



# More Filesystems!

## AFS, FILER, DFS, SAMBA

TechWeekStorage 2024

*Enrico BOCCHI, Sebastian BUKOWIEC*  
*CERN IT*

# Outline

- AFS, Filer
- Storage for Windows systems  
DFS, SMB

# AFS: Use cases

- **Current**

- \$HOME (interactive logon service) → user volumes (/afs/cern.ch/user)
- Filesystem interface for local BATCH access
  - preferably workspace volumes (/afs/cern.ch/work)
- Small Experiments data, projects → /afs/cern.ch/project, ...
- Software compilation, websites... → all (mainly workspaces)


- **Migrated into other services**

- Large Experiments data
- Most of Project spaces
- Archival of old data/projects (**first website in the world!**)
- Software distribution → now on a dedicated, RO system (CVMFS)



# AFS: Infrastructure and Operations

- **Fully virtualized on OpenStack VMs + Ceph storage:**
  - 69 Fileservers (shared, 8cores/15G/4partitions )
  - CEPH block storage
  - 4 DBs
  - Some FS+DB under critical power (\$HOME + Critical projects)
- **Software:**
  - OpenAFS 1.8.10
  - RHEL 9 (on-going migration)
- **Backup to Tape:**
  - Daily
  - Integrated with CERN Tape Archive

# AFS In Numbers



- 44487 volumes, 7k clients
- 487 TB (-4% than 2022)
- 4.4 billion files (+26% than 2022)
  
- Numbers stable, no big additions or removals
- AFS is global, but its usage at CERN is mostly local



# File

- **Virtualized NFS (v3 or v4) server on VM**

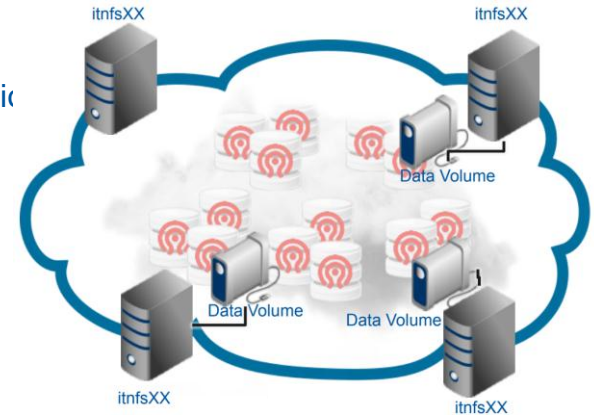
- Underlying storage is Ceph RBD
  - Plain and simple attached volume
  - QoS from volume type or aggregation of volumes
- ZFS as filesystem on top
  - Comes as 3<sup>rd</sup>-party Kernel module
  - Provides snapshots and replication capabilities (``zfs send | receive``)
  - Quota management (``.zfs.{user/group}quotas``), compression, deduplication



- **No HA, no Clustering, Failover is manual**

- We could do better (e.g., DRBD, Pacemaker, ...)
- Never manifested as lacking feature

- **Plans? Consolidate in favour of Ceph**



# DFS Service

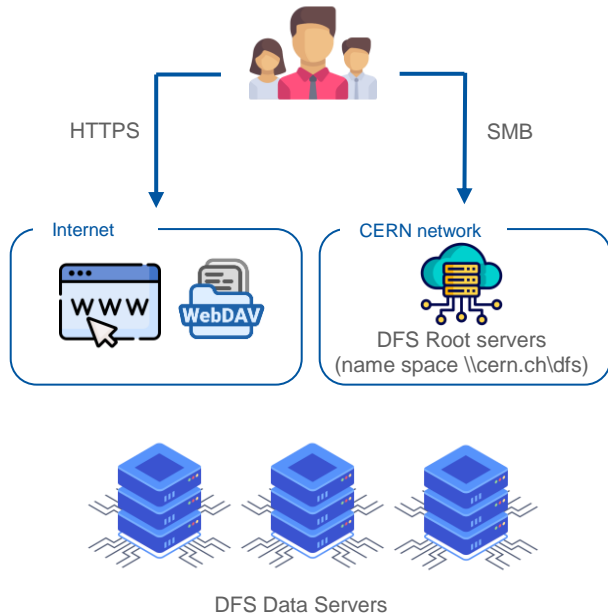
**DFS (Distributed File System)** is a role service in Windows Server that enables to group shared folders located on different servers into one or more logically structured namespaces. Provides a virtual view of shared folders, where a single path leads to files located on multiple servers.

**DFS** provides to all CERN users a centrally hosted file space. Primarily dedicated to the **Windows based ecosystem**.




# DFS overview

## Core components



## Auxiliary components



## DFS features

- Windows Network Share access
- Windows-native permissions management
- Web GUI + WebDAV access
- Quota management
- Snapshots (Previous versions)
- Backup
- Managed by Puppet, Group policies and CMF
- Alarms: icinga2
- Metrics: telegraf > influxDB > Grafana

# DFS Backup

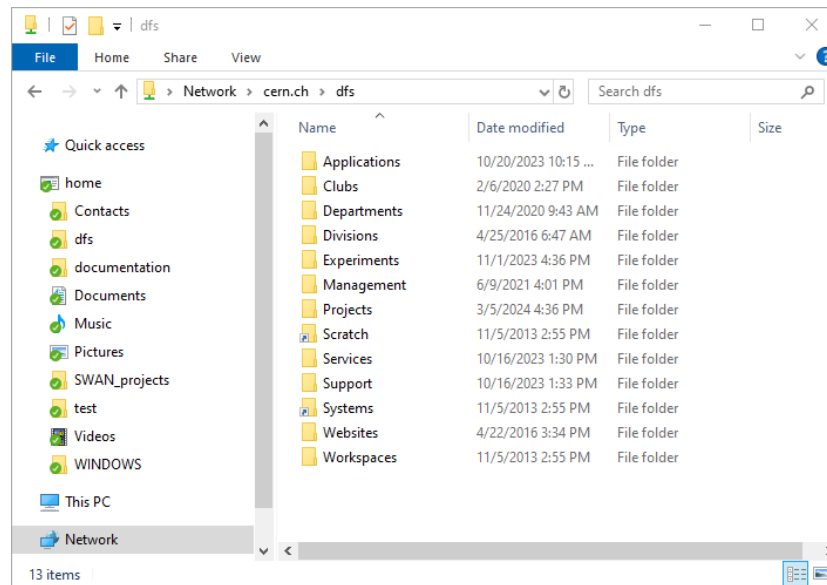
	DFS
<b>Technology</b>	TSM – IBM Spectrum Protect
<b>Schedule</b>	Daily invoked at 18:00
<b>Type</b>	- Incremental - Image every 4 weeks
<b>Retention</b>	Extra Versions: 300 days Only Version: 370 days
<b>Restore</b>	2 <sup>nd</sup> or 3 <sup>rd</sup> line support



IBM  
**Spectrum**  
**Protect Plus**

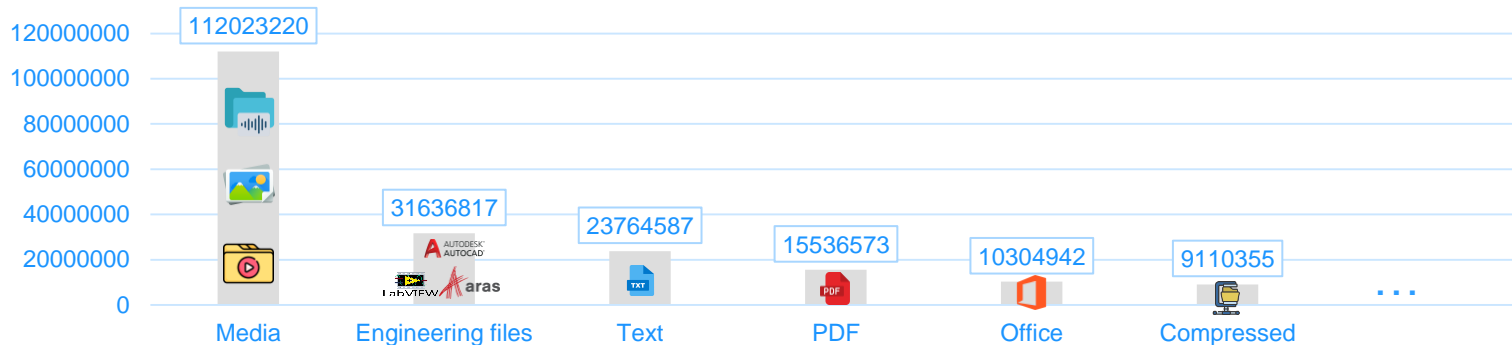
# Usage overview

- **Standard DFS projects**
- **UNOSAT dedicated servers**  
(5 servers, 227 TB)
- **MEDIA (CDS content)**  
(12 servers, 257 TB)
- **Engineering data (PLM/ARAS)**  
(5 servers, Storage Spaces Direct)
- **DFS Web hosting storage**  
(2 servers, 4155 websites)



# Examples of DFS usage

Total number of files stored on DFS: ~500 millions



- Media (CDS multimedia content, lecture recordings etc.)
- Engineering storage
- Documents
- Web for IIS websites
- Sources (images, drivers) for Windows installations (Diane) and CMF applications
- Antimatter experiments: ALPHA, BASE, ...
- Cryogenics, cooling and ventilation, power supply infrastructure data

# Size of the service

Number of “projects”	2 870 (**)
Number of websites	4 102 (**)
Users (unique)	6 338 (*)
Accounts (unique, users + web + computers)	15 984 (*)
Data servers	75
Namespace servers	3
Total capacity	~3 PB
Used capacity	~1.5 PB


(\*) data from 2024

(\*\*) data from 07/03/2024

# DFS plans

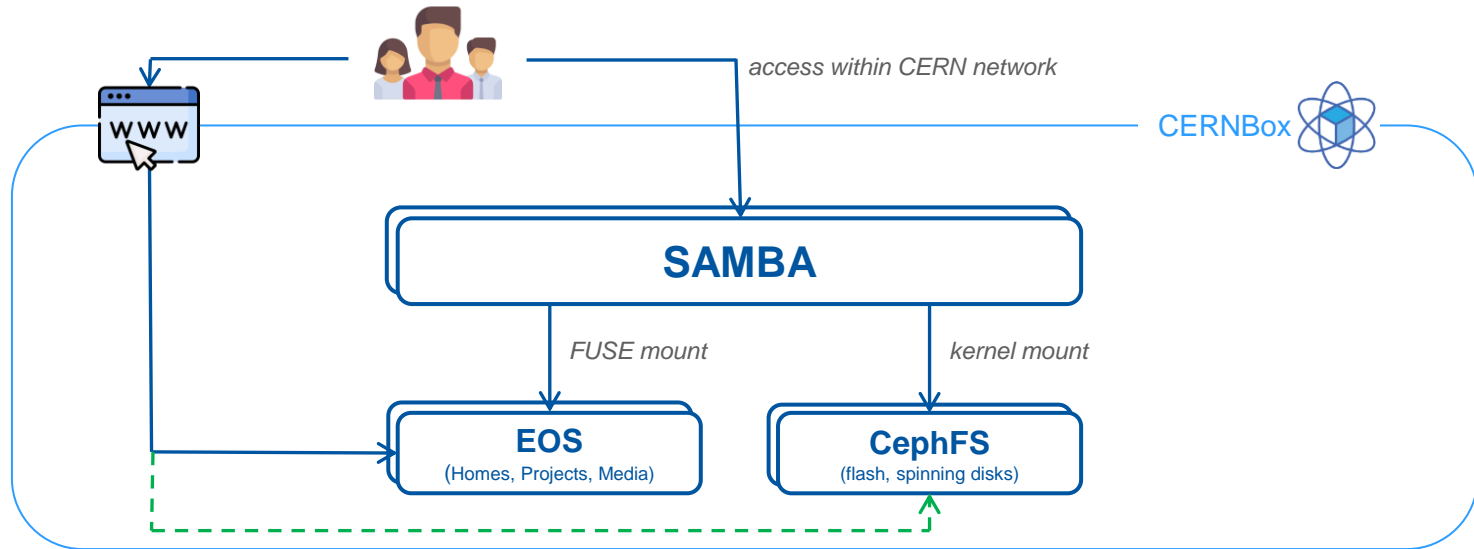
In the context of rationalisation of document management systems at CERN and consolidation of available storage systems the **DFS** is slowly **phased-out**.

**CERNBox** service replaces DFS delivering higher capacity, sync and share capabilities, modern web interface with applications and many more functionalities.

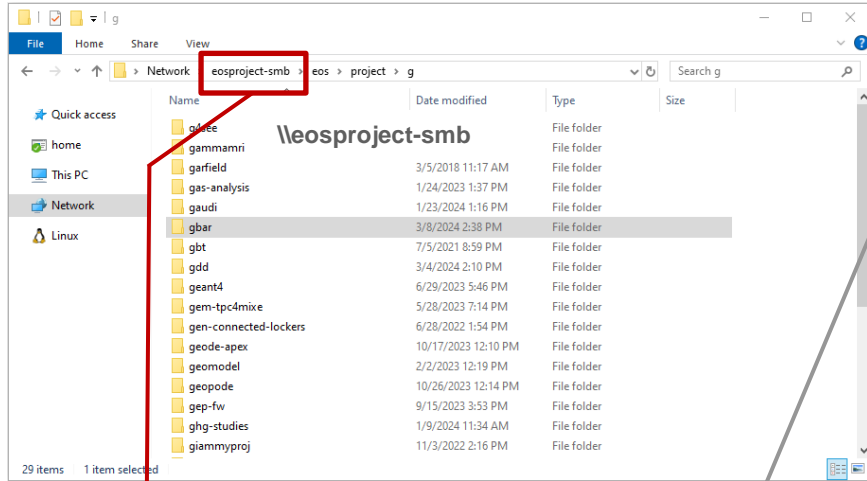


# SAMBA gateway in CERNBox

- Provides SAMBA access for Windows clients to CERNBox
- High-available solution implemented based on the CTDB daemon



# SAMBA architecture



## SAMBA server internals

- CTDB - a high-availability load-sharing CIFS server cluster (floating IPs)
- Bare-metal deployment (due to network constraints at CERN)
- EOS (fuse mount)
  - Homes, media, projects instances
- CephFS (kernel mount)
  - Multiple volumes
- Different ACLs management between EOS and CephFS



# Thank you!

