

Storage Decisions


Hosted by STORAGE 

Integrating Disk Into Backup for Faster Restores

Presented by
W. Curtis Preston
President
The Storage Group, Inc.




Storage Decisions

Hosted by STORAGE 

Tape Backups Aren't Cutting It


- High-speed tape drives in a library are the standard, but the cost of these units causes many people to cut corners elsewhere.
- Tape drives must be streamed, creating difficulty when creating network backups, incremental backups & offsite tapes.
- Tape drives also have reliability issues.
- Many people aren't utilizing their tape drives properly and are not getting all their backups done.
- Also, many are not creating offsite copies.

Storage Decisions


Hosted by STORAGE 

Am I Making This Up?

Let's find out.



Storage Decisions

Hosted by **STORAGE** 

What Percentage of Your Drives' Throughput are you Using? (e.g., If you have one LTO-2 capable of 30 MB/s, and you're getting 1.5:1 compression, your goal is 45 MB/s. If you've got 10 of them, your goal is 450 MB/s, or 1.6 TB/hr. How much of that are you using on a regular basis?)

- A. 75% or higher**
- B. 50% - 75%**
- C. 25% - 50%**
- D. 0% - 25%**
- E. We have no idea.**


Storage Decisions

Hosted by **STORAGE** 

How Often Does Tape Fail in Your Environment?

- A. Daily**
- B. Weekly**
- C. Monthly**
- D. Yearly**
- E. Never**

Storage Decisions

Hosted by **STORAGE** 

What Do You Send Offsite?

- A. We copy the originals and send the copy or original offsite.**
- B. We don't copy the originals and send them offsite.**
- C. We're not sending anything offsite. Please don't tell my boss.**

Tape Advantages

- High-speed, low-cost
- Good archival solution. Allows multiple copies without significant cost
- Lots of new tape drives on the market
 - 9940B (30/70 MB/s)
 - AIT-3 (15/30 MB/s)
 - LTO-2 (30/60 MB/s)
 - SDLT 320 (16/32 MB/s)
 - SAIT (30/60 MB/s)

Tape Challenges

Hard to Stream

- Tapes are now *too fast!*
 - **Must** use multiplexing to stream them during network backups.
 - Must use higher multiplexing values than ever before, hurting restore performance even more.
- Tape-to-tape copying takes time, and multiplexing increases that time, especially if you de-multiplex.
- Must perform regular full backups to reduce number of tapes required for restore
- Incremental backups do not supply enough data to stream a tape drive.

Tape Challenges

Flexibility & Reliability

- Cannot write to single tape drive from two shared servers simultaneously
- Single tape can cause large backup (or more importantly a restore) to fail.
- Tapes have more read/write errors than disk.
- *You never know if a tape is good until you really need it.*

Tape Challenges

Off-site Copies

- Assuming copy is same speed as backup, must buy at least twice as many drives to perform copies in one day
- If copy is not same speed, must accept longer copy window or buy more tape drives
- Additional drives cost a lot of money.
- **Result:** Many people still not making offsite copies

Solution: New Backup Media

- **Really inexpensive disk arrays**
 - IDE/ATA based
 - Addressable via Fibre Channel, SCSI, Firewire, NFS or CIFS
 - Prices from \$3/GB to \$20/GB
 - Disk array, NAS filer, Virtual Tape Library (VTL) and removable virtual tape (RVT) configurations

What Do These Things Look Like?

- **Disk Array**
 - Typical JBOD or RAID array that will need a filesystem & volume manager to be used for backups.
 - **Vendors:** EMC, IBM, HDS, NexSAN, STK, etc.
- **NAS Filer**
 - Mount huge disk array via NFS/CIFS and back up to it – no need for volume manager & filesystem.
 - **Vendors:** BlueArc, Data Domain, EMC, NetApp, OnStor, Procom, etc.

Storage Decisions

Hosted by STORAGE SearchStorage.com

What Do These Things Look Like?

- Virtual Tape Library (VTL)
 - Software running on an appliance (or in a smart switch) that makes a disk array look like one or more tape libraries.
 - Usually sold complete with storage, sometimes sold software only, or software and appliance only.
 - S/W Vendors: Alacritus, Diligent, FalconStor, NearTek, Quantum, Sepaton
 - H/W Vendors: ADIC, Copan, MaXXan

Storage Decisions

Hosted by STORAGE SearchStorage.com

What Do These Things Look Like? (2)

- Removable Virtual Tape
 - Disk drive managed by robot in a tape library; plugs into virtual tape drive
 - Acts like tape in every way, *can even be ejected*.
 - Vendors: Spectralogic

Storage Decisions

Hosted by STORAGE SearchStorage.com

What to do With Them?

- **Connect array to backup servers via Fibre Channel & SANs, or GbE & NFS/CIFS**
- **Back up to disk using your backup software of choice**
- **Copy disk backups to tape, send offsite**
- **D2d2T = Backups only on disk until they're put on tape; most restores come from tape.**
- **D2D2t = All onsite backups on disk; all restores come from disk except DR.**

```


graph LR
    A[Backup Client] --> B[Backup Server]
    B -- "NFS/CIFS/SAN" --> C[ATA Disk Array]
    C -- "Copy or second backup" --> D[Tape]
  
```

What to do With Them (2)

- Most backup systems do things that are not necessary when backing up to disk
 - Occasional full backups
 - Incremental backups of entire files
 - Backing up redundant files
- Some products are "getting it"
 - Forever incremental w/virtual full backups
 - Block level incremental backup
 - Redundant file & block elimination across data center

What to do With Them? (3)

- Replicate many clients to a central array, back up that array & send offsite
- Replication without the cost of traditional RAID arrays
- Could allow for very short backup windows, or centralized backup of remote sites



What to do With Them (4)

- Could also use software-based RAID to create additional mirror, and split mirror for backups
- Gives you BCV functionality for a fraction of the price!
- Back up large databases with no I/O overhead on server!

Why Are We Using Disk?

- Doesn't require constant stream.
- Disk is definitely more reliable than tape.
- Instantaneous loads, rewinds, FSF, FSR -- load hundreds of incrementals in the time it takes to load one full.
- Don't need to multiplex for performance reasons. (May want to multiplex for licensing reasons.)
- If you do multiplex, multiplexed restores may still be faster on disk.
- NFS/CIFS devices can be used simultaneously by many clients, without needing to stream each device.

Why Are We Using Disk? (2)

- Some protected via monitored RAID -- the loss of a single disk would be monitored and repaired, while the RAID group continued to protect the data.
- Disk-to-tape copies are easier than tape-to-tape copies.
- Incremental backups with little data will not hurt performance of other backups -- could perform infrequent full backups without increasing the chance of failure.
- Fewer full backups means
 - Media savings (buy less "tape")
 - Reduces impact of backup on servers & network
 - Backups finish faster

Disk Array/Filer or a VTL?

Disk Arrays & Filers

- Pay for disk, possibly volume mgr & filesystem
- Possibly pay per GB to backup s/w vendor
- Speed of filesystem
- Fragmentation
- Some backup s/w not fully filesystem aware
- Will have to change how some backups are done
- Can also be used as a replication destination

VTLs

- Pay for disk & value of VTL software
- Pay per slot, or drive to backup s/w vendor
- Speed of mult. raw disks
- Usually no fragmentation
- All backup s/w tape aware
- Should be faster
- Just like what you're using today
- Some VTL vendors offer dual-purpose VTL/filer units

Storage Decisions

Hosted by STORAGE SearchStorage.com

NAS Filer or a Disk Array?

Disk Arrays	Fileers
<ul style="list-style-type: none"> ● Pay for disk, possibly volume mgr & filesystem ● Usually partition to mount to mult. systems ● Pay for FC infrastructure ● FC/SCSI performance 	<ul style="list-style-type: none"> ● Pay for disk, and filer OS, NFS/CIFS software ● Easily mounted to multiple systems ● Pay for GbE infrastructure ● GbE/NFS/CIFS perform.

Storage Decisions

Hosted by STORAGE SearchStorage.com

Why Not do Everything to Disk?

- Archiving purpose of backups requires older backups to be available
- Tapes still much cheaper, allowing for multiple, stable copies to be put on "the shelf" onsite or offsite.
- Tapes not susceptible to filesystem corruption


Storage Decisions

Hosted by STORAGE SearchStorage.com

How do You Feel About Tape?

- Tape is the bomb. I just came to this talk to get a good seat for the next talk.
- I like D2d2T (little disk in front of big tape – disk used only as cache to tape)
- I like D2D2t (big disk in front of little tape – tape used only for offsite)
- If I could do it today, I'd rip out every tape drive I have, replace it with disk, and never use tape again.


Storage Decisions

Hosted by **STORAGE** 

Prerequisites: Who Needs Disk in Their Backup System?

- **Everyone. Make disk a part of any *new* system.**
- **If you're not ready to throw out the old already, you should meet one or more of the following prerequisites:**
 - Those needing a more reliable backup medium.
 - Those who have a high incremental/full ratio; incremental backups perform *much* better on disk.
 - Those who must reduce the amount of data being transferred across the network; switch to disk and perform full backups less frequently.
 - Those who perform frequent restores; disk restores happen in seconds, not minutes.
 - Those who are ready to switch to replication-based backup.


Storage Decisions

Hosted by **STORAGE** 

Who Doesn't Need Disk?

- **Just make sure that disk will solve the problem you have. It's not a cure-all.**
- **If you're about to throw out a non-depreciated backup system just to put in disk, get an independent opinion; most people just need a good tune-up.**
- **Just because a gun slinger can draw quick doesn't mean they'll be any good in the four-minute mile; *disk is not always faster than tape.***

Storage Decisions

Hosted by **STORAGE** 

In Short

- **Doing backups to inexpensive disk first allows for**
 - Faster, easier, more reliable backups – especially incremental backups
 - Easier creation of offsite tapes
 - Easier restores both on- and offsite
 - Many other features
- **Directories of these products:**
<http://www.storagemountain.com>
- **Questions to curtis@thestoragegroup.com**

More Questions?

Mr. Preston will be available for one-on-one questions at the "Ask the Experts" area of the Exhibit hall this evening from 5:00 - 6:00 PM.

Sponsored by CA
