

Samba: lessons learned

Online access to EOS for Windows-based users


Aritz Brosa Iartza, Giuseppe Lo Presti

CERN IT-ST

Status @ EOS Workshop 2020




- Windows DFS at CERN phase-out started, users being migrated to CERNBox
- Samba gateways from prototype to production
- High Available cluster in place since Sep 2019
 - Starting to gain some operational experience
 - Growing user base, production service at its inception
 - Still dealing with functionality issues, e.g. concurrent Microsoft Office editing was broken
- See [\[EOS Workshop 2020\]](#)




Today: significant usage growth



- Close to peaks of 3K connections, average has doubled compared to one year ago.



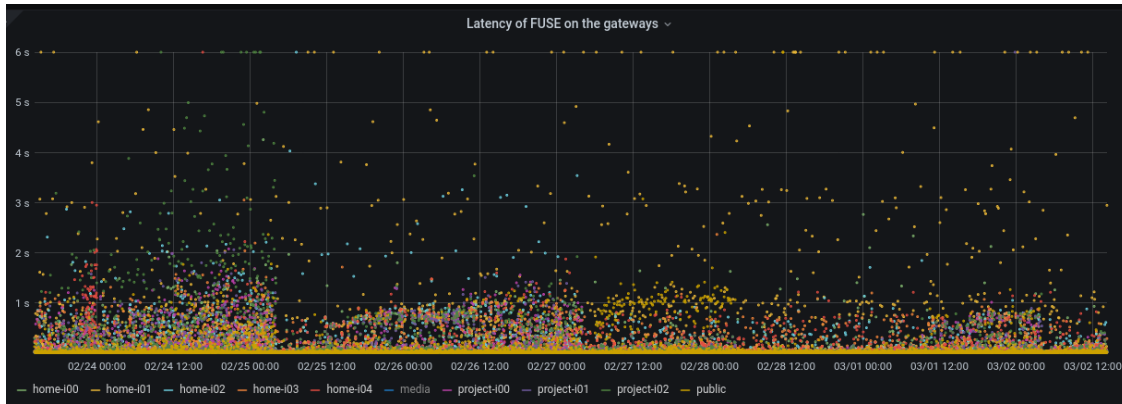
- Usage is significant also in terms of I/O ops sustained by FUSEX
 - ~10 kHz of IO ops on a regular basis



Monitoring and Alerting

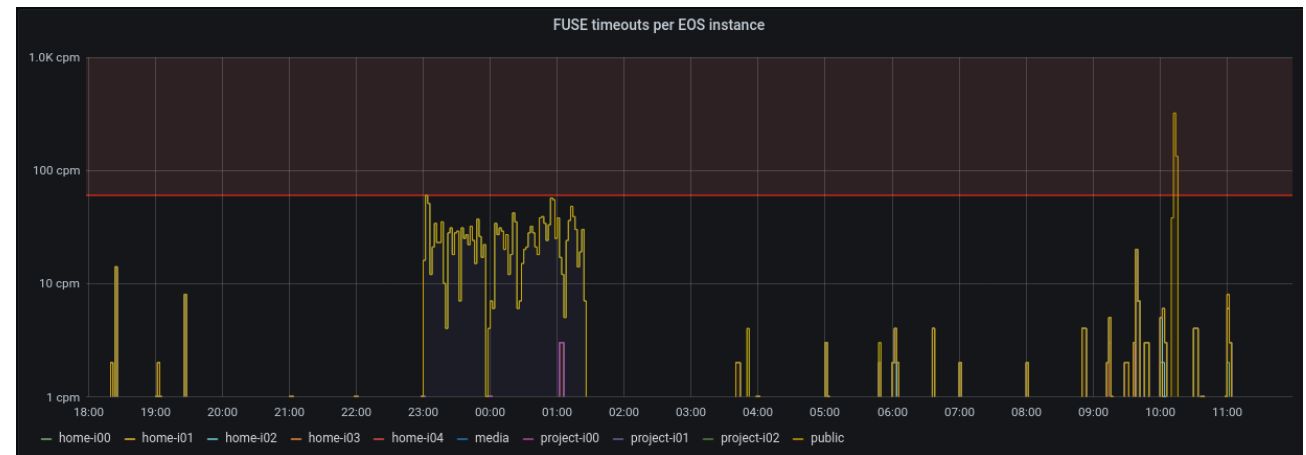


- Developed from scratch in Grafana
 - Aim to address issues before users are impacted



- This plot shows long run trend in terms of FUSEX latency.

- Faulty behavior in FUSEX logs can be alerted to administrators.




Experience and issues



- Samba over FUSEX stack is **extremely latency-sensitive**
 - Many issues fixed in EOS to address latency

Windows SMB mounts latency



Latency to stat the Samba mount from a Windows client (note the log scale)

Experience and issues: a few examples



Windows client's inefficiencies => fetch of all xattrs of all target files

- In POSIX jargon, `fstatat`, `listxattr`, and `getxattr` called for *each file* and *each extended attribute* of the folder
 - Aggressive caching + disabling of eos-specific system-level xattrs
- "Surprisingly", `smbclient` on Linux does NOT query all xattrs

Experience and issues: a few examples



Windows client's inefficiencies => fetch of all xattrs of all target files

- In POSIX jargon, `fstatat`, `listxattr`, and `getxattr` called for *each file* and *each extended attribute* of the folder
 - Aggressive caching + disabling of eos-specific system-level xattrs
- "Surprisingly", `smbclient` on Linux does NOT query all xattrs

Most recurrent filenames: **Thumbs.db**, **desktop.ini**

- Representing ~40-50% of the Samba traffic in terms of requests count!
- This is because of our strategy to provide home dirs through Samba
- Considering caching them in a separate SSD-based pool

Coming next




- **Usage growth** expected ahead:
about to migrate more Windows-based use cases
- Possibly need to isolate the most demanding use cases
in a separate EOS backend, optimized for low latency operations
- Looking forward to the EOS Windows native client from Comtrade!

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

Thank you for your attention!
Time for questions

