

Re: [SLE] Backups -- near-line disc and off-line tape

Source: <http://linux.derkeiler.com/Mailing-Lists/SuSE/2006-06/msg02742.html>

- *From:* "Bryan J. Smith" <b.j.smith@xxxxxxxx>
 - *Date:* Wed, 28 Jun 2006 22:35:03 -0400
-

On Wed, 2006-06-28 at 18:51 -0700, suse@xxxxxxxxxxxxxxxx wrote:

Looking for suggestions for backups:
Current situation:
3 data servers: daily tars copied to a 4th server with a tape backup running at night.

That's actually a good strategy -- you buffer to disc, and then disc to tape. I **NEVER** stream directly from a system over the network to tape. In fact, such real-time tape backup is why most people hate tape.

The ideal setup is to use disc as your near-line store (for quick backup with immediate restore), and tape as your off-line store (for taking backups off-site). In fact, using that strategy, you only really need to do a tape backup when you actually take backups off-site -- e.g., weekly.

The most ideal setup is when you have enough near-line storage to accommodate a *_full_* backup of *_all_* your servers. That way you only need to send incremental deltas every night from servers to that backup with near-line storage. From there, at your convenience, you can decide what to put from the near-line disc to off-line tape (or not).

You don't need SCSI drives for this backup system, since your discs will be idle most of the time. A 3Ware 4-channel SATA card with four (4) 320GB WD3200SD** (24x7 rated near-line enterprise commodity disc**) would give you 1 TB usable RAID-5 for \$600. You can then put that in a new \$500-600 backup server, or just use an existing server (that doesn't have much load at night for the deltas to near-line disc, and weekends for the write to near-line tape).

For a more complete discussion of the evolution of legacy real-time tape backup to buffered disc-to-tape backup to today's combination of near-line disc with off-line tape and the final evolution of the "virtual tape library," see 2005 September Sys Admin: "Dissecting Virtual Tape Libraries"
<http://www.samag.com/documents/sam0509a/>

Re: [SLE] Backups -- near-line disc and off-line tape

The article is purposely written from a `_sysadmin_` functional concept and viewpoint, and `_not_` a network/storage admin viewpoint.

[****NOTE:** Pay the \$20–30 "premium" for the Hitachi "T7K", Seagate "NL35" or Western Digital Caviar "RE" series. These are commodity discs that are `_rated_` for 24x7 "near-line" (sometimes, and quite `_in_` correctly, called "RAID" rated). In a nutshell, they are the Hitachi 7K, Seagate Barracuda ATA 7200.x and Western Digital Caviar SE products that test to `_higher_tolerances_` than their consumer/retail rated models. They are well worth it. I'm partial to the WD3200SD (107GB/platter, 320GB Caviar RE, `_not_` RE2) because it is very proven as a reliable model (even in its consumer WD3200JD version too). `_Never_` use "brand name" to base a decision -- as even hard drive vendors outsource these days. E.g., Seagate Barracuda ATA 7200.8/9/10 models are outsourced to Maxtor, and not as reliable as their prior 7200.7 series, so it didn't surprise me when Maxtor and Seagate merged just recently. Some of the newer Western Digital models (like the 500GB) aren't the same, sound design -- and Hitachi has gone back to using 5-platters in their 500GB disc (the last disc to do that was the IBM Deskstar 75GXP -- the Hitachi 500GB is also the Deskstar series ;-).

But at least with the Hitachi "T7K" (not the consumer 7K), Seagate NL35 (not the consumer Barracuda ATA 7200.x) and the Western Digital Caviar RE (not the consumer Caviar SE), they `_are_` tested to higher tolerances and `_rated_` for 24x7 "near-line" (i.e., not spinning 24x7, but still okay to be powered 24x7) instead of the standard 8x5 "consumer" tolerances.]

The local PCs are NOT backed up, nothing is supposed to be stored there.

With an intermediate backup server, you can always choose to `rsync` select directories to it.

BTW, a popular strategy for backup servers is to maintain "snapshots" of backups. Using kernel 2.4 with LVM**, you can take daily and weekly "snapshots" of the backup server, and then serve those "snapshots in time" back out (read-only) via NFS and SMB.

E.g., a typical strategy I personally use are the last 6 days and then the last 2 Fridays before those last 6 days (3 weeks back) -- 8 snapshots. Every night I do the "snapshot" of the filesystem `_before_` the servers and workstations do a `rsync` to send their deltas of what has changed in the last 24 hours. From that, users can go back and look at `_any_` backup via NFS or SMB from the last night, any of the previous 6 days, or the 2 Fridays before last.

**NOTE: kernel 2.6 with LVM2 and DeviceMapper is still developing

Re: [SLE] Backups -- near-line disc and off-line tape

snapshots (with write support and other advanced features). Red Hat has been shipping the feature in more recent kernels, but most (including myself) still don't trust it yet. So I'd stick with a kernel 2.4 distro.

The tape drive has died: a replacement is around \$3,000. The tapes were convenient because somebody could take them home next night and rotate.

What technology? AIT? DLT? LTO? Other? You'll want to stick with what technology you already had so you can read tapes from before.

If you don't need to read your old backup cartridges, then I suggest going with LTO -- a multi-vendor standard that is the best price/performance and also DLT-quality longevity (only IBM's proprietary formats are longer lasting).

The backups are right now 40Gb, but could easily be reduced to 10Gb.

Don't do that to yourself. If you need 40GB, backup 40GB. Don't add administrative burden to your job.

In reality, consider spending \$1,200 on a new backup server with 1TB of near-line (ATA-based) storage. If you don't need to read your old tapes, LTO-1 (100GB/20MBps native, 240GB/48MBps 2.4:1 typical) starts at \$800 and tapes are under \$25/each.

If you use a combination of near-line disc and off-line tape, you really only need enough disc to store full backups (2x if you want to use snapshots) and then enough tape cartridges to store 4 recent weeks (rotated) plus one per month for long-term retention (if you need it). That cuts down on your tape usage/need/costs.

What can you suggest so that somebody can take the backup home?

Since LTO-1 is so cheap, I recommend you stick with tape if you need portability. Use disc for near-line (local, on network) and only use tape for when you need to take it off-line (like off-site).

A common (and poor) assumption is that disc can take "swapping around." It can not. At the most, I'd by 2.5" discs and pad them like made in a tray because they can take 10x the stock of 3.5" discs. But I would not rely on 3.5" portable discs. Especially today's 3.5" commodity discs that are not designed for staying off-line long.

Re: [SLE] Backups -- near-line disc and off-line tape

Lastly, you have the "hot-swap" issue. SATA does not "hot-swap" well unless you put it on a 3Ware or other true hardware RAID controller in at least hardware RAID-1 (long story). Now a 3Ware Escalade 8006-2LP will only cost you \$125, plus another few hundred for a pair of discs or trays, but at that point, you're half-way to a LTO-1 drive anyway.

The key to having a complete recovery solution is to:

A. Realize that you not only can but should use disc backup today. Disc is a reliable mechanism, fast, random-access and can store "full copies" you can "delta" against (instead of having to pass full copies over the network every night).

B. Realize that you still need an off-line solution for off-site portability. Since this solution only goes off-site every week or month, it means it does not and often should not be your primary mechanism of backup. You should be using near-line disc, and then only commit from disc to off-line when you actually need to take stuff off-line/off-site.

C. Reality that disc is not portable and has poor tolerance to vibration, at least compared to tape cartridge. Which is why you should not be using disc as your off-line/off-site, but tape. Now you could opt to do Internet-based backup -- and you probably can do it with only 40GB. But it's also nice to have multiple off-line copies, and not just one. And that's where tape still has a great advantage.

Again, if you read the 2005 September article, it lays out all the aspects of not only "getting stuff backed up," but looking at the aspects of the recovery from a feasibility/reliability standpoint.

In a nutshell ...

Tape sucks compared to disc when it comes to immediate backup and recovery.

But disc sucks compared to tape when it comes to longevity, environment, tolerance and portability.

The ideal setup is to leverage the two, as they complement each other quite well.

--

Bryan J. Smith Professional, technical annoyance
<mailto:b.j.smith@xxxxxxxxx> <http://thebs413.blogspot.com>

The existence of Linux has far more to do with the breakup of AT&T's monopoly than anything Microsoft has ever done.

Re: [SLE] Backups -- near-line disc and off-line tape

--

Check the headers for your unsubscribe address

For additional commands send e-mail to suse-linux-e-help@xxxxxxxx

Also check the archives at <http://lists.suse.com>

Please read the FAQs: suse-linux-e-faq@xxxxxxxx