

EOS - Task #9306

Make EOS test instance functional

06/26/2015 11:02 AM - Gerard Bernabeu Altayo

Status:	Resolved	Start date:	06/26/2015
Priority:	Normal	Due date:	
Assignee:	Gerard Bernabeu Altayo	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Description			
In order to have a properly functional EOS test instance I have to:			
1. Reinstall Cmsstor150 as fst			
2. EOS keytab test must be different and take out keytab-test			
3. Wipe out test namespace			
4. Add some (dummy) data			
5. Test! !			
Subtasks:			
Task # 9307: Create separate eos keytab for cmseos-test			Resolved
Related issues:			
Precedes EOS - Task #9235: Test EOS version that supports links (0.3.119)		Resolved	06/29/2015 06/29/2015

History

#1 - 06/26/2015 11:05 AM - Gerard Bernabeu Altayo

- Precedes Task #9235: Test EOS version that supports links (0.3.119) added

#2 - 06/26/2015 11:20 AM - Gerard Bernabeu Altayo

I'm reinstalling cmsstor150 as EOS FST for the test instance:

```
-bash-4.1$ cms-shoot cmsstor150
removing host from rocks on cmsrocks51, if necessary
cmsstor24.fnal.gov: no host cmsstor150 to remove
Connection to cmsrocks51 closed.
removing host from rocks on cmsrocks52, if necessary
cmssrv26.fnal.gov: no host cmsstor150 to remove
Connection to cmsrocks52 closed.
stopping puppet on cmsstor150, if applicable
telling host to netboot on next boot
cmsstor150: netboot -> True
set 1 hosts to boot
1 system(s) updated
telling cmstpuppetca to remove host's cert, if present
cleaning cert for cmsstor150.fnal.gov
Notice: Revoked certificate with serial 1943
Notice: Removing file Puppet::SSL::Certificate cmsstor150.fnal.gov at '/var/lib/puppet/ssl/ca/signed/cmsstor150.fnal.gov.pem'
Notice: Removing file Puppet::SSL::Certificate cmsstor150.fnal.gov at '/var/lib/puppet/ssl/certs/cmsstor150.fnal.gov.pem'
telling cmstpuppetca to update autotoken information
when you're ready to start, run:
cmspower-powerit --action cycle --comment 'reinstalling' cmsstor150
don't forget to disable zabbix monitoring if applicable
-bash-4.1$ date;  cmspower-powerit --action cycle --comment 'reinstalling' cmsstor150
Fri Jun 26 11:01:02 CDT 2015
/usr/bin/ssh -l root cmsconsole cmspower-powerit --action cycle --comment '\gerard1: reinstalling' cmsstor150
Outlet state: OFF
Outlet state: ON === cmsstor150 ===
connecting to APC apccms1382-1, outlet 7
connecting to APC apccms1382-1, outlet 7

cmsstor150.fnal.gov - unconfigured/production (SLF 6.6)
8-core Xeon E5430 @ 2.66GHz (PowerEdge 1950); 15.57 GB RAM, 16.00 GB swap
WARNING: based on your node, cmsstor150.fnal.gov, ENSTORE_CONFIG_HOST has been set to conf-stken.fnal.gov
WARNING: If this is not correct; either reset ENSTORE_CONFIG_HOST by hand, set ENSTORE_USER_DEFINED_CONFIG_HOST by hand before
running setup or use a qualifier in your setup command!
```

```
[root@cmsstor150 ~]# touch GERARD.thisshouldnotbehere.file
[root@cmsstor150 ~]# Write failed: Broken pipe
```

```
[root@cmsconsole ~]# cmspower-cons cmsstor150
```

Right now installation is stuck with a blank console. Node is pingable but can't SSH in.... I will try to powercycle again.

#3 - 06/26/2015 11:32 AM - Gerard Bernabeu Altayo

cmsstor150 reinstalled all good in the end, I just needed to wait a bit more :)

```
[root@cmsstor150 ~]# uptime
11:31:47 up 7 min, 1 user, load average: 0.01, 0.11, 0.08
[root@cmsstor150 ~]# ll
total 228
-rw-r--r--. 1 root root 11193 Jun 26 11:10 anaconda-ks.cfg
-rw-r--r--. 1 root root 11366 Jun 26 11:20 cobbler.ks
-rw-r--r--. 1 root root 21370 Jun 26 11:10 install.log
-rw-r--r--. 1 root root 5557 Jun 26 11:09 install.log.syslog
-rw-r--r--. 1 root root 164257 Jun 26 11:20 ks-post.log
-rw-r--r--. 1 root root 3986 Jun 26 11:06 ks-pre.log
[root@cmsstor150 ~]# uname -a
Linux cmsstor150.fnal.gov 2.6.32-504.23.4.el6.x86_64 #1 SMP Tue Jun 9 11:55:03 CDT 2015 x86_64 x86_64 x86_64 GNU/Linux
[root@cmsstor150 ~]# fdisk -l
```

Disk /dev/sda: 250.0 GB, 250000000000 bytes
255 heads, 63 sectors/track, 30394 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000080

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	131	1048576	83	Linux
Partition 1 does not end on cylinder boundary.						
/dev/sda2		131	2220	16777216	82	Linux swap / Solaris
/dev/sda3		2220	30395	226313216	83	Linux

Disk /dev/sdb: 12001.7 GB, 12001652244480 bytes
255 heads, 63 sectors/track, 1459117 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 131072 bytes / 402653184 bytes
Disk identifier: 0x00000000

Disk /dev/sdc: 12001.7 GB, 12001652244480 bytes
255 heads, 63 sectors/track, 1459117 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 131072 bytes / 402653184 bytes
Disk identifier: 0x00000000

Disk /dev/sdd: 12001.7 GB, 12001652244480 bytes
255 heads, 63 sectors/track, 1459117 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 131072 bytes / 402653184 bytes
Disk identifier: 0x00000000

```
[root@cmsstor150 ~]#
```

Going to apply the (new) eos fst test role...

#4 - 06/26/2015 02:43 PM - Gerard Bernabeu Altayo

I've made a few changes in the modules to properly support 2 instances (test and prod).

I'm running puppet in cmsstor150 and so far not too bad, I found the issue with xrootd... Puppet could not install it, I manually did the following:

```
[root@cmsstor150 ~]# yum install xrootd --disablerepo=epel
```

Of course puppet ended with dependency errors because of this.

After a 2nd run there are still problems, eos-dsi is required but does not exist anymore... I'm thinking on what we can do about this, in EOS upstream this has been replaced by xrootd-dsi but of course that requires newer xrotd libraries :/

#5 - 06/26/2015 03:04 PM - Gerard Bernabeu Altayo

we still have eos-dsi in an old copy of the repo:

```
[root@cmssrv201 repo]# pwd
/srv/repo
./eos-slf6-x86_64.old-20150429/eos-dsi-0.2.8-1.x86_64.rpm
```

For now I'll just copy it to the repo...

```
[root@cmssrv201 repo]# cp ./eos-slf6-x86_64.old-20150429/eos-dsi-0.2.8-1.x86_64.rpm eos-slf6-x86_64/
[root@cmssrv201 repo]# cp /srv/repo/eos-slf6-x86_64.old-20150429/vdt_
vdt_compile_globus_core-VDT1.10.1_x86_64_rhap_5-1.x86_64.rpm vdt_globus_essentials-VDT1.10.1x86_64_rhap_5-4.x86_64.rpm
vdt_globus_data_server-VDT1.10.1x86_64_rhap_5-3.x86_64.rpm vdt_globus_sdk-VDT1.10.1x86_64_rhap_5-1.x86_64.rpm
[root@cmssrv201 repo]# cp /srv/repo/eos-slf6-x86_64.old-20150429/vdt_* /srv/repo/eos-slf6-x86_64/
[root@cmssrv201 repo]# cp /srv/repo/eos-slf6-x86_64.old-20150429/sparsehash-1.11-1.noarch.rpm
/srv/repo/eos-slf6-x86_64.old-20150429/gpt-3.2_4.0.8p1_x86_64_rhap_5-1.x86_64.rpm /srv/repo/eos-slf6-x86_64/
```

```
[root@cmssrv201 repo]# make eos_server
```

Now yum finds all it needs to install, and so does puppet, which now runs clean!

```
[root@cmsstor150 ~]# service eos status
xrootd for role: fst (pid 6679) is running ...
[root@cmsstor150 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       213G  2.4G 200G  2% /
tmpfs           7.8G   0 7.8G  0% /dev/shm
/dev/sda1       976M  35M 891M  4% /boot
/dev/sdd        11T  34M 11T  1% /storage/data3
/dev/sdb        11T  34M 11T  1% /storage/data1
/dev/sdc        11T  34M 11T  1% /storage/data2
```

#6 - 06/26/2015 05:03 PM - Gerard Bernabeu Altayo

OK, I'm moving forward with this. I did fix LOTS of code for EOS in my branch, and I'm getting it ready to have 2 working instances that do share code :D

Already done:

- 1. Reinstall Cmsstor150 as fst
2. EOS keytab test must be different and take out keytab-test

I will reinstall cmsstor150 to make sure I didn't cleanup too much, before reinstall it does show up good:

```
[root@cmssrv153 ~]# eos fs ls
```

#.....
1. host (#...) # id # path # schedgroup # geotag # boot # configstatus # drain # active
#.....
cmssrv153.fnal.gov (1095) 2 /mnt/eos default.0 rw nodrain

#.....
1. host (#...) # id # path # schedgroup # geotag # boot # configstatus # drain # active
#.....
cmsstor150.fnal.gov (1095) 3 /storage/data1 spare booted rw nodrain
cmsstor150.fnal.gov (1095) 4 /storage/data2 spare booted rw nodrain
cmsstor150.fnal.gov (1095) 5 /storage/data3 spare booted rw nodrain

```
[root@cmssrv153 ~]# eos node ls
#-----
2. type # hostport # geotag # status # status # txgw #gw-queued # gw-ntx #gw-rate # heartbeatdelta #nofs
#-----
nodesview cmssrv153.fnal.gov:1095 unknown on off 0 10 120 ~ 1
nodesview cmsstor150.fnal.gov:1095 online on off 0 10 120 2 3
[root@cmssrv153 ~]#
```

Also the process for adding nodes is being simplified :D

#7 - 06/26/2015 05:59 PM - Gerard Bernabeu Altayo

Node reshoot, the only thing I had to do for it to come back was:

```
yum install xrootd --disablerepo=epel
puppet agent -t
```

I can move fwd with next steps now (remove metadata and start from scratch!).

#8 - 07/10/2015 04:22 PM - Gerard Bernabeu Altayo

There seems to be no way to remove all the entries from the mgm other than just do RMs, or wiping out the metadata files... either way does not seem awesome to me so I'll go for the easier: rm -rf.

I've checked and the MGM had fsck disabled, just enabled:

```
[root@cmssrv151 ~]# eos fsck enable
success: enabled fsck
[root@cmssrv151 ~]# eos fsck stat
150710 16:19:39 1436563179.986752 started check
[root@cmssrv151 ~]# date
Fri Jul 10 16:21:46 CDT 2015
[root@cmssrv151 ~]#
```

It will detect that all files are unavailable.... Then I'll wipe them out! Now test machines have a keytab that does not allow them in the production instance, so it should be safe...

I could not remove files so change of plans, I learned a bit more of the MGM by wiping it out:

```
[root@cmssrv153 ~]# puppet agent --disable 'wiping out md'
[root@cmssrv153 ~]# /etc/init.d/eos
eos eosd eos-gridftp eossha eosslave eossync
[root@cmssrv153 ~]# service eosd stop; service eos stop; service eossync stop
Stopping eosd:
[ OK ]
Skipping fuse mount for instance produnmerged - no /etc/sysconfig/eos.produnmerged configuration file
[WARNING]
Stopping xrootd: mgm [ OK ]
Stopping xrootd: mq [ OK ]
Stopping xrootd: sync [ OK ]

Stopping eossync: [ OK ]
[root@cmssrv153 ~]# mv /var/eos/md/
directories.cmssrv153.fnal.gov.mdlog fmd.0002.sql iostat.cmssrv153.fnal.gov.dump so.fst.dump
files.cmssrv153.fnal.gov.mdlog iostat.cmssrv151.fnal.gov.dump old/ so.mgm.dump
[root@cmssrv153 ~]# mv /var/eos/md /var/eos/md.old
[root@cmssrv153 ~]# ll /var/eos/
total 48
drwx-----+ 2 daemon root 4096 Mar 26 2014 auth
drwx----- 5 daemon root 4096 Nov 12 2014 config
-rw-r--r- 1 daemon root 0 Jun 24 12:21 eos.mgm.rw
-rw-r--r- 1 daemon root 0 Jun 24 12:21 eos.mq.master
-rwxr-xr-x 1 daemon daemon 0 Jun 24 11:48 eos.mq.remote.up
drwx----- 2 daemon daemon 4096 Apr 27 18:02 html
drwx----- 3 daemon daemon 4096 Jul 10 16:33 md.old
drwxr-xr-x 3 daemon root 4096 Apr 2 2014 report
drwxr-xr-x 2 daemon root 4096 Mar 26 2014 stage
drwxr-xr-x 2 daemon root 20480 Jul 10 03:27 tx
[root@cmssrv153 ~]# mkdir /var/eos/md; chown daemon.daemon /var/eos/md; chmod 700 /var/eos/md
[root@cmssrv153 ~]# service eos start
```

Starting xrootd as mgm with -n mgm -c /etc/xrd.cf.mgm -m -l /var/log/eos/xrdlog.mgm -b -Rdaemon

[OK]

Starting xrootd as mq with -n mq -c /etc/xrd.cf.mq -l /var/log/eos/xrdlog.mq -b -Rdaemon

[OK]

Starting xrootd as sync - already started [FAILED]

```
[root@cmssrv153 ~]# service eossync start
```

Starting eossync:

```
FILE 0 => TARGET cmssrv153.fnal.gov:1096 [PASSED]
FILE 0 => TARGET cmssrv151.fnal.gov:1096 [PASSED]
FILE 1 => TARGET cmssrv153.fnal.gov:1096 [PASSED]
FILE 1 => TARGET cmssrv151.fnal.gov:1096 [PASSED]
FILE 2 => TARGET cmssrv153.fnal.gov:1096 [PASSED]
FILE 2 => TARGET cmssrv151.fnal.gov:1096 [PASSED]
CONF => TARGET cmssrv153.fnal.gov:1096
CONF => TARGET cmssrv151.fnal.gov:1096 [PASSED]
[PASSED]
```

```
[root@cmssrv153 ~]# ll /eos/
```

total 0

```
[root@cmssrv153 ~]# service eosd status
```

eosd is stopped

```
[root@cmssrv153 ~]# service eosd start
```

```
Starting eosd: [ OK ]
EOS_FUSE_DEBUG      : 0
EOS_FUSE_NOACCESS   : 1
EOS_FUSE_KERNELCACHE : 1
EOS_FUSE_DIRECTIO   : 0
EOS_FUSE_CACHE      : 1
EOS_FUSE_CACHE_SIZE : 300000000
EOS_FUSE_CACHE_WRITE : 0
EOS_FUSE_BIGWRITES  : 1
EOS_FUSE_EXEC       : 0
EOS_FUSE_LOCK_ENVIRONMENT : 0
EOS_FUSE_NO_MT      : 0
```

Skipping fuse mount for instance produnmerged - no /etc/sysconfig/eos.produnmerged configuration file
[WARNING]

```
[root@cmssrv153 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       213G  28G  175G  14% /
tmpfs           16G   0  16G   0% /dev/shm
/dev/sda1       976M  86M  839M  10% /boot
eosmain         33T  106M  33T   1% /eos
[root@cmssrv153 ~]# ll /eos
total 0
drwxr-xr-x 1 root root 1 Jul 10 16:43 cmseos-test.fnal.gov
drwxr-xr-x 1 root root 1 Jul 10 16:43 storetest
drwxr-xr-x 1 root root 2 Jul 10 16:43 uscms
[root@cmssrv153 ~]# tree /eos/
```

```
/eos/
├── cmseos-test.fnal.gov
│   └── proc
│       ├── archive
│       ├── conversion
│       ├── master
│       ├── quota
│       ├── reconnect
│       ├── recycle
│       ├── who
│       └── whoami
├── storetest
├── uscms
├── store
└── storetest
```

9 directories, 5 files
[root@cmssrv153 ~]#

On the backup machine I did the same.

#9 - 07/10/2015 05:05 PM - Gerard Bernabeu Altayo

I'm copying data in here:

```
[root@cmssrv153 ~]# eos mkdir /eos/gba.test
[root@cmssrv153 ~]# eos chmod 777 /eos/gba.test/
success: mode of file/directory /eos/gba.test/ is now '777'
[root@cmssrv153 ~]# ls -lah /eos/gba.test/
total 0
drwxrwxrwx 1 root root 0 Jul 10 16:51 .
drwxrwxr-x 1 root root 4 Jul 10 16:51 ..
[root@cmssrv153 ~]#
```

Unfortunately the quotas are still there and I can't really write in the FS... I think the best will be to reinstall the node and start from scratch because there is too much broken stuff in here I don't know how to fix:

```
[root@cmssrv153 ~]# eos space ls
#-----
-
1. type #      name #  groupsize #  groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota #balancing
# threshold # converter # ntx # active #intergroup
#-----
-----
spaceview     default     ?      ?      4      3      110.22 M      36.00 T      36.00 T      0 ? ?      0 ? 0 0 ?
[root@cmssrv153 ~]#
```

The slave MGM does not start either:

```

150710 16:58:27 time=1436565507.736325 func=BootNamespace level=NOTE logid=cb1e901a-274e-11e5-a2b0-782bcb3c8025 unit=
mgm@cmssrv151.fnal.gov:1094 tid=00007f95a9fe0740 source=Master:1889          tident=<service> sec=   uid=0 gid=0 name= geo="" eos
directory view configure started
150710 16:58:27 time=1436565507.736797 func=BootNamespace level=CRIT logid=cb1e901a-274e-11e5-a2b0-782bcb3c8025 unit=
mgm@cmssrv151.fnal.gov:1094 tid=00007f95a9fe0740 source=Master:1904          tident=<service> sec=   uid=0 gid=0 name= geo="" eos
view initialization failed after 0 seconds
150710 16:58:27 time=1436565507.736850 func=BootNamespace level=CRIT logid=cb1e901a-274e-11e5-a2b0-782bcb3c8025 unit=
mgm@cmssrv151.fnal.gov:1094 tid=00007f95a9fe0740 source=Master:1906          tident=<service> sec=   uid=0 gid=0 name= geo=""
initialization returned ec=9 Unable to write the record data at offset 0x156c; Bad file descriptor
150710 16:58:27 4133 XrootdConfig: Unable to create file system object via libXrdEosMgm.so
150710 16:58:27 4133 XrootdConfig: Unable to load file system.
----- xrootd protocol initialization failed.
150710 16:58:27 4133 XrdProtocol: Protocol xrootd could not be loaded
----- xrootd mgm@cmssrv151.fnal.gov:-1 initialization failed.

```

One other problem I have (puppet missconfig) is that the FST is mounting CERN's test MGM....

```

[root@cmsstor150 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       213G  2.6G  199G   2% /
tmpfs           7.8G   0 7.8G   0% /dev/shm
/dev/sda1       976M  35M  891M   4% /boot
/dev/sdd        11T   36M  11T   1% /storage/data3
/dev/sdb        11T   36M  11T   1% /storage/data1
/dev/sdc        11T   36M  11T   1% /storage/data2
eosmain        417T  852G  416T   1% /eos
[root@cmsstor150 ~]# mount
/dev/sda3 on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sda1 on /boot type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
/dev/sdd on /storage/data3 type xfs (rw,nobarrier,inode64)
/dev/sdb on /storage/data1 type xfs (rw,nobarrier,inode64)
/dev/sdc on /storage/data2 type xfs (rw,nobarrier,inode64)
eosmain on /eos type fuse (rw,nosuid,nodev,allow_other)

```

#10 - 07/14/2015 08:46 AM - Gerard Bernabeu Altayo

Reseted the config as recommended by Andreas:

```

[root@cmssrv153 ~]# eos space ls
#-----
-

1. type #      name #  groupsize #  groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota #balancing
# threshold # converter #  ntx # active #intergroup
#-----
-----
spaceview     default   ?    ?    4    3    110.22 M    36.00 T    36.00 T    0 ?    ?    0    ?    0    0    ?
[root@cmssrv153 ~]# eos config reset
success: configuration has been reset(cleaned)!
[root@cmssrv153 ~]# eos space ls
[root@cmssrv153 ~]# eos fs ls
[root@cmssrv153 ~]# eos ns

2. -----
3. Namespace Statistic
4. -----
ALL  Files                5 [booted] (0s)
ALL  Directories           30

5. -----
ALL  Compactification      status=off waitstart=0 interval=0 ratio-file=0.0:1 ratio-dir=0.0:1

6. -----
ALL  Replication           mode=master-rw state=master-rw master=cmssrv153.fnal.gov configdir=/var/eos/config/cmssrv153.fnal.gov/
config=default active=true mgm:cmssrv151.fnal.gov=ok mgm:mode=slave-ro mq:cmssrv151.fnal.gov:1097=ok

7. -----
ALL  File Changelog Size    580 B
ALL  Dir  Changelog Size    6116 B

8. -----
ALL  avg. File Entry Size   116 B
ALL  avg. Dir  Entry Size   203 B

9. -----
ALL  memory virtual         645.51 MB

```

```

ALL   memory resident      298.87 MB
ALL   memory share         9.58 MB
ALL   memory growths      190.99 MB
ALL   threads              45
ALL   uptime               55066

```

10. -----

```

[root@cmssrv153 ~]# eos config help
Usage: config ls|dump|load|save|diff|changelog|reset|autosave [OPTIONS]
[eos] config provides the configuration interface to EOS.

```

Options:

```

config ls [--backup|-b] :
list existing configurations
--backup|-b : show also backup & autosave files
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
config save [-f] [<name>] [--comment|-c "<comment>"] :
save config (optionally under name)
-f : overwrite existing config name and create a timestamped backup
=> if no name is specified the current config file is overwritten

```

```

config load <name> :
load config (optionally with name)
config diff :
show changes since last load/save operation
config changelog [-#lines] :
show the last <#> lines from the changelog - default is -10
config reset :
reset all configuration to empty state
config autosave [on|off] :
without on/off just prints the state otherwise set's autosave to on or off

```

```

[root@cmssrv153 ~]# eos config changelog
2015-07-14 08:42:54 del config fs:/eos/cmsstor150.fnal.gov:1095/fst/storage/data1
2015-07-14 08:42:54 del config fs:/eos/cmsstor150.fnal.gov:1095/fst/storage/data2
2015-07-14 08:42:54 del config fs:/eos/cmsstor150.fnal.gov:1095/fst/storage/data3
2015-07-14 08:42:54 set config global:/config/cmseos-test.fnal.gov/mgm/#nextfsid => 0
2015-07-14 08:42:54 del config fs:/eos/cmssrv153.fnal.gov:1095/fst/mnt/eos
[root@cmssrv153 ~]# eos fs ls
[root@cmssrv153 ~]#

```

This will clear quotas, etc. The FS was cleared before. Info from Andreas on this:

Hi Gerard,
the best is to do:

```
eos config reset
```

This wipes the whole configuration to scratch including the filesystems ... but not sure you want that.

If you want just to keep the filesystems, erase from the configuration file on disk (default.eoscf) all the keys you want to remove like lines starting with "quota:", mappings "vid:", space/node settings "global:" if you want to get rid of the quota nodes, the fastest is to remove the namespace files "/var/eos/md/*.mdlog" and recreate the basic directory structure or to delete all the quota nodes one by one with "quota rmnode" ... but this is sort of painful.

Do you think we want an "instance reset" command where we just keep all the filesystem configuration but reinitialize the namespace, mappings, quotas? If yes, this should have a proper security question ;-) before execution ...

Cheers Andreas.

#11 - 07/14/2015 10:49 AM - Gerard Bernabeu Altayo

OK, so after wiping config:

```

[root@cmssrv153 ~]# eos config reset
success: configuration has been reset(cleaned)!

```

```

[root@cmssrv153 ~]# eos vid enable sss

```

```
[root@cmsstor150 ~]# rm -f /storage/data*/.eosfs*
[root@cmsstor150 ~]# eosfstregister `grep '^ *all.manager' /etc/xrd.cf.fst | awk '{print $2}'` /storage/data spare:`mount | grep -c /storage/data`
#####
```

```
1. <eosfstregister> v1.0.0 #####
/storage/data1 : uuid=c1124330-245c-4090-87db-c4c89ee4bfb3 fsid=undef
success: mapped 'c1124330-245c-4090-87db-c4c89ee4bfb3' <=> fsid=1
/storage/data2 : uuid=61f98b7f-50b6-4c41-94f3-35c77b2275cb fsid=undef
success: mapped '61f98b7f-50b6-4c41-94f3-35c77b2275cb' <=> fsid=2
/storage/data3 : uuid=6fcd9e57-1a43-46d8-9ee8-c7c74a86143d fsid=undef
success: mapped '6fcd9e57-1a43-46d8-9ee8-c7c74a86143d' <=> fsid=3
```

Now removing the old quotas still left:

```
[root@cmssrv153 cmssrv153.fnal.gov]# XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/storetest/lisa/
Do you really want to delete the quota node under path /eos/uscms/storetest/lisa/ ?
Confirm the deletion by typing => 8102483610
=> 8102483610
```

```
Deletion confirmed
success: removed quota node /eos/uscms/storetest/lisa/ (errc=0) (Success)
[root@cmssrv153 cmssrv153.fnal.gov]#
```

...

It looks like now we're good:

```
[root@cmssrv153 cmssrv153.fnal.gov]# eos space ls
#-----
-

1. type #      name # groupsize # groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota #balancing
# threshold # converter # ntx # active #intergroup
#-----
spaceview      spare      0      0      3      0      110.21 M      36.00 T      0      0 off off 20 off 2 0
off
[root@cmssrv153 cmssrv153.fnal.gov]# eos fs ls
```

```
#-----
1. host (#...) # id #      path # schedgroup # geotag # boot # configstatus # drain # active
#-----
cmsstor150.fnal.gov (1095) 1 /storage/data1 spare booted rw nodrain
cmsstor150.fnal.gov (1095) 2 /storage/data2 spare booted rw nodrain
cmsstor150.fnal.gov (1095) 3 /storage/data3 spare booted rw nodrain
[root@cmssrv153 cmssrv153.fnal.gov]# eos space ls ↓
```

```
#-----
2. type #      name # groupsize # groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota
#-----
spaceview      spare      0      0      3      0      110.21 M      36.00 T      0      0 off
```

```
#-----
3. host #port # id #      uuid #      path # schedgroup # headroom # boot # configstatus # drain # active# scaninterval
#-----
cmsstor150.fnal.gov 1095 1 c1124330-245c-4090-87db-c4c89ee4bfb3 /storage/data1 spare 0.00 booted rw
nodrain 604800
cmsstor150.fnal.gov 1095 2 61f98b7f-50b6-4c41-94f3-35c77b2275cb /storage/data2 spare 0.00 booted rw
nodrain 604800
cmsstor150.fnal.gov 1095 3 6fcd9e57-1a43-46d8-9ee8-c7c74a86143d /storage/data3 spare 0.00 booted rw
nodrain 604800
[root@cmssrv153 cmssrv153.fnal.gov]#
```

```
[root@cmsstor150 ~]# mgm=cmssrv153.fnal.gov #For production: cmssrv222.fnal.gov
[root@cmsstor150 ~]# ssh $mgm eos node set ${HOSTNAME}:1095 on
[root@cmsstor150 ~]# ssh $mgm eos vid add gateway ${HOSTNAME}
success: set vid [ eos.rgid=0 eos.ruid=0 mgm.cmd=vid mgm.subcmd=set mgm.vid.auth=tident mgm.vid.cmd=map mgm.vid.gid=0
mgm.vid.key=<key> mgm.vid.pattern="*@cmsstor150.fnal.gov" mgm.vid.uid=0 ]
[root@cmsstor150 ~]# fsids=`ssh $mgm eos fs ls -m $HOSTNAME | awk '{print $3}' | grep id= | cut -d= -f2`
[root@cmsstor150 ~]# if [ $(('echo $HOSTNAME | grep -o '[0-9]+' % 2 )) -eq 0 ]; then
```

```
dest=0
else
dest=2
fi
```



```
[root@cmsstor150 ~]# for i in $fsids; do
```

```
    cmd="eos fs mv $i default.$dest"
    echo $cmd
    ssh $mgm $cmd
    let dest++
done
```

```
eos fs mv 1 default.0
success: moved filesystem 1 into space default.0
eos fs mv 2 default.1
success: moved filesystem 2 into space default.1
eos fs mv 3 default.2
success: moved filesystem 3 into space default.2
[root@cmsstor150 ~]#
```

Now it shows all good:

```
[root@cmssrv153 cmssrv153.fnal.gov]# eos space ls
```

```
#-----
```

```
1. type #      name # groupsize # groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota #balancing
# threshold # converter # ntx # active #intergroup
#-----
```

```
-----
spaceview      default      0      0  3      0      110.21 M      36.00 T      0      0 off  off  20      off  2      0
off
```

```
[root@cmssrv153 cmssrv153.fnal.gov]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda3       213G  24G  179G  12% /
tmpfs           16G   0  16G   0% /dev/shm
/dev/sda1       976M  86M  839M  10% /boot
eosmain         33T  106M  33T   1% /eos
[root@cmssrv153 cmssrv153.fnal.gov]#
```

```
[root@cmssrv153 cmssrv153.fnal.gov]# eos quota ls | grep 'Quota Node:' | awk '{print "XrdSecPROTOCOL=unix eos quota rmnode -p " $5}'
```

```
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/gba.test/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/storetest/catalind/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/burt/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/burt/feb6/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/burt/test5/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/burttest/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/burttest4/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/catalind/a/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/lisa/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/lpcbtag/
XrdSecPROTOCOL=unix eos quota rmnode -p /eos/uscms/store/user/lpcmuon/
```

#12 - 07/17/2015 11:26 AM - Gerard Bernabeu Altayo

Finally the procedure to follow that actually gets rid of the quotas and everything on one simple step is:

```
eos config reset
eos vid enable sss
eos config save -f
eos config autosave on
```

This persists across restarts... Then I had to add the FST manually as per

https://cmsweb.fnal.gov/bin/view/Storage/EOSOperationalProcedures#Install_a_new_EOS_FST_node

I had to restart the FST for the space to be available in the just reseted EOS instance. I also had to restart the slave MGM.

Now it works:

```
[root@cmssrv153 ~]# su - gerard1
-bash-4.1$ cp /etc/passwd /eos/gba.test/
-bash-4.1$ ll /eos/gba.test/
total 2
-rw-r--r-- 1 daemon daemon 1682 Jul 17 11:22 passwd
-bash-4.1$
[root@cmssrv153 ~]# wc -l //eos/gba.test/passwd
```

```
33 //eos/gba.test/passwd
```

Still, there was considerable FS corruption coming from the quotas so I wiped out the FS again, and then created a directory where everyone can write, so I can test :)

And basic links seem to work!

```
[root@cmssrv153 ~]# eos mkdir /eos/gerard
[root@cmssrv153 ~]# eos chmod 777 /eos/gerard
success: mode of file/directory /eos/gerard is now '777'
[root@cmssrv153 ~]# cp /etc/passwd /eos/gerard/
[root@cmssrv153 ~]# cd /eos/gerard/
[root@cmssrv153 gerard]# ll
total 2
-rw-r--r-- 1 daemon daemon 1682 Jul 17 12:25 passwd
[root@cmssrv153 gerard]# ln -s passwd passwd.link
[root@cmssrv153 gerard]# ll
total 2
-rw-r--r-- 1 daemon daemon 1682 Jul 17 12:25 passwd
lrwxrwxrwx 1 daemon daemon    0 Jul 17 12:26 passwd.link -> passwd
[root@cmssrv153 gerard]#
```

My test is going to be to build the linux kernel on EOS, and then I'll do it on the local FS and let's see what's faster :)

Will follow <http://kernelnewbies.org/KernelBuild>

#13 - 07/17/2015 12:45 PM - Gerard Bernabeu Altayo

Looks like we're not there yet, it may not be related to links but a simple 'git clone' has trouble to run, not sure if the errors are fatal yet:

```
[root@cmssrv153 gerard]# git clone git://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable.git
error: could not commit config file /eos/gerard/linux-stable/.git/config
error: could not commit config file /eos/gerard/linux-stable/.git/config
error: could not commit config file /eos/gerard/linux-stable/.git/config
Initialized empty Git repository in /eos/gerard/linux-stable/.git/
error: could not commit config file /eos/gerard/linux-stable/.git/config
error: could not commit config file /eos/gerard/linux-stable/.git/config
remote: Counting objects: 4611038, done.
remote: Compressing objects: 100% (365800/365800), done.
remote: Total 4611038 (delta 294560), reused 0 (delta 0)
Receiving objects: 100% (4611038/4611038), 1.05 GiB | 3.40 MiB/s, done.
Resolving deltas: 4% (160510/3796463)
```

Tried with another repository from another node and it shows the same errors:

```
[root@cmssrv151 ~]# cd /eos/
[root@cmssrv151 eos]# ls
cmseos-test.fnal.gov gerard
[root@cmssrv151 eos]# cd gerard/
[root@cmssrv151 gerard]# ls
linux-stable passwd passwd.link
[root@cmssrv151 gerard]# mkdir clone.from.151
[root@cmssrv151 gerard]# cd clone.from.151/
[root@cmssrv151 clone.from.151]# ls
[root@cmssrv151 clone.from.151]# git clone git://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable.git
fatal: could not create work tree dir 'linux-stable': Device or resource busy
[root@cmssrv151 clone.from.151]# ll
total 0
drwxrwxrwx 1 daemon daemon 0 Jul 17 12:40 linux-stable
[root@cmssrv151 clone.from.151]# pwd
/eos/gerard/clone.from.151
[root@cmssrv151 clone.from.151]# ll linux-stable/
total 0
[root@cmssrv151 clone.from.151]# rmdir linux-stable/
[root@cmssrv151 clone.from.151]# cp /etc/passwd .
[root@cmssrv151 clone.from.151]# wc -passwd
32 passwd
[root@cmssrv151 clone.from.151]# git clone git://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable.git
/eos/gerard/clone.from.151/linux-stable/.git/refs: Device or resource busy
[root@cmssrv151 clone.from.151]# mount
/dev/sda3 on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
```

```

/dev/sda1 on /boot type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
eosmain on /eos type fuse (rw,nosuid,nodev,allow_other)
[root@cmssrv151 clone.from.151]# rpm -q eos-client
eos-client-0.3.123-aquamarine.sle6.x86_64
[root@cmssrv151 clone.from.151]# rpm -q eos
package eos is not installed
[root@cmssrv151 clone.from.151]# rpm -q eos-server
eos-server-0.3.123-aquamarine.sle6.x86_64
[root@cmssrv151 clone.from.151]# git clone https://github.com/xrootd/xrootd.git
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
Initialized empty Git repository in /eos/gerard/clone.from.151/xrootd/.git/
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
remote: Counting objects: 46795, done.
remote: Compressing objects: 100% (51/51), done.
remote: Total 46795 (delta 22), reused 0 (delta 0), pack-reused 46744
Receiving objects: 100% (46795/46795), 27.02 MiB | 3.90 MiB/s, done.
Resolving deltas: 100% (35864/35864), done.
error: Unable to create /eos/gerard/clone.from.151/xrootd/.git/HEAD
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
error: could not commit config file /eos/gerard/clone.from.151/xrootd/.git/config
[root@cmssrv151 clone.from.151]# ll xrootd/
total 67
drwxrwxrwx 1 daemon daemon 1 Jul 17 12:42 bindings
drwxrwxrwx 1 daemon daemon 0 Jul 17 12:42 cmake
-rw-r--r- 1 daemon daemon 1702 Jul 17 12:42 CMakeLists.txt
-rw-r--r- 1 daemon daemon 35147 Jul 17 12:42 COPYING
-rw-r--r- 1 daemon daemon 2811 Jul 17 12:42 COPYING.BSD
-rw-r--r- 1 daemon daemon 7651 Jul 17 12:42 COPYING.LGPL
drwxrwxrwx 1 daemon daemon 1 Jul 17 12:42 docs
-rw-r--r- 1 daemon daemon 11119 Jul 17 12:42 Doxyfile
-rwxr-xr-x 1 daemon daemon 7621 Jul 17 12:42 genversion.sh
-rw-r--r- 1 daemon daemon 1432 Jul 17 12:42 LICENSE
drwxrwxrwx 1 daemon daemon 3 Jul 17 12:42 packaging
-rw-r--r- 1 daemon daemon 3268 Jul 17 12:42 README
drwxrwxrwx 1 daemon daemon 36 Jul 17 12:42 src
drwxrwxrwx 1 daemon daemon 3 Jul 17 12:42 tests
drwxrwxrwx 1 daemon daemon 0 Jul 17 12:42 utils
-rw-r--r- 1 daemon daemon 49 Jul 17 12:42 VERSION_INFO
[root@cmssrv151 clone.from.151]# cd xrootd/
[root@cmssrv151 xrootd]# git branch -a

```

- master
 - remotes/origin/HEAD -> origin/master
 - remotes/origin/master
 - remotes/origin/stable-3.3.6-x
 - remotes/origin/stable-3.3.x
 - remotes/origin/stable-4.0.x
 - remotes/origin/stable-4.1.x
 - remotes/origin/stable-4.1.x-cern
 - remotes/origin/stable-4.2.x
 - remotes/origin/xrdposixcl
 - remotes/origin/xrdssi
- ```

[root@cmssrv151 xrootd]# cd
[root@cmssrv151 ~]# git clone https://github.com/xrootd/xrootd.git
Initialized empty Git repository in /root/xrootd/.git/
remote: Counting objects: 46795, done.
remote: Compressing objects: 100% (51/51), done.
remote: Total 46795 (delta 22), reused 0 (delta 0), pack-reused 46744
Receiving objects: 100% (46795/46795), 27.02 MiB | 17.96 MiB/s, done.
Resolving deltas: 100% (35864/35864), done.
[root@cmssrv151 ~]#

```

#### #14 - 07/17/2015 12:55 PM - Gerard Bernabeu Altayo

I updated to the latest version and this errors went away (with 0.3.125):

```

[root@cmssrv151 ~]# cd /eos/gerard/
[root@cmssrv151 gerard]# ls
clone.from.151 linux-stable passwd passwd.link
[root@cmssrv151 gerard]# cd clone.from.151/
[root@cmssrv151 clone.from.151]# ls

```

```
passwd xrootd
[root@cmssrv151 clone.from.151]# rm -rf xrootd
[root@cmssrv151 clone.from.151]# git clone https://github.com/xrootd/xrootd.git
Initialized empty Git repository in /eos/gerard/clone.from.151/xrootd/.git/
remote: Counting objects: 46795, done.
remote: Compressing objects: 100% (51/51), done.
remote: Total 46795 (delta 22), reused 0 (delta 0), pack reused 46744
Receiving objects: 100% (46795/46795), 27.02 MiB | 3.60 MiB/s, done.
Resolving deltas: 100% (35864/35864), done.
[root@cmssrv151 clone.from.151]# ll
total 2
-rw-r--r- 1 daemon daemon 1630 Jul 17 12:40 passwd
drwxrwxrwx 1 daemon daemon 8 Jul 17 12:52 xrootd
```

Will try to compile xrootd, if works will go for linux kernel!

#### #15 - 07/17/2015 02:21 PM - Gerard Bernabeu Altayo

I've detected a few issues checking out a the Kernel GIT repo, opened <https://its.cern.ch/jira/browse/EOS-1196>

I also detect general slowness when compiling xrootd (over 2h to do it!), however this may not be a relevant usecase! Functionality wise it works :)

The 'make install' did not take that long:

```
[root@cmssrv151 clone.from.151]# time make install
real 0m46.817s
user 0m1.120s
sys 0m1.363s
```

#### #16 - 07/20/2015 11:56 AM - Gerard Bernabeu Altayo

- Status changed from Assigned to Resolved

eos test is functional, closing this ticket. Will do performance testing elsewhere.