Mamoon Ansari

From: Oliver Lemke
Sent: Friday, July 10, 2015 12:50 PM
To: DL-Service - SW Support
Cc: German Krackow; Ken Hall
Subject: Useful notes on troubleshooting Replication issues

This issue was per SR3559030 with missing replication policies, but offers a lot of good checks and commands for other replication related issues. See summary at end.

Worked by Steve Cole, after being escalated through SES > SUS

Issue: 10 Folders under FS: /Video, each had a replication policy, but they disappeared around 6/16/15

Troubleshooting Methodology:
1) Determine what is actually replicating by command line
- Note I have lines prefixed with 'Source MDC:' or 'Target MDC:' to clarify where commands are run from.

Checking some of the folders, looking for BLOCKLET:
Source MDC: Run dm_info for folder /Producers
dm_info PRODUCERS
- output shows: 'eventlist: 0x10000000 BLOCKLET' indicates replication

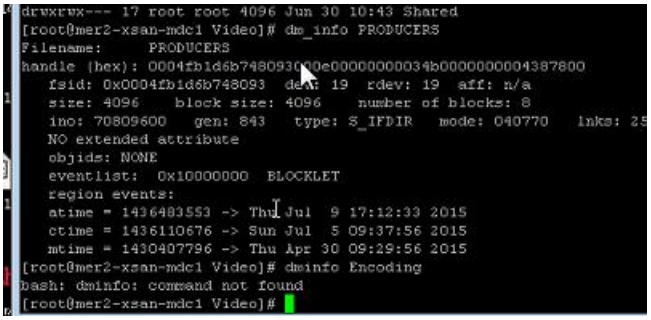

Same under folder 'Encoding':
dm_info Encoding
- output shows: 'eventlist: 0x10000000 BLOCKLET' indicates replication

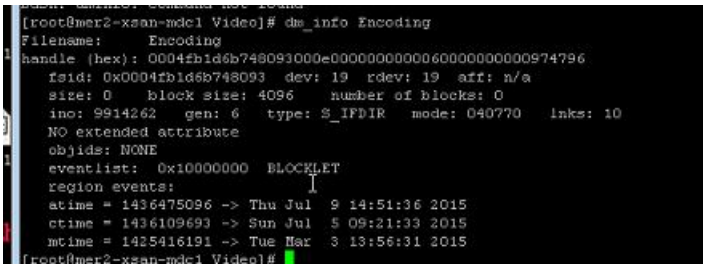

Checking some of the folders, looking for existence of the /.rep_private/config folder, then looking for the replication key:
Source MDC: Under /stornext/<FS>/.rep_private/config folder, key and folder info shows for working replication on FS /Creative:

cd /stornext/<FS>/.rep_private/config – where <FS> is the FileSystem name
ls –la
- Note key_xxxxxxxxx – this number is actually the folder's inode.

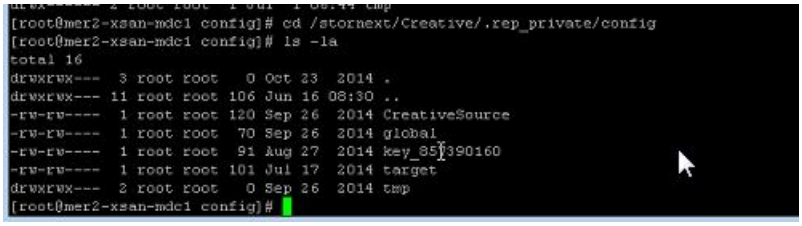

Use 'more' command on folder to show replication info:
more <FS> – where <FS> is the FileSystem name
- shows rep_output=true

Use 'more' command on key* to show all keys:
more key*
- Note here there should be 10 keys for each of the 10 folders to be replicated, but there is only 1 showing

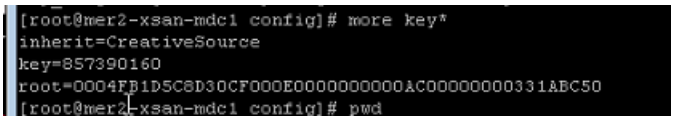

Source MDC: Now a listing of the folders under /Video which has missing replication policies:

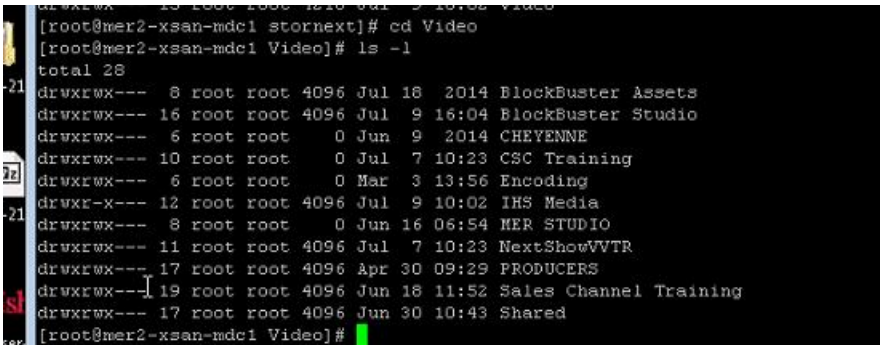

- All these folders had no replication policies and in turn no /.rep_private/config folders
(sorry no screenshot)

##

2) We have confirmed that replication policies are missing and that this is not just a GUI issue.
So what's the fix to getting the replication policies back?
Just recreate the replication policy again – no concerns over having to match keys, the keys are inherited from the folder's inode
Under GUI:

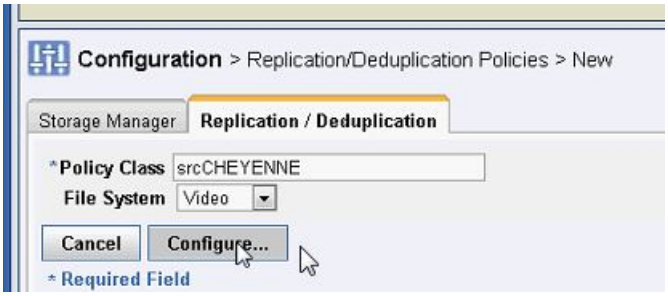

- Take note, CAPS are not allowed, see screenshot below. Policy Class was just recreated in just lower case.
- Also take note that the naming convention for the 'Policy Class' is just a name and can be called whatever you want. Most likely it reflects the name of what the user is replicating.

- Make sure 'Outbound Replication' is checked
- Note option 'Copies to Keep on Target' (referenced later in document)

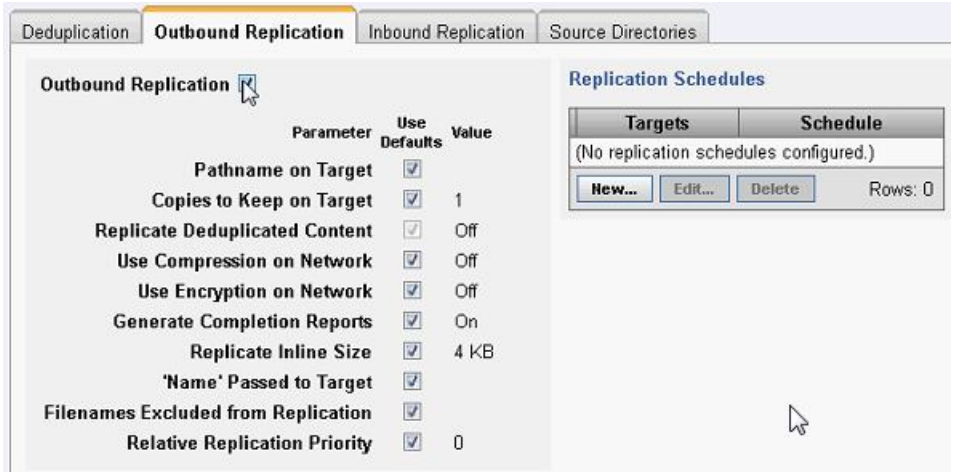

- Now to configure the Target and Schedules
- Select Available Target:

- Define schedule. This one is set for all days, once a day. Good idea to stagger the schedules an hour apart, if there are multiples. For this option select All Weekdays, All Months, All Days (of month), but only 1 of the hourly option – note Cron spec shows below.

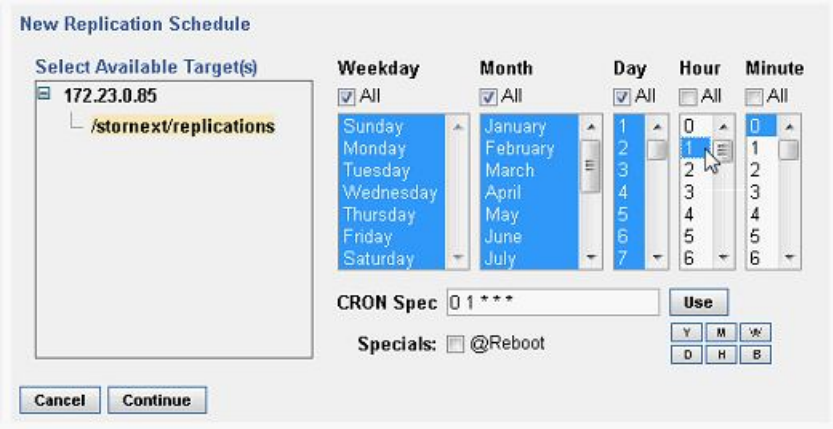

- Now the schedule appears, along with target info:

- Now you have to define the folder or folders within the <FS> that you wish to replicate, in this case it was /Video/CHEYENNE:

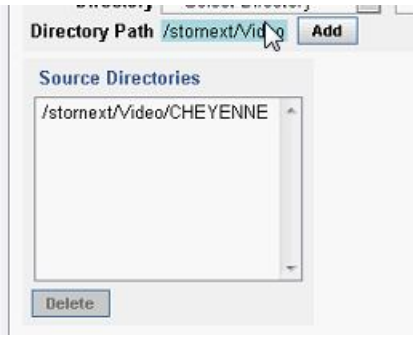

- Associate the Policy Class name:

##

3) Back to the CLI to check the newly configured Replication Policy

Source MDC: Now you will see under /stornext/Video/.rep_private/config, the key and details for Policy Class: srccheyenne:

cd /stornext/<FS>/.rep_private/config – where <FS> is the FileSystem name

ls –l

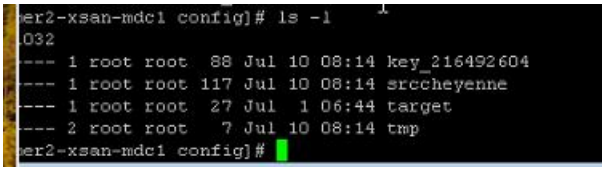

- Confirm that snpolicyd is running:

ps –ef | grep snpolicyd

```
[root@mer2-xsan-mdc1 config]# ps -ef | grep snpolicyd
root     15360 15058  0 Jul01 ?        00:04:07 /usr/cvfs/bin/snpolicyd -H mer2-xsan-mdc1.node-1 p
rocess/snpolicyd mer2-xsan-mdc1.node-1 0
root     46538 49101  0 08:16 pts/1    00:00:00 grep snpolicyd
[root@mer2-xsan-mdc1 config]#
```

The only log that shows replication activity is 'snpolicy.out'
tail –f /usr/cvfs/debug/snpolicy.out
- This log snippet shows a missing policy for key 9914261:

```
[root@mer2-xsan-mdc1 config]# tail -f /usr/cvfs/debug/snpolicy.out

(E)[20150710 08:16:44.453784 47346] Policy file for key 9914261 does not exist.

(W)[20150710 08:16:44.453794 47346] No policy for key 9914261, dropped inode 326179094
(E)[20150710 08:16:44.453803 47346] Policy file for key 9914261 does not exist.

(E)[20150710 08:16:44.453810 47346] Policy file for key 9914261 does not exist.

 I [20150710 08:16:44.453814 47346] /stornext/Video: queued 0 files for ingest, deferred 0
 I [20150710 08:16:44.453819 47346] /stornext/Video: Ingested 0 files 0 Kbytes in 0.000048 secs
```

General: Name of policy (Policy Class) does not matter, just where it is going to and from

Under the snpolicy CLI command, there is a '–repcleanup' command used to clean up issues. Note be cautious of using the –allcopies option, it WILL remove ALL copies which the customer may not want.

```
    -repcancel=pathname                 # cancel pending replication
    -repcleanup=mount_path -key=key [-copy=copy|-allcopies] [-path=path]
                                        # cleanup rep target for key
                                        # make copy N policy root or
                                        # remove all copies
    -rephistory=pathname [-count=n]     # list replication completion info
```

##

Target MDC: Note replications transferred:

```
root@chy-xsan-mdc1:/stornext/replications
drwxrwx--- 17 root root 4096 Jun 11 04:28 PRODUCERS.4
drwxr-x--- 17 root root 4096 Jun 10 04:02 PRODUCERS.5
drwxr-x--- 17 root root 4096 Jun  6 04:03 PRODUCERS.6
drwxr-x--- 17 root root 4096 Jun  5 04:02 PRODUCERS.7
drwxr-x--- 17 root root 4096 Jun  4 04:02 PRODUCERS.8
drwxr-x--- 17 root root 4096 Jun  3 04:02 PRODUCERS.9
drwxrwx--- 19 root root 4096 Jun 16 05:02 Sales Channel Training
drwxrwx--- 19 root root 4096 Jun 13 05:02 Sales Channel Training.1
drwxr-x--- 19 root root 4096 Jun  2 05:10 Sales Channel Training.10
drwxr-x--- 19 root root 4096 May 27 05:02 Sales Channel Training.11
drwxr-x--- 19 root root 4096 May 23 05:02 Sales Channel Training.12
drwxr-x--- 19 root root 4096 May 22 05:02 Sales Channel Training.13
drwxr-x--- 19 root root 4096 May 21 05:02 Sales Channel Training.14
drwxr-x--- 19 root root 4096 May 20 05:02 Sales Channel Training.15
drwxrwx--- 19 root root 4096 Jun 12 05:15 Sales Channel Training.2
drwxrwx--- 19 root root 4096 Jun 11 05:15 Sales Channel Training.3
drwxr-x--- 19 root root 4096 Jun 10 05:02 Sales Channel Training.4
drwxr-x--- 19 root root 4096 Jun  9 05:02 Sales Channel Training.5
drwxr-x--- 19 root root 4096 Jun  6 05:02 Sales Channel Training.6
drwxr-x--- 19 root root 4096 Jun  5 05:02 Sales Channel Training.7
drwxr-x--- 19 root root 4096 Jun  4 05:02 Sales Channel Training.8
drwxr-x--- 19 root root 4096 Jun  3 05:02 Sales Channel Training.9
drwxrwxr-x  2 root root    0 Jun 11  2014 target01
[root@chy-xsan-mdc1 replications]#
```

- To cleanup, you can also can manually remove by hand,oyu don't have to use the 'snpolicyd –repcleanup' option.

##

4) Process checks after triggering a manual replication run from the GUI.

Grep for 'Realized'

Grep for "realized" to confirm replication completions:

- If 'realized' is shown the transfer of replicated data was complete!

grep "realized" /usr/cvfs/debug/snpolicy.out

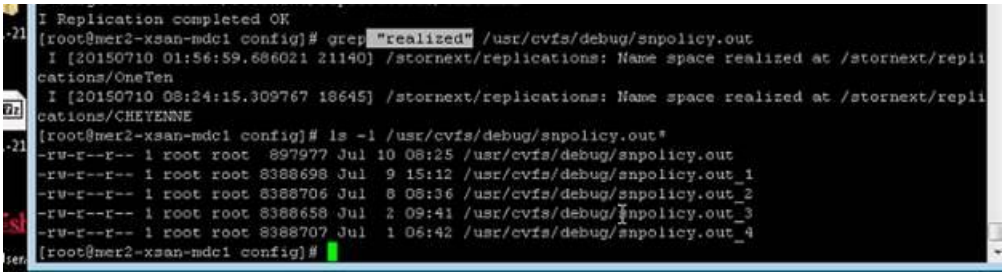

Snpolicy.out is the only place where replications are logged

##

Comparing Keys between Source and Target:

Source MDC: Note running ls –lid on the <FS>/Folder will show the folder inode (highlighted)

ls –lid <foldername> - Assuming you are already at the parent level folder or FS.

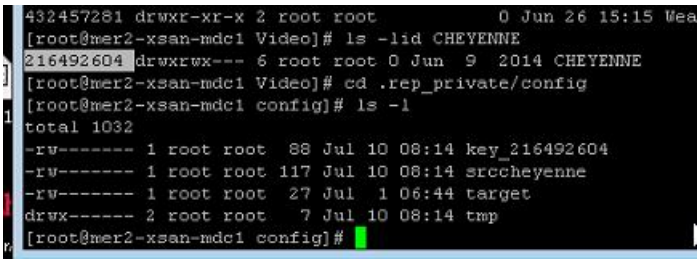

...and the inode is actually the Key!

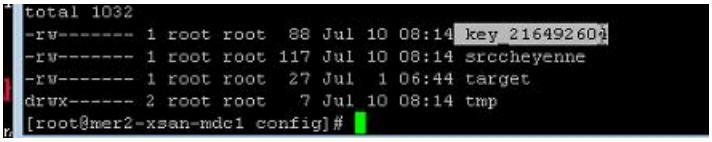

Source MDC: Data also relates to the Superblock ID of the <FS>. Run the following to get the <FS> Superblock info:

cvsfdb <FS>

cvsfdb> show sb

- Epoch # when file system was created:

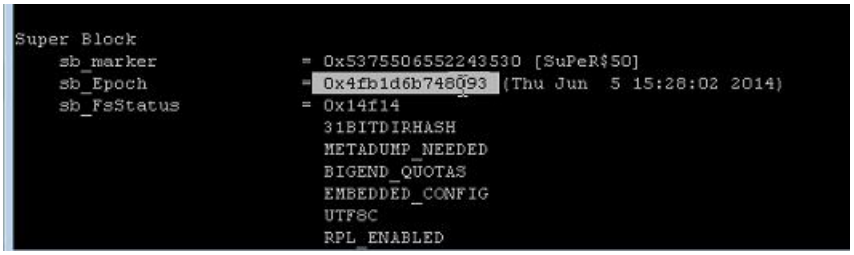

Target MDC: You will notice on the target under the defined replications folder /.rep_private/<superblock epoch #>

Target MDC: Under the superblock ID (Epoch #) on the Target MDC, you will see the inode ref for each folder, in this case 211105, then directory copies, then the replicated files:

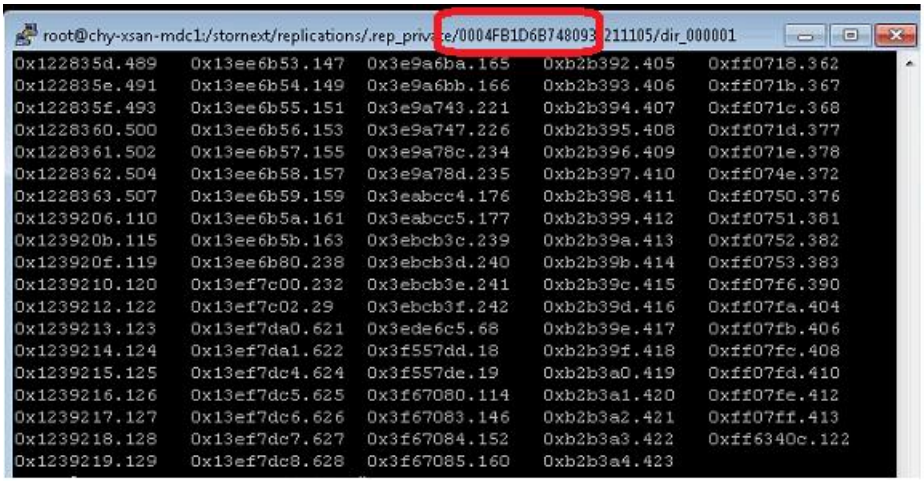

On Target Superblock ID:

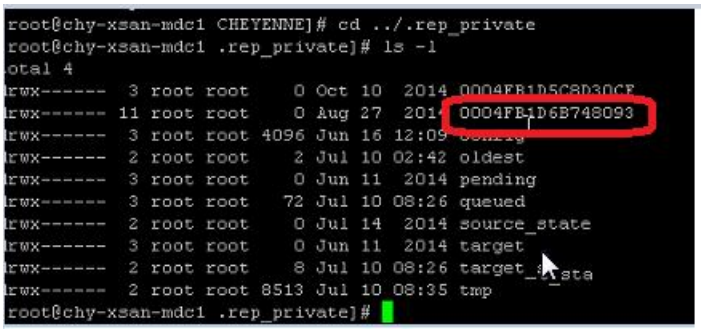

…and inode:

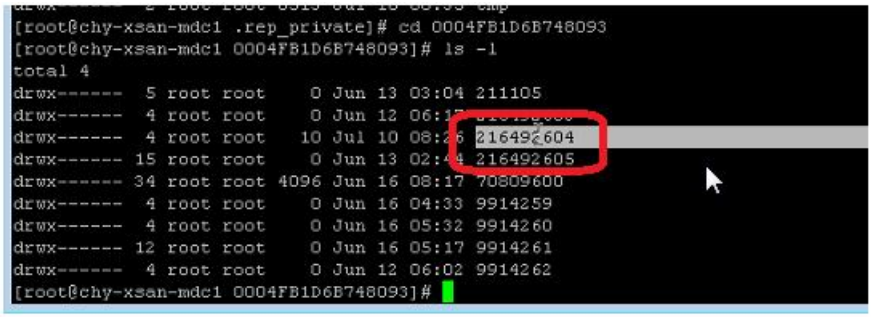

##

Reviewing Link Counts

Before being "realized" the Link Count on a file (on the Source MDC) will be 1, after it is realized it will be 2

In this example, the customer had a file copy count of 16, but as it was realized, then the Link count is shown as 17…. for 16 copies of the file and 1 for the realization:

```
drwxrwx--- 75 root root    4096 Jun 16 04:03 zach.hendrix
-rw-rw---- 17 root root 1800860 Oct  1  2014 zach.hendrix alias
[root@chy-xsan-mdc1 PRODUCERS]#
```

##

5) Summary

There are a number of ways to check that replication is working:
1) Check for the existence of a /.rep_private folder under the folder of the <FS>, and corresponding /config folder under that.
2) Check for the existence of the key_xxxxxxx file and confirm its number against the inode of the folder with ls –lid on the folder.
3) Run a 'more <FS>' to confirm replication is set to true.
4) Correlate Superblock ID (Epoch #) and folder inode info on Source with Target folders to confirm with customer that the files are present on the Target MDC
5) Confirm under Source GUI that the replication schedule is correctly setup to run non the correct days/times. Avoid triggering schedules to run at same time (not sure if there is any real impact of them running at same time, except for load on MDCs, network bandwidth).
6) Use link counts on files, comparing against file copies +1 to confirm successful replication.
7) Search snpolicy.out for 'realized' to confirm successful transfer
8) …and if the replication policies do go missing, just recreate them. As the key is pulled from the folder's inode , it will always be unique and the same. The policy name would not have to be the same, just the Source and Target folders and of course MDCs.
If a folder is removed and recreated with the same name, I would assume this to be a very different story, as the inode info will be different.

##

RCA on missing replication policies
- Policies seemed to disappear 6/16, the same time the target was updated from 5.0.1 to 5.2.1, but still this was inconclusive to be the root cause. Essentially there was no RCA here, so we just had to move on and recreate the policies via the source GUI.
- Actions taken: Recreated one policy and let the customer go ahead and recreate all the other himself via the Source GUI.

~~~~~~~~~ END ~~~~~~~~~~~~

Kind Regards
Oliver

Oliver Lemke *Software Product Support (SPS)*
Working Hours: 8am to 5pm MT
719.536.5642 | Oliver.Lemke@Quantum.com | Quantum.com
~~~~~~~~~