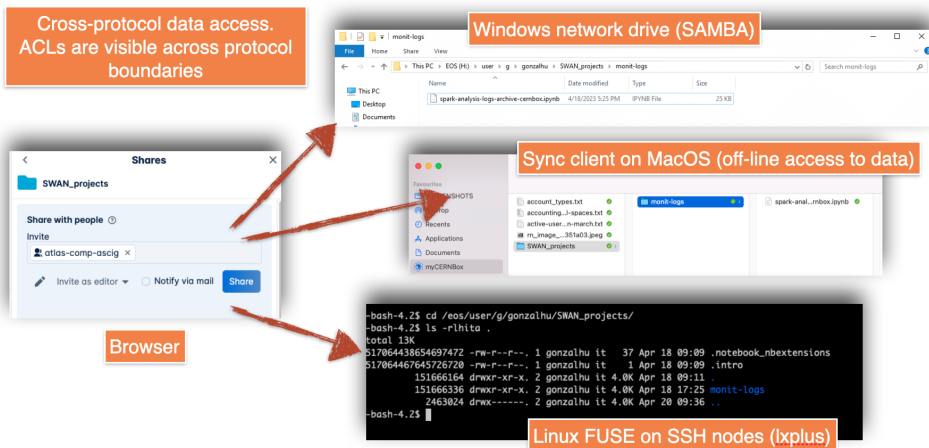


Figure 4: ACLs are consistent across multiple access protocols

#### 3.1 Consolidation Efforts

Within CERN's IT department, CERNBox has played a pivotal role in consolidation efforts, helping to reduce the cost of operating multiple services and consolidating into one platform. Figure 5 highlights these efforts.

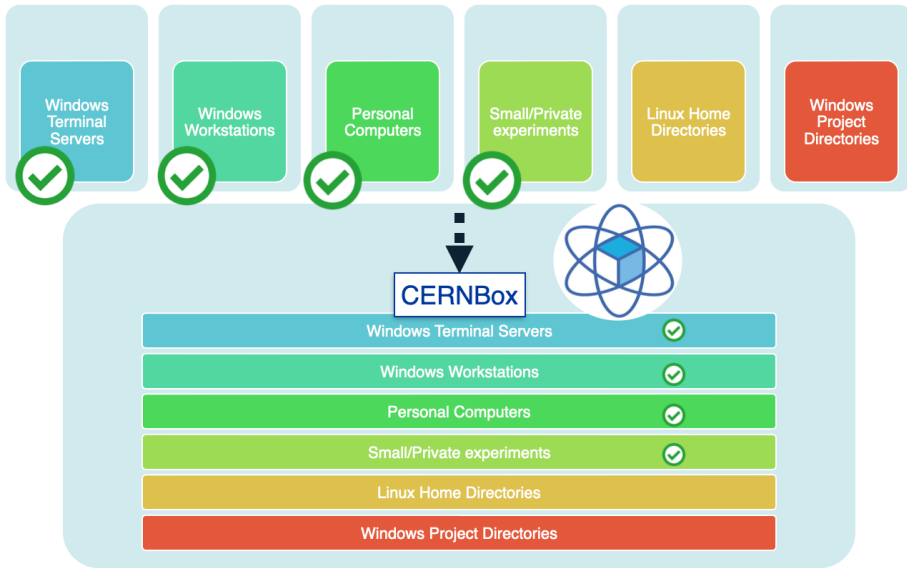


Figure 5: This figure shows the consolidated platforms in CERNBox and the challenges to face in the near future: Linux home directories and Windows workspaces

## 4 Architecture

CERNBox’s architecture is a testament to its versatility and adaptability. It integrates with various storage systems and components, ensuring efficient data handling. Key components of CERNBox’s architecture can be visualized in Figure 6 alongside a subset of its use-cases.

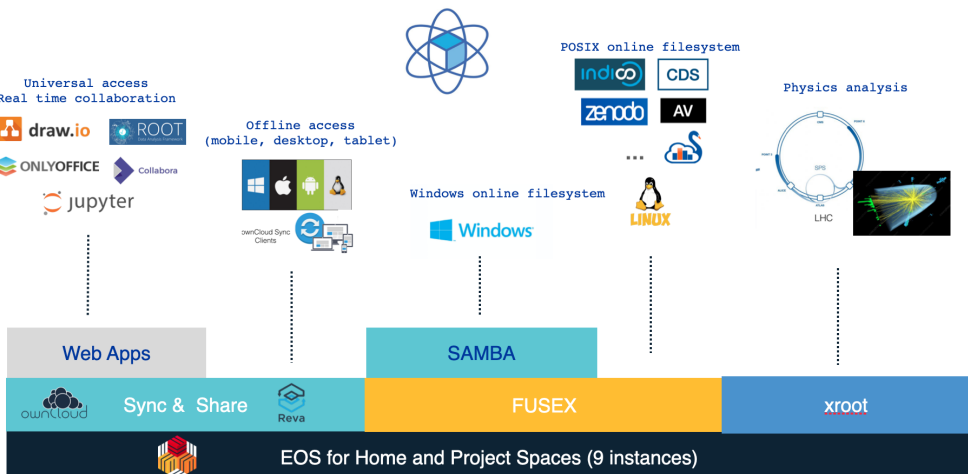


Figure 6: General architecture of CERNBox

The four main elements are explained in the following sections.

## 4.1 Web

CERNBox provides a modern web interface structured as a Single Page Application.. It employs OIDC (OpenID Connect) with CERN Central Authentication SSO (Single Sign-On) for secure access. The Web component is based on ownCloud's new Web platform, code named OCIS Web[3]. This new platform is a long-standing desired improvement over the old PHP-based MVC web, which was highly tied to the PHP OC 10 server. The new Web relies on the open WebDAV protocol for data transfers and in ownCloud's own APIs (OCS) for data sharing. The software stack of this front-end is VueJS, CSS and HTML. Recent efforts from ownCloud included migration from Vue2 to Vue3.

## 4.2 Sync&Share server backend

The core of CERNBox is a high-performance gRPC and HTTP server code named Reva. This component is responsible for serving the APIs required by the Web and Sync clients to access the underlying storage systems. Reva, while originally created at CERN, is now a community driven effort with active participation from industry (namely ownCloud,, which incorporates this component as part of their new flagship product, OCIS) and contributions from EU projects, like ScienceMesh[11] (where Reva is the core component to provide interoperability across multiple nodes, independently of the vendor). Reva provides a flexible framework to add new services and new APIs to be consumed by a variety of clients.

## 4.3 EOS

EOS is an open source distributed disk storage system in production at CERN since 2011 . The development focus has been on low-latency analysis use cases for LHC and non-LHC experiments and life-cycle management using JBOD2 hardware for multi PB storage installations. EOS has been the preferred storage back-end for CERNBox and both platforms have co-evolved together since early 2014. EOS provides all the asafety guarantees needed by the ownCloud synchronization protocol. When CERNBox was created, EOS was the only available storage platform supporting these requirements. With time, other file-system gained attraction, like CephFS, that matured from a R&D project to a production storage solution across multiple industries and its integration in CERNBox is currently being explored (more in Section 5.2).

# 5 Current Work

CERNBox continues to adapt to the needs of the community with several key initiatives currently in progress.

## 5.1 Facilitating Inter-site Collaboration with ScienceMesh

CERNBox is pioneer in fostering collaboration within the scientific community. It actively contributes to addressing the challenges associated with isolated collaboration efforts through initiatives like OpenCloudMesh (OCM)[13]. One noteworthy feature of CERNBox is its provision of the OCM API, which enables users to securely receive shared data from external institutions.

Presently, our focus is on streamlining the adoption of the ScienceMesh[11] federated infrastructure. This initiative involves active participation from CERN, alongside numerous other National Research and Education Networks (NRENs) across Europe. ScienceMesh serves as a promising platform for advancing inter-site collaboration and knowledge sharing.

## 5.2 CephFS integration

To facilitate the ideal transition of project areas within the DFS service (Windows-based proprietary distributed storage system), CERNBox is introducing access to CephFS alongside EOS. The adoption of CephFS is primarily motivated by the need for high Input/Output Operations Per Second (IOPS) and minimal latency, especially for Windows-based computers.

Ceph has already established itself as a cornerstone of storage systems at CERN. The integration of CephFS incorporates a novel driver that harnesses the capabilities of the go-ceph library. This library represents a set of Application Programming Interface (API) bindings that enable the direct utilization of native Ceph APIs from within "Go" code. These bindings leverage Go's cgo feature to create a robust, efficient, and secure implementation. This implementation not only ensures speed and reliability but also facilitates updates and security enhancements without imposing any unnecessary burden on our team.

Furthermore, the extension of Reva provides an abstraction layer for storage back-ends. This layer enables the smooth translation of high-level user interactions, such as sharing and access control, into low-level file system operations, particularly for managing POSIX Access Control Lists (ACLs). In doing so, Reva enhances the overall user experience presenting a single gateway to various CERN storage systems.

### 5.3 Notification System

After the release of the updated CERNBox web platform in fall 2022, a feedback campaign was conducted. Users frequently requested notifications for various system actions. Responding to these requests, CERNBox implemented a notifications system. This system is now widely utilized within the CERN community to inform peers and third parties about activities on their shared storage. It enhances collaboration by enabling both reactive and asynchronous workflows. To create this system, CERNBox integrated NATS technology as depicted in Figure 7, providing fast and lightweight message delivery, ideal for notification generation and communication. NATS is known for its suitability in adaptive edge and distributed systems like CERNBox.




Figure 7: Notification's service flow for an email notification. The process starts with Alice (bottom center) sharing a folder with Bob, which triggers a 'Create share' action and registers a new notification on the service using NATS. Afterwards, Bob uploads a file, which triggers the sending of this notification.

## 6 Outlook

CERNBox has garnered interest from various educational and scientific institutions. However, its current software stack is intricately tied to CERN's infrastructure, making it challenging to use elsewhere without significant modifications. To address this limitation and enable easier deployment, we are currently refactoring Reva. The goal of this refactoring effort is to create a plugin portal within Reva. This portal will allow users to specify customizations without affecting the core functionality of the system. By doing so, we aim to make CERNBox more adaptable and accessible for a broader range of institutions. Ultimately, this initiative will pave the way for other organizations to deploy CERNBox while tailoring it to their specific needs. Our aim is to make CERNBox upstream compatible, ensuring that it can easily integrate into various environments and serve as a valuable resource for a wider audience in the Education and Scientific sectors.

## References

- [1] NATS, <https://nats.io/> (access time: 22/09/2023)
- [2] CERNBox BoF, <https://indico.cern.ch/event/1230126/sessions/492063/#20230506>
- [3] ownCloud, <https://owncloud.com> (access time: 22/09/2023)
- [4] Mascetti, Luca and Labrador, H Gonzalez and Lamanna, M and Mościcki, JT and Peters, AJ, *Journal of Physics: Conference Series* **664**, 062037 (2015)
- [5] H. G. Labrador, *CERNBox: Petabyte-Scale Cloud Synchronisation and Sharing Platform* (University of Vigo, Ourense, 2015) EI15/16-02
- [6] H. G. Labrador, *CERNBox BoF*, (WLCG-HSF Workshop 2023, Norfolk, US)
- [7] Peters, AJ and Sindrilaru, EA and Adde, G, *Journal of Physics: Conference Series* **664**, 062037 (2015)
- [8] Peters, Andreas J and Janyst, Lukasz *Journal of Physics: Conference Series* **331**, 052015 (2011).
- [9] E. Tejedor et al, (these proceedings) **23rd International Conference on High Energy and Nuclear Physics**, Facilitating Collaborative Analysis in SWAN (2018)
- [10] D. Piparo et al, **Future Generation Computer Systems Volume 78**, SWAN: A service for interactive analysis in the cloud, 1071-1078 (2018)
- [11] Arora, I., Sainz, S., Ferreira, P., Labrador, H. & Moscicki, J. Enabling interoperable data and application services in a federated ScienceMesh. *EPJ Web Of Conferences*. **251** pp. 02041 (2021)
- [12] Lo Presti, G., Brosa Iartza, A. & Bukowiec, S. Samba and CERNBox: Providing online access to Windows-based users at CERN. *European Physical Journal Web Of Conferences*. **251** pp. e02024 (2021,5)
- [13] Open Cloud Mesh, <https://cs3org.github.io/OCM-API/docs.html>