

On master01, media02, media03, media04 and media5 servers we have 9 shared tape drives for the SLPs. Most of the time only 3 or 4 would be working. Every day I would start in the evening with 8 or 9 working tape drives and wake up in the morning to 2 or 3 working drives.

On a daily basis I would try up the drives, vary them in the robot, get Vendor Support to check them for errors, complain to the storage team about paths going missing and have them check and re-check them, and go through all the troubleshooting steps I could think of – reboot, restart netbackup, rescan, release allocations, release drives and on and on. This had been going on for at least a year or two in various forms, and gradually getting worse. The backlog on duplications was running into dozens of terabytes taking up days to finish after any large backup completed, and causing a backlog onto the disk pools and putting every backup at risk of failing due to insufficient free space.

Attaching the tapes to new fibre (about two years ago) gave some hope that things would improve, but it didn't seem to. In fact, it may have made things worse – some drives now had multiple paths which would appear and disappear which made things even more complicated.

It is possible there's a reason joining to the new fabric may have made things worse though... SCSI Reservation.

I finally found out that if the entire end-to-end fabric, including switches, ISLs and any device in-between is not fully enabled and capable of handling SCSI reservation, the reservation may not complete properly, may corrupt, or might not release.



“it is imperative that SCSI traffic passes through the network/SAN infrastructure unmolested and without corruption, otherwise issues such as SCSI reservations can arise due to devices not releasing their reservations correctly and/or being unable to acquire them”

https://www.veritas.com/support/en_US/article.100027655

`/usr/openv/netbackup/db/media/error` log file showed page after page of RESERVATION_CONFLICT errors so perhaps this was something to look into.

With the drives dedicated to this backup server set only, Netbackup itself is actually capable of managing drive allocation without SCSI reservation, so I disabled it completely in Netbackup for all my master and media servers, by unchecking “Enable SCSI reserve”. (Restart netbackup, etc)

Things seemed to still work, but I still had at least 4 drives down.

```
vmopr cmd -crawlreleasebyname <drive_name>
```

Running this command released all the stale or broken reservations and within 10 minutes **all** drives were UP and in ACS.

Since this change on Saturday night the servers have caught up the entire backlog of 300 SLP jobs (120-odd TB). All new backup jobs complete with plenty of space and their duplications kick off within 10 minutes and finish relatively soon after.

There is no longer any backlog on SLPs for the first time in years.

Only one drive has gone down, but this is a confirmed hard failure in the robot. All other drives are performing beautifully and are able to be used on any of the media servers as netbackup desires.

The lesson seems to be that at least in this case the fibre infrastructure does not support SCSI reservation properly.

If your drives are DOWNing, you may want to check this out.

So far so good - let's hope this level of performance from the drives continues.

From:

<https://wiki.dewberry.co.za/> - **Shaun's Wiki**

Permanent link:

https://wiki.dewberry.co.za/doku.php?id=tape_drives_from_hell

Last update: **2020/08/13 09:43**

