



EOS site report of the Joint Research Centre

EOS workshop 14-15 March 2024

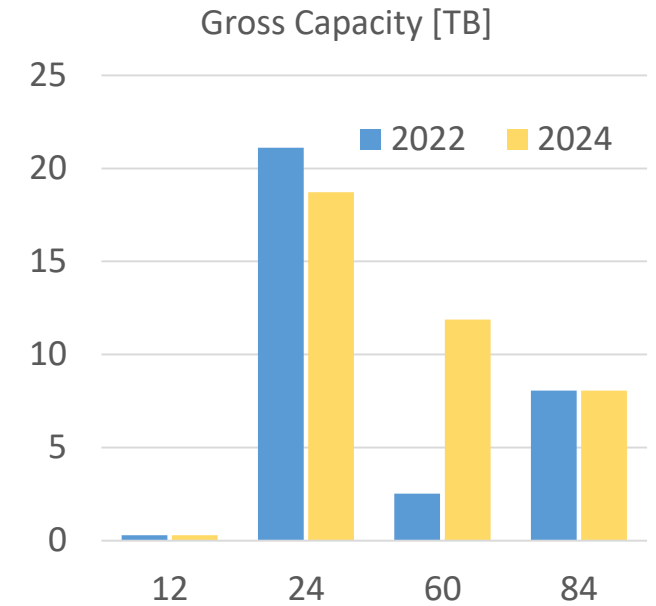
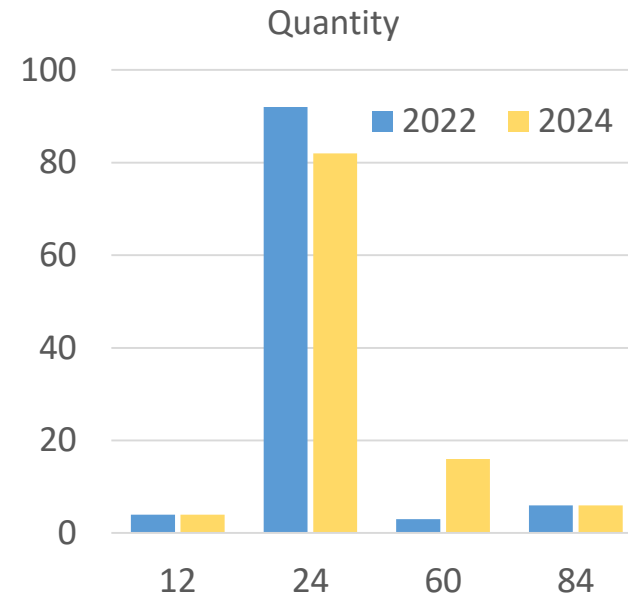
Armin Burger, Franck Eyraud

EOS instance at JRC – Overview

- Main storage backend of the JRC Big Data Analytics Platform (BDAP)
- In use since 2016
- EOS versions
 - MGM: 5.1.22 FST: 5.1.20
 - Processing & service nodes: mainly 5.1.x and 5.2.x, some 4.8.x
- Access via EOSX FUSE client
- Replica 2
- Server OS: CentOS 7

EOS instance at JRC – Hardware

- Heterogeneous hardware
- JBODs
 - Disks per JBOD: 12, 24, 60, 84
 - Disk sizes: 6, 10, 14, 16, 18 TB
 - Replacement after ~7 years
 - New JBODs: 60-bay, 18 TB disks
- FST servers
 - 12-20 cores, 128-256 GB RAM, 1-2 TB space for logs



EOS instance at JRC – Current status

- 56 configured FSTs with 1 or 2 JBODs connected
- 30 PiB gross capacity (36 PiB soon)
- 93% used
- 1.9 G files
- 265 M directories

EOS activities in 2023 at JRC

- Upgraded EOS MGM from 4.8 to 5.1
 - Overall worked fine
 - Initial issues of MGM crashes have been very quickly solved by CERN EOS team with new provided version
 - Solved various issues
 - fsck was usable again
- Migrated from LevelDB to extended attributes
 - Solved issues of blocked FSTs due to insufficient memory

EOS activities in 2023 at JRC [2]

- Updated fsck status
 - Use the repair feature to lower numbers of reported files with problems
 - Still some files that can't be repaired
- Decommissioned 10 old nodes (24x6TB)
 - Needed to balance between groups to free space on the groups concerned by drain
 - Identified additional corrupted files with respect to the ones reported by fsck
- Setting up encryption on EOS, using native XRootD functionality
 - For project folders with sensitive data
 - Decryption via fuse client, key on mount configuration

EOS issues experienced

- Crashes of FUSEX client can still happen
 - Mainly on read-write mounts
 - Usage of side-car containers and multiple mount points reduces impact on users
 - Might be reduced by using latest EOS clients
- Some client crashes due to filled up local file cache
- A few cases with 0-byte or corrupted files when writing to EOS
- Arbitrary read errors on TIFF and NetCDF files
- File missing replica leading to a *read only file system* error at file updating

Some nice-to-have features for EOS

- Permission handling like POSIX style
 - Permissions of parent directory effective also for child directory
 - Configurable at directory level
- S3 native support
 - Ideally with a possible configuration as gateway
=> no MGM/FST access of client necessary

EOS related activities planned for 2024

- Upgrade EOS to latest 5.2.x version
- Migrate all EOS nodes to Alma Linux 9
- Extend capacity with additional 60-disk JBODs
- Evaluate S3 gateway setup towards EOS
- Investigate in suitable backup solution for project data

Many thanks to
the CERN EOS development and
management teams for the support !



© European Union 2024

Unless otherwise noted the reuse of this presentation is authorised under the [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/) license. For any use or reproduction of elements that are not owned by the EU, permission may need to be sought directly from the respective right holders.

