



**EOS installation and maintenance from site
admin perspective
Adrian Sevcenco, ISS, RO**

Planning

- <https://eos-docs.web.cern.ch/>
- <https://eos-community.web.cern.ch/>
- <https://gitlab.cern.ch/dss/eos>
- <https://gitlab.cern.ch/dss/eos/-/tree/master/doc/configuration>

Plan the installation:

- Decide on the needs and scale
- Ask “how to begin” advices on forum
 - EOS is distributed FS !! → Metadata is separate from data
- Metadata management planning (MGM)-> Plan it carefully
 - QuarkDB - <https://eos-docs.web.cern.ch/configuration/quarkdb.html>
 - number of machines
 - configuration location and maybe service customizations
 - password line vs password file in configuration
 - **N.B.!!! Upgrading EOS4->EOS5 assumes that MGM have it's own configuration in QuarkDB !! see the docs**

- MGM <https://eos-docs.web.cern.ch/quickstart/admin/configure.html#setup-mgm>
 - Decide the level of redundancy
 - N.B. on configuration and upgrade:
 - EOS_GEOTAG needs 2 strings splitted by “::”

METADATA IS THE MOST IMPORTANT PART OF DATA!!! PLAN FOR DISASTERS!

- **Data storage planning**

- Data sits on File Storage (FST) type node
- Decide the economics of robustness vs IO capacity
 - single or multiple replica
 - single physical block device or mdraid one
 - Compute the potential bottlenecks:
 - take into account the data path from HDD interface through switch(es) to the worker node for both sequential and random IO
 - ALICE have heavy IO profile
 - streaming data
 - any bulk access becomes random when there are $O(100-1000)^3$

Configuration

- I skip the installation part:
 - straight forward
 - small set of packages
 - but ask on the forum about the versions!
 - for EOS5 the rpms are located in “diopside/tag/testing”
 - N.B.!!! make backup of “/etc/grid-security/xrootd/TkAuthz.Authorization”
 - some version of alicetokenacc package (there could be a new one by now) overwrote the existing TkAuthz.Authorization
- Configuration
 - It has 3 components
 - /etc/sysconfig/eos_env
 - configuration of quarkDB service(s): /etc/xrd.cf.quarkdb
 - for MGM: /etc/xrd.cf.mq /etc/xrd.cf.mgm
 - for FST: /etc/xrd.cf.fst



Configuration cont.



- Assuming the services were started successfully → EOS configuration
 - N.B.!! The following is not (yet) vetted with the EOS experts
 - actually a part of the following `_is_` but i do not know which part

```
eos -b vid enable sss  
eos -b vid enable https
```

```
# adds a host as a (fuse) gateway with 'su' priviledges  
eos -b vid add gateway mgm
```

```
# Create local uid/gid corresponding to aliproduct:alice definition from cern ldap  
groupadd -g 1395 alice  
useradd -m -g alice -u 10367 aliproduct
```



Configuration cont.



```
# create main storage path structures and map them to eos location
```

```
ALICE_GRID_PATH="/eos/alice/grid"
```

```
eos -b mkdir ${ALICE_GRID_PATH}
```

```
for name in $(seq -w 00 15); do
```

```
echo -e "mkdir ${ALICE_GRID_PATH}/${name}\nmap link /${name}/
```

```
${ALICE_GRID_PATH}/${name}/" | eos -b
```

```
done
```

```
echo -e "chown -r aliproduct:alice ${ALICE_GRID_PATH}\nchmod 777
```

```
${ALICE_GRID_PATH}" | eos -b
```

```
# everybody authenticated via UNIX will be mapped to the ALICE user
```

```
10367(aliproduct) and the corresponding group
```

```
# Andreas dixit :) ←
```

So, i do have some
proofs of validity :D

```
vid set map -unix <pwd> void:10367 vgid:1395
```

```
N.B!! to be done from EOS cli! the initial below one is not escaped properly
```

```
>>> eos -b 'vid set map -unix \\<pwd\\> void:10367 vgid:1395' <<<
```

```
eos -b vid set membership 2 +sudo
```

Configuration cont.

- the vid subject can be pointy

eos vid ls

if tident is shown then remove it with:

vid rm tident:"*@mgm":gid

vid rm tident:"*@mgm":uid

Currently, working on EOS5 i have:

https:"<pwd>":gid => root

https:"<pwd>":uid => root

publicaccesslevel: => 1024

sss:"<pwd>":gid => root

sss:"<pwd>":uid => root

sudoer => uids(daemon)

unix:"<pwd>":gid => alice

unix:"<pwd>":uid => aliproduct

Configuration cont.

- **Config the space (named “default” in my case)**

```
eos -b space config default space.scaninterval=1814400  
eos -b space config default space.autorepair=on  
eos -b space config default space.graceperiod=3600  
eos -b space config default space.drainperiod=86400  
eos -b space config default space.headroom=5.1G
```

```
eos -b space config default fs.scaninterval=1814400  
eos -b space config default fs.graceperiod=3600  
eos -b space config default fs.drainperiod=86400  
eos -b space config default fs.headroom=5.1G
```




Configuration cont.



- Enable converter and group balancer to equalize space between groups

<https://eos-docs.web.cern.ch/configuration/converter.html>

<https://eos-docs.web.cern.ch/configuration/groupbalancer.html>

eos space config default space.converter=on

eos space config default space.converter.ntx=32

eos space config default space.converter.ntx=128

eos space config default space.groupbalancer=on

eos space config default space.groupbalancer.ntx=16

eos space config default space.groupbalancer.engine=std

Configuration cont.

- FST node preparation

```
eos -b node config ${FST}:1095 gw.rate=2000
```

```
eos -b node config ${FST}:1095 gw.ntx=100
```

```
eos -b node txgw ${FST}:1095 on
```

```
eos -b node set ${FST} on
```

- Map a node FS to the FST

```
eos fs add -m ${FSID} $(uuidgen -r) ${FST}:1095 ${MNT_DIR} "${SPACE}.${FSID}"  
off
```

```
eos group set "${SPACE}.${FSID}" on
```

Now inspect the output of eos fs ls and enable the fs with

```
eos fs config ${FSID} configstatus=rw
```

- **eos-config-inspect**
 - **dump (legacy file-based config) and import to QDB an MGM configuration**
 - **relocate-filesystem: change the FST to which a filesystem belongs to**
- **eos ns stat**
 - **NS statistics**
- **eos io stat -x**
 - **Transfers table per application → check groupbalancer**
- **eos group ls --io | eos fs ls --io**
 - **usage metrics per group or per file system**
- **eos io enable -r -p**
 - **enable collection of io statistics**
- **eos io ns [-b] -100**
 - **print hot directories (unfortunately no hot files)**
- **eos file replicate [<path>|fid:<fid-dec>|fxid:<fid-hex>] <fsid1> <fsid2>**
 - **replicate file <path> part on <fsid1> to <fsid2>**



Misc info



- The **tokenauthz** and **xrootd-alicetokenacc** packages were merged under the new project **alicetokenacc**
 - A very important new feature!! → **alicetokenacc.multiprocess N**
 - Threaded processing of envelopes (capped to 128)

- **XRootD plain installation**
 - detailed documentation: https://jalien.docs.cern.ch/site/xrootd_rpms/



Support Info



Adrian.Sevcenco@cern.ch
@asevcenc on Mattermost

THANK YOU !