

# Samba: service evolution and experience with bind mounts

*Aritz Brosa Iartza, Giuseppe Lo Presti*

**CERN IT-ST**

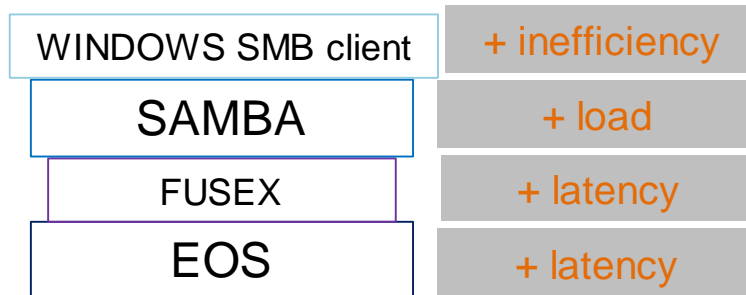
# Status @ EOS Workshop 2021



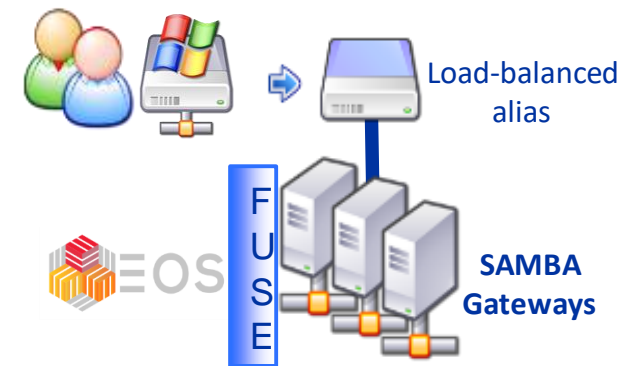
Overview of the architecture of the service:

- Quick recap from last workshop:  
Important service growth since then (more use cases being migrated gradually)

Suffering from user-visible latency



- See [\[EOS Workshop 2021\]](#)



# Split load across clusters




- There has been a recent upgrade on the layout of the hardware:

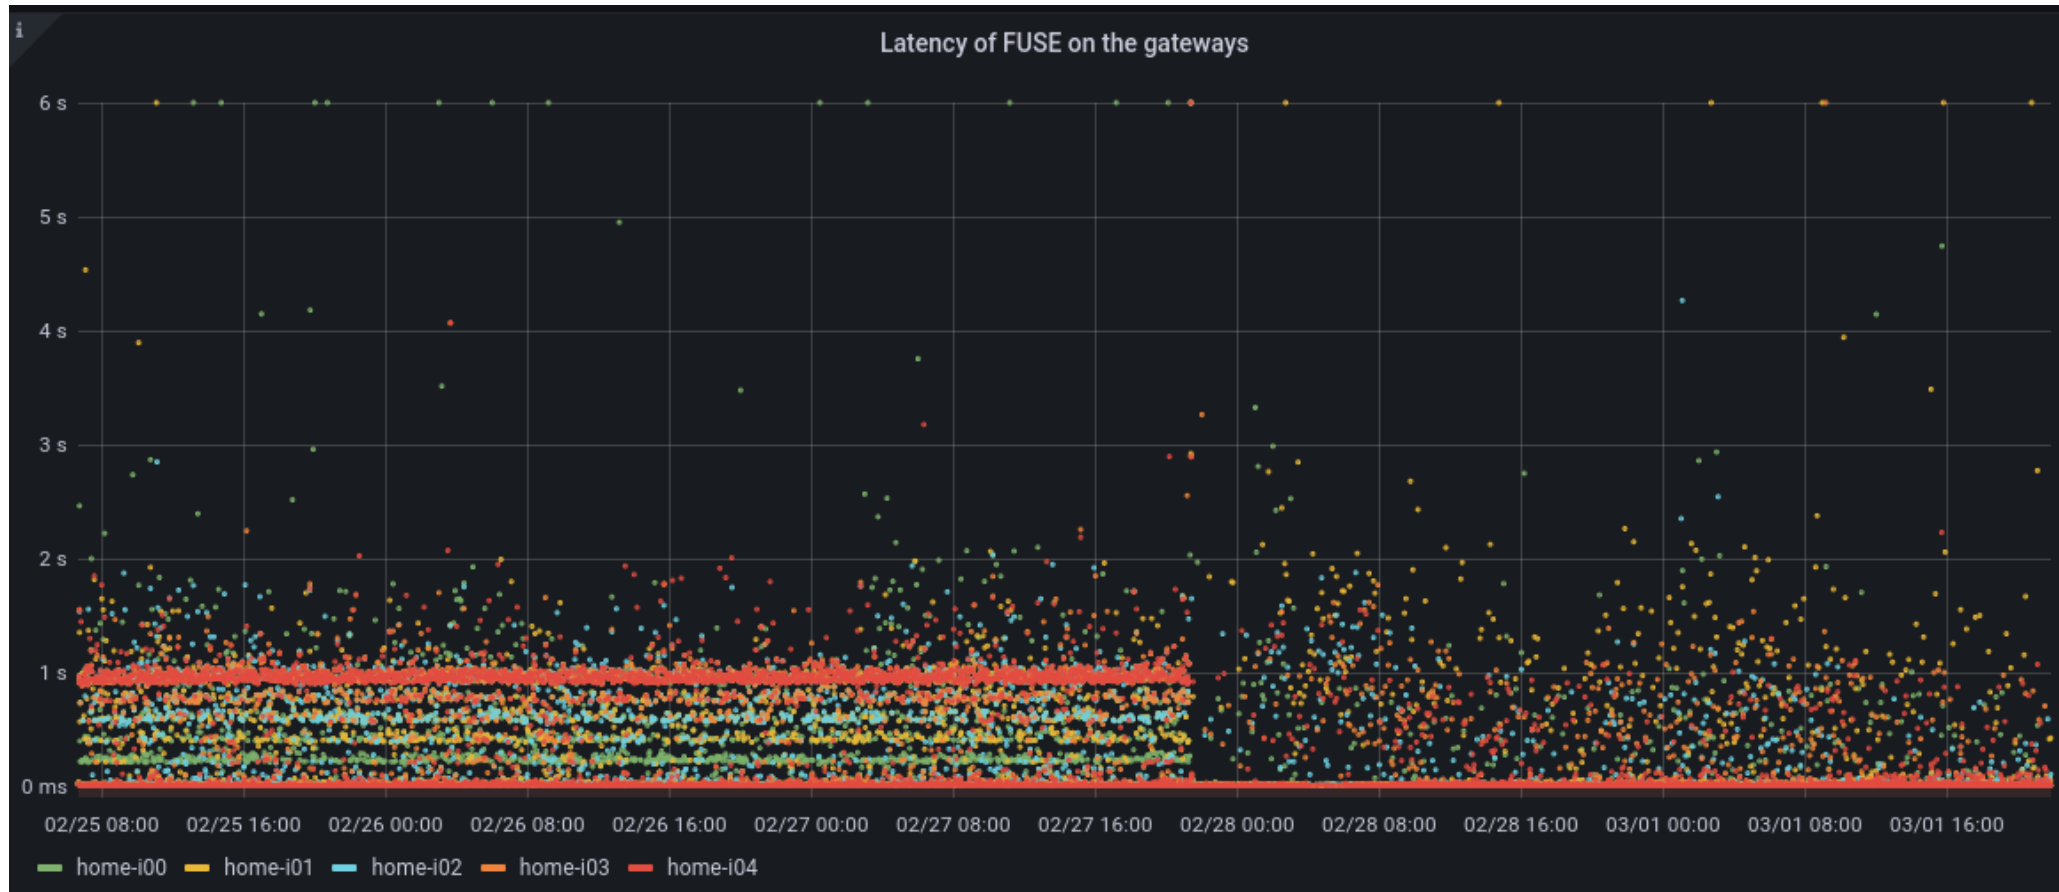
This way, diminishing the load on each cluster, we have seen an improvement on latency.

The use of bind mounts helps dividing more conveniently the mounts, reducing the amount of eos fuse mounts.

Upgraded to EOS 5.0.13.



# Split load across clusters (II)

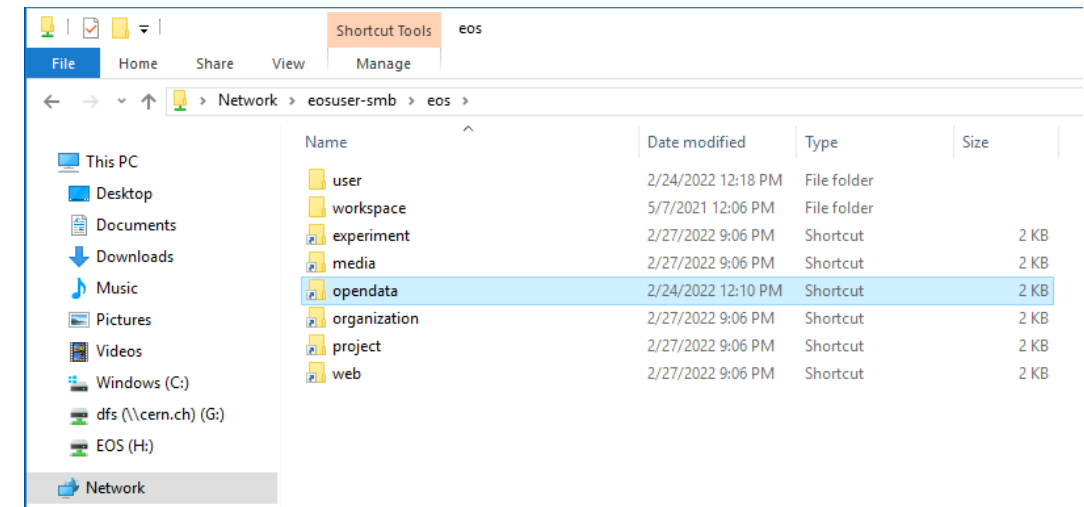


# Bind mounts: a handy asset





**Using bind mounts also helps customizing the directory tree to the user, enabling a more coherent UX**

- Along with the use of DNS aliases and Windows shortcuts (.lnk files)
- This opens up options for further enhancements:
  - Creation of an organigram-based view of all project spaces (already in prod)
  - Creation of a virtual mount to expose the user's "My Shares" and "My Projects" web views



# Further improvements



- Sustained high rate of requests observed in the system.
- After a thorough analysis, it turned out the most requested files were unnecessary files.
- The cleanup intervention drastically reduced the request rate.

Cleanup intervention:  
19 Nov 2021

# Future prospects



- Many DFS projects are being migrated, but some require a level of performance we are not yet able to provide

↳ Certain apps' workflow for interactive use-cases are particularly bad.

- Promising preliminary results in prototyping a low-latency optimized EOS backend
  - See e.g. [\*EOS on CephFS presentation\*](#)
- An assessment of all clients' impact on EOS homes/projects (including internal ones, such as the backup system) is being carried out
- The new multi-cluster deployment with bind mounts provides margin for further scaling out the service

Thank you for your attention!  
Time for questions

