

## Re: Backup tools too slow for LTO-3

---

*Source:* <http://unix.derkeiler.com/Newsgroups/comp.unix.solaris/2006-05/msg00798.html>

---

- *From:* Robert Lawhead <[news0000.5.unixguru@xxxxxxxxxxxxxxxxxxx](mailto:news0000.5.unixguru@xxxxxxxxxxxxxxxxxxx)>
  - *Date:* Fri, 12 May 2006 09:56:55 -0700
- 

Stefan Jankowski wrote:

Remember the days when the tape drive was the primary I/O bottleneck in backup operations? Well those days are gone: with that shiny new HP Ultrium 3 attached to my trusty V240[1] the limiting factor seem to be filesystem read rather than tape write operations.

The Ultrium 3 drive needs at least 27 MB/s to be able to stream (times the effective compression factor that is). The 3510 RAID as such is able to deliver data at a sufficient rate, provided it can be read in large chunks, as the following table shows

```
sdd if=/dev/md/rdisk/MYSLICE -onull bs=256k => 80-90 MB/s [2][3]
sdd if=/dev/md/rdisk/MYSLICE of=/dev/rmt/0n bs=256k => 55-60 MB/s
sdd if=/dev/md/rdisk/MYSLICE of=/dev/rmt/0n bs=512k => 60-67 MB/s
sdd if=/dev/md/rdisk/MYSLICE of=/dev/rmt/0n bs=1024k => 65-70 MB/s
sdd if=/dev/md/rdisk/MYSLICE of=/dev/rmt/0n bs=2048k => 65-70 MB/s
```

However, none of the onboard backup tools (or common replacements) appears to be able to deliver a sufficient /sustained/ transfer rate. I have tried:

1. ufsdump with blocking factors of up to the max allowed of 512 (= 256k blocks)
2. Solaris tar with blocking factors up to 2048 (= 1024k blocks)
3. star[4] with up to 1024k blocks and 512M FIFO (bs=1024k fs=512M)

All of these methods perform well for some regions of the file system and abysmally (20 MB/s and less with the tape stopping frequently) for others; the problem seems to be backing up many small files.

At this point, I am stuck. As I currently lack the disk space for the obvious solution (create tar archive on disk and dump to tape later) I'd like to ask:

## Re: Backup tools too slow for LTO-3

Are there any other optimizations, tweaks, tools that I've overlooked?

Should I look at commercial backup software?

How do other users of LTO-3 drives do it?

Thank you,  
Stefan

[1] System setup:

- V240 with 2x 1.5Ghz USIIIi, 2GB RAM
- SE 3510 with 12x 300GB/10000RPM, 2 logical drives (RAID-5, 6 and 4 disks), random optimization, ld0 attached to the V240 via a 2GB FC adapter (SG-XPCI2FC-QF2) in the 66 MHz PCI slot
- Overland Neo 2000 w/ 1x HP Ultrium 3, attached to the V240 via U320 SCSI (SSG-XPCI2SCSI-LM320) in one of the remaining 33MHz slots (slot slower than card, I know, but the 66Mhz slot is taken...)

[2] Approx. throughput measured by observing the output of 'iostat -xM 5' for a few minutes

[3] `sdd` is a replacement for `dd` available at <ftp://ftp.berlios.de/pub/sdd/>

[4] `star` is a replacement for `tar` available at <http://cdrecord.berlios.de/old/private/star.html>

You only need enough disk to buffer the difference between the rates at which data is produced and consumed.

Write yourself a buffer program (think TIVO) that provides two endpoints (named pipes `stdin/stdout`) that buffers its input to a large "ring file". Manage the flow in and out of the ring file such that reads on the read endpoint are blocked until some high watermark is reached, but once reached will satisfy reads until a low watermark (empty) is reached. At that point, block reads until the ring is refilled to some high watermark reached, and repeat...

Now, if you now "dump" to the write-endpoint, and "dd" from the read-endpoint to tape, and the average rate of the dump process exceeds the rate that the tape drive requires to stream, (and the ring file is big enough), life will be good. If not, at least you can limit the number of times the tape has to start/stop streaming.