

## CMS dCache - Support #9896

### Upgrade dCache-tape to dcache-2.2.29-1

08/25/2015 04:09 PM - Gerard Bernabeu Altayo

<b>Status:</b>	Resolved	<b>Start date:</b>	08/25/2015
<b>Priority:</b>	Normal	<b>Due date:</b>	08/26/2015
<b>Assignee:</b>	Gerard Bernabeu Altayo	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	4.00 hours
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>component:</b>	base	<b>Co-Assignees:</b>	
<b>Stakeholders:</b>			

#### Description

This ticket will contain the instructions and logbook for:

- move to openjdk
- the dCache-2.2.29-1 minor version upgrade on dCache TAPE instance.
- update kernel and reboot

All we need to do is:

1. Upgrade the repo in [http://cms-install.fnal.gov/cobbler/repo\\_mirror/uscmst1-el6-x86\\_64](http://cms-install.fnal.gov/cobbler/repo_mirror/uscmst1-el6-x86_64)
2. Standard 'yum clean all; yum update -y; reboot'

#### History

##### #1 - 08/25/2015 04:23 PM - Gerard Bernabeu Altayo

I've done the 'offline' steps to upgrade the RPM:

```
ssh -lroot cmsinstall.fnal.gov
/srv/repo/sl6-x86_64
[root@cmssrv201 sl6-x86_64]# ls dcache-2.2.*
dcache-2.2.17-1.noarch.rpm dcache-2.2.23-1.noarch.rpm dcache-2.2.28-1.noarch.rpm
[root@cmssrv201 sl6-x86_64]# wget --no-check-certificate https://srm.fnal.gov/wiki/pub/DcacheCorner/CmsDcache/dcache-2.2.29-1.noarch.rpm
```

##### #2 - 08/25/2015 05:03 PM - Gerard Bernabeu Altayo

I have declared the downtime for tomorrow in check\_mk already.

Got the list of Tape Servers and Pools by doing:

```
ssh cmsadmin1.fnal.gov
-bash-4.1$ cd GIT/CMS/enc/hosts
-bash-4.1$ git up
-bash-4.1$ grep role::dcache::pool::tape: * | cut -d. -f1,2,3 > ~/tape-pool-nodes
-bash-4.1$ echo -e 'cmsdcacheadmin.fnal.gov\ncmssrmtmp.fnal.gov\ncmschimera.fnal.gov\ncmschimera.backup.fnal.gov' > ~/tape-server-nodes #Some day we will get the dcache servers in their own ENC category and be able to g rep...
```

Also running the Kernel update today so that it's all faster tomorrow:

```
-bash-4.1$ pssh -h tape-server-nodes -l root -t 0 -p 50 -o upgr_srvs.log -e upgr_srvs.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; yum update -y; rpm -q kernel-${kernelversion}'
-bash-4.1$ pssh -h tape-pool-nodes -l root -t 0 -p 50 -o upgr_pools.log -e upgr_pools.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; yum update -y; rpm -q kernel-${kernelversion}'
```

Tomorrow I will have to:

```
root cmsinstall
cd /srv/repo
make uscmst1
exit
```

Follow [https://cmsweb.fnal.gov/bin/view/Storage/DCache22Procedures#Rebooting\\_whole\\_dCache\\_instance](https://cmsweb.fnal.gov/bin/view/Storage/DCache22Procedures#Rebooting_whole_dCache_instance)

The adapted commands are:

```
ssh cmsadmin1
cd ~
pssh -h tape-server-nodes -l root -t 0 -p 50 -o upgr_srvs.log -e upgr_srvs.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; extra_rpm_check="dcache-2.2.29-1.noarch"; puppet agent --disable; service dcache-server stop; umount -l /pnfs; yum clean all; yum update -y; rpm -q kernel-${kernelversion} ${extra_rpm_check} && puppet agent --enable && reboot'
```

```
pssh -h tape-server-nodes -l root -t 60 -p 50 -o check_srvs.log -e check_srvs.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion && (df /pnfs; service dcache-server status | grep -v DOMAIN | grep -v running)'
```

```
pssh -h tape-pool-nodes -l root -t 0 -p 50 -o upgr_pools.log -e upgr_pools.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; extra_rpm_check="dcache-2.2.29-1.noarch"; puppet agent --disable 'rebooting'; service dcache-server stop; umount -l /pnfs; yum clean all; yum update -y; rpm -q kernel-${kernelversion} ${extra_rpm_check} && puppet agent --enable && reboot'
```

```
pssh -h tape-pool-nodes -l root -t 60 -p 50 -o check_pools.log -e check_pools.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion && (df /pnfs; service dcache-server status | grep -v DOMAIN | grep -v running)'
```

### #3 - 08/26/2015 01:23 PM - Gerard Bernabeu Altayo

Starting the downtime...

I've just updated the repo and told Natalia (I'm in her office), now starting the outage below...

```
bash-4.1$ date
Wed Aug 26 13:03:03 CDT 2015
-bash-4.1$ pssh -h tape-server-nodes -l root -t 0 -p 50 -o upgr_srvs.log -e upgr_srvs.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; extra_rpm_check="dcache-2.2.29-1.noarch"; puppet agent --disable; service dcache-server stop; umount -l /pnfs; yum clean all; yum update -y; rpm -q kernel-${kernelversion} ${extra_rpm_check} && puppet agent --enable && reboot'
[1] 13:03:49 [SUCCESS] cmschimera.backup.fnal.gov
[2] 13:03:52 [FAILURE] cmssrmtmp.fnal.gov Exited with error code 1
[3] 13:03:56 [SUCCESS] cmschimera.fnal.gov
[4] 13:03:58 [SUCCESS] cmsdcacheadmin.fnal.gov
-bash-4.1$
```

Fixed cmssrmtmp by doing 'yum remove dcache; yum install dcache', it was running the 'SNAPSHOT' version... This was properly detected during testing by Natalia and I was aware, so easy to fix :)

It took a while for cmschimera to boot...

```
-bash-4.1$ date; pssh -h tape-server-nodes -l root -t 60 -p 50 -o check_srvs.log -e check_srvs.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion && (df /pnfs; service dcache-server status | grep -v DOMAIN | grep -v running) '
Wed Aug 26 13:10:06 CDT 2015
[1] 13:10:08 [SUCCESS] cmschimera.backup.fnal.gov
[2] 13:10:08 [SUCCESS] cmssrmtmp.fnal.gov
[3] 13:10:08 [FAILURE] cmschimera.fnal.gov Exited with error code 1
[4] 13:10:09 [FAILURE] cmsdcacheadmin.fnal.gov Exited with error code 1
-bash-4.1$
```

Checking this 2 FAILED nodes... It looks it's all good but our check is not... Some more info to debug later:

```
[root@cmschimera ~]# uptime; uname -a
13:10:57 up 3 min, 1 user, load average: 1.87, 0.84, 0.31
Linux cmschimera.fnal.gov 2.6.32-573.3.1.el6.x86_64 #1 SMP Thu Aug 13 12:55:33 CDT 2015 x86_64 x86_64 x86_64 GNU/Linux
[root@cmschimera ~]# kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion
13:11:03 up 3 min, 1 user, load average: 2.11, 0.93, 0.34
Linux cmschimera.fnal.gov 2.6.32-573.3.1.el6.x86_64 #1 SMP Thu Aug 13 12:55:33 CDT 2015 x86_64 x86_64 x86_64 GNU/Linux
[root@cmschimera ~]# (df /pnfs; service dcache-server status | grep -v DOMAIN | grep -v running)
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/sda3        270681640 150988740 105936332  59% /
[root@cmschimera ~]# service dcache-server status
DOMAIN          STATUS  PID  USER
namespaceDomain running 8827 root
nfsDomain       running 8923 root
[root@cmschimera ~]# kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion && (df /pnfs; service dcache-server status | grep -v DOMAIN | grep -v running)
13:11:34 up 4 min, 1 user, load average: 2.01, 1.03, 0.39
Linux cmschimera.fnal.gov 2.6.32-573.3.1.el6.x86_64 #1 SMP Thu Aug 13 12:55:33 CDT 2015 x86_64 x86_64 x86_64 GNU/Linux
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/sda3        270681640 150988740 105936332  59% /
```

```
[root@cmschimera ~]# echo $?
1
[root@cmschimera ~]# service dcache-server status
DOMAIN          STATUS PID USER
namespaceDomain running 8827 root
nfsDomain       running 8923 root
[root@cmschimera ~]# rpm -q dcache
dcache-2.2.29-1.noarch
[root@cmschimera ~
```

On a node where this worked, this line is just bad! I'm looking for dcache to detect non-good status!

```
-bash-4.1$ cat check_srvs.log/cmsstrtemp.fnal.gov
13:10:06 up 1 min, 0 users, load average: 0.87, 0.35, 0.12
Linux cmsstrtemp.fnal.gov 2.6.32-573.3.1.el6.x86_64 #1 SMP Thu Aug 13 12:55:33 CDT 2015 x86_64 x86_64 x86_64 GNU/Linux
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sda3 270681640 30103760 226821312 12% /
dCacheDomain stopped dcache
-bash-4.1$
```

FIXED ON THE PROCEDURE AS OF 1:21, notified Natalia, the right one is:

```
kernelversion=2.6.32-573.3.1.el6.x86_64; uptime; uname -a | grep $kernelversion && (df /pnfs; (service dcache-server status | grep -v DOMAIN | grep -v running) || service dcache-server status)
```

Carrying on:

```
bash-4.1$ date
Wed Aug 26 13:13:13 CDT 2015
-bash-4.1$ pssh -h tape-pool-nodes -l root -t 0 -p 50 -o upgr_pools.log -e upgr_pools.err 'kernelversion=2.6.32-573.3.1.el6.x86_64;
extra_rpm_check="dcache-2.2.29-1.noarch"; puppet agent --disable --rebooting; service dcache-server stop; umount /pnfs; yum clean all; yum
update -y; rpm -q kernel${kernelversion} ${extra_rpm_check} && puppet agent --enable && reboot'
```

Now running (the fixed) verification step for pools:

```
-bash-4.1$ pssh -h tape-pool-nodes -l root -t 60 -p 50 -o check_pools.log -e check_pools.err 'kernelversion=2.6.32-573.3.1.el6.x86_64; uptime;
uname -a | grep $kernelversion && (df /pnfs; (service dcache-server status | grep -v DOMAIN | grep -v running) || service dcache-server status)'
[1] 13:23:14 [SUCCESS] cmsstor117.fnal.gov
[2] 13:23:14 [SUCCESS] cmsstor113.fnal.gov
[3] 13:23:15 [SUCCESS] cmsstor114.fnal.gov
[4] 13:23:15 [SUCCESS] cmsstor111.fnal.gov
[5] 13:23:15 [SUCCESS] cmsstor141.fnal.gov
[6] 13:23:15 [SUCCESS] cmsstor115.fnal.gov
[7] 13:23:15 [SUCCESS] cmsstor146.fnal.gov
[8] 13:23:15 [SUCCESS] cmsstor148.fnal.gov
[9] 13:23:15 [SUCCESS] cmsstor130.fnal.gov
[10] 13:23:15 [SUCCESS] cmsstor142.fnal.gov
[11] 13:23:15 [SUCCESS] cmsstor139.fnal.gov
[12] 13:23:15 [SUCCESS] cmsstor134.fnal.gov
[13] 13:23:15 [SUCCESS] cmsstor143.fnal.gov
[14] 13:23:15 [SUCCESS] cmsstor137.fnal.gov
[15] 13:23:15 [SUCCESS] cmsstor128.fnal.gov
[16] 13:23:15 [SUCCESS] cmsstor132.fnal.gov
[17] 13:23:15 [SUCCESS] cmsstor136.fnal.gov
[18] 13:23:15 [SUCCESS] cmsstor135.fnal.gov
[19] 13:23:15 [SUCCESS] cmsstor162.fnal.gov
[20] 13:23:15 [SUCCESS] cmsstor145.fnal.gov
[21] 13:23:15 [SUCCESS] cmsstor129.fnal.gov
[22] 13:23:15 [SUCCESS] cmsstor133.fnal.gov
[23] 13:23:15 [SUCCESS] cmsstor147.fnal.gov
[24] 13:23:15 [SUCCESS] cmsstor144.fnal.gov
[25] 13:23:15 [SUCCESS] cmsstor140.fnal.gov
[26] 13:23:15 [SUCCESS] cmsstor167.fnal.gov
[27] 13:23:15 [SUCCESS] cmsstor138.fnal.gov
[28] 13:23:15 [SUCCESS] cmsstor131.fnal.gov
[29] 13:23:15 [SUCCESS] cmsstor163.fnal.gov
[30] 13:23:15 [SUCCESS] cmsstor165.fnal.gov
[31] 13:23:15 [SUCCESS] cmsstor166.fnal.gov
[32] 13:23:15 [SUCCESS] cmsstor156.fnal.gov
-bash-4.1$ date
Wed Aug 26 13:23:31 CDT 2015
-bash-4.1$
```

Downtime took less than 22 minutes :)

#### #4 - 08/26/2015 01:33 PM - Gerard Bernabeu Altayo

Apparently Natalia decided to deviate from the agreed procedure...

Looking at what she did I see she added something in the verification phase which made my bug not affect her procedure, but she was not verifying properly either... I've fixed my bug and added her version verification so that we'll have more information for forensics if needed, the end of the verification steps now has:

(service dcache-server status && dcache version)

Check\_mk is not completely happy as some nodes came back with half of the network cards without link... Doing ifdown eth0; ifup eth0 I fixed it on cmsstor113 but there are others... I'll wait a bit to see if it stabilizes on its own, the error status was:

```
[root@cmsstor113 ~]# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: adaptive load balancing
Primary Slave: None
Currently Active Slave: eth0
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: eth0
MII Status: up
Speed: Unknown
Duplex: Unknown
Link Failure Count: 0
Permanent HW addr: 00:26:18:63:04:62
Slave queue ID: 0
```

```
Slave Interface: eth1
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:26:18:63:04:ee
Slave queue ID: 0
```

```
[root@cmsstor113 ~]# ethtool eth0
```

```
Settings for eth0:
Supported ports: [ TP ]
Supported link modes:  10baseT/Half 10baseT/Full
100baseT/Half 100baseT/Full
1000baseT/Half 1000baseT/Full
Supported pause frame use: No
Supports auto-negotiation: Yes
Advertised link modes:  10baseT/Half 10baseT/Full
100baseT/Half 100baseT/Full
1000baseT/Half 1000baseT/Full
Advertised pause frame use: Symmetric
Advertised auto-negotiation: Yes
Link partner advertised link modes:  10baseT/Half 10baseT/Full
100baseT/Half 100baseT/Full
1000baseT/Full
Link partner advertised pause frame use: Symmetric Receive-only
Link partner advertised auto-negotiation: Yes
Speed: 1000Mb/s
Duplex: Full
Port: Twisted Pair
PHYAD: 1
Transceiver: internal
Auto-negotiation: on
MDI-X: off
Supports Wake-on: g
Wake-on: g
Current message level: 0x000000ff (255)
drv probe link timer ifdown ifup rx_err tx_err
Link detected: yes
```

```
[root@cmsstor113 ~]# ifdown eth0
```

```
[root@cmsstor113 ~]# ifup eth0
```

```
[root@cmsstor113 ~]# ethtool eth0
```

```
Settings for eth0:
Supported ports: [ TP ]
Supported link modes:  10baseT/Half 10baseT/Full
```

100baseT/Half 100baseT/Full  
1000baseT/Half 1000baseT/Full  
Supported pause frame use: No  
Supports auto-negotiation: Yes  
Advertised link modes: 10baseT/Half 10baseT/Full  
100baseT/Half 100baseT/Full  
1000baseT/Half 1000baseT/Full  
Advertised pause frame use: Symmetric  
Advertised auto-negotiation: Yes  
Link partner advertised link modes: 10baseT/Half 10baseT/Full  
100baseT/Half 100baseT/Full  
1000baseT/Full  
Link partner advertised pause frame use: Symmetric Receive-only  
Link partner advertised auto-negotiation: Yes  
Speed: 1000Mb/s  
Duplex: Full  
Port: Twisted Pair  
PHYAD: 1  
Transceiver: internal  
Auto-negotiation: on  
MDI-X: off  
Supports Wake-on: g  
Wake-on: g  
Current message level: 0x000000ff (255)  
drv probe link timer ifdown ifup rx\_err tx\_err  
Link detected: yes  
[root@cmsstor113 ~]# cat /proc/net/bonding/bond0  
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: adaptive load balancing  
Primary Slave: None  
Currently Active Slave: eth1  
MII Status: up  
MII Polling Interval (ms): 100  
Up Delay (ms): 0  
Down Delay (ms): 0

Slave Interface: eth1  
MII Status: up  
Speed: 1000 Mbps  
Duplex: full  
Link Failure Count: 0  
Permanent HW addr: 00:26:18:63:04:ee  
Slave queue ID: 0

Slave Interface: eth0  
MII Status: up  
Speed: 1000 Mbps  
Duplex: full  
Link Failure Count: 0  
Permanent HW addr: 00:26:18:63:04:62  
Slave queue ID: 0  
[root@cmsstor113 ~]#

It does not really make much sense...

**#5 - 08/26/2015 05:18 PM - Gerard Bernabeu Altayo**

- *Status changed from New to Resolved*

The upgrade is complete, closing the ticket.