



# EOS installation notes

Miguel Martinez Pedreira



# First notes

- There are several ways to install it, mainly done for **SLC5/6**
  - eos-deploy
    - This script asks for some information in a wizard, and then does all the work for us
      - Firewall
      - Self-signed host certs if not existing
      - Download LCG CAs, configures EOS repo
      - Masks xrootd (and other) from EPEL and download RPMs from YUM repository
      - Creates configuration for the different components
        - And register/map the filesystems under `/var/eos/fs/<n>...`
      - Takes care of ALICE specifics
        - Token plugin
      - Run the services
      - Standard space/fs configuration settings (for repair, balance...)
      - Creates the ALICE home directory linked to aliproduct account
    - Manual RPM installation and configurations
      - Not very sure if the docs contain 'all' the necessary information...

# + Firewall

- eos-deploy open the ports for the installation/configuration
  - But admin needs to make that permanent

1094: XRootD MGM port (only on MGMs)

1095: XRootD FST port (only on FSTs)

1096: XRootD SYNC port (only on MGMs)

1097: XRootD MQ port (only on MGMs)

443: https X509 port (only on HTTPS gateways or MGM)

8443: https KRB5 port (only on HTTPS gateways or MGM)

8000: http port (only on MGMs)

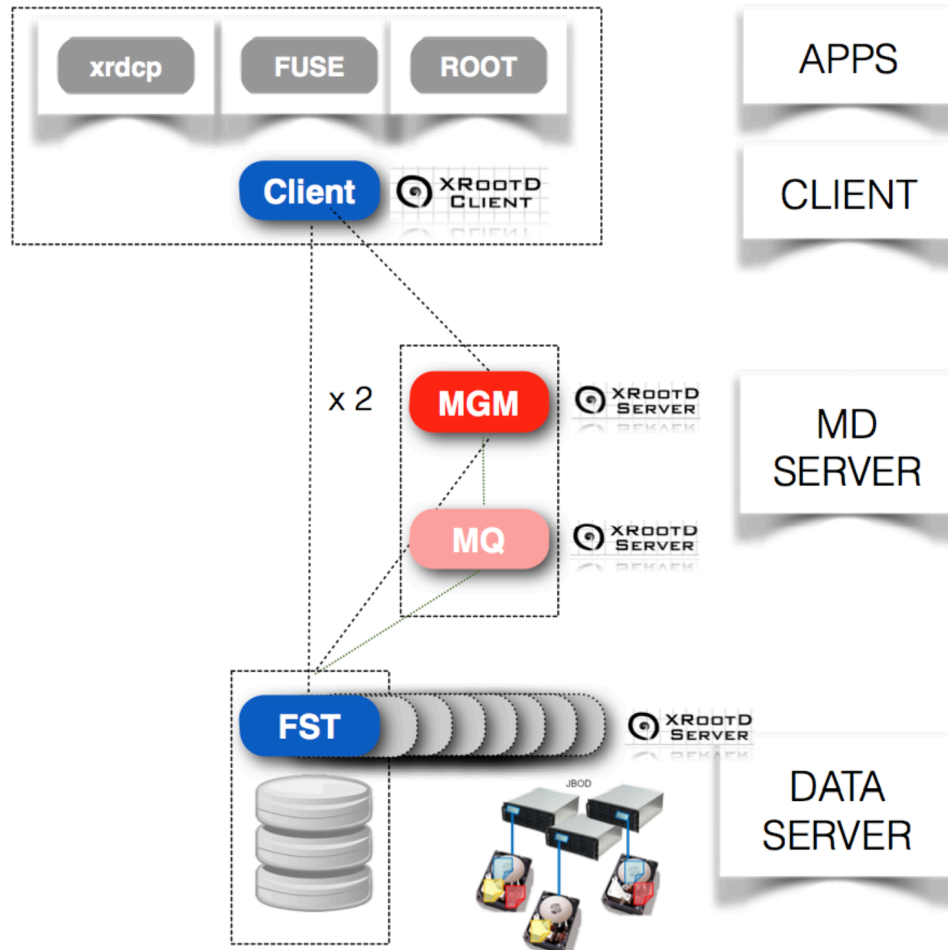
8001: http port (only on FSTs)



# More notes

- SSH access
  - To use the script, ran from the manager, we need ssh access to the FSTs
    - It would be interesting to make it password-less
  - In my case: ssh keys + ssh-agent
    - Asks for the pw once and keeps it active for the session
- When asks for regexp, is based on the `df -h` on the nodes
  - E.g. We have `/data<n>` as partitions, we type 'data'
- Number of filesystems
  - As many as partitions/disks, `<n>` in the previous point
- Careful with the repo files...

# + EOS architecture





# + LBL case

- 5 servers with 60 disks each
- Option 1: OS based RAID6 as usual
- Option 2: RAIN6 using groups of 'two rows', so 10 disks per group, setting the number of stripes to 10. Equivalent to option 1.
- Option 3: archive layout with 20 groups (3 disks per server in the same group) and 15 stripes
  - Adds more redundancy: if one machine pops out, they can still read & write data
  - If you use the raid6 layout you always need all 5 machines available to be able to write
  - The downside of this configuration is, that if a disk breaks it will probably take more than a week to recreate the lost disk, but since you can lose now three at the same time, this is probably a minor issue (Andreas)

name	redundancy	algorithm	description
raidp	4+2	dual parity raid	can loose 2 disks without dataloss
raid6	N+2	Erasure Code (Jerasure library)	can loose 2 disks without dataloss
archive	N+3	Erasure Code (Jerasure library)	can loose 3 disks without dataloss

# + Final notes and personal doubts

- Possibility to use puppet
  - Done in Subatech, but their configuration is not reusable by others
    - I have some of the manifest in case they have some EOS info
      - Plan to do a common knowledge base and wiki to help admins install easily
- Xrootd 4 + IPv6
  - EOS comes with xrootd 3.x (different xrootd versions for different EOS releases)
  - They plan to release with xrootd 4 in next release
  - Easy to upgrade? Apparently yes (will follow on this since will be done soon by Subatech)
- Xrootd to EOS and roll-back EOS
  - Andreas suggest emptying the existing storage, start EOS from scratch and then populate it with transfers
  - And the opposite...? I guess always better starting from clean installations...
- Don't replicate too much: RAID + eos replication
- After installation
  - There are commands to ask disks, fsts, etc and register them in the manager
  - Also able to add a second manager, switch them as master/slave





# Test server

- I have used eos-deploy to set up an EOS instance
- Dual MGM (master-slave) and 1 FSTs
- Using SLC6 CERN with standard machines
  - 8GB RAM, 2-4 cores
- We can use it for whatever purpose
  - Testing
  - Explore commands
  - Apply layouts
  - Replication options
- Is not trivial to me finding the right commands...
- **Andreas offered to give 'live' help for us while doing this 😊 !**