

Samba and CERNBox:

Providing online access to Windows-based users at CERN

Giuseppe Lo Presti, Aritz Brosa Iartza, Sebastian Bukowiec

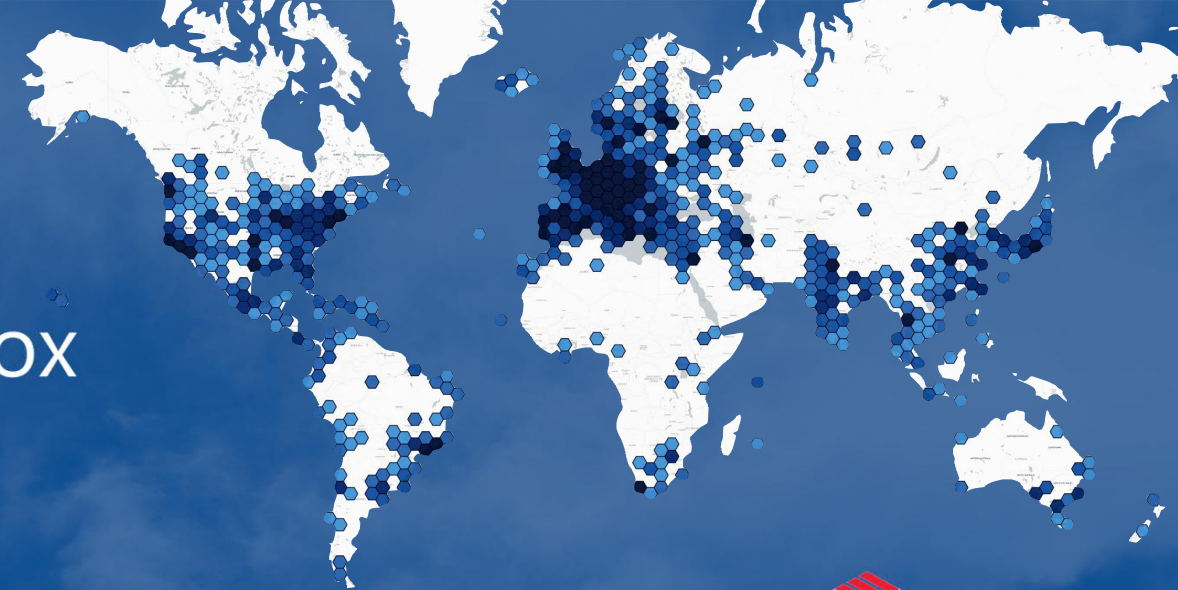
CERN, IT Dept.

Storage solutions for the HEP Community



- 8 years of dev & ops
 - 5K+ monthly active users, 37K users in total
 - 6PB+ data, 1.7B+ files, 110K+ shares
- Sync & Share + online access
- Consolidating “home directories” into CERNBox
 - Onboarding DFS-based use cases
 - **Samba gateways instrumental to support them**
- Central Hub for CERN Data and Apps

<https://cernbox.web.cern.ch>



Powered by EOS



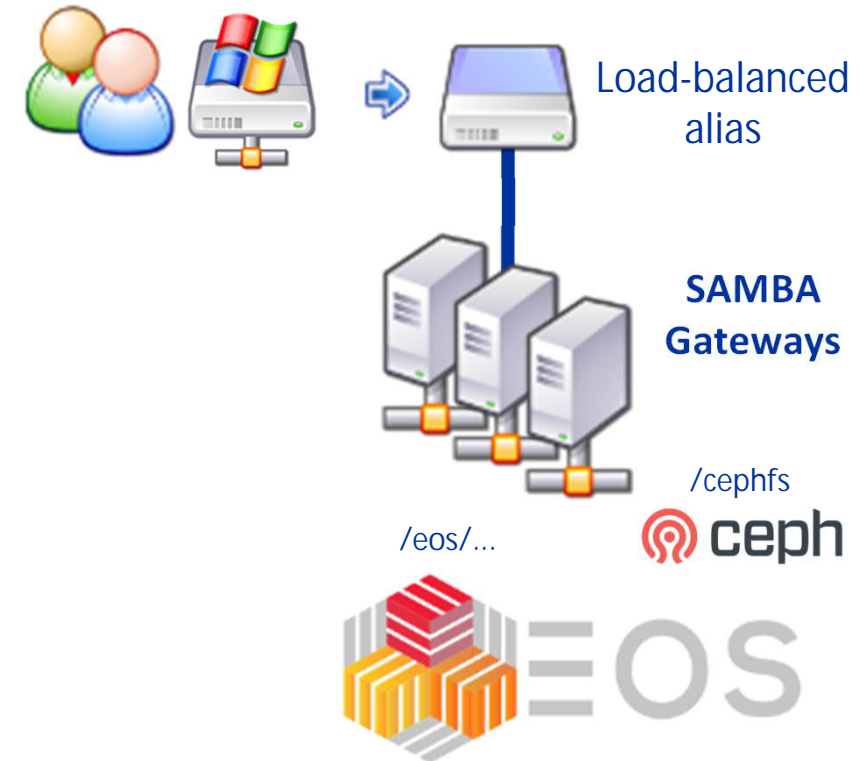
- Open Source in-house storage solution
- 10+ years of dev & ops
- Serving the LHC storage and throughput requirements
 - 100s of PBs
 - Disk & Tape

<https://eos.web.cern.ch>


Samba for CERNBox



- Cluster of 4 nodes, High-Available setup
 - 192 GB RAM, 25 Gbps NIC, *floating IPs* as required by ctdb
 - **Samba 4.11** on CentOS 8.3
 - A small /cephfs mount in place to share the state
- Distributed Storage (EOS) is FUSE-mounted
 - Multiple instances all exposed via \\cernbox-smb\eos\...
- Windows Domain (AD) joined in **dedicated keytab** mode
 - Authc performed by winbind, Authz performed by EOS
- **File locking supported across all gateways**
 - A must to support Office concurrent usage notifications




Timeline




- **Production service as of September 2019**
 - **Samba 4.8** on CentOS 7, 4 nodes with 64 GB RAM
 - Windows Terminal Servers configured to use it for roaming profiles
- **Usage growth in Q2 2020**
 - Upgraded to **4.10**, then **reverted** because of too much pressure on our underlying FUSE-mounted storage
- **New cluster commissioned in October 2020**
 - Samba **4.11**, improved EOS FUSE implementation, **very stable service**
- **“Coming Soon”**
 - Upgrade to Samba **4.13.latest** + deployment of a **VFS module to support RichACL-based permissions**


Usage Evolution




- Close to peaks of 3K connections, average has doubled compared to ~ a year ago




- Significant usage also in terms of I/O ops sustained by FUSE
 - Rates of ~10 kHz seen on a regular basis
 - Windows clients often insist on some files!



Monitoring and Alerting




- Custom probe pushing data to InfluxDB, Grafana dashboard
 - Aim to address issues before users are impacted: log parsing + active probing




- Long run trends in terms of FUSE latency can be analyzed

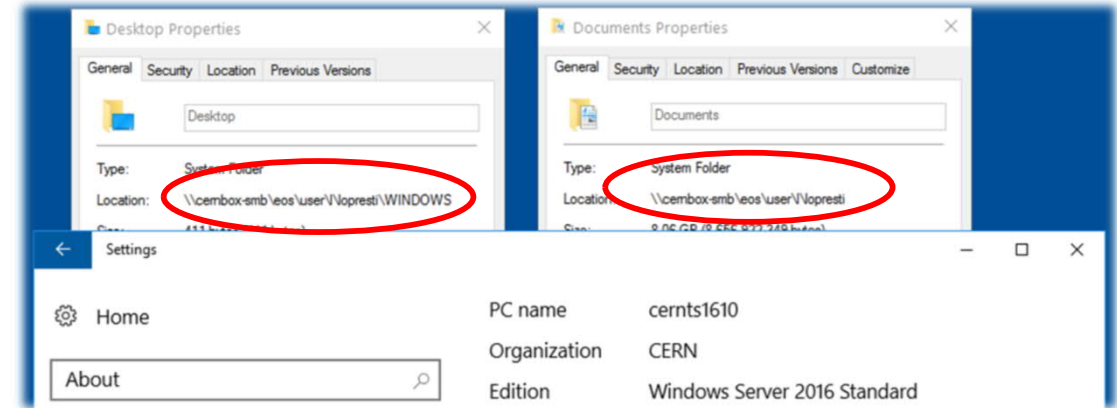
- Faulty behavior in FUSE logs can be alerted to administrators




Monitoring and Alerting



- A Samba over FUSE stack is **extremely latency-sensitive**
 - Substantial efforts invested in our storage to address latency
 - Extensive usage of tracing tools (strace, wireshark) to identify and fix bottlenecks



Time to stat each SMB mount from a Windows Server, daily



Coming next

- **Further usage growth expected ahead:**
about to migrate more Windows-based use cases, in particular concerning **shared project areas used by engineering applications**
 - Possibly need to isolate the most demanding use cases in a separate storage backend, optimized for low latency operations
- Getting more engaged with the **upstream Samba community**:
possible synergies emerged at the recent *SambaXP 2021* Conference
 - <https://sambaxp.org/archive>, held online, May 3rd-5th 2021

Conclusions: the bigger picture

- Multiple access paths . . .

- Windows Desktop/Documents/... system folders
 - Either synchronized, or mapped to Samba

- Samba became a first-class citizen among the available access methods to CERNBox
 - Significant usage, critical service for many workflows in the user community

- . . . aiming at a coherent cross-platform UX

