

## EOS - Feature #11694

### Retire cmseos50-63 (will be added to dCache disk)

02/09/2016 01:31 PM - Gerard Bernabeu Altayo

<b>Status:</b>	Resolved	<b>Start date:</b>	02/09/2016
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Gerard Bernabeu Altayo	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>Description</b>			
As per David Fagan's requests:			
Gerard,			
Based on our brief conversation the right machines are from cmseos50 up, as you mentioned one at a time and see what happens.			
Please let me know how things go, I'll watch the net lot of this will be local to a switch but I don't know if the space is totally all over the space or specific to nodes..			
If I see a network top out I'll let you know, we are pushing but they aren't bouncing off the ceiling.			
thanks, David.			
This is similar to <a href="https://cdcvs.fnal.gov/redmine/issues/10763">https://cdcvs.fnal.gov/redmine/issues/10763</a>			

### History

#### #1 - 02/09/2016 01:36 PM - Gerard Bernabeu Altayo

Draining all this nodes means removing ~2PB of disk from EOS:

71\*2\*14  
1988

Following the drain documentation from <http://eos.readthedocs.org/en/latest/configuration/draining.html>.

```
[root@cmsrv222 ~]# for i in `seq 166 193`; do eos fs config $i configstatus=drain; done
```

```
[root@cmsrv222 ~]# eos fs ls
```

```
#.....  
#.....  
#          host (#...) #   id #           path #       schedgroup #       geotag #       boot # confi  
gstatus #       drain # active  
#.....  
#.....  
      cmseos12.fnal.gov (1095)   41   /storage/data3       default.2       booted  
rw      nodrain   online  
      cmseos12.fnal.gov (1095)   42   /storage/data1       default.3       booted  
rw      nodrain   online  
      cmseos12.fnal.gov (1095)   43   /storage/data2       default.1       booted  
rw      nodrain   online  
      cmseos11.fnal.gov (1095)   44   /storage/data1       default.0       booted  
rw      nodrain   online  
      cmseos11.fnal.gov (1095)   45   /storage/data2       default.3       booted  
rw      nodrain   online  
      cmseos11.fnal.gov (1095)   46   /storage/data3       default.2       booted  
rw      nodrain   online  
      cmseos13.fnal.gov (1095)   47   /storage/data1       default.0       booted  
rw      nodrain   online  
      cmseos13.fnal.gov (1095)   48   /storage/data2       default.1       booted  
rw      nodrain   online  
      cmseos13.fnal.gov (1095)   49   /storage/data3       default.2       booted  
rw      nodrain   online  
      cmseos14.fnal.gov (1095)   50   /storage/data1       default.0       booted  
rw      nodrain   online
```

	cmseos14.fnal.gov (1095)	51	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos14.fnal.gov (1095)	52	/storage/data3	default.3	booted
rw	nodrain online				
	cmseos15.fnal.gov (1095)	53	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos15.fnal.gov (1095)	54	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos15.fnal.gov (1095)	55	/storage/data3	default.3	booted
rw	nodrain online				
	cmseos16.fnal.gov (1095)	56	/storage/data1	default.3	booted
rw	nodrain online				
	cmseos16.fnal.gov (1095)	57	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos16.fnal.gov (1095)	58	/storage/data3	default.2	booted
rw	nodrain online				
	cmseos17.fnal.gov (1095)	59	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos17.fnal.gov (1095)	60	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos17.fnal.gov (1095)	61	/storage/data3	default.2	booted
rw	nodrain online				
	cmseos1.fnal.gov (1095)	62	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos1.fnal.gov (1095)	63	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos1.fnal.gov (1095)	64	/storage/data3	default.2	booted
rw	nodrain online				
	cmseos5.fnal.gov (1095)	65	/storage/data1	default.0	booted
drain	stalling online				
	cmseos5.fnal.gov (1095)	66	/storage/data2	default.1	booted
drain	stalling online				
	cmseos5.fnal.gov (1095)	67	/storage/data3	default.2	booted
drain	stalling online				
	cmseos4.fnal.gov (1095)	68	/storage/data1	default.0	booted
ro	nodrain online				
	cmseos4.fnal.gov (1095)	69	/storage/data2	default.1	booted
ro	nodrain online				
	cmseos4.fnal.gov (1095)	70	/storage/data3	default.2	booted
ro	nodrain online				
	cmseos3.fnal.gov (1095)	71	/storage/data1	default.0	booted
ro	nodrain online				
	cmseos3.fnal.gov (1095)	72	/storage/data2	default.2	booted
ro	nodrain online				
	cmseos3.fnal.gov (1095)	73	/storage/data3	default.1	booted
ro	nodrain online				
	cmseos2.fnal.gov (1095)	74	/storage/data1	default.0	booted
ro	nodrain online				
	cmseos2.fnal.gov (1095)	75	/storage/data2	default.1	booted
ro	nodrain online				
	cmseos2.fnal.gov (1095)	76	/storage/data3	default.3	booted
ro	nodrain online				
	cmseos25.fnal.gov (1095)	103	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos25.fnal.gov (1095)	104	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos26.fnal.gov (1095)	105	/storage/data1	default.2	booted
rw	nodrain online				
	cmseos26.fnal.gov (1095)	106	/storage/data2	default.3	booted
rw	nodrain online				
	cmseos27.fnal.gov (1095)	107	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos27.fnal.gov (1095)	108	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos27.fnal.gov (1095)	109	/storage/data3	default.2	booted
rw	nodrain online				
	cmseos28.fnal.gov (1095)	110	/storage/data1	default.0	booted
rw	nodrain online				
	cmseos28.fnal.gov (1095)	111	/storage/data2	default.1	booted
rw	nodrain online				
	cmseos29.fnal.gov (1095)	112	/storage/data1	default.2	booted
rw	nodrain online				
	cmseos29.fnal.gov (1095)	113	/storage/data2	default.3	booted
rw	nodrain online				
	cmseos30.fnal.gov (1095)	116	/storage/data1	default.0	booted

rw	nodrain	online	cmseos30.fnal.gov (1095)	117	/storage/data2	default.1	booted
rw	nodrain	online	cmseos31.fnal.gov (1095)	119	/storage/data1	default.2	booted
rw	nodrain	online	cmseos31.fnal.gov (1095)	120	/storage/data2	default.3	booted
rw	nodrain	online	cmseos32.fnal.gov (1095)	122	/storage/data1	default.0	booted
rw	nodrain	online	cmseos32.fnal.gov (1095)	123	/storage/data2	default.1	booted
rw	nodrain	online	cmseos33.fnal.gov (1095)	125	/storage/data1	default.2	booted
rw	nodrain	online	cmseos33.fnal.gov (1095)	126	/storage/data2	default.3	booted
rw	nodrain	online	cmseos39.fnal.gov (1095)	134	/storage/data1	default.2	booted
rw	nodrain	online	cmseos39.fnal.gov (1095)	135	/storage/data2	default.3	booted
rw	nodrain	online	cmseos34.fnal.gov (1095)	136	/storage/data1	default.0	booted
rw	nodrain	online	cmseos36.fnal.gov (1095)	137	/storage/data1	default.0	booted
rw	nodrain	online	cmseos36.fnal.gov (1095)	138	/storage/data2	default.1	booted
rw	nodrain	online	cmseos34.fnal.gov (1095)	139	/storage/data2	default.1	booted
rw	nodrain	online	cmseos38.fnal.gov (1095)	140	/storage/data1	default.0	booted
rw	nodrain	online	cmseos38.fnal.gov (1095)	141	/storage/data2	default.1	booted
rw	nodrain	online	cmseos37.fnal.gov (1095)	142	/storage/data1	default.2	booted
rw	nodrain	online	cmseos37.fnal.gov (1095)	143	/storage/data2	default.3	booted
rw	nodrain	online	cmseos35.fnal.gov (1095)	144	/storage/data1	default.2	booted
rw	nodrain	online	cmseos35.fnal.gov (1095)	145	/storage/data2	default.3	booted
rw	nodrain	online	cmseos42.fnal.gov (1095)	146	/storage/data1	default.0	booted
rw	nodrain	online	cmseos42.fnal.gov (1095)	147	/storage/data2	default.1	booted
rw	nodrain	online	cmseos40.fnal.gov (1095)	148	/storage/data1	default.0	booted
rw	nodrain	online	cmseos40.fnal.gov (1095)	149	/storage/data2	default.1	booted
rw	nodrain	online	cmseos41.fnal.gov (1095)	150	/storage/data1	default.2	booted
rw	nodrain	online	cmseos41.fnal.gov (1095)	151	/storage/data2	default.3	booted
rw	nodrain	online	cmseos43.fnal.gov (1095)	152	/storage/data1	default.2	booted
rw	nodrain	online	cmseos43.fnal.gov (1095)	153	/storage/data2	default.3	booted
rw	nodrain	online	cmseos44.fnal.gov (1095)	154	/storage/data1	default.0	booted
rw	nodrain	online	cmseos44.fnal.gov (1095)	155	/storage/data2	default.1	booted
rw	nodrain	online	cmseos45.fnal.gov (1095)	156	/storage/data1	default.2	booted
rw	nodrain	online	cmseos45.fnal.gov (1095)	157	/storage/data2	default.3	booted
rw	nodrain	online	cmseos46.fnal.gov (1095)	158	/storage/data1	default.0	booted
rw	nodrain	online	cmseos46.fnal.gov (1095)	159	/storage/data2	default.1	booted
rw	nodrain	online	cmseos47.fnal.gov (1095)	160	/storage/data1	default.2	booted
rw	nodrain	online	cmseos47.fnal.gov (1095)	161	/storage/data2	default.3	booted
rw	nodrain	online	cmseos49.fnal.gov (1095)	162	/storage/data1	default.2	booted
rw	nodrain	online	cmseos49.fnal.gov (1095)	163	/storage/data2	default.3	booted
rw	nodrain	online					

```

    cmseos48.fnal.gov (1095) 164 /storage/data1 default.0 booted
  rw    nodrain  online
    cmseos48.fnal.gov (1095) 165 /storage/data2 default.1 booted
  rw    nodrain  online
    cmseos50.fnal.gov (1095) 166 /storage/data1 default.0 booted
drain  prepare  online
    cmseos50.fnal.gov (1095) 167 /storage/data2 default.1 booted
drain  prepare  online
    cmseos51.fnal.gov (1095) 168 /storage/data1 default.2 booted
drain  prepare  online
    cmseos51.fnal.gov (1095) 169 /storage/data2 default.3 booted
drain  prepare  online
    cmseos52.fnal.gov (1095) 170 /storage/data1 default.0 booted
drain  prepare  online
    cmseos52.fnal.gov (1095) 171 /storage/data2 default.1 booted
drain  prepare  online
    cmseos53.fnal.gov (1095) 172 /storage/data1 default.2 booted
drain  prepare  online
    cmseos53.fnal.gov (1095) 173 /storage/data2 default.3 booted
drain  prepare  online
    cmseos56.fnal.gov (1095) 174 /storage/data1 default.0 booted
drain  prepare  online
    cmseos56.fnal.gov (1095) 175 /storage/data2 default.1 booted
drain  prepare  online
    cmseos54.fnal.gov (1095) 176 /storage/data1 default.0 booted
drain  prepare  online
    cmseos54.fnal.gov (1095) 177 /storage/data2 default.1 booted
drain  prepare  online
    cmseos55.fnal.gov (1095) 178 /storage/data1 default.2 booted
drain  prepare  online
    cmseos55.fnal.gov (1095) 179 /storage/data2 default.3 booted
drain  prepare  online
    cmseos57.fnal.gov (1095) 180 /storage/data1 default.2 booted
drain  prepare  online
    cmseos57.fnal.gov (1095) 181 /storage/data2 default.3 booted
drain  prepare  online
    cmseos58.fnal.gov (1095) 182 /storage/data1 default.0 booted
drain  prepare  online
    cmseos58.fnal.gov (1095) 183 /storage/data2 default.1 booted
drain  prepare  online
    cmseos59.fnal.gov (1095) 184 /storage/data1 default.2 booted
drain  prepare  online
    cmseos59.fnal.gov (1095) 185 /storage/data2 default.3 booted
drain  prepare  online
    cmseos60.fnal.gov (1095) 186 /storage/data1 default.0 booted
drain  prepare  online
    cmseos60.fnal.gov (1095) 187 /storage/data2 default.1 booted
drain  prepare  online
    cmseos61.fnal.gov (1095) 188 /storage/data1 default.2 booted
drain  prepare  online
    cmseos61.fnal.gov (1095) 189 /storage/data2 default.3 booted
drain  prepare  online
    cmseos62.fnal.gov (1095) 190 /storage/data1 default.0 booted
drain  prepare  online
    cmseos62.fnal.gov (1095) 191 /storage/data2 default.1 booted
drain  prepare  online
    cmseos63.fnal.gov (1095) 192 /storage/data1 default.2 booted
drain  prepare  online
    cmseos63.fnal.gov (1095) 193 /storage/data2 default.3 booted
drain  prepare  online
[root@cmsrv222 ~]#

```

## #2 - 02/16/2016 04:04 PM - Gerard Bernabeu Altayo

So the drains are not working very well, will now set all this disks as Read Only and (re)start the drain only for one of the pools:

```

[root@cmsrv222 ~]# eos fs ls -d
#.....
#.....
#          host (#...) #   id #          path #   drain #   progress #   files # bytes-left #
timeleft #retry #wopen
#.....
#.....
    cmseos5.fnal.gov (1095) 65 /storage/data1   expired           0       188.00   284.18 MB 99

```

999999999	0	0								
	cmseos5.fnal.gov	(1095)	66	/storage/data2	expired	0	249.00	3.49 GB	99	
999999999	0	0								
	cmseos5.fnal.gov	(1095)	67	/storage/data3	expired	0	239.00	2.84 GB	99	
999999999	0	0								
	cmseos4.fnal.gov	(1095)	68	/storage/data1	expired	0	69.72 k	4.10 TB	99	
999999999	0	0								
	cmseos4.fnal.gov	(1095)	69	/storage/data2	expired	0	68.17 k	3.98 TB	99	
999999999	0	0								
	cmseos4.fnal.gov	(1095)	70	/storage/data3	expired	0	74.04 k	4.40 TB	99	
999999999	0	0								
	cmseos3.fnal.gov	(1095)	71	/storage/data1	expired	0	69.38 k	4.03 TB	99	
999999999	0	0								
	cmseos3.fnal.gov	(1095)	72	/storage/data2	expired	0	103.29 k	6.27 TB	99	
999999999	0	0								
	cmseos3.fnal.gov	(1095)	73	/storage/data3	expired	0	65.91 k	4.02 TB	99	
999999999	0	0								
	cmseos2.fnal.gov	(1095)	74	/storage/data1	expired	0	98.74 k	6.05 TB	99	
999999999	0	0								
	cmseos2.fnal.gov	(1095)	75	/storage/data2	expired	0	98.90 k	6.06 TB	99	
999999999	0	0								
	cmseos2.fnal.gov	(1095)	76	/storage/data3	expired	0	316.00	281.15 MB	99	
999999999	0	0								
	cmseos50.fnal.gov	(1095)	166	/storage/data1	draining	0	427.76 k	34.65 TB		
4726	0	6								
	cmseos50.fnal.gov	(1095)	167	/storage/data2	draining	0	420.56 k	34.18 TB		
4726	0	2								
	cmseos51.fnal.gov	(1095)	168	/storage/data1	draining	0	438.24 k	35.81 TB		
4726	0	5								
	cmseos51.fnal.gov	(1095)	169	/storage/data2	stalling	0	1.00	965.53 MB		
4726	0	5								
	cmseos52.fnal.gov	(1095)	170	/storage/data1	stalling	0	437.25 k	34.83 TB		
4726	0	4								
	cmseos52.fnal.gov	(1095)	171	/storage/data2	stalling	0	418.86 k	34.22 TB		
4725	0	4								
	cmseos53.fnal.gov	(1095)	172	/storage/data1	stalling	0	440.64 k	35.84 TB		
4726	0	2								
	cmseos53.fnal.gov	(1095)	173	/storage/data2	stalling	0	467.83 k	37.93 TB		
4726	0	2								
	cmseos56.fnal.gov	(1095)	174	/storage/data1	stalling	0	435.39 k	34.92 TB		
4725	0	4								
	cmseos56.fnal.gov	(1095)	175	/storage/data2	stalling	0	418.34 k	34.62 TB		
4726	0	6								
	cmseos54.fnal.gov	(1095)	176	/storage/data1	stalling	0	427.91 k	35.75 TB		
4726	0	6								
	cmseos54.fnal.gov	(1095)	177	/storage/data2	stalling	0	422.56 k	34.38 TB		
4727	0	3								
	cmseos55.fnal.gov	(1095)	178	/storage/data1	stalling	0	436.75 k	35.99 TB		
4727	0	0								
	cmseos55.fnal.gov	(1095)	179	/storage/data2	stalling	0	462.48 k	37.66 TB		
4727	0	0								
	cmseos57.fnal.gov	(1095)	180	/storage/data1	stalling	0	437.23 k	36.10 TB		
4727	0	9								
	cmseos57.fnal.gov	(1095)	181	/storage/data2	stalling	0	460.50 k	37.95 TB		
4727	0	6								
	cmseos58.fnal.gov	(1095)	182	/storage/data1	stalling	0	431.55 k	35.09 TB		
4727	0	4								
	cmseos58.fnal.gov	(1095)	183	/storage/data2	stalling	0	418.25 k	34.73 TB		
4727	0	9								
	cmseos59.fnal.gov	(1095)	184	/storage/data1	stalling	0	442.29 k	36.11 TB		
4727	0	9								
	cmseos59.fnal.gov	(1095)	185	/storage/data2	stalling	0	466.13 k	37.85 TB		
4727	0	4								
	cmseos60.fnal.gov	(1095)	186	/storage/data1	stalling	0	434.35 k	35.36 TB		
4727	0	4								
	cmseos60.fnal.gov	(1095)	187	/storage/data2	stalling	0	415.03 k	34.95 TB		
4727	0	4								
	cmseos61.fnal.gov	(1095)	188	/storage/data1	stalling	0	434.57 k	35.97 TB		
4727	0	8								
	cmseos61.fnal.gov	(1095)	189	/storage/data2	stalling	0	470.36 k	37.45 TB		
4728	0	7								
	cmseos62.fnal.gov	(1095)	190	/storage/data1	stalling	0	431.21 k	35.71 TB		
4728	0	5								
	cmseos62.fnal.gov	(1095)	191	/storage/data2	stalling	0	414.52 k	34.50 TB		
4728	0	1								

```

cmseos63.fnal.gov (1095) 192 /storage/data1 stalling 0 435.64 k 35.83 TB
69876 0 0
cmseos63.fnal.gov (1095) 193 /storage/data2 stalling 0 460.88 k 37.47 TB
69876 0 0
[root@cmsrv222 ~]# for i in `seq 65 76` `seq 166 193`; do eos fs config $i configstatus=ro; done
[root@cmsrv222 ~]# for i in 168; do eos fs config $i configstatus=drain; done
[root@cmsrv222 ~]#
[root@cmsrv222 ~]# eos fs ls -d

#.....
#.....
# host (#...) # id # path # drain # progress # files # bytes-left #
timeleft #retry #wopen
#.....
#.....
cmseos51.fnal.gov (1095) 168 /storage/data1 draining 0 438.09 k 35.71 TB
86276 0 5
[root@cmsrv222 ~]#

```

And the good news is that this seems to be working now! Data is coming down and the box is draining at ~500MB/s!

```

[root@cmsrv222 ~]# eos fs ls -d

#.....
#.....
# host (#...) # id # path # drain # progress # files # bytes-left #
timeleft #retry #wopen
#.....
#.....
cmseos51.fnal.gov (1095) 168 /storage/data1 draining 0 436.43 k 35.60 TB
85785 0 5
[root@cmsrv222 ~]#

```

My new theory is that drains stall whenever they find files in error state, which means I'll have to go dig out all what's wrong in the FS:

```

[root@cmsrv222 ~]# eos fsck stat
160216 15:35:32 1455658532.284256 started check
160216 15:35:32 1455658532.284286 Filesystems to check: 115
160216 15:35:42 1455658542.322971 d_cx_diff : 14 (14)
160216 15:35:42 1455658542.322992 d_mem_sz_diff : 13729 (16516)
160216 15:35:42 1455658542.322997 orphans_n : 2 (2)
160216 15:35:42 1455658542.323000 rep_diff_n : 251 (298)
160216 15:35:42 1455658542.323005 rep_offline : 0 (0)
160216 15:35:42 1455658542.323009 unreg_n : 3 (3)
160216 15:35:42 1455658542.323013 zero_replica : 44 (44)
160216 15:35:43 1455658543.937750 stopping check
160216 15:35:43 1455658543.937786 => next run in 30 minutes
[root@cmsrv222 ~]#

```

### #3 - 02/16/2016 05:06 PM - Gerard Bernabeu Altayo

I've noticed that the migration speed with a single FST was gradually decreasing speed. Since there are sooo many small files, I've increased the parallelism per FST migration from 2 to 20 transfers at a time, this boosted the transfer bandwidth back up :)

### #4 - 02/16/2016 05:25 PM - Gerard Bernabeu Altayo

With 20 the load increased very high in many FSTs, lowered it to 5...

### #5 - 02/17/2016 09:47 AM - Gerard Bernabeu Altayo

The node is almost drained:

```

[root@cmsrv222 ~]# eos fs ls -d

#.....
#.....
# host (#...) # id # path # drain # progress # files # bytes-left #
timeleft #retry #wopen
#.....
#.....
cmseos51.fnal.gov (1095) 168 /storage/data1 draining 99 98.00 2.50 GB
22028 0 5
[root@cmsrv222 ~]#

```

Starting another node migration:

```
[root@cmsrv222 ~]# for i in 166 167 169; do eos fs config $i configstatus=drain; done
```

#### #6 - 02/22/2016 01:31 PM - Gerard Bernabeu Altayo

The drain is working pretty well going FST by FST (it will just take a while), starting a new one and this afternoon I'll retire 2 nodes:

```
[root@cmsrv222 ~]# eos fs ls -d
```

```
#.....
#.....
#          host (#...) #    id #          path #          drain #    progress #          files # bytes-left #
timeleft #retry #wopen
#.....
.....
          cmseos50.fnal.gov (1095)    166    /storage/data1    expired          99          21.00    396.48 MB 99
999999999 0    6
          cmseos50.fnal.gov (1095)    167    /storage/data2    expired          99          6.00     1.75 GB 99
999999999 0    2
          cmseos51.fnal.gov (1095)    168    /storage/data1    expired          99          30.00    596.66 MB 99
999999999 0    5
          cmseos51.fnal.gov (1095)    169    /storage/data2    expired          0          1.00    965.53 MB 99
999999999 0    5
[root@cmsrv222 ~]# for i in 170 171; do eos fs config $i configstatus=drain; done
[root@cmsrv222 ~]#
```

In order to clear the running transfers, restarting cmseos50 & cmseos51 and then I'll look at each individual FST to clear them up.

#### #7 - 02/23/2016 05:31 PM - Gerard Bernabeu Altayo

Nodes keep draining, now it is time to look at the few remaining files and solve the individual issues, for example:

```
[root@cmsrv222 ~]# for i in `eos fs dumpmd 166 -path | cut -d= -f2`; do eos file info $i; done
File: '/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_24/cache/quant-toys:0-wp3000-c04209
a4ad.db' Flags: 0644
Size: 61440
Modify: Wed Oct 14 05:54:21 2015 Timestamp: 1444820061.136974000
Change: Wed Oct 14 05:54:02 2015 Timestamp: 1444820042.80888040
CUid: 47452 CGid: 5063 Fxid: 081c8ce6 Fid: 136088806 Pid: 2434882 Pxid: 00252742
XStype: adler XS: 32 43 c4 88 ETAG: 36531060695105536:3243c488
replica Stripes: 2 Blocksize: 4k LayoutId: 00100112
#Rep: 2
# fs-id #.....
#.....
#          host #          schedgroup #          path #          boot # configstatus #          drain
# active #          geotag
#.....
0    166    cmseos50.fnal.gov    default.0    /storage/data1    booted          ro    nodrain
online
1    136    cmseos34.fnal.gov    default.0    /storage/data1    booted          rw    nodrain
online
*****
File: '/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_MT_3fb/Muon_MT_3fb_16/cache/nll_der-toys:0.0-wp3000-04
f3872ac8.db' Flags: 0644
Size: 693248
Modify: Wed Oct 14 05:55:31 2015 Timestamp: 1444820131.970225000
Change: Wed Oct 14 05:55:28 2015 Timestamp: 1444820128.408139999
CUid: 47452 CGid: 5063 Fxid: 081ca568 Fid: 136095080 Pid: 2435048 Pxid: 002527e8
XStype: adler XS: 18 26 2b 73 ETAG: 36532744859156480:18262b73
replica Stripes: 2 Blocksize: 4k LayoutId: 00100112
#Rep: 2
# fs-id #.....
#.....
#          host #          schedgroup #          path #          boot # configstatus #          drain
# active #          geotag
#.....
0    166    cmseos50.fnal.gov    default.0    /storage/data1    booted          ro    nodrain
online
1    158    cmseos46.fnal.gov    default.0    /storage/data1    booted          rw    nodrain
online
*****
.... <OUTPUT CUT>
```

#### #8 - 02/24/2016 02:41 PM - Gerard Bernabeu Altayo

I've managed to empty one of the FSTs:

```
[root@cmsrv222 ~]# for i in `eos fs dumpmd 166 -path | cut -d= -f2`; do eos file move $i 166 45; done
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
success: scheduled move from source fs=166 => target fs=45
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_21/nll_der-toys:0.0-wp2000-ee81029424.cfg
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_MT_3fb/electron_gelbtags_IDMTJet250Lep150MET150_discovery.json
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_MT_3fb/electron_gelbtags_IDMTJet250Lep150MET150_discovery.json
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Muon_ML_3fb/Muon_ML_3fb_20/model_summary_general.thtml
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]# eos fs dumpmd 166 -path
path=/eos/uscms/store/user/drinkin/Theta_Oct13/Elec_ML_3fb/Elec_ML_3fb_16/quant-toys:0-wp3000-4160ae2cb8.cfg
[root@cmsrv222 ~]#
```

#### #9 - 02/24/2016 03:05 PM - Gerard Bernabeu Altayo

now just following [https://cmsweb.fnal.gov/bin/view/Storage/EOSOperationalProcedures#Decommission\\_an\\_EOS\\_FST\\_node](https://cmsweb.fnal.gov/bin/view/Storage/EOSOperationalProcedures#Decommission_an_EOS_FST_node)

Will do cmseos50 first:

```
-bash-4.1$ /srv/admin/bin/cis-please-retire cmseos50
you (DCSO) are in charge of removing the host from zabbix and the ENC
Primary Ticket Information
Number:            INC000000669206
Summary:          Retire host: cmseos50
Status:           Assigned
Submitted:        2016-02-24 14:59:48 CST
Urgency:          4 - Low
Priority:          3 - Medium
Service Type:     Server

Requestor Info
Name:             Gerard Bernabeu Altayo
Email:           gerardl@fnal.gov
Created By:       cd-srv-cms-snow

Assignee Info
Group:           ECF-CIS
Name:            (none)
Last Modified:   2016-02-24 14:59:53 CST

User-Provided Description
Please retire host: cmseos50
-bash-4.1$
```

#### #10 - 02/24/2016 03:12 PM - Gerard Bernabeu Altayo

Now doing cmseos51, it has some 'sticky' files:



```

[root@cmsrv222 ~]# for i in `eos fs dumpmd 168 -path | cut -d= -f2`; do eos file move $i 168 46; done
success: scheduled move from source fs=168 => target fs=46
[root@cmsrv222 ~]# eos file info /eos/uscms/store/user/kreis/Hbb/20150608_wide_fa3zz_nohup/log_Wh_CMS_float_0
bserverd_minuit_0p01_tries10_strategy2.log
File: '/eos/uscms/store/user/kreis/Hbb/20150608_wide_fa3zz_nohup/log_Wh_CMS_float_observed_minuit_0p01_tries
10_strategy2.log' Flags: 0644
Size: 6602829
Modify: Tue Jun 9 14:42:16 2015 Timestamp: 1433878936.112778000
Change: Tue Jun 9 13:58:57 2015 Timestamp: 1433876337.474202243
Cuid: 44456 CGid: 5063 Fxid: 0608092a Fid: 101189930 Pid: 1895919 Pxid: 001cedef
XStype: adler XS: 4e 20 7e 0d ETAG: 27162965002158080:4e207e0d
replica Stripes: 2 Blocksize: 4k LayoutId: 00100112
#Rep: 2
# fs-id #.....
.....
# active # host # schedgroup # path # boot # configstatus # drain
# geotag
#.....
0 61 cmseos17.fnal.gov default.2 /storage/data3 booted rw nodrain
online
1 168 cmseos51.fnal.gov default.2 /storage/data1 booted ro nodrain
online
*****
[root@cmsrv222 ~]# eos file drop /eos/uscms/store/user/kreis/Hbb/20150608_wide_fa3zz_nohup/log_Wh_CMS_float_0
bserverd_minuit_0p01_tries10_strategy2.log 168
success: dropped stripe on fs=168
[root@cmsrv222 ~]#

[root@cmseos51 ~]# for i in $fsids; do ssh $mgm eos fs dumpmd $i -path; done
path=/eos/uscms/store/user/kreis/Hbb/20150608_wide_fa3zz_nohup/log_Wh_CMS_float_observed_minuit_0p01_tries10_s
strategy2.log
[root@cmseos51 ~]# for i in $fsids; do ssh $mgm eos fs dumpmd $i -path; done
[root@cmseos51 ~]#

```

So now I can move forward with standard procedure:

```

[root@cmseos51 ~]# fsids=`ssh $mgm eos fs ls -m $HOSTNAME | awk '{print $3}' | grep id= | cut -d= -f2`
[root@cmseos51 ~]# for i in $fsids; do
> ssh $mgm eos fs rm $i
> done
error: you can only remove file systems which are in 'empty' status (errc=22) (Invalid argument)
error: you can only remove file systems which are in 'empty' status (errc=22) (Invalid argument)
[root@cmseos51 ~]# ssh $mgm eos vid remove gateway ${HOSTNAME}
success: rm vid [ eos.rgid=0 eos.ruid=0 mgm.cmd=vid mgm.subcmd=rm mgm.vid.cmd=unmap mgm.vid.key=tident:"*@cms
eos51.fnal.gov":uid]
success: rm vid [ eos.rgid=0 eos.ruid=0 mgm.cmd=vid mgm.subcmd=rm mgm.vid.cmd=unmap mgm.vid.key=tident:"*@cms
eos51.fnal.gov":gid]
[root@cmseos51 ~]# puppet agent --disable 'This node needs to be reshoot'
[root@cmseos51 ~]# service eosd stop
Stopping eosd:
[ OK ]

[root@cmseos51 ~]# service eos stop
Stopping xrootd: fst
[ OK ]

[root@cmseos51 ~]# chkconfig eos off
[root@cmseos51 ~]# chkconfig eosd off
[root@cmseos51 ~]# chkconfig eos-gridftp off
[root@cmseos51 ~]# sleep 5
[root@cmseos51 ~]# ssh $mgm eos node rm ${HOSTNAME}:1095
error: unable to remove node '/eos/cmseos51.fnal.gov:1095/fst' - filesystems are not all in empty state - try
to drain them or: node config <name> configstatus=empty
(errc=16) (Device or resource busy)
[root@cmseos51 ~]# fsids=`ssh $mgm eos fs ls -m $HOSTNAME | awk '{print $3}' | grep id= | cut -d= -f2`
[root@cmseos51 ~]# for i in $fsids; do
> ssh $mgm eos fs config $i configstatus=empty
> done
[root@cmseos51 ~]# ssh $mgm eos node rm ${HOSTNAME}:1095
success: removed node '/eos/cmseos51.fnal.gov:1095/fst'
[root@cmseos51 ~]#
[root@cmseos51 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2       40G   6.1G   32G   17% /
tmpfs           32G     0    32G    0% /dev/shm
/dev/sda1      1008M  138M   820M   15% /boot

```

```
/dev/sda5      857G   72M  814G   1% /storage/local/data1
/dev/sdb       71T   169M   71T   1% /storage/data1
/dev/sdc       71T   213M   71T   1% /storage/data2
[root@cmseos51 ~]# poweroff
```

```
Broadcast message from root@cmseos51.fnal.gov
(/dev/pts/0) at 15:09 ...
```

```
The system is going down for power off NOW!
[root@cmseos51 ~]# Connection to cmseos51 closed by remote host.
Connection to cmseos51 closed.
-bash-4.1$
```

```
-bash-4.1$ /srv/admin/bin/cis-please-retire cmseos51
you (DCSO) are in charge of removing the host from zabbix and the ENC
```

#### Primary Ticket Information

```
Number:          INC000000669215
Summary:         Retire host: cmseos51
Status:          Assigned
Submitted:       2016-02-24 15:10:51 CST
Urgency:         4 - Low
Priority:         3 - Medium
Service Type:    Server
```

#### Requestor Info

```
Name:           Gerard Bernabeu Altayo
Email:          gerard1@fnal.gov
Created By:     cd-srv-cms-snow
```

#### Assignee Info

```
Group:          ECF-CIS
Name:           (none)
Last Modified:  2016-02-24 15:10:56 CST
```

#### User-Provided Description

```
Please retire host: cmseos51
-bash-4.1$
```

```
comp-4:hosts gerard1$ HOST=cmseos51; git rm $HOST.fnal.gov.yaml; git commit -m "retiring $HOST"; git push
rm 'hosts/cmseos51.fnal.gov.yaml'
[master 228e747] retiring cmseos51
 1 file changed, 12 deletions(-)
 delete mode 100644 hosts/cmseos51.fnal.gov.yaml
Counting objects: 7, done.
Delta compression using up to 4 threads.
Compressing objects: 100% (3/3), done.
Writing objects: 100% (3/3), 296 bytes | 0 bytes/s, done.
Total 3 (delta 2), reused 0 (delta 0)
remote:
remote: diff-tree:
remote: :100644 000000 23e5d116ff7eff80ed1f02a4204375ed5824b282 0000000000000000000000000000000000000000 D
hosts/cmseos51.fnal.gov.yaml
remote: fatal: Path 'hosts/cmseos51.fnal.gov.yaml' does not exist in '228e7477063198cc7ab6764e1e92c4bcc35c0886
'
remote: omd-host-crud update cmseos51 --role UNSET --instance UNSET --extra 'UNSET'
remote: cmseos51 - host updated, will now inventory...
remote: cmseos51 - unknown error
remote: omd-host-crud delete cmseos51
remote: cmseos51 - host deleted
remote: Recieved from stdin:
remote: oldrev: 218f475579d7028a4a2808ce0064c1ca708b7a3b
remote: newrev: 228e7477063198cc7ab6764e1e92c4bcc35c0886
remote: refname: refs/heads/master
remote: Derived Configuration:
remote: REPO: puppet@cms-git:/var/lib/puppet/enc.git
remote: BRANCH: master
remote: BRANCH_DIR: /srv/puppet/enc
remote: PUPPET_SERVERS: puppet@cmssrv166.fnal.gov puppet@cmspuppet2.fnal.gov puppet@cmspuppet1.fnal.gov
remote: Updating remote branch /srv/puppet/enc/master on puppet@cmssrv166.fnal.gov
remote: From cms-git:/var/lib/puppet/enc
remote: * branch          master      -> FETCH_HEAD
remote: Updating 218f475..228e747
remote: Fast-forward
remote:  hosts/cmseos51.fnal.gov.yaml | 12 -----
remote:  1 files changed, 0 insertions(+), 12 deletions(-)
```

```

remote: delete mode 100644 hosts/cmseos51.fnal.gov.yaml
remote: Updating remote branch /srv/puppet/enc/master on puppet@cmspuppet2.fnal.gov
remote: From cms-git:/var/lib/puppet/enc
remote: * branch                master      -> FETCH_HEAD
remote: Updating 218f475..228e747
remote: Fast-forward
remote: hosts/cmseos51.fnal.gov.yaml | 12 -----
remote: 1 files changed, 0 insertions(+), 12 deletions(-)
remote: delete mode 100644 hosts/cmseos51.fnal.gov.yaml
remote: Updating remote branch /srv/puppet/enc/master on puppet@cmspuppet1.fnal.gov
remote: From cms-git:/var/lib/puppet/enc
remote: * branch                master      -> FETCH_HEAD
remote: Updating 218f475..228e747
remote: Fast-forward
remote: hosts/cmseos51.fnal.gov.yaml | 12 -----
remote: 1 files changed, 0 insertions(+), 12 deletions(-)
remote: delete mode 100644 hosts/cmseos51.fnal.gov.yaml
To puppet@cms-git.fnal.gov:enc
 218f475..228e747 master -> master
comp-4:hosts gerardl$

```

**#11 - 02/24/2016 03:27 PM - Gerard Bernabeu Altayo**

cmseos52 retired too.

**#12 - 02/24/2016 03:56 PM - Gerard Bernabeu Altayo**

cmseos53 retired, cmseos56 started drain :)

**#13 - 03/07/2016 04:55 PM - Gerard Bernabeu Altayo**

cmseos56 retired :)

All the others keep draining...

**#14 - 03/11/2016 04:44 PM - Gerard Bernabeu Altayo**

only 2 more nodes to go :D

[root@cmssrv222 ~]# eos fs ls -d

#.....

1. host (#...)	# id #	path #	drain #	progress #	files #	bytes-left #	timeleft #	retry #	wopen
cmseos61.fnal.gov (1095)	188	/storage/data1	drained	100	0.00	0.00 B	0	0	0
cmseos61.fnal.gov (1095)	189	/storage/data2	draining	85	59.71 k	5.67 TB	587624	0	0
cmseos63.fnal.gov (1095)	192	/storage/data1	draining	25	283.88 k	26.64 TB	587624	0	0
cmseos63.fnal.gov (1095)	193	/storage/data2	stalling	2	391.70 k	35.78 TB	587625	0	0

[root@cmssrv222 ~]#

**#15 - 03/14/2016 12:00 PM - Gerard Bernabeu Altayo**

- Status changed from New to Resolved

all nodes have been retired!