



# **EOS Operations: 101 Admins guide Installation and Operations**

**Luca Mascetti  
CERN - IT Storage**

**[luca.mascetti@cern.ch](mailto:luca.mascetti@cern.ch)**





EOS

**EOS Open Storage**



# ABOUT EOS

## Elastic, Adaptable and Scalable

EOS is a simple and scalable open source software solution for central data recording, user analysis and data processing.

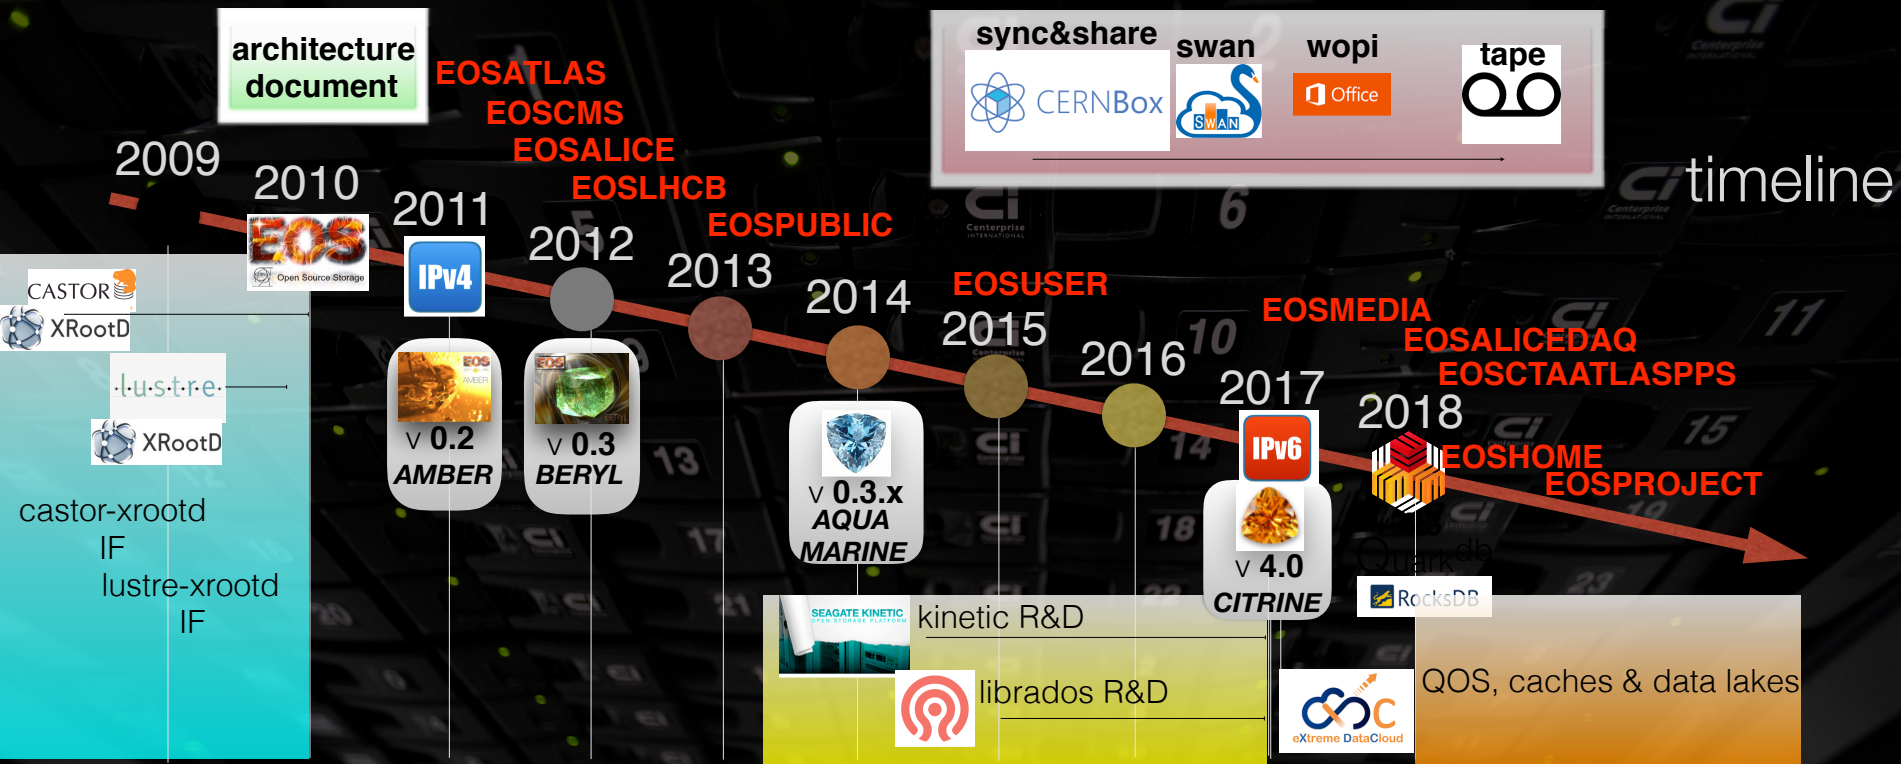
EOS supports thousands of clients with random remote I/O patterns with multiprotocol support and tunable QoS.

**HTTP, WebDAV, CIFS, FUSE, XRoot, gsiFTP**


EOS offers a variety of authentication methods and user/project quotas.

**KRB5, X509, Shared Secret and unix**



# EOS Project History



timeline



# EOS Architecture




Client

MGM

MQ

```

eos-check-bl
build.ma
checksum
  Adle
  Chec
  crc3
  crc3
cmake_cl
CXX.incl
DependIn
depend.i
depend.m
    
```



namespace

FST



data



## EOS Production Releases

**phasing out**

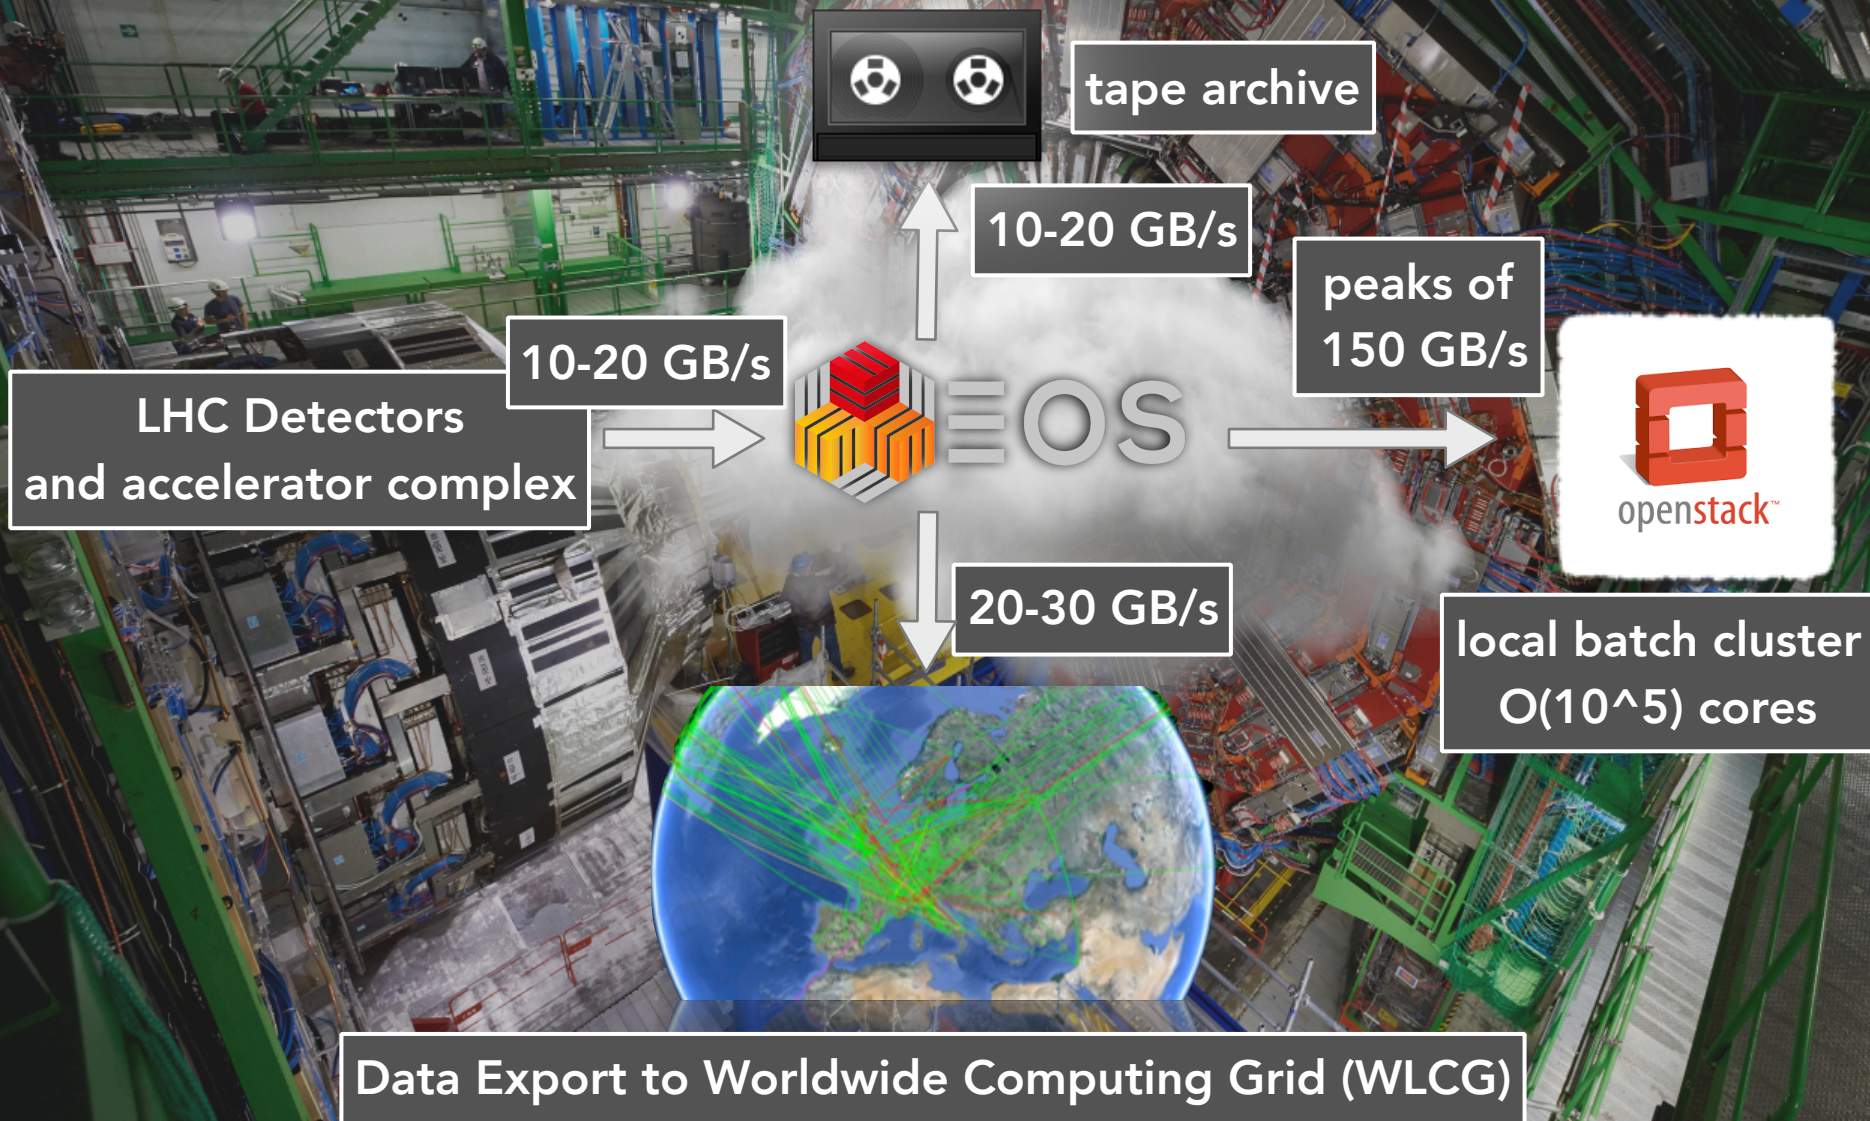
**Aquamarine**  
**V 0.3.X**

**Citrine**  
**V 4.X**

<p><b>IPV4</b></p> <p>namespace in-memory data on attached disks</p>	<p><b>IPV6</b></p> <p>plugins for meta data &amp; data persistency</p>
<p>XRootD V3</p>	<p>XRootD V4</p>

# How is it used?

CERN's mainstream usecase





# EOS Production instances @ CERN



Total Space  
**340 PB**

Files Stored  
**4.92 Bil**





# EOS Production instances @ CERN

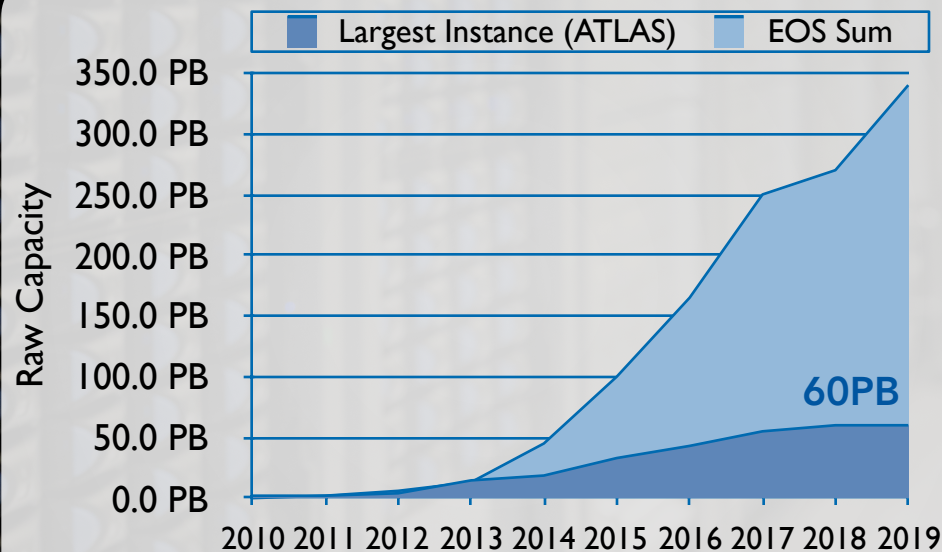


## EOS instances:

- 5 for the LHC experiments
- 9 CERNBox (EOSUSER + 5 EOSHOME + 3 EOSPROJECT)
- EOSMEDIA (photo/audio/video)
- EOSPUBLIC (Open Data and non-LHC experiments)
- EOSBACKUP (backup for CERNBox)
- 5 for various tests


~1500 storage nodes

~60k disks



# Hardware evolution

- Profiting from economy of scale
  - minimise price per GB
- System Unit:
  - 8 physical cores (16 virtual) 64-128GB RAM
  - disk-tray of 24x 4-6-10-12TB HDDs



- Running different generations
  - 2 trays per system unit
  - 4 trays per system unit
  - 8 trays per system unit





EOS

Architecture

# In-memory namespace

Client

Client

Client

## Management Server

Pluggable Namespace, Quota  
Strong Authentication  
Capability Engine  
File Placement  
File Location

## Message Queue

Service State Messages  
File Transaction Reports  
Shared Objects (queue+hash)

## File Storage

File & File Meta Data Store  
Capability Authorization  
Check-summing & Verification  
Disk Error Detection (Scrubbing)

MGM

MQ

sync

namespace

MGM

MQ

sync

namespace

FST

FST

FST



data




data



data

```
eos-check-b1
  build.ma
  checksum
    Ad1e
    Chec
    crc3
    crc3
  cmake_c1
  CXX_incl
  DependIn
  depend.i
  depend.m
  ...
```



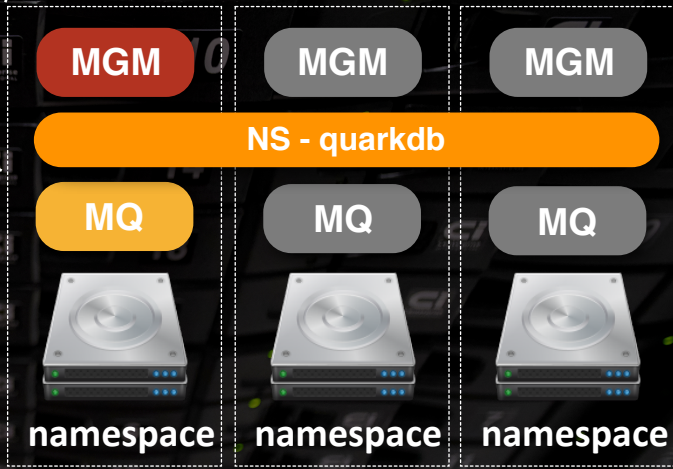
# QuarkDB namespace




**Management Server**  
Pluggable Namespace, Quota  
Strong Authentication  
Capability Engine  
File Placement  
File Location

**Message Queue**  
Service State Messages  
File Transaction Reports  
Shared Objects (queue+hash)

**File Storage**  
File & File Meta Data Store  
Capability Authorization  
Check-summing & Verification  
Disk Error Detection (Scrubbing)



```
eos-check-bl
  build.ma
  checksum
  Ad1e
  Chec
  crc3
  crc3
  cmake_cl
  CXX.incl
  DependIn
  depend.i
  depend.m
```



# QuarkDB namespace and multi-FST

Client

Client

Client

## Management Server

- Pluggable Namespace, Quota
- Strong Authentication
- Capability Engine
- File Placement
- File Location

## Message Queue

- Service State Messages
- File Transaction Reports
- Shared Objects (queue+hash)

## File Storage

- File & File Meta Data Store
- Capability Authorization
- Check-summing & Verification
- Disk Error Detection (Scrubbing)

MGM

MGM

MGM

NS - quarkdb

MQ

MQ


MQ




namespace

namespace


namespace



data




data

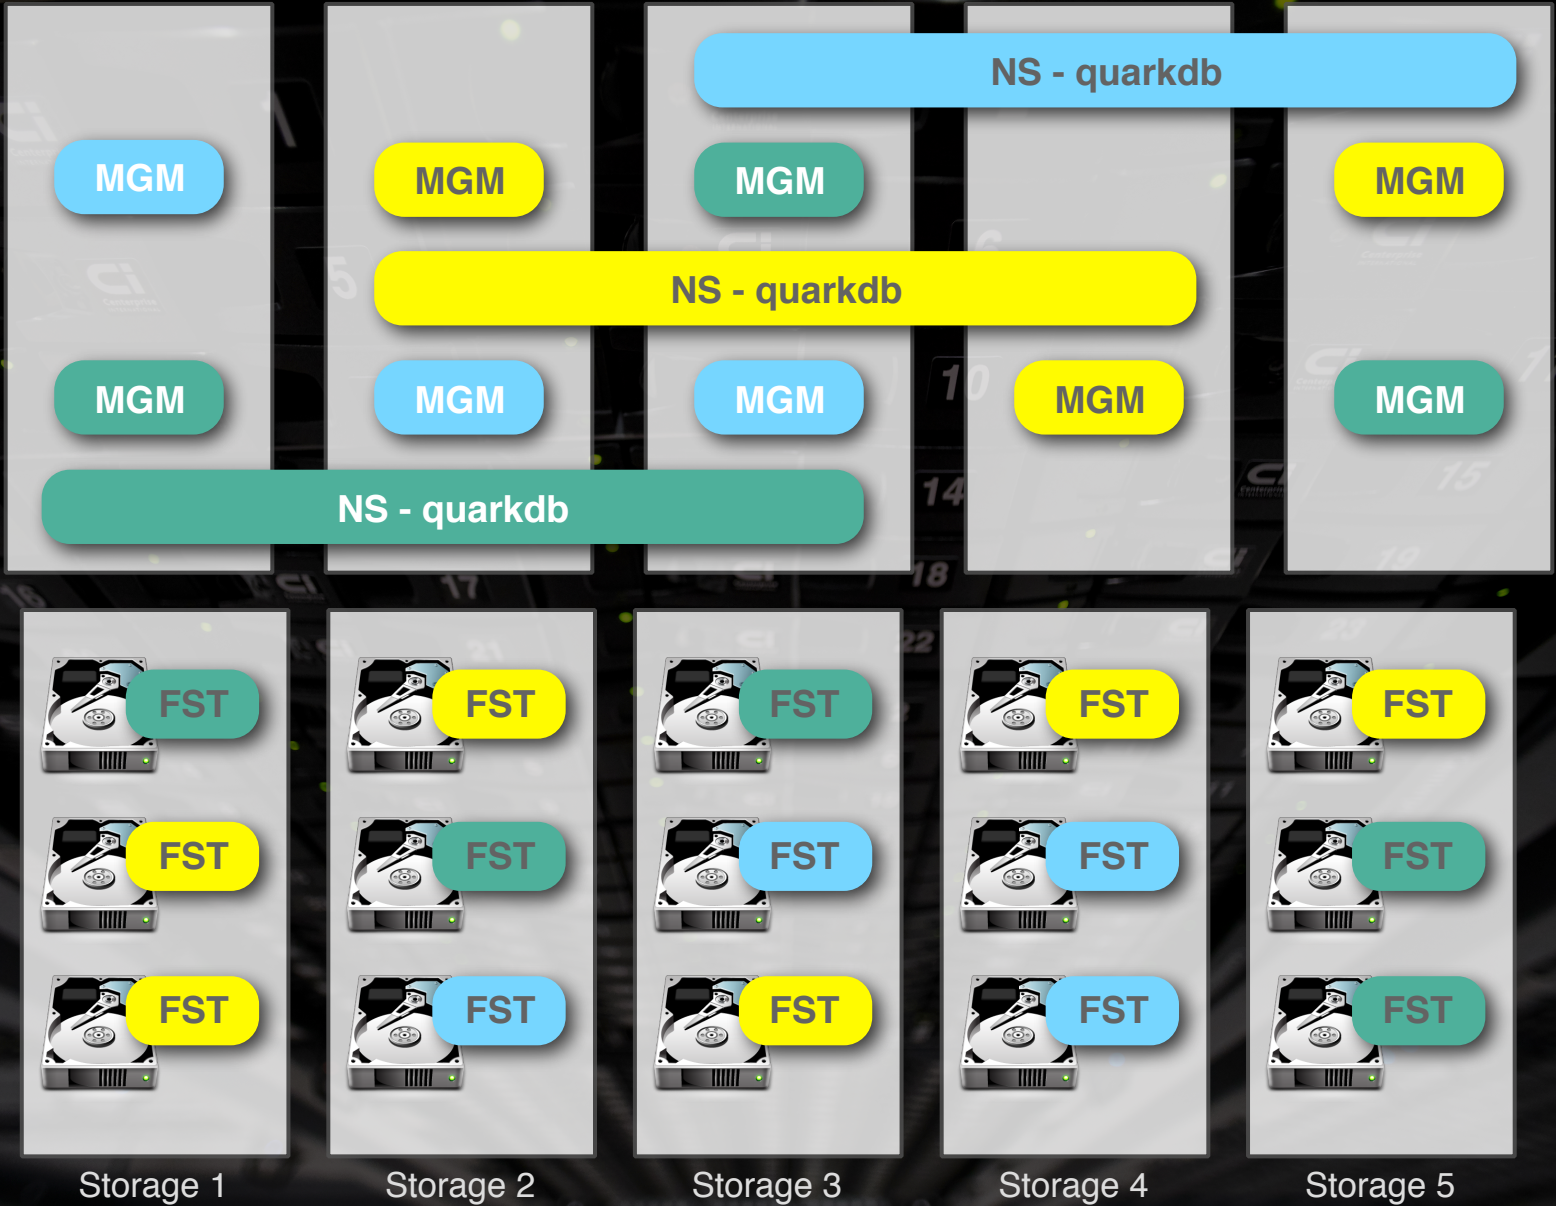


data

```
eos-check-b1
  build.ma
  checksum
  Ad1e
  Chec
  crc3
  crc3
  cmake_c1
  CXX.inc1
  DependIn
  depend.i
  depend.m
  23
```



# Hardware optimisation





EOS

**Cluster Creation**



# EOS Software and Dependencies

## Software Repositories

- **Production Repo**

- Citrine <https://storage-ci.web.cern.ch/storage-ci/eos/citrine/tag/>
- Citrine-Dependencies <https://storage-ci.web.cern.ch/storage-ci/eos/citrine-depend/>
- QuarkDB <https://storage-ci.web.cern.ch/storage-ci/quarkdb/tag/>

- **Pre-production Repo**

- Citrine-Testing <https://storage-ci.web.cern.ch/storage-ci/eos/citrine/tag/testing/>

- **Development Branch**

- Citrine-Dev <https://storage-ci.web.cern.ch/storage-ci/eos/citrine/commit/>

## Code

- **Gitlab**

<https://gitlab.cern.ch/dss/eos.git>

- **Github**

<https://github.com/cern-eos/eos>

# EOS Documentations and help

**EOS main website:**

<https://eos.web.cern.ch/>

**EOS Documentation:**

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

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

<https://quarkdb.web.cern.ch/quarkdb/docs/master/>

**EOS Community website:**

<https://eos-community.web.cern.ch/>

**OpenLab Collaboration Project with COMTRADE**

**EOS Whitepaper and EOS beginner setup:**

<https://cernbox.cern.ch/cernbox/desktop/index.php/s/Nbpi3hGZYYqHN93>

# Setup MGM

Location	Significance
/etc/xrd.cf.mgm	xrootd server configuration file
/etc/sysconfig/eos_env	instance configuration file
/var/eos/md/files.<hostname>.mdlog	changelog file cont. file meta data of the name space
/var/eos/md/directories.<hostname>.mdlog	changelog file cont. directory meta data of the name space
/var/eos/config/<hostname> to be backed up	directory cont. configuration files and configuration changelog
/etc/eos.keytab	keytab file for 'sss' authentication (shared secret)
/var/log/eos/mgm/xrdlog.mgm	eos/xrootd MGM log file
/var/eos/md/so.mgm.dump	Dump of the MGM shared object hash/queues

- How to generate a keytab
  - xrdsssadmin -k <eosinstance> -u daemon -g daemon add /etc/eos.keytab
- RPM needed
  - yum install xrootd-client-libs xrootd-client xrootd-debuginfo xrootd xrootd-libs xrootd-server-libs
  - yum install leveldb-1.10.0-0.el6.x86\_64 zeromq eos-server eos-client eos-debuginfo eos-xrootd libmicrohttpd-0.9.38-eos.wves libmicrohttpd-debuginfo-0.9.38-eos.wves

# Setup MQ and SYNC (in-memory)

Location	Significance
<code>/etc/xrd.cf.mq</code>	xrootd server configuration file
<code>/etc/sysconfig/eos_env</code>	instance configuration file
<code>/var/log/eos/mq/xrdlog.mq</code>	eos/xrootd MQ log file
<code>/var/log/eos/mq/proc/stats</code>	current broker statistics file

## Sync Daemon

keep in sync the ns files between headnodes

`/var/eos/md/files.<servername>.mdlog`

`/var/eos/md/directories.<servername>.mdlog`

keep in sync the configuration of the instance

`/var/eos/config/<servername>/*.eoscf`

# EOS headnode setup

## Headnode Daemons

```
systemctl start eos@master
```

```
systemctl start eos@*
```

## IPTABLES

```
-A INPUT -p tcp -m multiport --dports 1094 -m comment --comment "101 allow xroot to EOS-MGM" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 1097 -m comment --comment "101 allow xroot to EOS-MQ" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 1096 -m comment --comment "101 allow xroot to EOS-SYNC" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 443 -m comment --comment "102 allow GSI-HTTP to EOS-MGM" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 8000 -m comment --comment "102 allow HTTP to EOS-MGM" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 8443 -m comment --comment "102 allow KRB5-HTTP to EOS-MGM" -j ACCEPT  
  
-A INPUT -p tcp -m multiport --dports 7777 -m comment --comment "101 allow xroot to EOS-QDB" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 1100 -m comment --comment "102 allow HTTP to EOS-FUSEX" -j ACCEPT
```

# EOS headnode setup (in-memory ns)

```
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$ df
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/sda2       20642428    2845112   16748740   15% /
tmpfs           264585112         0   264585112    0% /dev/shm
/dev/sda1       1032088     137596    842064    15% /boot
/dev/sda3       20642428     178756   19415096    1% /tmp
/dev/sda4       2064208     129256    1830096    7% /usr/vice/cache
/dev/sda5      2839366856  411696920 2283438412   16% /var
/dev/sdb1       922459044  376630264  498970592   44% /var/eos/md
AFS              9000000         0    9000000    0% /afs

[root@eosuser-srv-b2 (mgm:master mq:master) ~]$
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$ free -m
              total          used          free      shared    buffers     cache
Mem:           516767      508532           8235           4         738      2300
-/+ buffers/cache: 277188      239579
Swap:            0             0             0

[root@eosuser-srv-b2 (mgm:master mq:master) ~]$
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$
```

RAID-1 HDD

RAID-1 SSD

Faster boot in case of daemon restart and the ns files are not in the server cache

RAM provision ~ 1KB per NS entry

# EOS headnode setup (quarkdb ns)

```
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$  
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$ df  
Filesystem      1K-blocks      Used Available Use% Mounted on  
/dev/sda2       20511312    4461720  14984632  23% /  
devtmpfs        131916536     0  131916536  0% /dev  
tmpfs           131927196     0  131927196  0% /dev/shm  
tmpfs           131927196    2312600  129614596  2% /run  
tmpfs           131927196     0  131927196  0% /sys/fs/cgroup  
/dev/sda1        999320     232692   697816  26% /boot  
/dev/sda3        20511312     47560  19398792  1% /tmp  
/dev/sda5        1998672      3140   1890676  1% /usr/vice/cache  
/dev/sda6       1801424464  916559724  793334244  54% /var  
/dev/sdb1       1874457348  53551720  1820905628  3% /var/lib/quarkdb  
tmpfs           26385440     0  26385440  0% /run/user/0  
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$  
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$ free -m  
Mem:      total      used      free      shared  buff/cache  
Swap:     0          0          0          2258    122918  
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$  
[root@eoshome-ns-i00-00 (mgm:master mq:master) ~]$
```

SSD

SSD

## RAM provision

NS cache, full NS stored in quarkdb  
on SSD /var/lib/quarkdb

# EOS headnode setup - QuarkDB NS

```
yum install quarkdb quarkdb-debuginfo
```

**Config file:** /etc/xrootd/xrootd-quarkdb.cfg

```
xrd.port 7777
xrd.protocol redis:7777 libXrdQuarkDB.so
redis.mode raft
redis.database /var/lib/quarkdb/eosns
redis.myself qdb-test-1.cern.ch:7777
redis.password_file /etc/eos.keytab
```

```
quarkdb-create --path /var/lib/quarkdb/eosns --clusterID your-cluster-id --
nodes qdb-test-1.cern.ch:7777,qdb-test-2.cern.ch:7777,qdb-
test-3.cern.ch:7777
```

```
systemctl start xrootd@quarkdb
```



# EOS storage node setup

Location	Significance
<code>/etc/xrd.cf.fst</code>	xrootd server configuration file
<code>/etc/sysconfig/eos_env</code>	instance configuration file
<code>/var/log/eos/fmd.&lt;unix-tst&gt;.&lt;fsid&gt;.mdlog</code>	FST meta data changelog file
<code>/var/log/eos/so.fst.dump</code>	Dump of the FST shared object hash/queues

FST Daemon

```
systemctl start eos@fst
```

IPTABLES

```
-A INPUT -p tcp -m multiport --dports 1095 -m comment --comment "101 allow xroot to EOS-FST" -j ACCEPT  
-A INPUT -p tcp -m multiport --dports 8001 -m comment --comment "102 allow HTTP to EOS-FST" -j ACCEPT
```

Register the FST back to the EOS instance

from the MGM

```
eos fs add `uuidgen` <host.domain>:1095 <disk> <pool>
```

from the FST

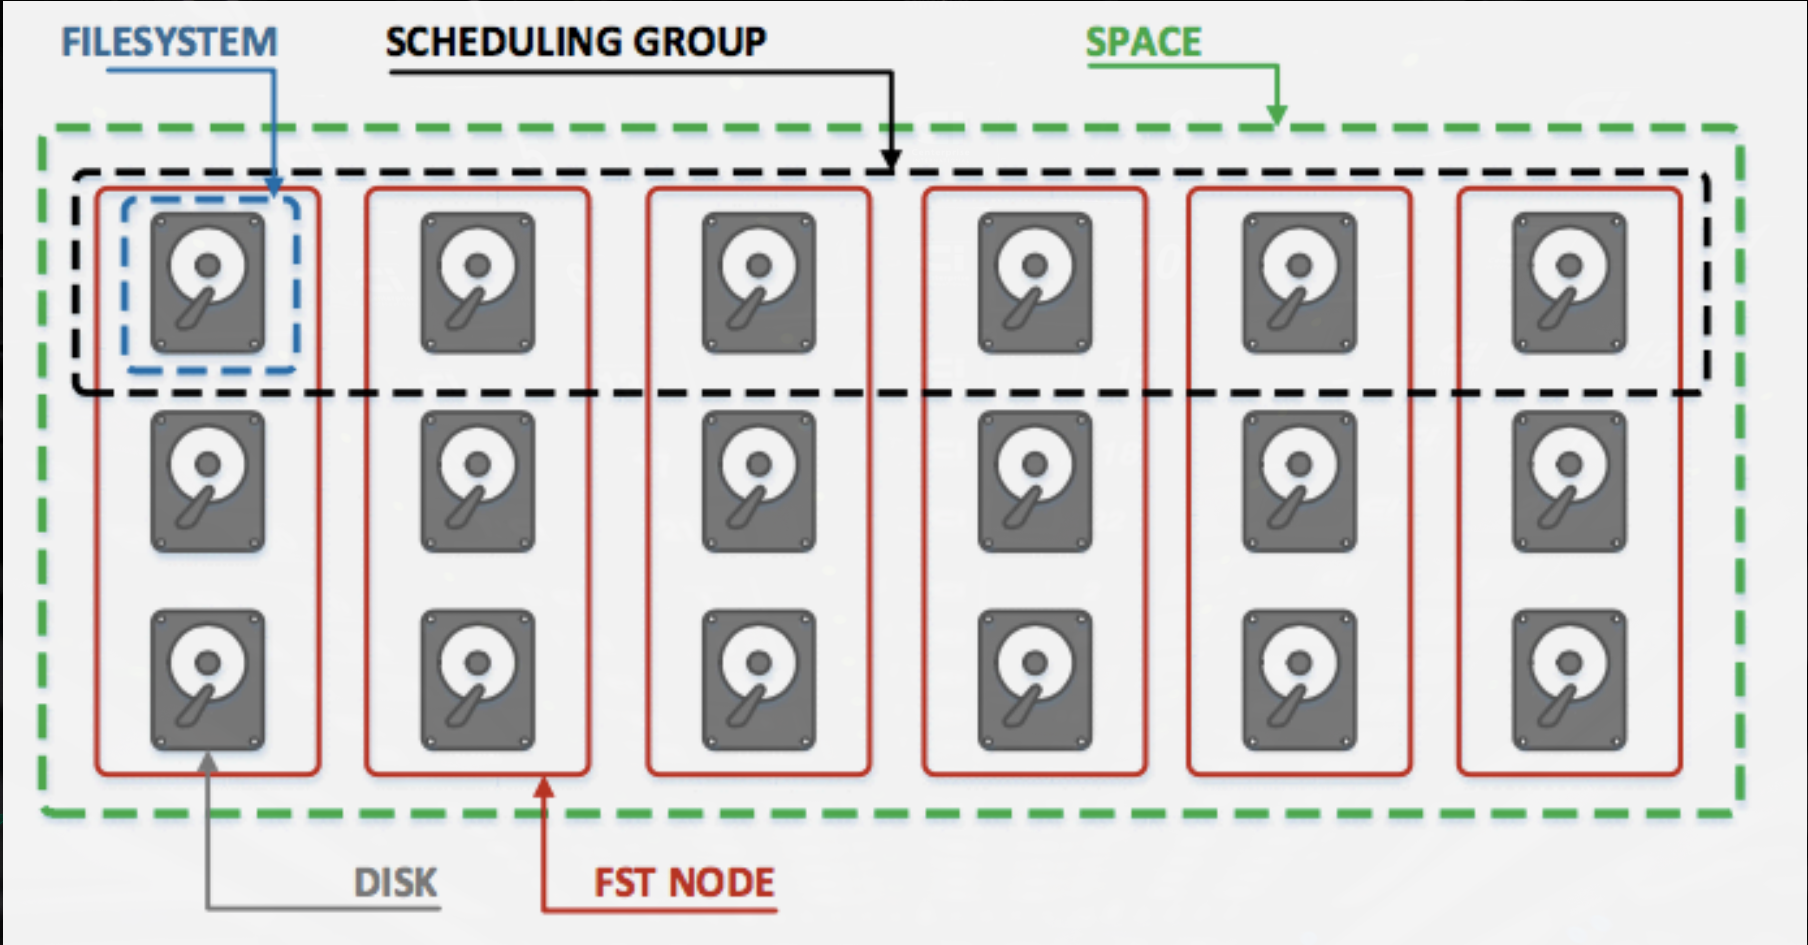
```
eosfstregister <disk> <pool>
```



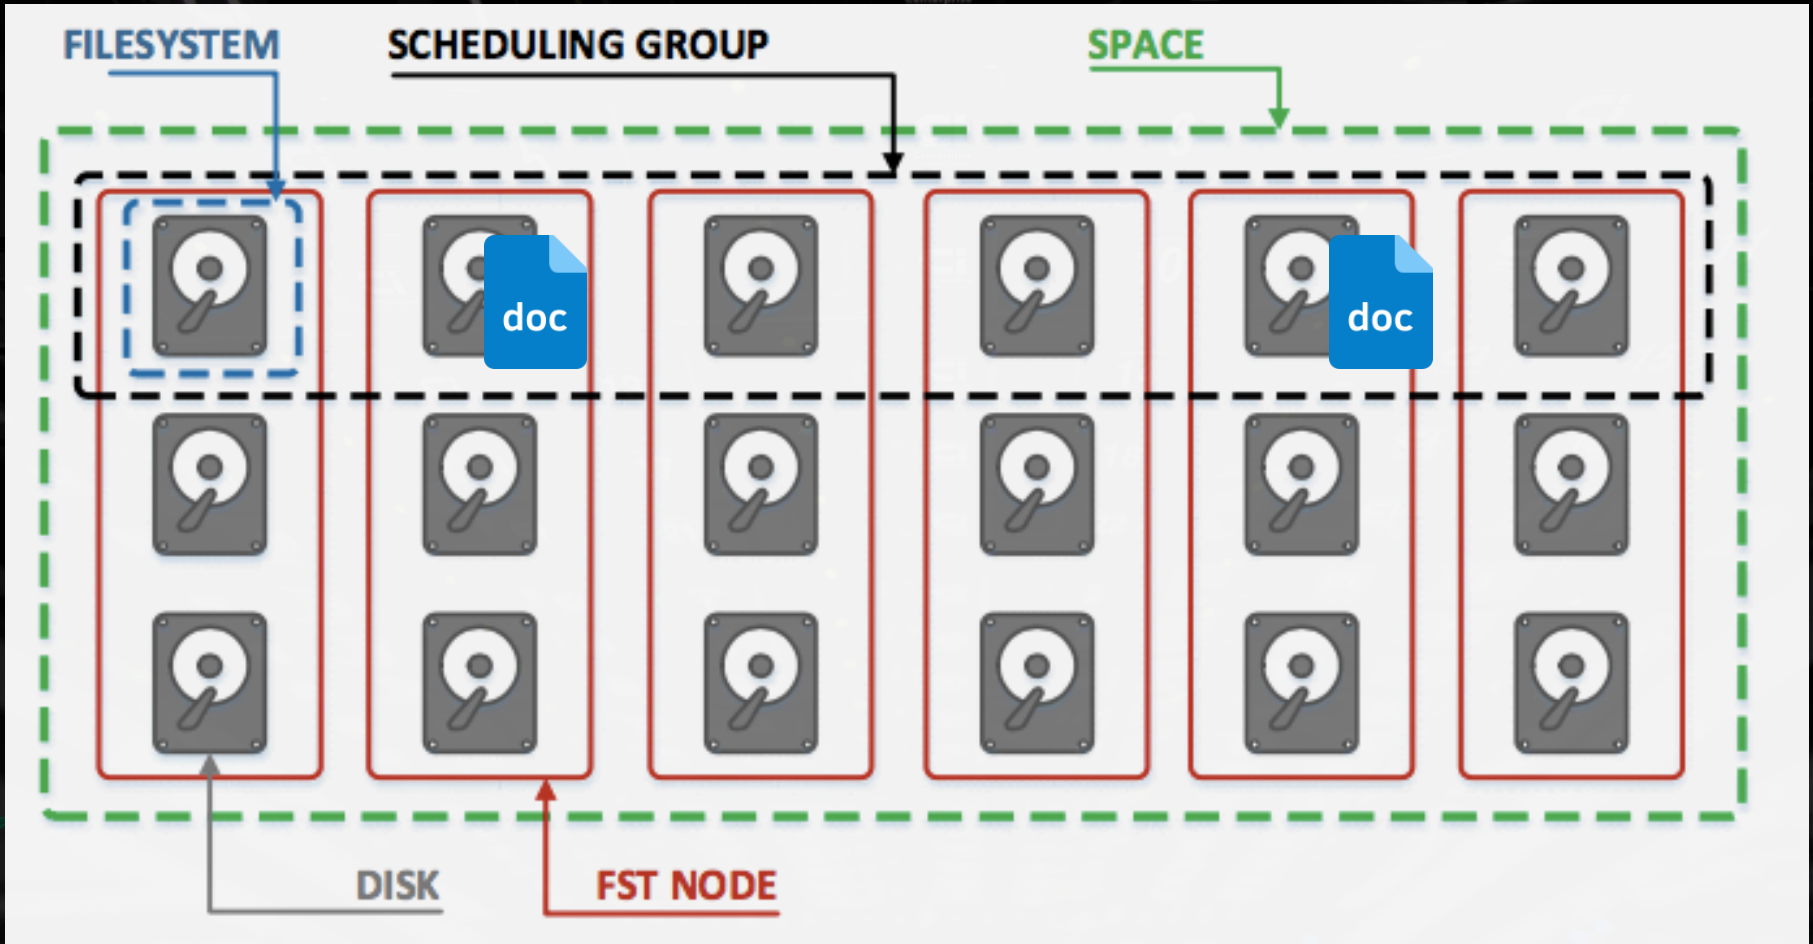
EOS

**Functionality**

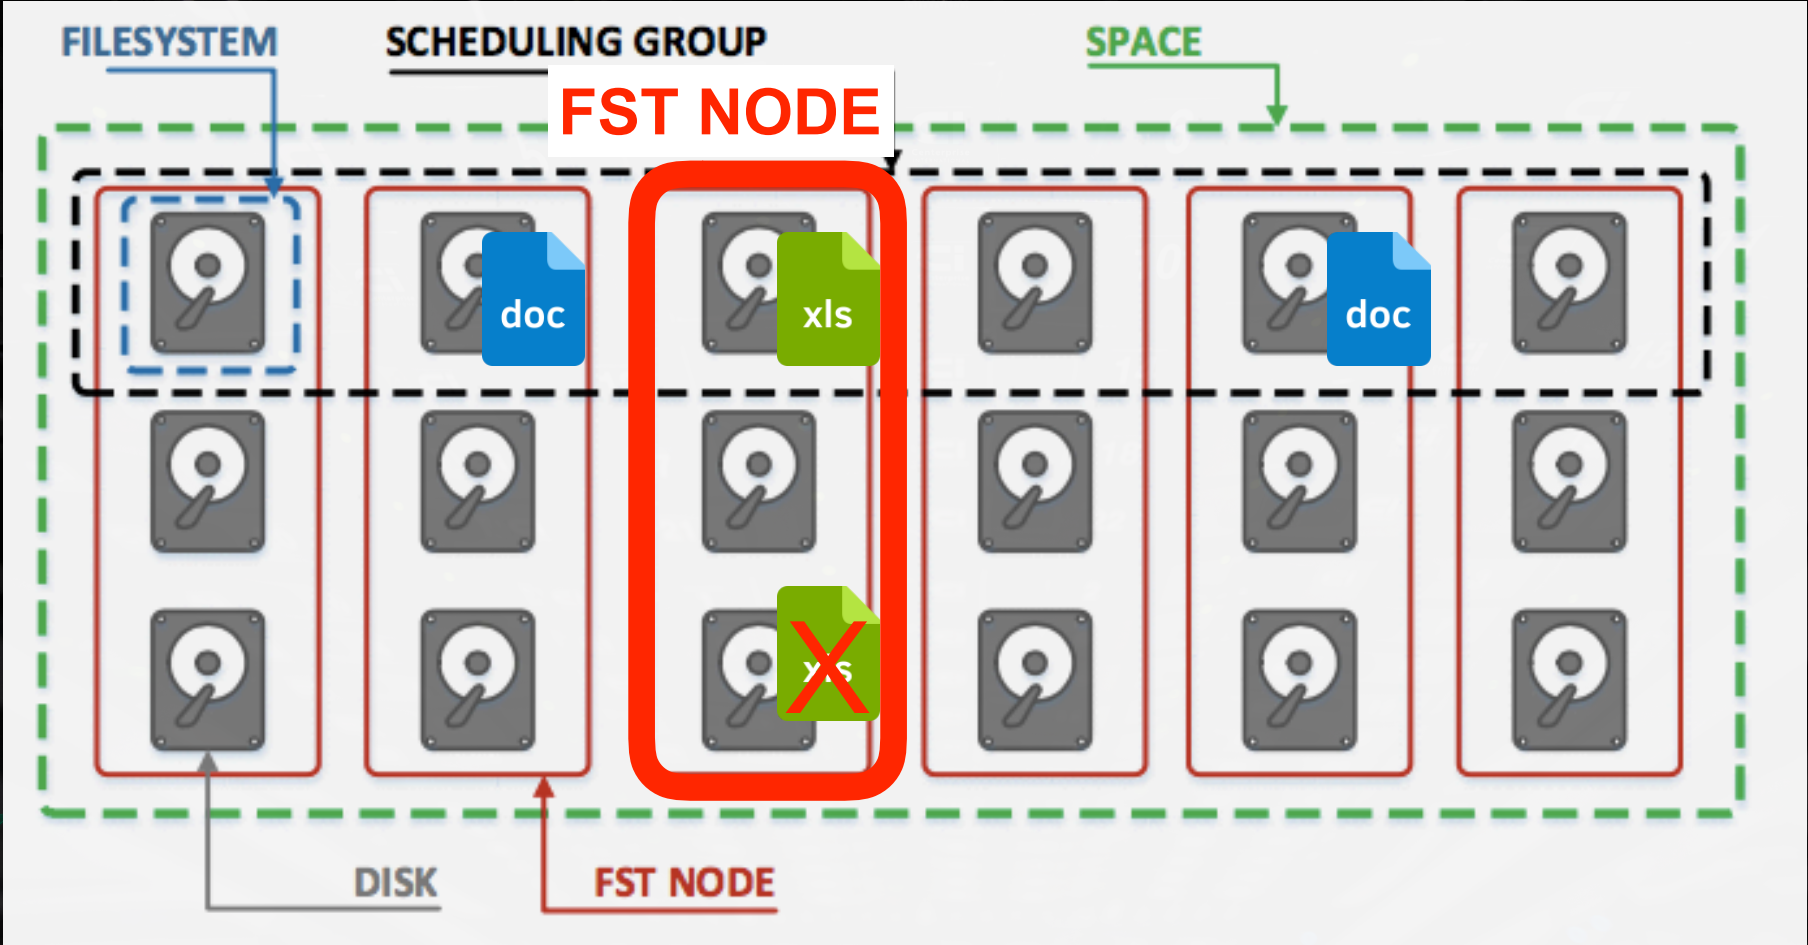
# Data Placement



# Data Placement

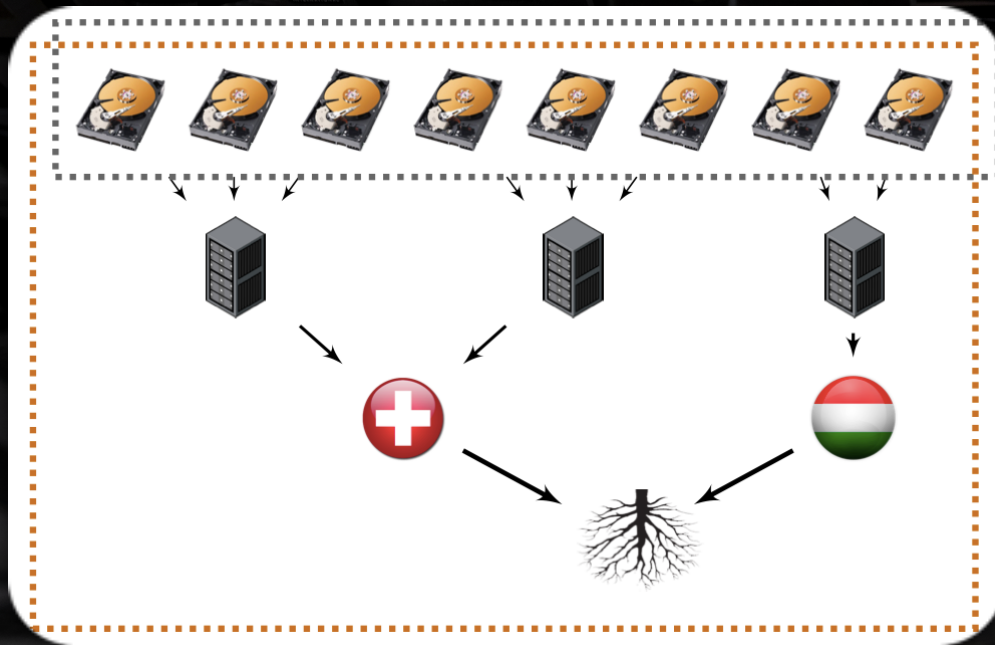


# Data Placement



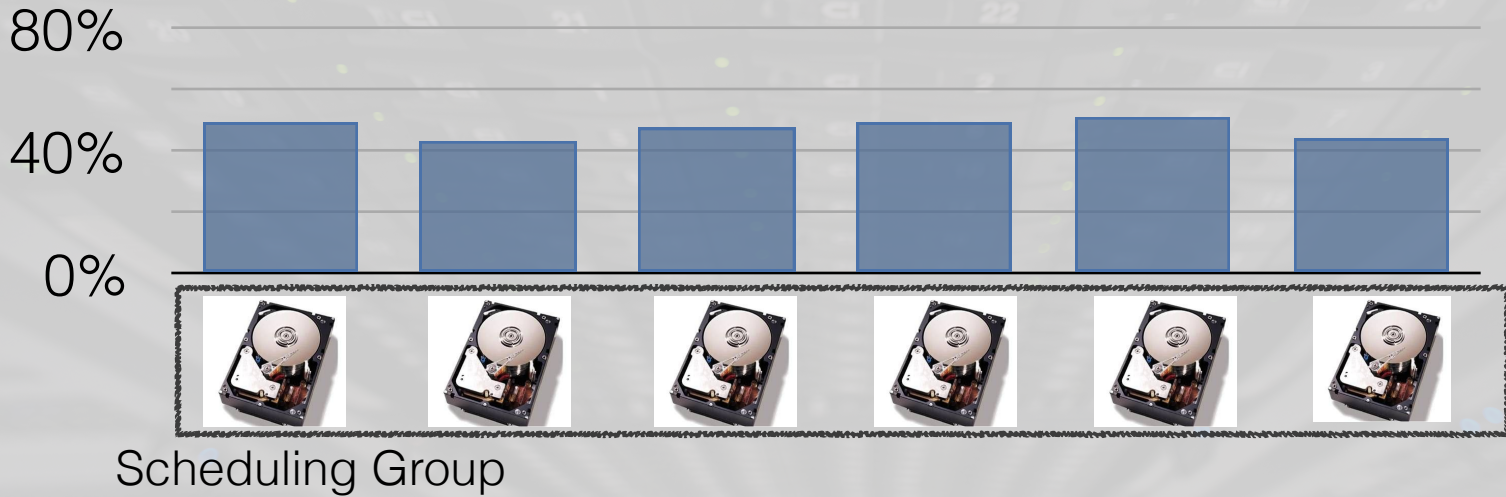
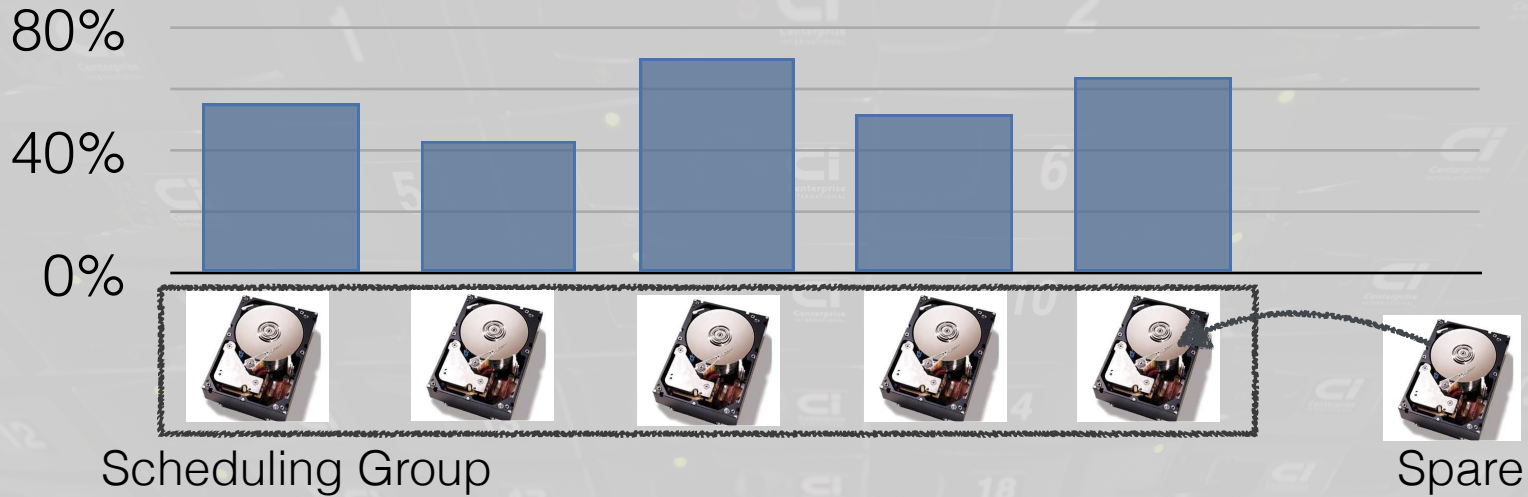
# EOS Infrastructure awareness

- Tree scheduling
- stateful
- tree efficient data structures
- multiple scheduling policies available



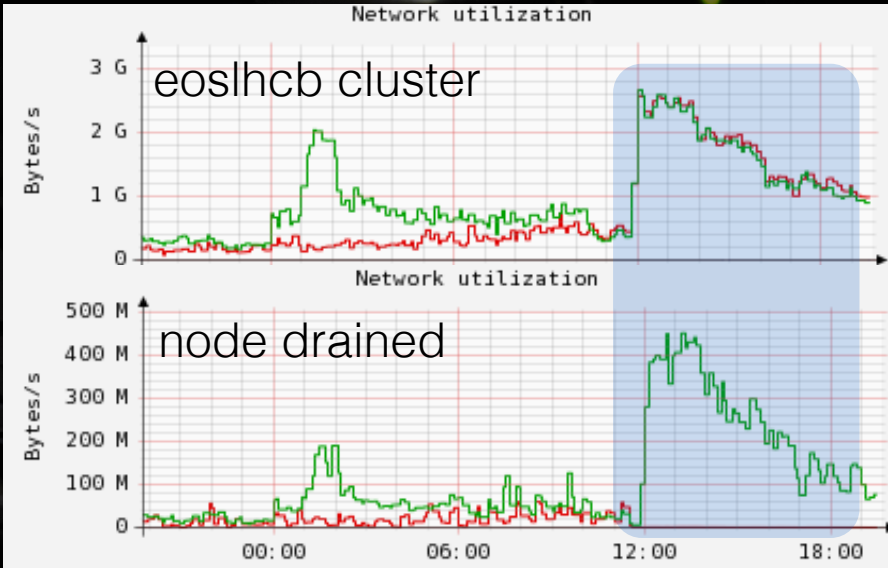
# Balancing

```
/> eos space config <space> space.balancer=on
```



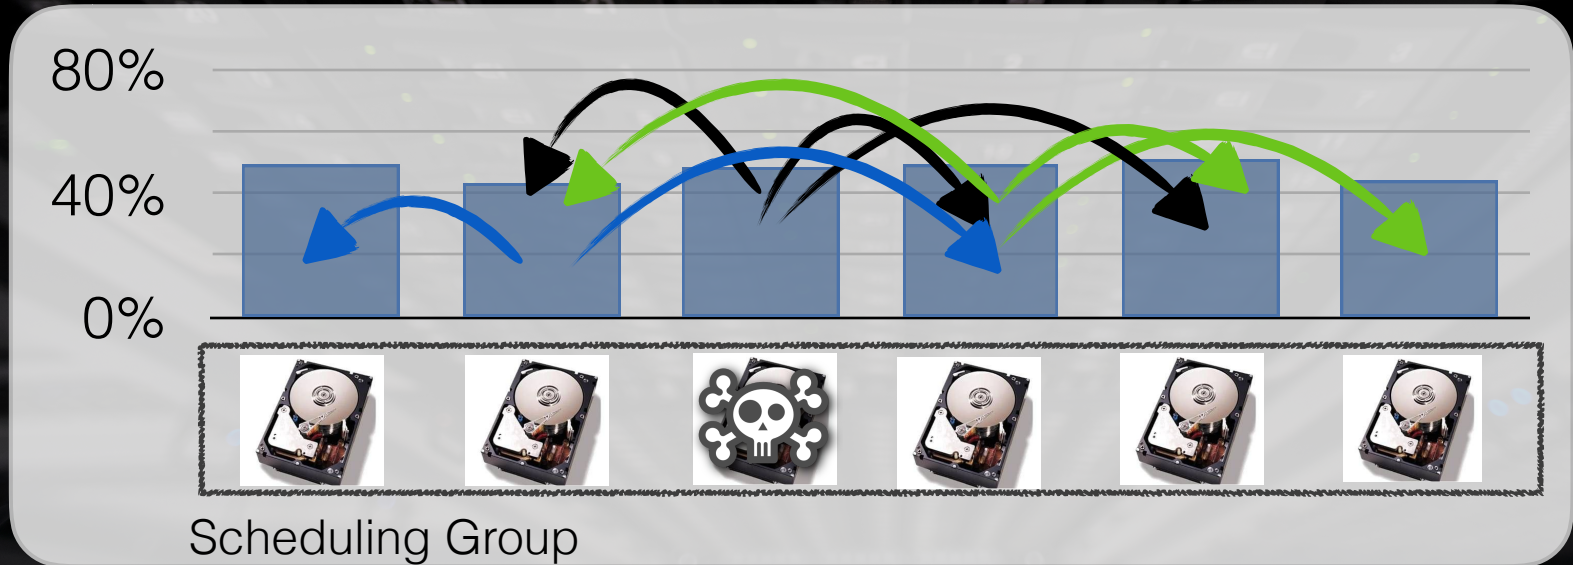
# Draining

```
/> eos fs config <fsid> configstatus=drain
```



```
[root@eoslhcb-srv-m2 (mgm:master mq:master) ~]$
[root@eoslhcb-srv-m2 (mgm:master mq:master) ~]$ eos fs ls -d
```

#	host (#...)	#	id #	path #	drain #
#	p05153074058585.cern.ch	(1095)	3536	/data16	drained
#	p05153074065039.cern.ch	(1095)	3621	/data11	drained
#	p05153074065039.cern.ch	(1095)	3658	/data21	drained
#	p05153074065039.cern.ch	(1095)	3669	/data25	drained
#	lxfsrd46c01.cern.ch	(1095)	3858	/data01	drained
#	lxfsrd46c01.cern.ch	(1095)	3859	/data02	drained
#	lxfsrd46c01.cern.ch	(1095)	3860	/data03	drained
#	lxfsrd46c01.cern.ch	(1095)	3861	/data04	drained
#	lxfsrd46c01.cern.ch	(1095)	3862	/data05	drained
#	lxfsrd46c01.cern.ch	(1095)	3863	/data06	drained
#	lxfsrd46c01.cern.ch	(1095)	3864	/data07	drained
#	lxfsrd46c01.cern.ch	(1095)	3865	/data08	drained
#	lxfsrd46c01.cern.ch	(1095)	3866	/data09	drained
#	lxfsrd46c01.cern.ch	(1095)	3867	/data10	drained
#	lxfsrd46c01.cern.ch	(1095)	3868	/data11	drained
#	lxfsrd46c01.cern.ch	(1095)	3869	/data12	drained
#	lxfsrd46c01.cern.ch	(1095)	3870	/data13	drained
#	lxfsrd46c01.cern.ch	(1095)	3871	/data14	drained





# Draining and Balancing

## Settings per space and per storage node

```
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$ eos space help
usage: space ls                                     : list spaces
usage: space ls [-s] [-m|-l|--io|--fsck] [<space>] : list in all spaces or select only <space>. <s
-s : silent mode
-m : monitoring key=value output format
-l : long output - list also file systems after ea
--io : print IO statistics
--fsck : print filesystem check statistics

space config <space-name> space.nominalsize=<value> : configure the nominal size for this space
space config <space-name> space.balancer=on|off      : enable/disable the space balancer [default=of
space config <space-name> space.balancer.threshold=<percent> : configure the used bytes deviation which trig
space config <space-name> space.balancer.node.rate=<MB/s>   : configure the nominal transfer bandwidth per r
space config <space-name> space.balancer.node.ntx=<#>       : configure the number of parallel balancing tr
space config <space-name> space.converter=on|off          : enable/disable the space converter [default=o
space config <space-name> space.converter.ntx=<#>         : configure the number of parallel conversions
space config <space-name> space.drainer.node.rate=<MB/s >   : configure the nominal transfer bandwidth per r
space config <space-name> space.drainer.node.ntx=<#>       : configure the number of parallel draining tra
space config <space-name> space.lru=on|off               : enable/disable the LRU policy engine [default
space config <space-name> space.lru.interval=<sec>       : configure the default lru scan interval
space config <space-name> space.headroom=<size>          : configure the default disk headroom if not de
space config <space-name> space.scaninterval=<sec>       : configure the default scan interval if not de
space config <space-name> space.drainperiod=<sec>        : configure the default drain period if not de
space config <space-name> space.graceperiod=<sec>        : configure the default grace period if not de
space config <space-name> space.autorepair=on|off        : enable auto-repair of faulty replica's/files
space config <space-name> space.geo.access.policy.write.exact=on|off : if 'on' use exact matching geo repli
space config <space-name> space.geo.access.policy.read.exact=on|off : if 'on' use exact matching geo repli

space config <space-name> fs.<key>=<value>             : configure file system parameters for each fil

space define <space-name> [<groupsize> [<groupmod>]]   : define how many filesystems can end up in one
=> <groupsize>=0 means, that no groups are bu
=> <groupmod> defines the maximum number of f

space reset <space-name> [--egroup|mapping|drain|scheduledrain|schedulebalance|ns|nsfilesystemview|nsfilemap
: reset a space e.g. recompute the drain state
space status <space-name> [-m]                       : print's all defined variables for space
space set <space-name> on|off                         : enables/disables all groups under that space
space rm <space-name>                                : remove space

space quota <space-name> on|off                      : enable/disable quota
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$
```



EOS

**Configuration**

# Layouts (replication)

```
.....  
..... Layouts ...  
.....  
- set 2 replica as standard layout ...  
  |eos> attr set default=replica /eos/instance/2-replica  
-----  
- set RAID-6 4+2 as standard layout ...  
  |eos> attr set default=raid6 /eos/instance/raid-6  
-----  
- set ARCHIVE 5+3 as standard layout ...  
  |eos> attr set default=archive /eos/instance/archive  
-----  
- re-configure a layout for different number of stripes (e.g. 10)  
  |eos> attr set sys.forced.stripes=10 /eos/instance/archive
```

```
sys.acl="u:lmascett:rwX!m"  
sys.allow.oc.sync="1"  
sys.forced.atomic="1"  
sys.forced.blockchecksum="crc32c"  
sys.forced.blocksize="4k"  
sys.forced.checksum="adler"  
sys.forced.layout="replica"  
sys.forced.maximumsize="10000000000"  
sys.forced.maxsize="10000000000"  
sys.forced.nstripes="2"  
sys.forced.space="default"  
sys.mask="700"  
sys.mtime.propagation="1"  
sys.owner.auth="*"  
sys.recycle="/eos/user/proc/recycle/"  
sys.versioning="10"
```

Layout settings are per directories

Single files can be changed with a converter

```
eos file convert <path> [<layout>:<stripes>] [targetspace]
```

# Authentication and Virtual IDs

```
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$ eos vid help
usage: vid ls [-u] [-g] [-s] [-U] [-G] [-g] [-a] [-l] [-n]           : list configured policies
        -u : show only user role mappings
        -g : show only group role mappings
        -s : show list of sudoers
        -U : show user alias mapping
        -G : show group alias mapping
        -y : show configured gateways
        -a : show authentication
        -l : show geo location mapping
        -n : show numerical ids instead of user/group names

vid set membership <uid> -uids [<uid1>,<uid2>,...]
vid set membership <uid> -gids [<gid1>,<gid2>,...]
vid rm membership <uid>           : delete the membership entries for <uid>.
vid set membership <uid> [+|-]sudo
vid set map -krb5|-gsi|-https|-sss|-unix|-tident|-voms <pattern> [vuid:<uid>] [vgid:<gid>]
                                           -voms <pattern> : <pattern> is <group>:<role>

vid set geotag <IP-prefix> <geotag> : add to all IP's matching the prefix <prefix> the geo location tag <geotag>
                                           N.B. specify the default assumption via 'vid set geotag default <default-tag>'
vid rm <key>                           : remove configured vid with name key - hint: use config dump to see the key names of vid rules

vid enable|disable krb5|gsi|sss|unix|https
                                           : enable/disables the default mapping via password database

vid add|remove gateway <hostname> [krb5|gsi|sss|unix|https]
                                           : adds/removes a host as a (fuse) gateway with 'su' privileges
                                           [<prot>] restricts the gateway role change to the specified authentication method
```

```
[root@eosuser-srv-b2 (mgm:master mq:master) ~]$
```

```
config dump [--fs|-f] [--vid|-v] [--quota|-q] [--policy|-p] [--comment|-c] [--global|-g] [--access|-a] [<name>] [--map|-m] :
                                           dump current configuration or configuration with name <name>

-f : dump only file system config
-v : dump only virtual id config
-q : dump only quota config
-p : dump only policy config
-g : dump only global config
-a : dump only access config
-m : dump only mapping config
```

# Disk Pools and Scheduling Groups

```
eos space define [groupsize] [groupmod]
```

```
eos space ls
```

```
eos group set <group> on/off
```

```
eos group ls
```

```
eos node ls
```

```
eos node ls --sys
```

```
eos node ls --io
```

```
eos node set <node> on/off
```

```
eos fs ls
```

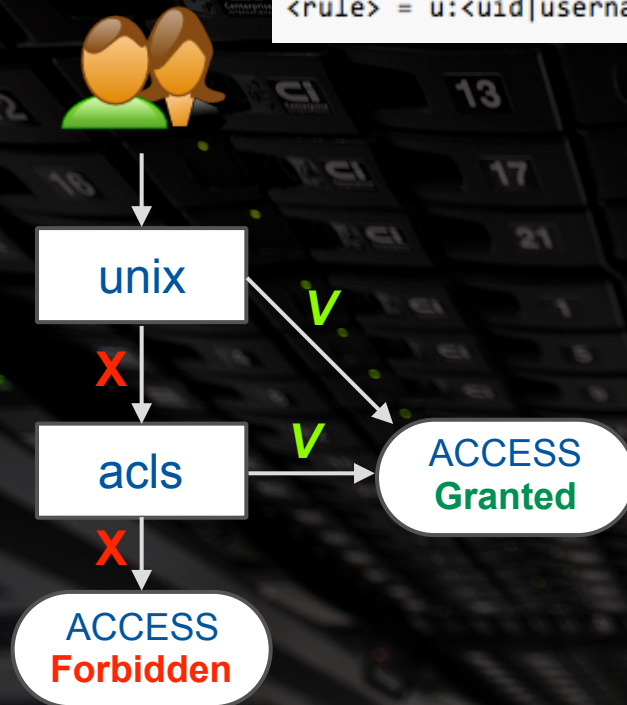
```
eos fs ls -d
```

```
eos fs ls --io
```

# Permission Systems

Permission are verified at directory level

2-step permissions:  
unix bits + ACLs



```
<rule> = u:<uid|username>|g:<gid|groupname>|egroup:<name>:{rwxomqc(!d)(+d)(!u)(+u)}
```

tag	definition
r	grant read permission
w	grant write permissions
x	grant browsing permission
m	grant change mode permission
!m	forbid change mode operations
!d	forbid deletion of files and directories
+d	overwrite a '!d' rule and allow deletion of files and directories
!u	forbid update of files
+u	overwrite a '!u' rule and allow updates for files
q	grant 'set quota' permissions on a quota node
c	grant 'change owner' permission on directory children

# Permission Systems

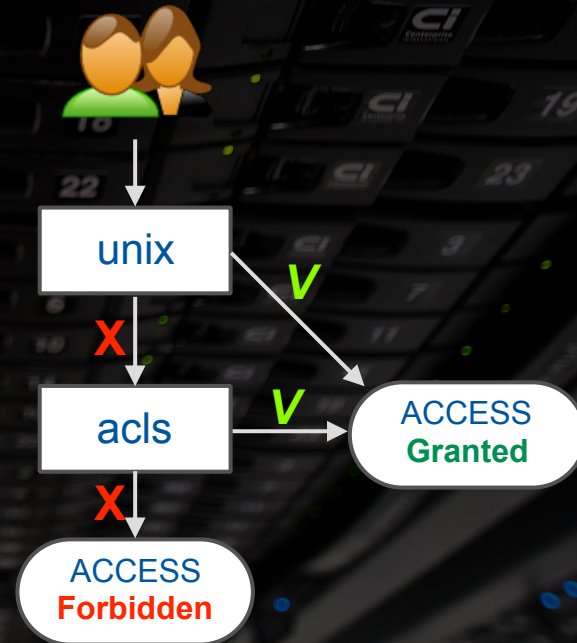
Admin and Users ACLs

```
sys.eval.useracl=true
```

sys and user acls are concatenated and evaluated

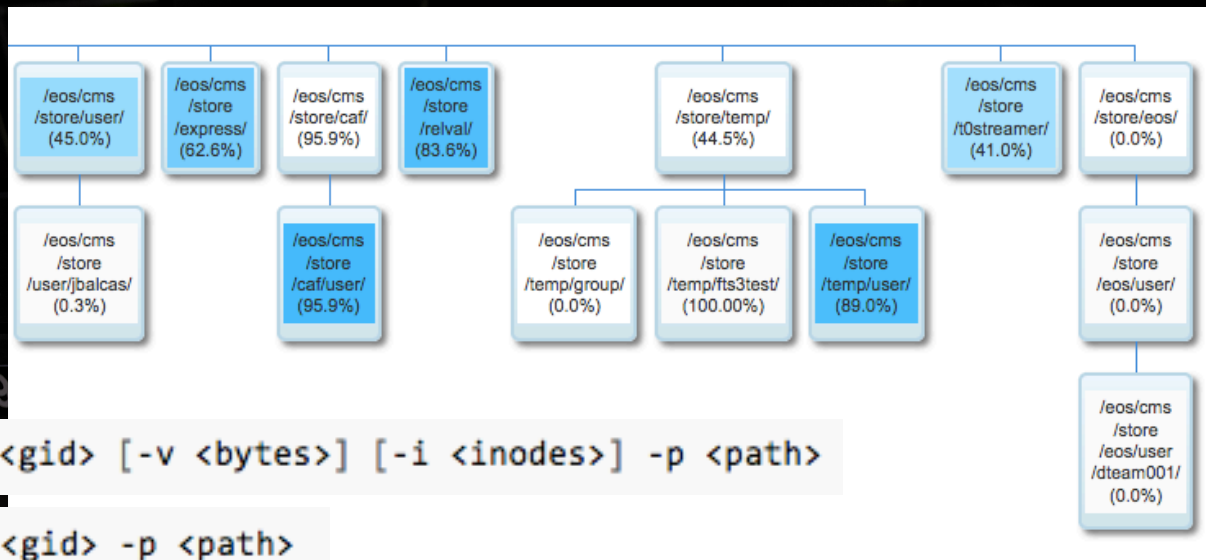
```
sys.acl=<acllist>  
user.acl=<acllist>
```

```
<acllist> = <rule1>,<rule2>...<ruleN>
```



# Quota System


## User Group Project



```
eos quota set -u <uid>|-g <gid> [-v <bytes>] [-i <inodes>] -p <path>
```

```
eos quota rm -u <uid> |-g <gid> -p <path>
```

```
eos quota rmnode -p <path>
```





# Recycle Bin

FIFO policy for delayed file deletion  
time and volume based

```
EOS Console [root://localhost] | /eos/> recycle config --lifetime 86400
```

```
EOS Console [root://localhost] | /eos/> recycle config --size 100G
```

```
EOS Console [root://localhost] | /eos/dev/2rep/subnode/> recycle config --add-bin /eos/dev/
```

```
EOS Console [root://localhost] | /eos/> recycle
```

```
#  
# used 0.00 B out of 100.00 GB (0.00% volume / 0.00% inodes used) Object-Lifetime 86400 [s]  
#
```

# Recycle Bin

## Listing recycle entries

```
EOS Console [root://localhost] | /eos/dev/2rep/subnode/> recycle ls
# Deletion Time          UID      GID      TYPE      RESTORE-KEY      RESTORE-PATH
# =====
Thu Mar 21 23:02:22 2013  apeters  z2       recursive-dir 00000000000007cf /eos/dev/2rep/subnode/tree
```

## restoring recycle entries

```
EOS Console [root://localhost] | /eos/dev/2rep/subnode/> recycle restore -f 00000000000007cf
warning: renamed restore path=/eos/dev/2rep/subnode/tree to backup-path=/eos/dev/2rep/subnode/tree.00000000000007cf
success: restored path=/eos/dev/2rep/subnode/tree
```

## emptying completely recycle entries

```
EOS Console [root://localhost] | /eos/dev/2rep/subnode/> recycle purge
success: purged 1 bulk deletions and 0 individual files from the recycle bin!
```

# EOS Features



rate limiting

recycle bin

user/host ban

https/S3 extensions

workflow engine

accounting

fuse optimisation

intergroup data balancing

versioning

geo scheduling policies

id mappings

io monitoring

fsck

stat monitoring

sticky ownership

transfers engine

rich ACLs

sharing

gateways

and many more...



EOS

# Summary

# Summary



**EOS** open storage provides a **very flexible** platform for large communities

- storage technology used to store LHC data - 340PB deployed
- more than 17k users storing data today via **CERNBox**

**EOS** is **deployed** in several sites around the world

- Aarnet (Australia), IHEP (China), FermiLab (US), EC-JRC (Italy), INFN Trieste (Italy), ASGC (Taiwan), KI, JINR (Russia), SPbSU and many others...
- multiple ALICE grid sites

Ongoing collaboration with **COMTRADE** within CERN OpenLab to provide commercial support to companies

**EOS workshops (yearly)**

- **Last one in February 2019**
  - over 80 registered participants from over 25 institutions
  - 32 oral contributions during the 2 days workshop

# References & Documentation



**EOS main website:**

<https://eos.web.cern.ch/>

**EOS Documentation:**

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

<https://quarkdb.web.cern.ch/quarkdb/docs/master/>


**EOS Community website:**

<https://eos-community.web.cern.ch/>



*Thanks for the attention!*

# CERNBox: the CERN cloud storage driven by EOS



CERNBox

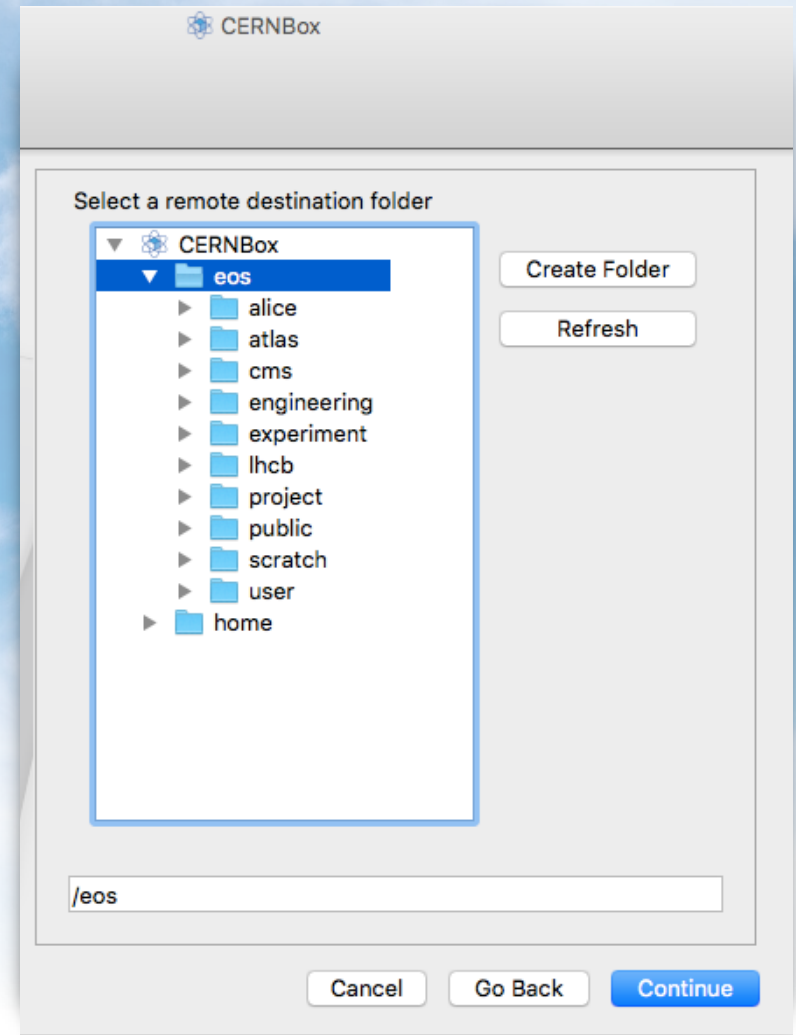
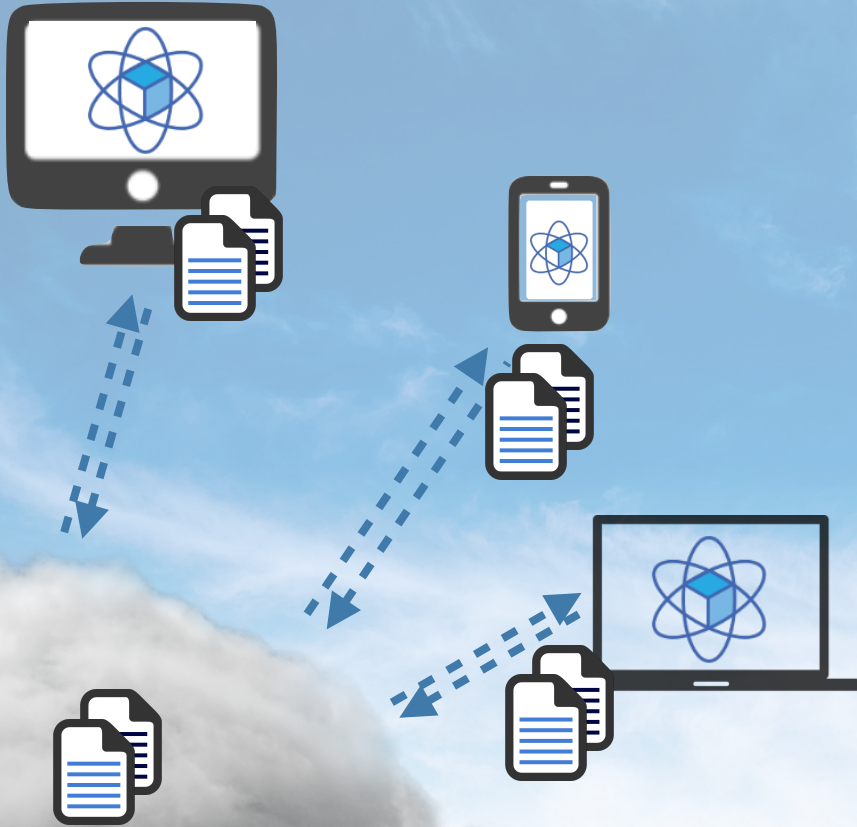
powered by



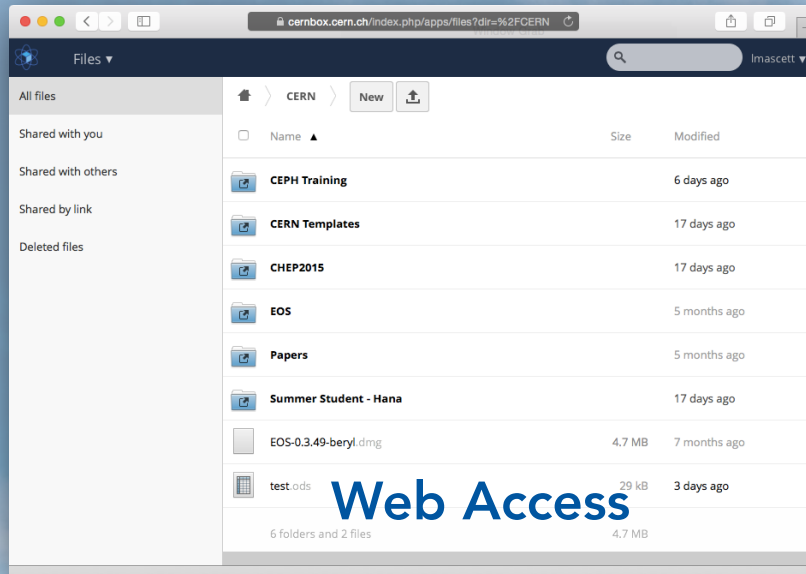
EOS



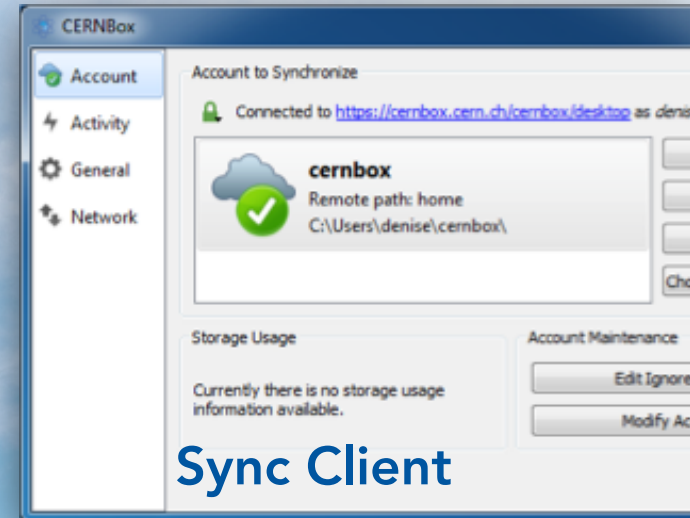
# Bring data closer to our users: CERNBox



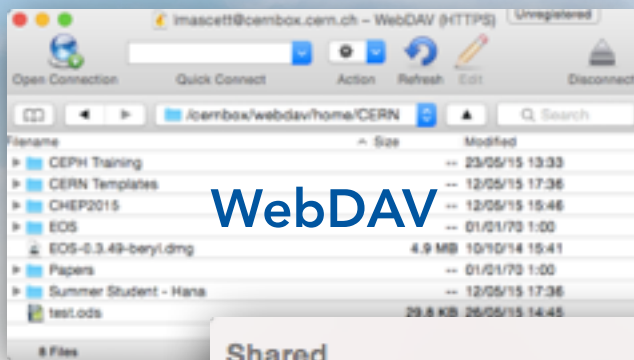
# Available Access Methods



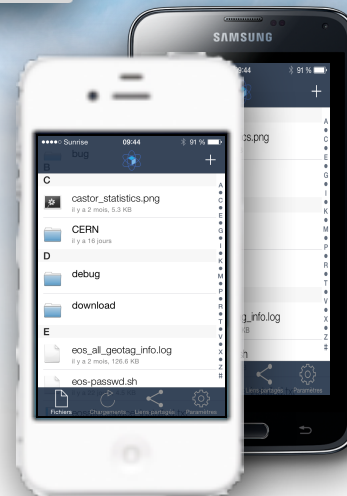

Web Access



Sync Client



WebDAV

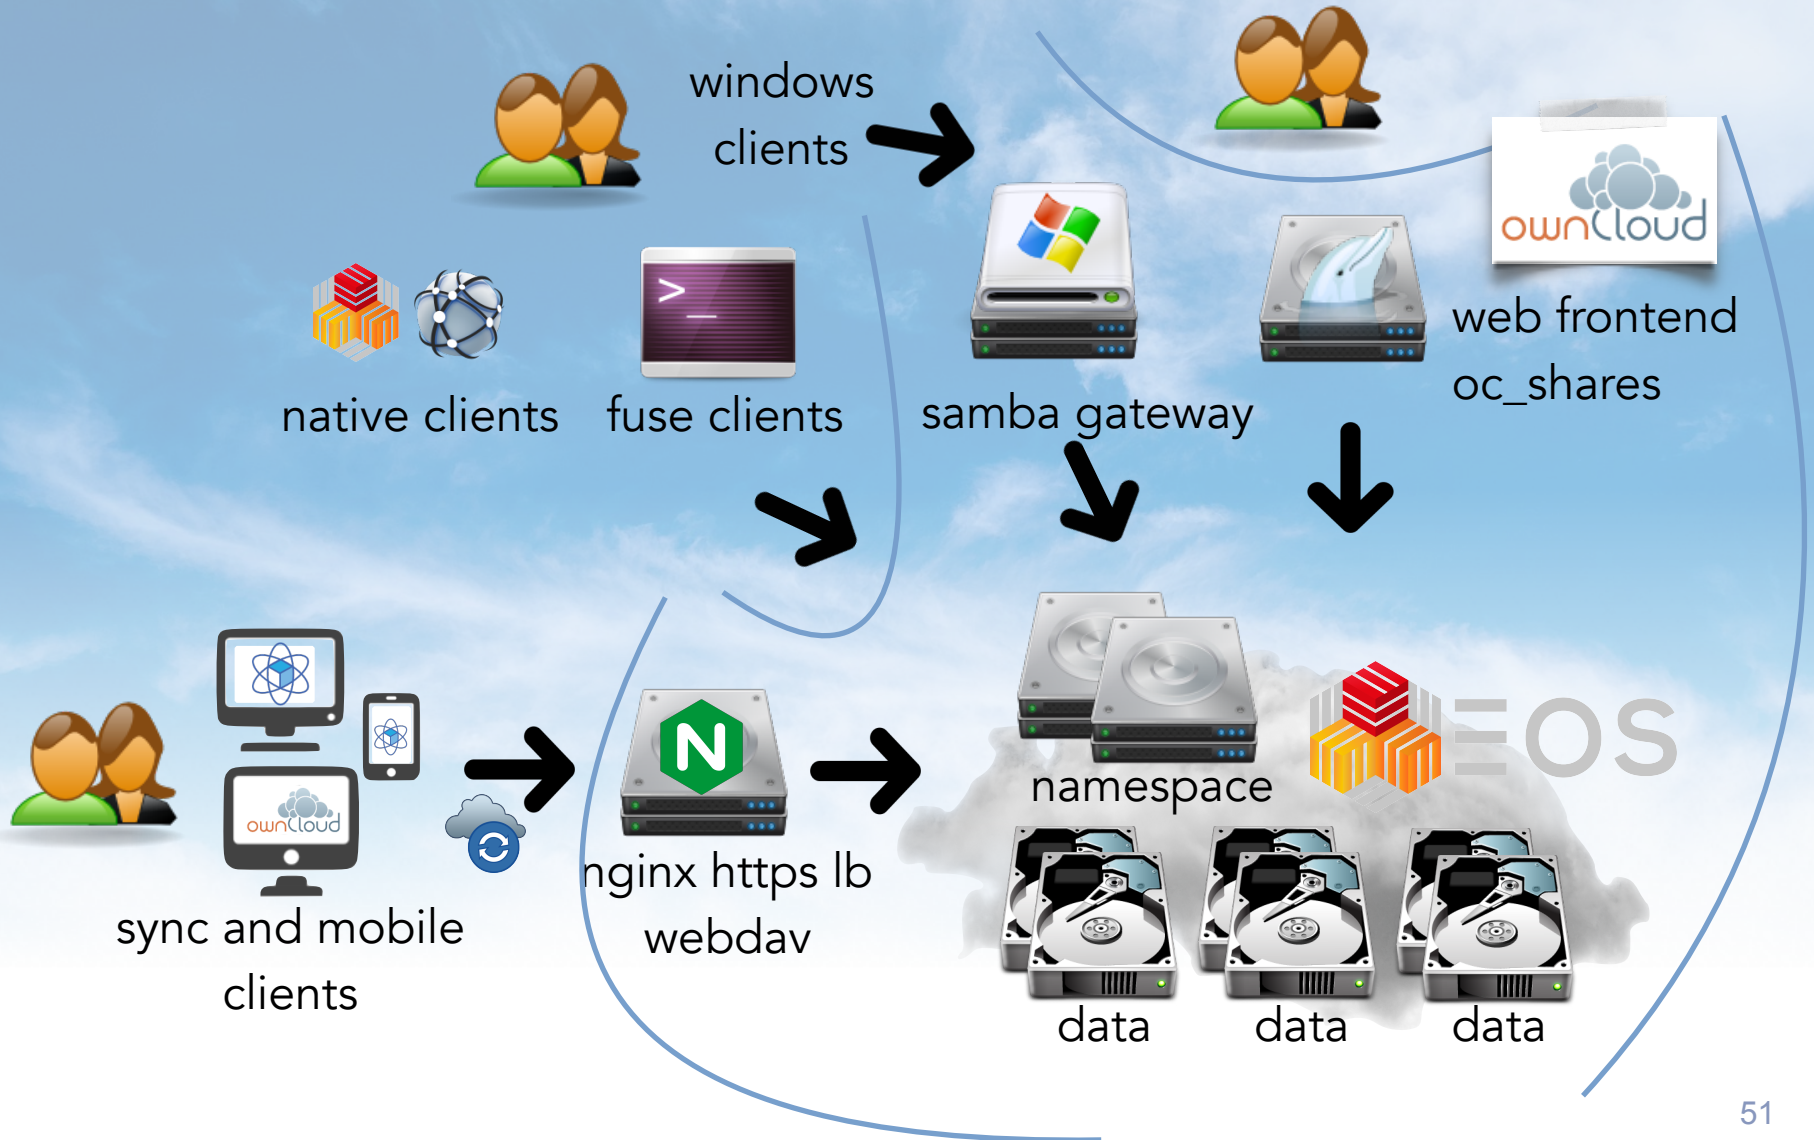



Mobile App




Directly from the storage backend  
**EOSUSER**  
(xroot, http, s3, ...)

# EOS: the CERNBox backend





CERNBox





[www.cern.ch](http://www.cern.ch)

