EOS deployment at GRIF

Vamvakopoulos Emmanouil

On behalf of Technical Committee at GRIF

LCG FRANCE meeting

Nov 2025 -IRFU CEA

EOS (large scale distributed storage)
key features

Core Architecture

- Built on XRootD framework
- QuarkDB key-value store for namespace metadata (KV store on RocksDB)
- MGM (Management): Namespace management, runs with QuarkDB
- FST (File Storage): JBOD-based storage servers with lightweight config

Authentication Methods

- WLCG tokens/SciTokens
- Kerberos, X.509 certificates, Pre-shared keys
- Local tokens and Macaroons

High Availability & Data Protection

- Master-slave MGM failover with minimal DNS dependencies
- File replication across multiple storage nodes
- Erasure coding (software RAID)
- In-memory namespace cache with persistent Key-Value backend

Access protocol/method

- root, https, https/TCP, grpc, fuse

- **Advanced Features**

- Space, group, filesystem (fs) organization
- Geo-balancing
- Intelligent geographical data placement
- ACL like nfs4 and quota subsystem

GRIF is a distributed site made of four (4) different subsites,
in different locations of the Paris region:
LLR, LPNHE, IJCLAB et IRFU/CEA.

- **IRFU, LLR** and **IJCLAB** are interconnected with 100Gbit link, LPHNE with 20 Gbit
- The worst network latency between the subsites is within 2-4 msec
- Four (4) independent DPM instances
- Total Pledges Capacity ~20PBytes
- Supports four (4) WLCG VOs: **Alice, Atlas, CMS and LHCb + several EGI VOs**
- Hardware configuration is mainly storage servers with 10Gbit nics (or more) with direct attached sas disks
- Data protection based on RAID-6 done by server's controller
- Quite heterogeneous hardware layout and hard drive sizes between the sites and servers' generations

EOS@GRIF

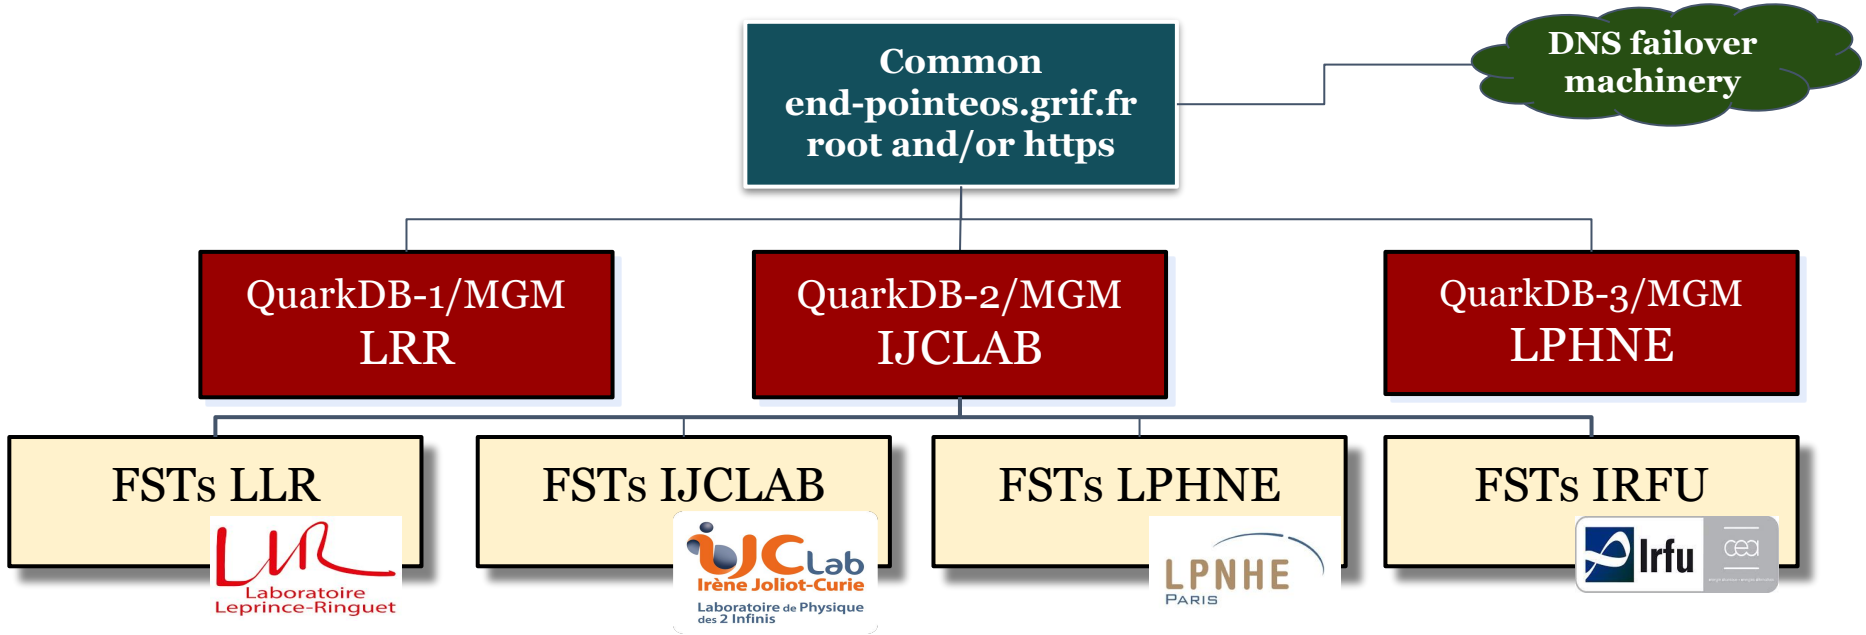

- Quarkdb (and MGMs) cluster with three (3) nodes
- FST nodes will span over four (4) sites
- Storage accounting

Milestone of the Project : 2021-2024

Phase A: Evaluation and Selection (Nov 2020 – Nov 2021)

- **1. Resilience Design:** Defined a multi-site architecture
- **2. Intensive Testing:** Evaluated EOS and dCache,
- **3. EOS Selection:** Selected **EOS**

Phase B: Installation and Validation (Nov 2021 – June 2022)

- **1. Hardware Constraint:** Deploy EOS the initial (RAID6) configuration
- **2. Integration & Access test of protocols (**https, root) with various authentication methods (X509, macaroons, tokens) for all VOs.
- **3. Migration Planning:** Defined and prepared the **VO-by-VO data migration plan**

Phase C: Transition and Production (Jun 2022 – Dec 2024)

- **1. (Zero) Downtime Migration:** Executed the data migration by leveraging **experiment-specific data management tools (e.g., Rucio)**
- **2. Hardware Recycling:** Systematically **reorganized data on DPM** to free up old servers and progressively add them to the new EOS configuration, fulfilling the hardware reuse constraint.

EOS version and infrastructure

- Running EOS 5.3.9 on mixed RH8 and RH9 Linux distributions
- **LPNHE:** Rocky Linux 8 and 9 (new installations)
- **IRFU/CEA:** Rocky Linux 8
- **IJLCAB and LLR:** Upgrading from CentOS Stream to Alma Linux 8, with some new Alma Linux 9 installations
- **Planned upgrade** to EOS 5.3.27 (!)
- Total ~100 FST nodes + 3 MGM

Supported Virtual Organizations (VOs) and IAM

- **WLCG VO**
  - alice (alice token)
  - atlas (IAM)
  - ops
  - wlcg (IAM)
  - cms (IMA)
  - lhcb(IMA)
- **EGI VOs**
  - complex
  - belle II
  - VO based on Dirac WMS
    - CTA (IAM)
    - HESS
  - France-Grilles (IAM)
  - clas12

Distribution of storage capacity

- **We have heterogeneous distribution of storage capacity over the four (4) sites which depends from**
  - Difference of funding streams of each subsite
  - Internal network architecture and cooling capabilities differ at each subsite
  - Different hardware layout due to different purchases campaigns
- Keep the data protection under raid6 and split large (~120-192TB) raid6 volumes on several partitions smaller (FS) partitions almost equal ~20-25TB

Lab	INSTALL (TB) Defaut Space (Production)	(%)
LLR	3614	**16**
IJCLAB	6024	**26**
LPHNE	3907	**17**
IRFU/CEA	9491	**41**
Total	**23036**	100%
Used	**18532**	~110M files

EOS@GRIF Default Space Organization

Space Structure and Access

- All VOs have access to all default space elements (group and fs) - the only constraint is the quota limit
- Default space organized into groups (10), each group contains ~90 Filesystems
- Each FS has almost shame same size (20-25TB) - each group has the same number of FS

Round-Robin Distribution and Load Balancing

- Round-robin group selection makes all group equivalent in the scheduling algorithm
- For each group the only bias weight is the number of FS per site
- Plus Internal metrics of fs/node load are taken into account
- Sites with more FSs accept proportionally more data

Benefits

- Uniform utilization of network resources for all VOs
- Uniform write and read patterns
- High redundancy and resiliency on write transactions
- Less work for balancer - try to reduced inter-site traffic
- Aggressive balancer use when adding fresh FS (not an issue - one-time operation)

LHC VO	Plegdes (TB) Q3-2025
ALICE	2799
ATLAS	8078
CMS	5648
LHCB	1772
Total	**18297**

Daily operations with EOS

- **Filesystem Management**
- Add new filesystems (FSs) and remove old ones
- Drain (FSs) before decommissioning
- Deal with offline FSs or nodes
- **Data Distribution and Balancing**
- Balance data inside a group and between groups
- Multi replicas for sensitive files

Access Control and Security

- Authentication issues (user/group mapping)
- ACLs (Access Control Lists) management
- Quota level enforcement
- **Data Integrity and Monitoring**
- Consistency checks (e.g. ATLAS, LHCb experiments)
- Check file status and state
- Verify replica health and availability

Incidents with EOS

Issues:

- ACL/mapping problems due to CA change to Harica (DN changes)
- FST offline (each FST: ~140TB, 0.5M files)
- FST hardware/disk failures (e.g., IJCLAB)
- DN changes
- File quota limits exceeded
- MGM components race condition bug during shutdown/initialization of MGM

Further Steps

- Monitoring and logging
  - Prometheus metrics monitoring (e.g/ eos_exporter)
  - ccnetlab, renater and geant monitoring
  - Scitags and packet marking
  - Elastic search for the MGM logs
- Test LRU deletion for temporary areas in namespace (e.g. cms temp dir)
- Understand better the namespace structure, fsck and durability process
- Incorporate wlcg token (e.g. for CMS) for root protocol with ztn

Acknowledgements

Many thanks to EOS developers team for the discussions and the recommendations

Many thanks for yours attention
Questions and Comments ?

BACKUP slides

An Ideal Matrix: N server by K Filesystem (of same size)

We do not have N with k drives per server

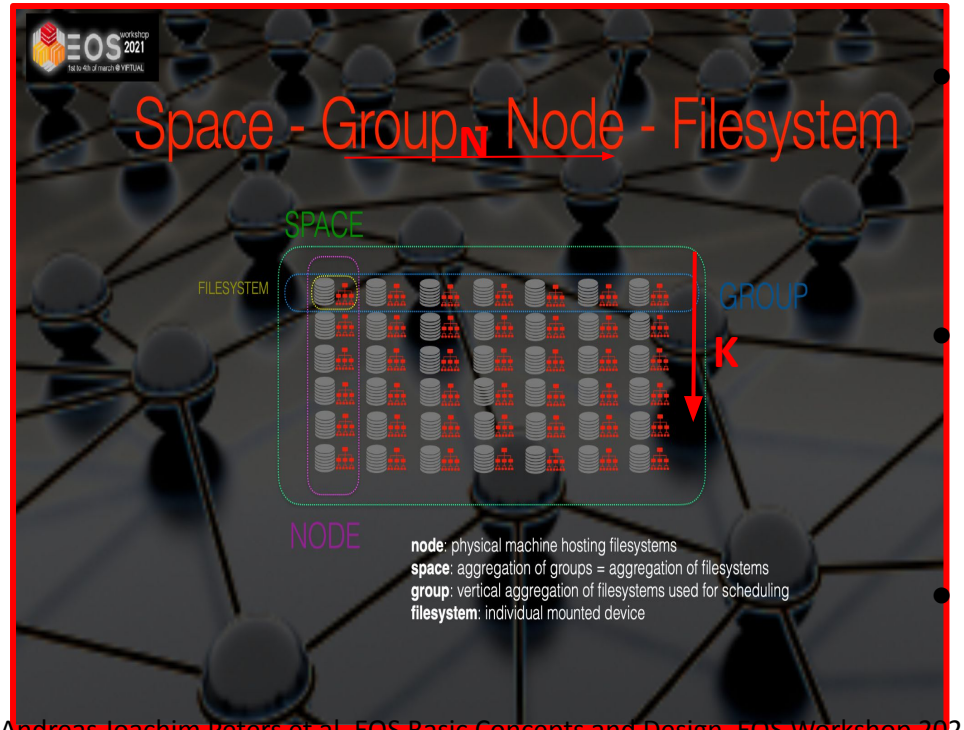

- On Ideal case we have:
- N servers with **K** individual FS on each server
- (of the same size)
- Thus we have **K** groups with N filesystem on each group (from N different servers)
- Easy to add a new server of same size (of K individual FS)

Andreas Joachim Peters et al. EOS Basic Concepts and Design, EOS Workshop 2021