



CSIRO



www.csiro.au

Generational Change

Peter Edwards



Overview

In IT, newer products are continually superseding previous ones.

In the case of tape architecture, the time it takes to copy media from an old format to a newer one means that generational upgrades may be in progress all the time.

Indeed, there may be more than one upgrade underway at a time.

Perhaps this should be allowed for when sizing a configuration?

Example

For example, at CSIRO ASC in the last three years we have had five generations with up to four in use concurrently.

As at Oct 2011:

T9840C	T9940B	T10000A	T10000B	T10000C
gone	2/3 gone	gone	production	Being installed

T9840C -> T9940B 29 weeks

T9940B -> T10kB 18 weeks to reach 2/3 done
(estimate 27 weeks total)

T10kA -> T10kB 25 weeks

When to start a transition

Keeping old tape drives as long as possible is frequently a good thing to do, as they provide increased parallelism which you don't give up lightly.

Reasons to do so may include:

- Cost of old drives (media availability/price, maintenance charges/effort, reliability)
- Compatibility with a new library (eg: T9840C unsupported in SL8500)
- Unable to install more new drives without removing old (FC ports, drive/tape slots in library, power supplies)
- Relative speeds (eg: a T10k that you have to wait for may still give faster recalls than a T9940 that can be used immediately)

Can new drives re-use old media?

If new drives can read/overwrite media written by old ones, then why not:

1. Flag old-format tapes with data as read-only using *dmvoladm*.
2. Remove or disable old drives.
3. Merge tapes by setting "hsparse" at a rate that suits you, choosing socket or disk merges as you prefer. Use of *run_merge_manager* is recommended. (Natural attrition is probably too slow.)
4. As tapes become free, clear the read-only flag. You may need to destroy existing tape labels with *tmlabel*, *dd* or similar.

Advantages: No new DGs or VGs, so no config changes
No load on the DMF daemon

Disadvantage: Monopolises drives (merging works at the volume level)

(CSIRO ASC went from T10kA to T10kB this way.)

Which tapes to convert first?

- To maintain drive parallelism for as long as possible, convert least active tapes first.
 - Can you get access statistics out of your tape subsystem? Even just as a one-off?
 - Use the DMF VOL database's *wfdate* field which gives time of last append (not ideal).
 - Convert a tape parasitically when next mounted due to user activity, though this tends to convert the most active tapes first and does so the most during busy periods.
- But if freeing up tape slots is a priority, convert emptiest tapes first.

Tape-to-tape directly, or via disk?

- Tape-to-tape ties up two drives (one of the old type and one of the new) for the duration.
- Tape-disk followed by disk-tape causes disk I/O and is slower. (Beware of IS220 disk, which handles concurrency poorly.)

Convert entire tape at once?

- For high capacity tapes, this will involve dedicating drive(s) for a lengthy periods of time.

For the old drives, this becomes less important as the conversion proceeds, and vice versa for the new ones. But presumably the new ones are faster.

- Alternatively, convert in batches of one or two zones, which takes longer due to extra tape mounts but interleaves better with production traffic.

How hard do you push?

run_merge_manager is intended to perform tape merges only when there are idle tape drives. At CSIRO ASC, we modified this to perform media conversions when the system was less busy.

If converting a zone or two at a time, this may be overkill.
(We don't bother with it any more.)

Beware of pushing too hard - drive availability may not be your bottleneck.

Perhaps only convert tapes during non-prime time?

The scripts

- A bit like Seymour Cray's sailboats...

Thanks

Thanks to Channel Ten and Google Images for the “*Talkin' 'bout your Generation*” image used on the cover slide.

Advanced Scientific Computing

Peter Edwards

Systems Support Manager

Email: peter.edwards@csiro.au

www.csiro.au

Thank you

Contact Us

Phone: 1300 363 400 or +61 3 9545 2176

Email: enquiries@csiro.au Web: www.csiro.au

