EOS Deployment at GRIF

Vamvakopoulos Emmanouil

On behalf of Technical Committee at GRIF

Journées LCG-France at IPHC

6-7 june 2022

Strasbourg

Outline

- Brief description of GRIF's current storage system
- Motivation for changes
- Context diagram of EOS services at GRIF
- Comments on installation and configuration
- Organization of EOS filesystems and data protection
- Progress up to the moments
- LHC VOs migration

Storage@GRIF for LHC/EGI VOs

- **GRIF** is a distributed site made of four (4) different subsites, in different locations of the Paris region.
- **IRFU, LLR** and **IJCLAB** are interconnected with 100Gb links.
- The worst network latency between the subsites is within 2-4 msec
- Four (4) independent DPM instances
- Total Pledges Capacity ~10 PBytes
- Supports four (4) WLCG VOs: **Alice, Atlas, CMS and LHCb** + several EGI VOs
- Hardware configuration is mainly storage servers with 10+ Gbit NIC and with direct attached sata 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

Motivation for changes

- DPM is reaching its end of life soon as a WLCG/EGI service
- GRIF represents a total of ~10 PB but is seen as 4 medium-size sites
  - Avoid duplication of data amongst the subsites
  - (depending on the VO's DDM workflow, e.g. atlas secondary replicas)
  - Optimum usage of storage resources in a common pool
- Eos common pool makes makes GRIF configuration appropriate from datalakes perspective
  - Grif Has the potential to be a major player in a French datalake if it can expose one GRIF endpoint for each VO
- We can share experience/tools but each subsite has to be managed independently for optimal management
- We need to consolidate our efforts amongst the four subsites with the same people workpower
- In addition, work started on a distributed Ceph instance could open the way for more things in common

EOS@GRIF

final

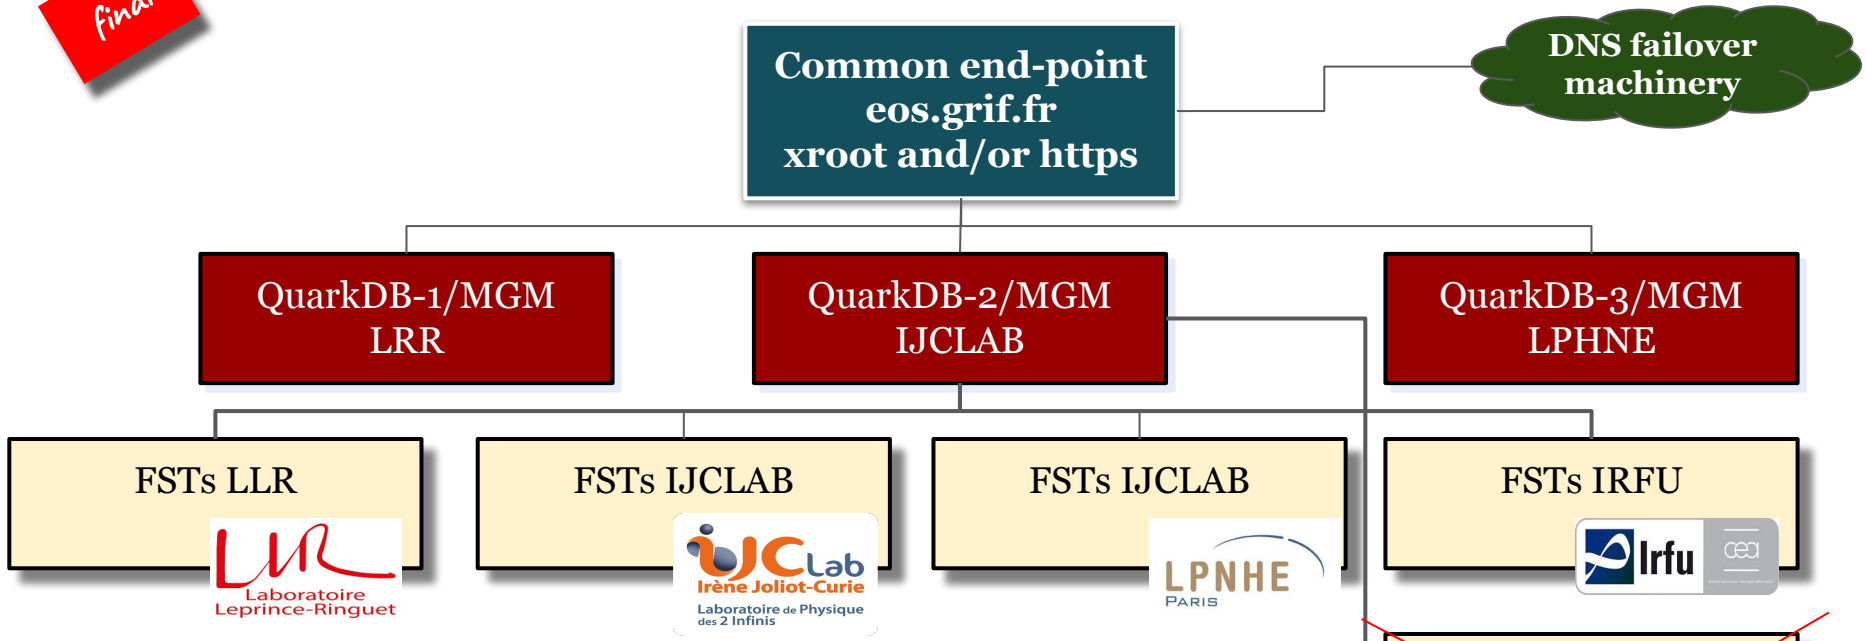

- Quarkdb (and MGMs) cluster: 3 nodes distributed on 3 sites
- FST nodes (disk servers) will be distributed in the 4 GRIF subsites
- ~~Some xrootd PSS gateways for xroot TPC with delegated proxies~~

EOS Installation & OS version

- EOS Diopside (5.0.x) installation on Centos Stream 8 and Rocky Linux
  - OS distribution based on GRIF subsite preference
  - Still in testing phase (repositories)
- EOS 5.0.x: 2 different repositories for CS8 and RHEL8/derivatives
  - CS8: always built with the latest CS8 RPMs
    - not working with RHEL8 and derivatives
  - RHEL8: built with official RHEL8 container image
    - https://hub.docker.com/r/redhat/ubi8

Installation and Configuration

- Automated deployment and configuration
  - Quattor (3 subsites), Puppet (1)
  - All the EOS/xrootd configuration files managed with the configuration tools
- IPV4, IPv6 public network
- Keytab secrets and macaroons
- MGM endpoint: alias failover through DynDNS managed by a script
  - Update DNS alias based on which instance is the MGM master (quick failover at the MGM level)
  - Used successfully for years for services like BDII
  - Latency: at least a minute (depends on cron job frequency)
- Freeze/Production version 5.0.18 (in-depth testing for auth/z and access protocol)
  - Still we have some issue with Alice and alicetokenacc xrootd auth/z plug-in

alicetokenacc xrootd auth/z plug-in

- Configure TkAuthz.Authorization
  - `EXPORT PATH:/ VO:* ACCESS:ALLOW CERT:*`
  - `RULE PATH:/eos/grif/alice/ AUTHZ:delete|read|write|write-once| NOAUTHZ:| VO:*| CERT:IGNORE`
  - `KEY VO:* PRIVKEY:/etc/grid-security/xrootd/privkey.pem PUBKEY:/etc/grid-security/xrootd/pubkey.pem`
- Maps default unix account/group to be alis:alice for unix authentication
- tpc https it works for token/gsi authentication
- Plain root access for non alice vo do not work for gsi auth (a atlas dn maps to alis:alice)
- We force the xrootd client to use gsi but do not work
- Need further careful checks …

GRIF deployment and Erasure Coding (EC)

Erasure coding considered but abandoned for EOS@GRIF for two major reason:

- Hardware profile at GRIF amongst the four sites is not uniform
  - Different number of servers, different number and size disks due to different local capacity plan
  - Continuous procurement process: no chance to buy the same HW config 2 years, not a volume large enough to build a new homogeneous FS group each year
- Achieving a global resiliency comparable to DPM would require a storage overhead ~30%
  - Failure at one site should not impact more that the data stored uniquely at the site
  - Current RAID6 diskservers have an average overhead of 12%: no budget for increasing it

EOS and Space Organization

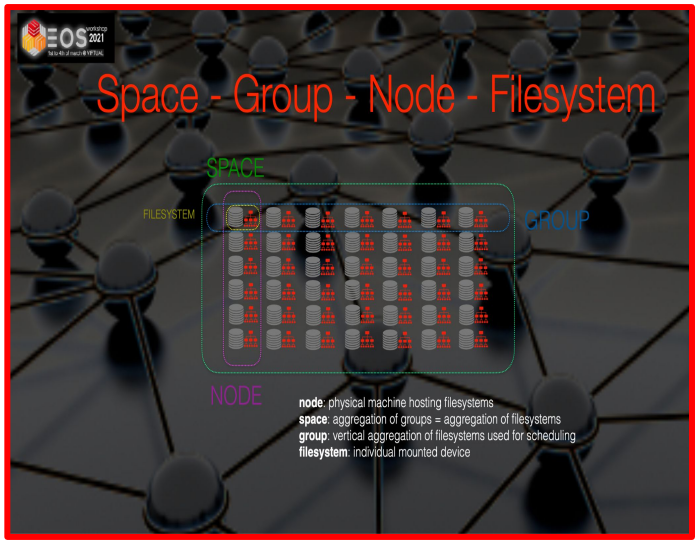

- **One EOS space all the VOs**
  - Uniform utilization of the capacity and the server bandwidth (disks and network) as much we can
  - One default space for pledge resources and one space for unpledge and local resource
- **Scheduling groups** with filesystems (FS) from every site, resilience at the FS level as no erasure coding
  - A file is in one FS only: losing a file system will impact neither files not stored in this file system nor ability to write to the FS group
  - Some VOs may be restricted to some subsites using FS geotags
- **EOS FS:** no strict requirement of adopting a unique size as EC is not used
  - EOS provide a balancer to ensure that the usage level of each volume is “the same” with various policies
  - Most FS wil be RAID6 volumes, typically in 14+2 configuration, sometimes splitted in several (equal) partitions.
  - It is not bad to standardize approximately FS size to limit the rebalancing
  - Large ceph volumes (500 TB) could as backend storage device, at least for the transition/migration period

Distribution of Used “space” for migration

FEB 22

	IRFU	IJCLAB	LLR	LPNHE	Total
ALICE	0 TB	966TB	0	0	1,4PB
ATLAS	1.9PB	1.3PB	0	1.3PB	4,5PB
CMS	1.5PB	0	1.8PB	0	3,3PB
LHCB	0	156TB	0	113TB	289TB

- We have servers with total attach capacity (from 100TB, 160TB 240TB up to 760TB)
- Number of servers per subsite: 4 server on LPNHE, 11 on LLR, 14 on IJCLAB, 32 on IRFU

LHC VOs first dialogue and data migration plan

- **Atlas**
  - Atlas DDM group will perform the data migration via rucio and FTS (on progress)
  - Atlas rucio and panda machinery are having capabilities for seamless transition
  - Thanks to existence of cached/secondaries replicas at GRIF sites the total amount of ATLAS data to migrate is only 3.3PB
- **CMS**
  - Data management group of CMS will perform the data migration via rucio and FTS
  - We need to assess with CMS people the capabilities for seamless transition
  - Not clear yet if we are required to migrate all the data (cache, secondary replicas…)
- **Alice**
  - Data management group will perform the data migration
  - The data migration will be done offline (maximum 2 steps)
- **LHCb**
  - Still in discussion, LHCb would prefer that GRIF handles the migration but we are not really ready to do it
  - Small volume of data compared to other VOs (~300 TB)

In all the cases, we need a decent amount (~1PB) of extra capacity for the initial data migration before releasing the first DPMs servers

Réalisations jusqu'à présent

- **Fulfill the milestone to have a working instance mid-mai**
  - **xrootd access**
  - **https TPC**
  - **Add new FSTs**
  - **Setup internal access, quota, space and groups and Accounting report**
- **Starts ATLAS the migration of LPNHE and LAL on 23/5/2022**
  - ATLAS LPNHE (O.5 PB) → ~ done
  - ATLAS ex-LAL (0.6 PB) → ~ almost done (50 TB left)
  - Atlas Panda queues for ex-LAL and LPNHE queue use the eos.grif.fr as primary storage endpoint

EOS@GRIF current configuration and volumetric

- **MGM 5.0.18-1**
  - ~160 FS on production
  - meta-data: 1,7M files et 800K directories
  - Size of quarkdb flat files is 10GB x 3
- **20 nodes FST online in total**
- **Spaces:**
  - Default (pledged resources) — 3.44 PB total / ~ 1 PB used
  - Localgroup (no pledged resources) — 720 PB total / ~70TB used
  - Spare (a temporary space for unallocated space) — 0.25 PB total

LPHNE AND LAL DATADISK to GRIF_EOS

23/5 to 6/6

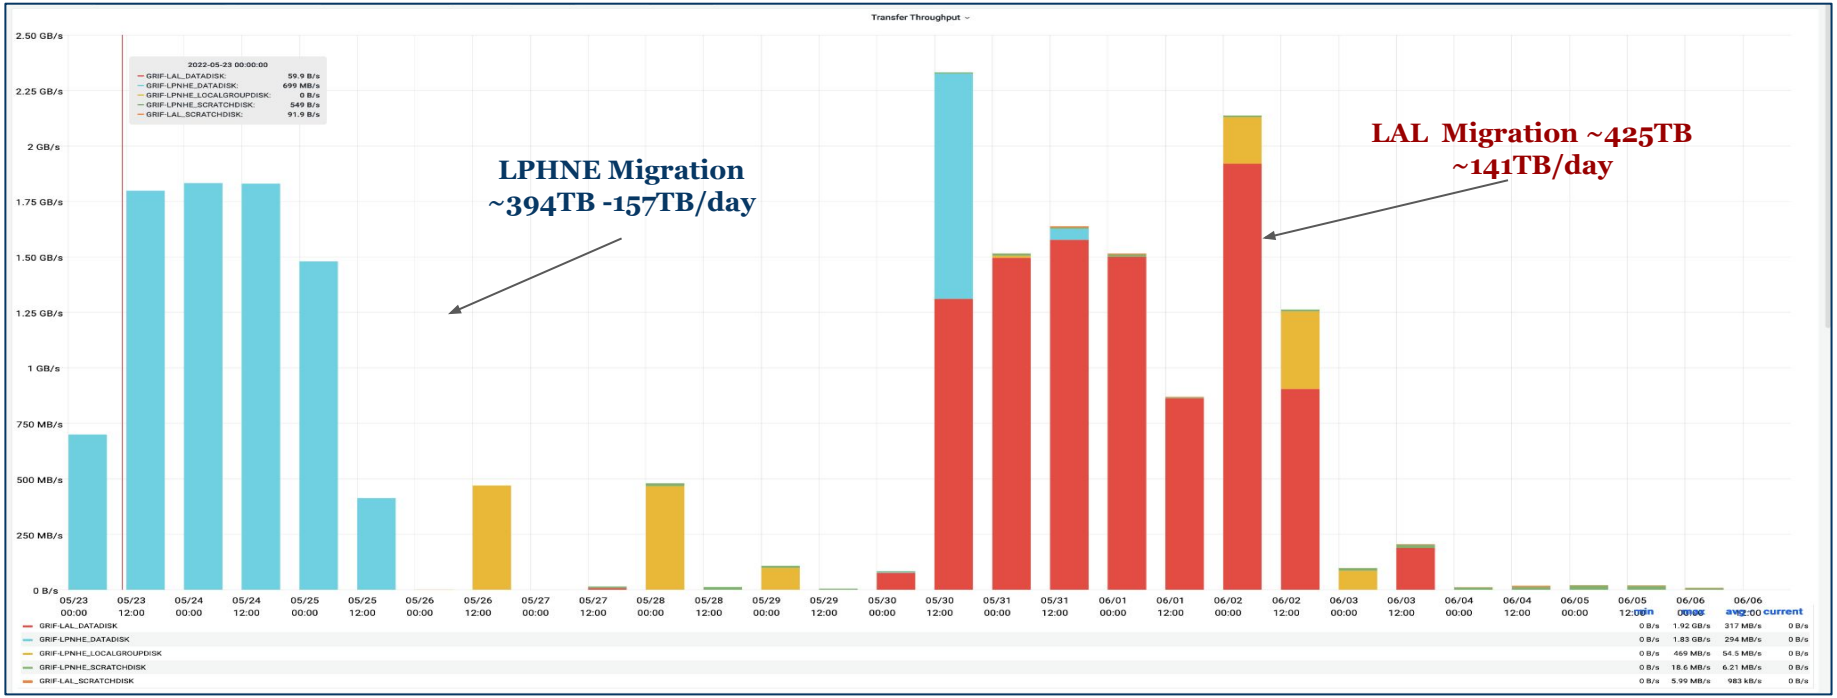

Total Aggregate bandwidth from Alice apmond tool

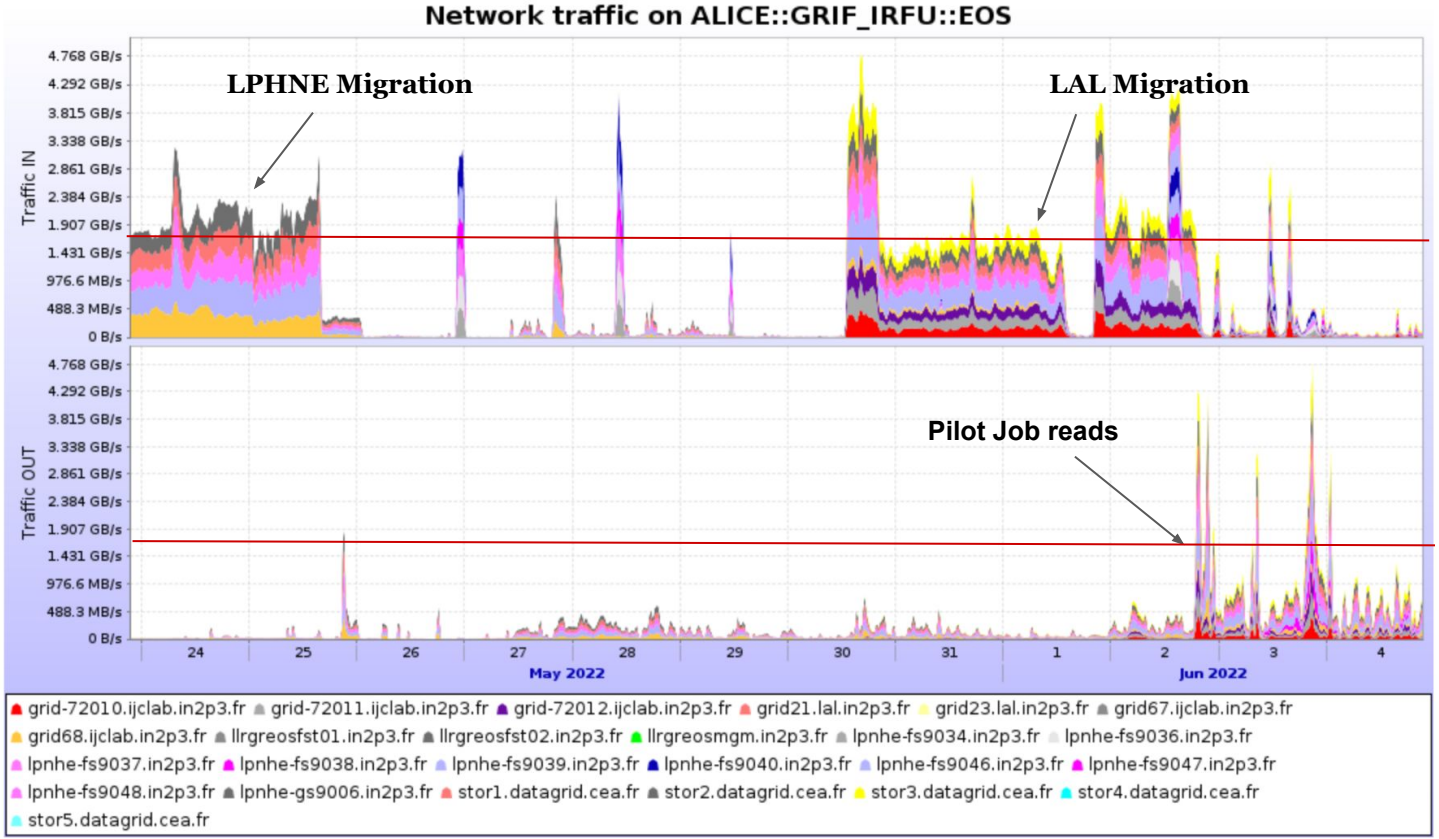

Updated Migration Roadmap

- **Phase 1 preparation and validation**
  - **Mid Mai- July (~10 weeks):**
    - ATLAS LPNHE (O.5 PB) → done
    - ATLAS ex-LAL (0.6 PB) → almost done
    - ALICE IPNO (1 PB) → Waiting, Still we have some issue with Alice and alicetokenacc xrootd auth/z plug-in
- **Phase 2 Massive migration**
  - **September - October (~6-8 weeks)**:
    - ATLAS Irfu (1.9 PB) + CMS LLR (1.8 PB)
  - **December - January (~6-8 weeks)**: CMS Irfu (1.9 PB)
- **Phase 3 Massive migration**
  - **March - June: LHCb** + Non LHC VOS

Acknowledgements

Many thanks to EOS developers team and LHC VOs technical representatives for the discussions

and for the recommendations

Many thanks for yours attention

Questions and Comments ?

BACKUP slides

Migration roadmap ~1 year

- **Remark**: no attempt to balance the data between sites during the migration
- **May 15:** final setup of the EOS instance, configuration of SAM tests for the 4 LHC VOs → Done
- **May 23 - June 19 (5 weeks):** ATLAS LPNHE (O.7 PB) → Done
  - Hope to migrate 150 TB/week (2 Gb/s); to be validated
  - Need to start draining DPM probably before the end of the migration, after file deletion
- **Mai 30 - July 31 (7 weeks)**: ATLAS LAL (0.7 PB)
- Requires new HW to be delivered and installed

  June 13 - July 31 (7 weeks): ALICE IPNO (1 PB)
  - Requires new HW to be delivered and installed
  - Overlap VO migrations. DPM drain and EOS space addition in parallel with migration
  - Also requires new HW to be delivered and installed
- **August: according to VO and local site availabilities:** not clear if we'll do more that completing the previous steps

Risks and mitigations

- **Setup of the EOS instance**
  - Time consuming task to test new eos releases
  - Important to avoid delays: summer period will be less favorable to start the the initial work with VO
  - Subject to new HW delivery delays leading to insufficient temporary space in EOS
- **DPM drain longer than expected,** introducing delays in the roadmap before we can start a new migration phase
  - Normally, should not be too long as we move forward as migrated files will be deleted from DPM
- **Underestimation of migration time**: clearly difficult to assess before we started
  - We think our current numbers are very conservative: 2x less that what we observe in real production… but production will continue in
  - Non LHC VOs migration may take time as not much contact
  - They also have no real tools to do it

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

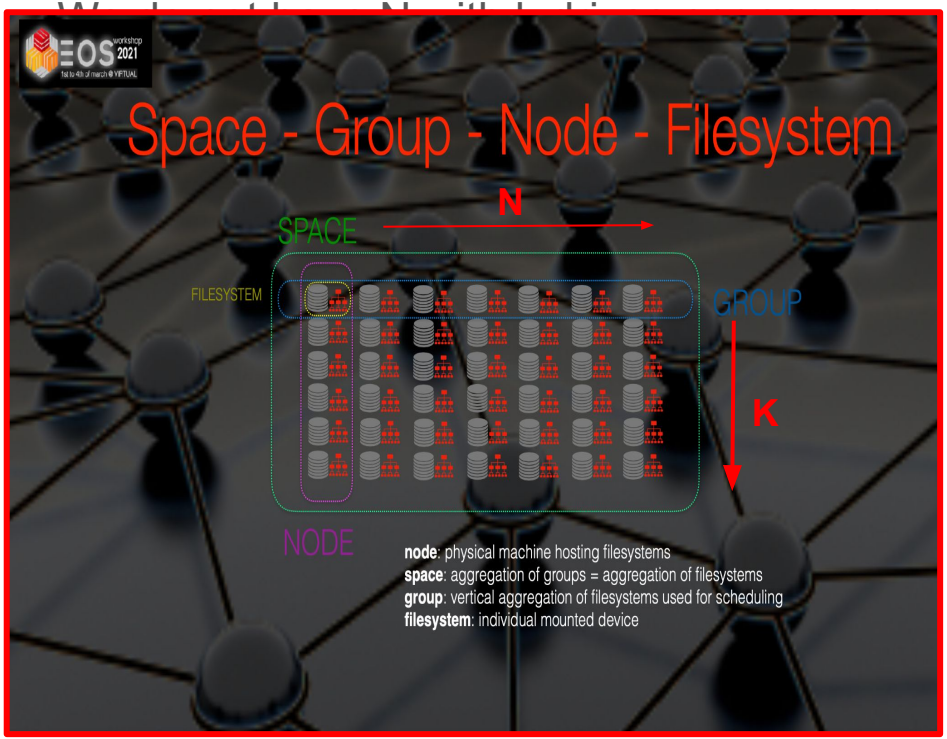

- On Ideal case we have:
- N servers with **K** individual FS on each server (of the same size)
- Thus we have **K** groups with N filesystem on each group (from N different servers)
- Easy to add a new server of same size (of K individual FS)

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

A non-uniform example of EOS FS Organization

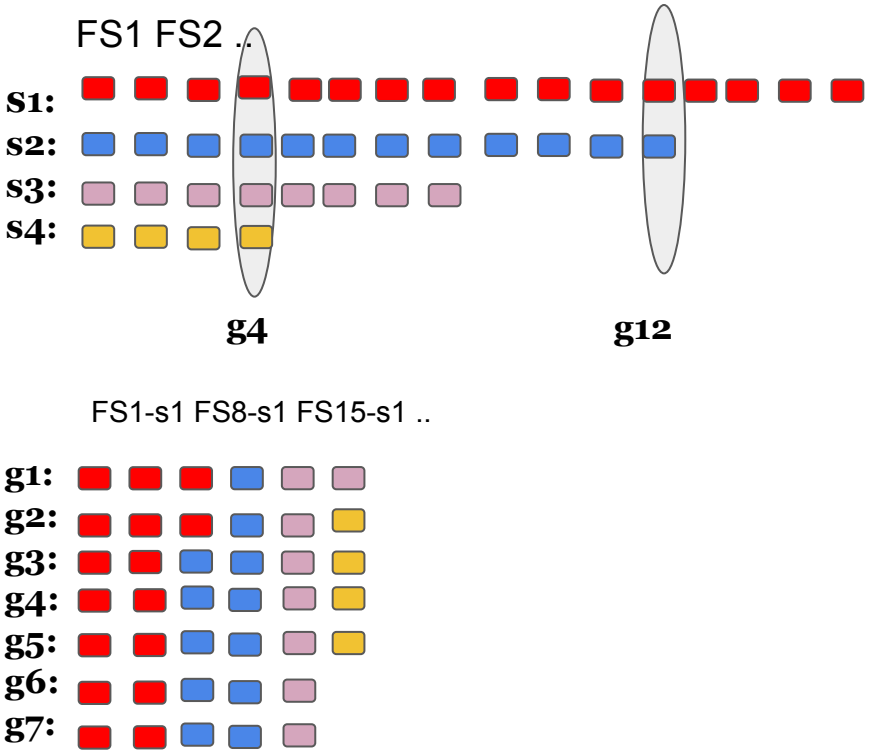

- Let's imagine 4 servers with 16,12, 8, and 4 FS of the same size
- The original organization of FS can not be deployed as we are going to have a group with a non-uniform number of FS
- in total, We have 40 groups
- *k=int(sqrt(40))+1 = 7 (a rule of thumb)*
- Sort the server by the # of filesystems
- Take the server with the largest number of FS and fill cyclically the group table
- And continue to the next one
- At the end, we have a matrix **of k group x k fs** which looks more uniform than the initial one
- We have as much as the minimum # of FS from the same server for each group
- We expect that with a larger number of server/fs this will converge better (more uniform groups)
- This procedure is easy to deploy when we add a new FST
- This procedure is not unique

Configuration details

- EOS 5.0.x
  - Mixing nodes with Centos 7 and Centos 8 flavors
- Identical gridmap file along the sites
- Identical pool unix accounts for the VOs
  - Logically we need 2-3 accounts (depending on VO internal DN/proxies usage)
  - VOs, which give access to each user can drive to a large gridmapfile
  - We are not sure if we need the VOMS extension matching or not (?)
  - **e.g. http.secxtractor /opt/eos/xrootd/lib64/libXrdVoms.so -vomsfunparms:certfmt=pem\|vos=atlas,dteam\|grps=/atlas,/dteam,/dteam/france\|grpopt=10\|dbg**
  - **Plus the vid mapping: DN/voms role→User**
- Usage of native http(s) xrootd interface only on specific ports
  - Do not use microhttpd interface - under decommission
  - EOS_MGM_HTTP_PORT=9000 and EOS_FST_HTTP_PORT=9001
- Looking forward for the redirection from Slave to Master MGM (for xroot and http(s))

EOS@MGM

- sec.protparm gsi -vomsfun:/opt/eos/xrootd/lib64/libXrdSecgsiVOMS.so -vomsfunparms:**certfmt=pem|vos=atlas,dteam|grps=/atlas,/dteam,/dteam/france|grpopt=10|dbg**
- sec.protocol gsi -crl:3 -cert:/etc/grid-security/daemon/hostcert.pem -key:/etc/grid-security/daemon/hostkey.pem -gridmap:/etc/grid-security/grid-mapfile -d:4 **-gmapopt:11** -vomsat:1 -moninfo:1 -gmapto:1

...

- http.cadir /etc/grid-security/certificates/
- http.cert /etc/grid-security/daemon/hostcert.pem
- http.key /etc/grid-security/daemon/hostkey.pem
- http.gridmap /etc/grid-security/grid-mapfile
- http.secxtractor /opt/eos/xrootd/lib64/libXrdVoms.so -vomsfunparms:certfmt=pem|**vos=atlas,dteam|grps=/atlas,/dteam,/dteam/france**|grpopt=10|dbg
- http.trace all
- http.exthandler xrdtpc /opt/eos/xrootd/lib64/libXrdHttpTPC.so
- http.exthandler EosMgmHttp /usr/lib64/libEosMgmHttp.so **eos::mgm::http::redirect-to-https=1**

…

- mgmofs.cfgtype quarkdb
- mgmofs.nslib /usr/lib64/libEosNsQuarkdb.so
- Mgmofs.qdbpassword mystrongsecret