NAS Backup Strategies

Jeff Bain
Hewlett-Packard

Network Storage Solutions Organization
Automated Storage
700 71st Ave.
Greeley, CO 80634

jeff_bain@hp.com

NAS Backup Strategies

- NAS: A Changing Definition
- Pros and Cons of NAS
- NAS Backup Strategies
- Technology Survey
- Products
- Wrap-up / Questions

NAS: A Changing Definition

- Not just “filers”
- Disk / Tape / Optical storage on the LAN
- “Storage Over IP”, “Storage Over Ethernet”
- Protocols include: NDMP, iSCSI, DAFS, backup *appliances* with pre-configured software agents
- Related technologies: Gigabit Ethernet (GbE), TCP Offload Engines (TOE) on NICs

Pros and Cons of NAS

- ☺ Works with your current infrastructure
- ☺ It's Ethernet, "You're soaking in it right now!"
- ☺ Reliable, easy to scale
- ☺ Plug-and-Play "intelligent appliance" operation (generally not available in fibre channel)

Pros and Cons of NAS

- ☹ 10/100Base LANs are too slow
- ☹ The TCP stack is too fat and slow
- ☹ File systems (NFS/CIFS) are too slow
- ☹ Technology is too new

NAS Backup Strategies

Local with tape drives

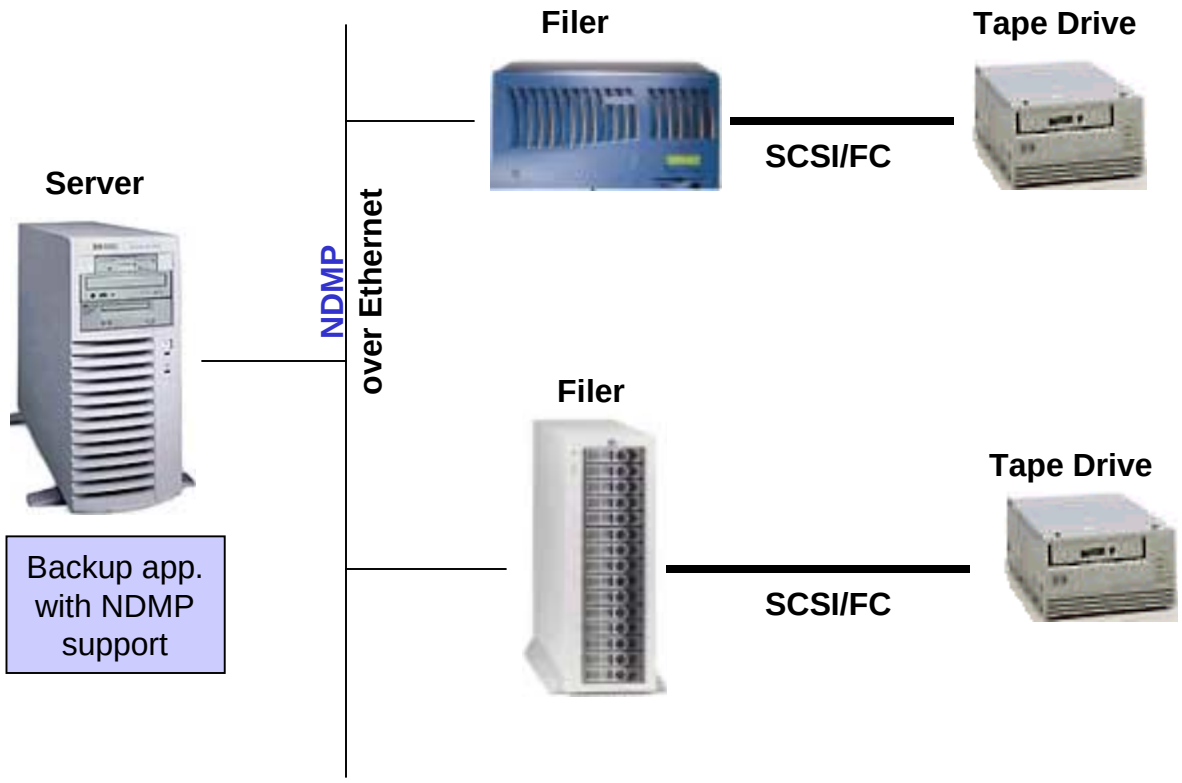

Relatively inexpensive, but requires manual tape pickup and replacement

NAS Backup Strategies

Local with tape library

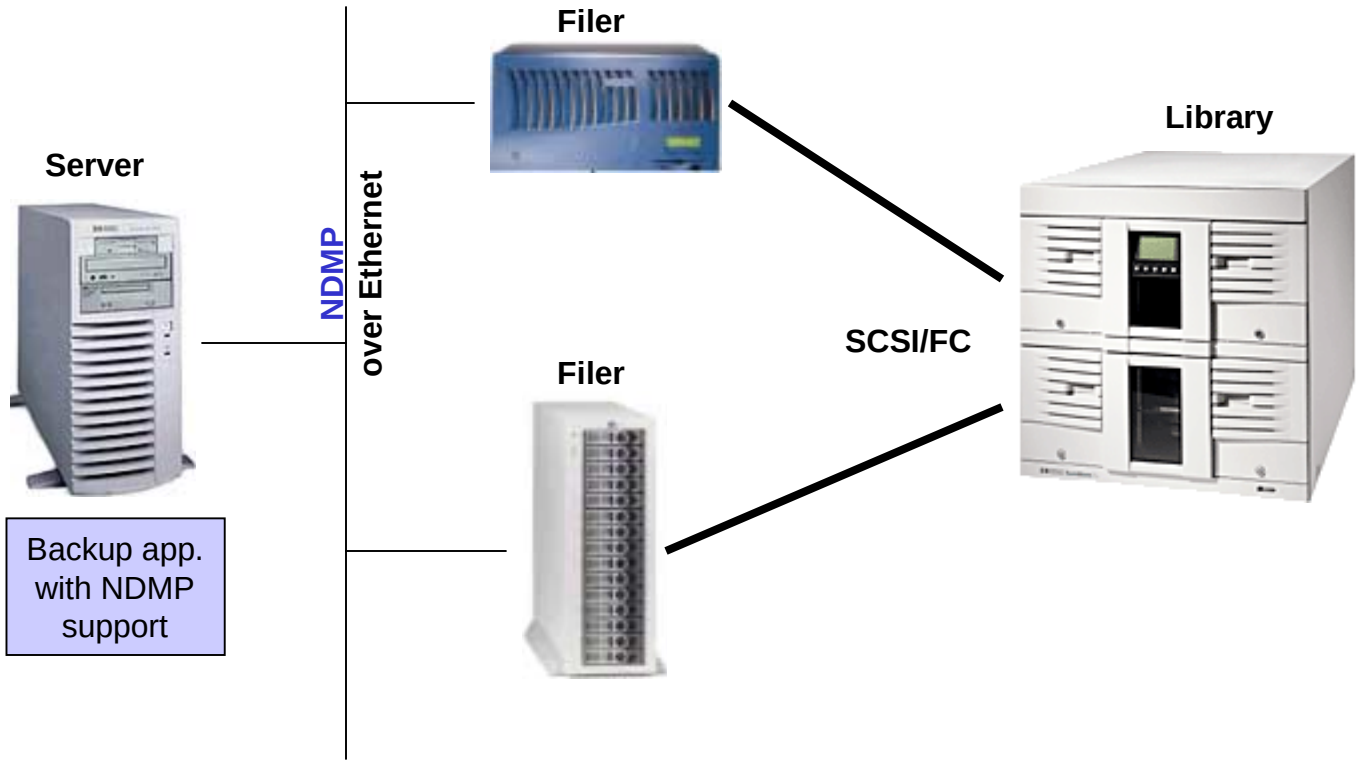

More automated: tape replacement + centralized storage

NAS Backup Strategies

Filer to Filer

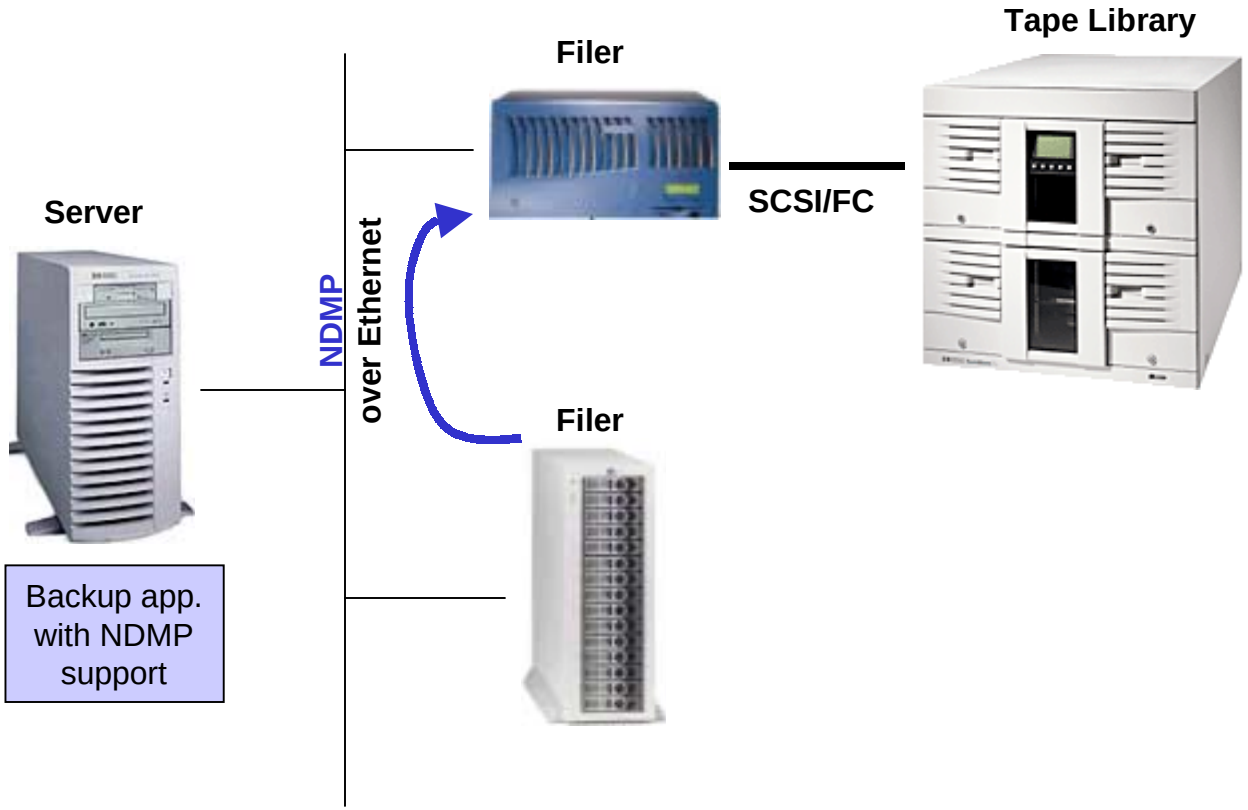

Removes backup loading from single filer using **snapshots** and **snap-mirror** technology

NAS Backup Strategies

Filer to Server (NDMP)

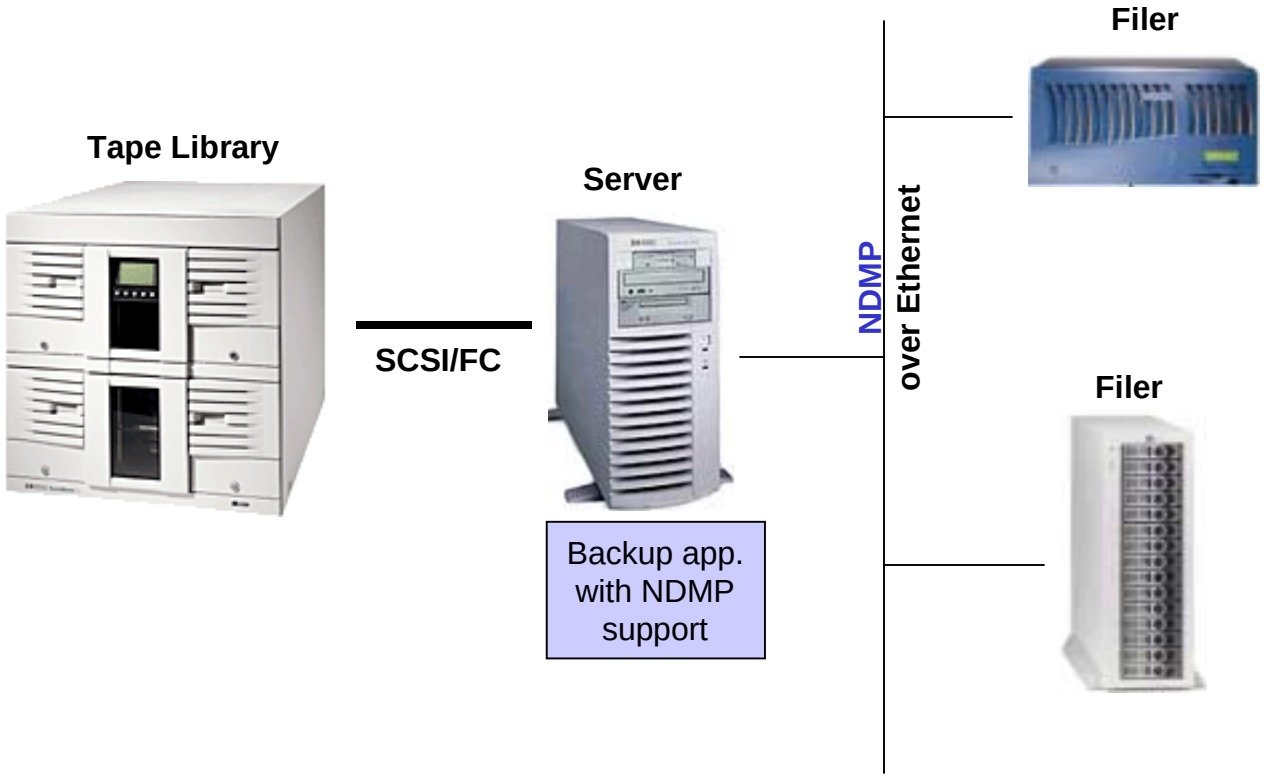

Great for larger data centers. Allows centralized backup without needing a SAN to connect filers to library.

NAS Backup Strategies

Filer to Server (NFS/CIFS) ☹

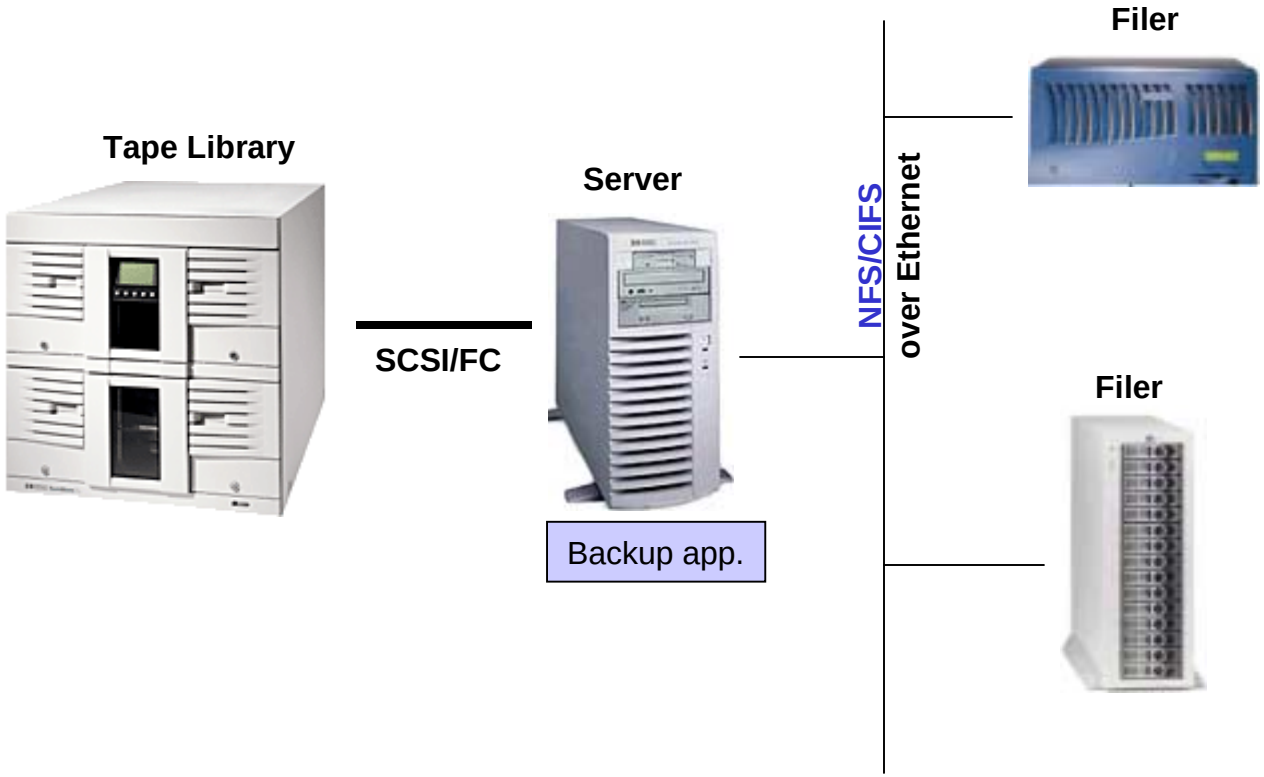

Doesn't require NDMP, but NFS/CIFS over TCP is very slow!
DAFS will help...

Technology Survey

Getting past the 'Cons'

- 10/100Base LANs are too slow?
  - **Solution: Gigabit Ethernet, 10-GbE**
- The TCP stack is too fat and slow?
  - **Solution: TCP Offload Engines on NICs**
- File systems (NFS/CIFS) are too slow?
  - **Solution: NDMP, iSCSI, and DAFS**
- Technology is too new?
  - **Solutions are maturing...**

Technology Survey

Gigabit Ethernet (GbE), 10-GbE

- Fully backward compatible ethernet link
- 125 MBytes / sec line speed
- Available as an option in most filers
- Many server NICs available
- Requires CPU horsepower for full advantage:
  - 1GHz CPU, 64/66 PCI bus
- 10 Gbit products available, but not many yet...

Technology Survey

TCP Offload Engines (TOE)

- Handles TCP checksums and copying in NIC hardware
- Takes the load off the server CPU
- Available in almost every GbE NIC today
- Look for even smarter devices in the future
- Virtual Interface Architecture (VIA) in hardware, used by DAFS

Technology Survey

NDMP, iSCSI, DAFS

- NDMP: Network Data Management Protocol
  - Block-level, sockets-based protocol to transfer data between filers, servers, libraries
  - Available in most filers, some backup apps, a few libraries
- iSCSI: SCSI over Ethernet (SoE)
  - Products announced by Cisco, IBM, EMC, HP
  - Look for iSCSI support for disk arrays and libraries by mid-2002
  - Also look for Nishan's SCSI over IP (SoIP)
- DAFS: Direct Access File System
  - Uses the Virtual Interface Architecture (VIA) to move data between devices without OS, kernel, or TCP stack overhead.
  - Like DMA over the network
  - Predicted availability in 2003

Technology Survey

iSCSI

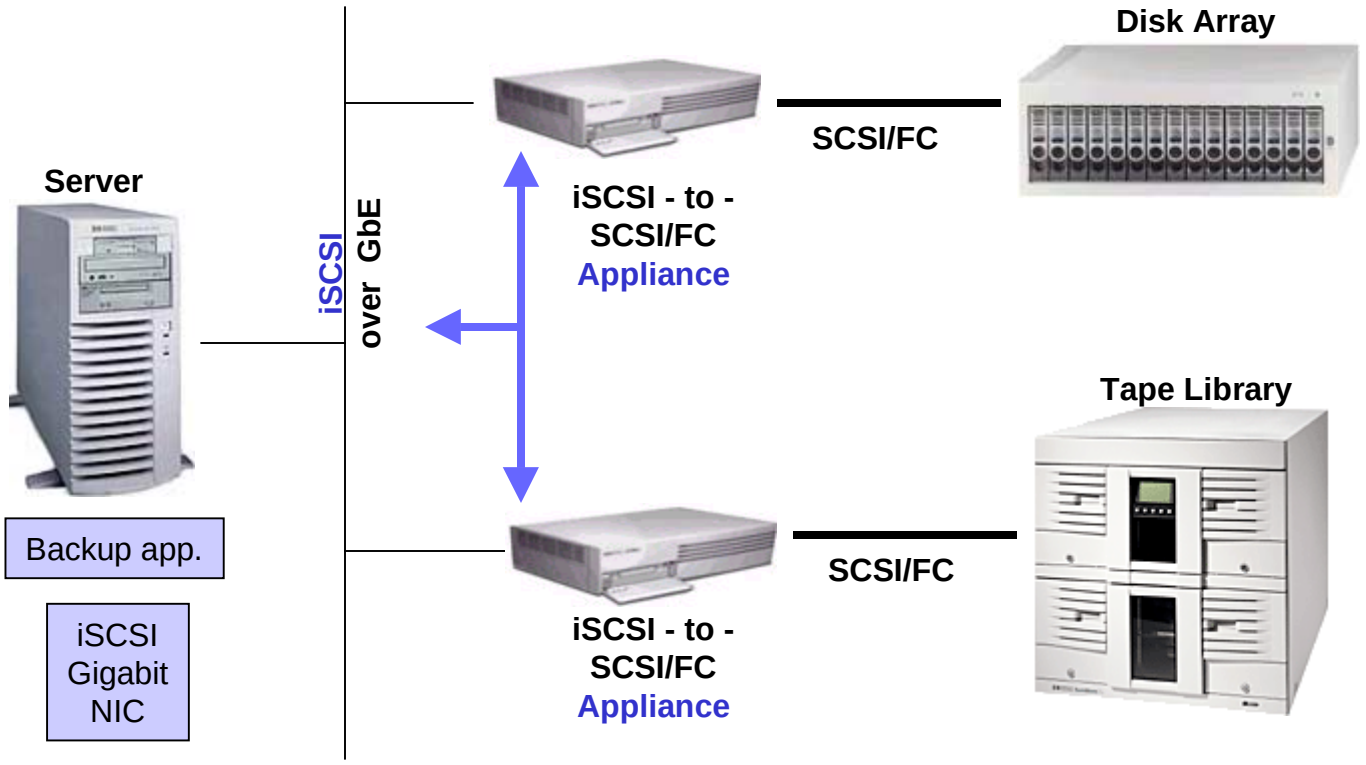

Uses your current infrastructure for disk and tape access.
Intelligent Appliances for easy installation and management

Products

- NAS Disks and Arrays
  - HP's AutoBackup, NAS-VA, and NAS-XP series
  - NetApp's 700 and 800 series
- Gigabit Ethernet NICs
  - HP, Intel, 3Com, NetGear, SysConnect, more...
- Disk Arrays
  - HP's VA and XP series
- Tape Libraries
  - HP's library family: 1-20 drives, 9-700 slots
- Tape Drives
  - HP's Ultrium 230 (LTO) and DLT8000 and drives
- iSCSI Appliances: *coming soon!*

Wrap-Up

- Hardware: "NAS isn't just filers anymore!"
- Protocols: NDMP (now), iSCSI (soon), DAFS (future)
- Storage Appliances:
  - any SCSI device can soon be connected over the LAN
  - Intelligent appliances mean easy manageability and lower cost of ownership
- Infrastructure: GbE switches and NICs available today

BACKUP SLIDES

(no pun intended)

- Why NAS? -
and
- NAS Evolution -

Why NAS?

My legacy environment:

Each workstation has its own disk and take drive.

Eventually, I'll need to add more servers, more disks, more backup…not very scaleable!

Legacy environment

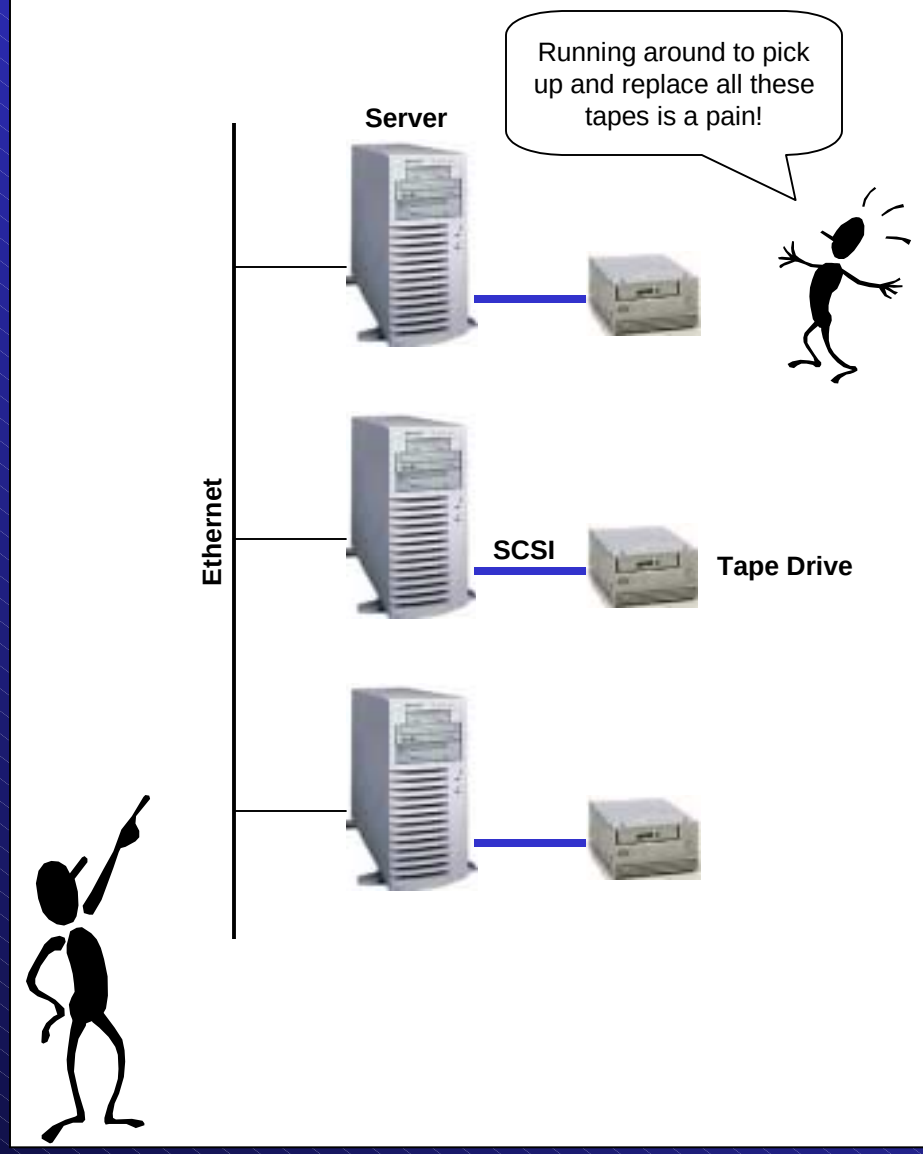

Why NAS?

I can add a disk array and a centralized "file server" that all my workstations can mount to for more storage: much more cost effective (for disks)!

I could save even more time if I didn't have to build my own file server and set up the disk array...

Centralized file server

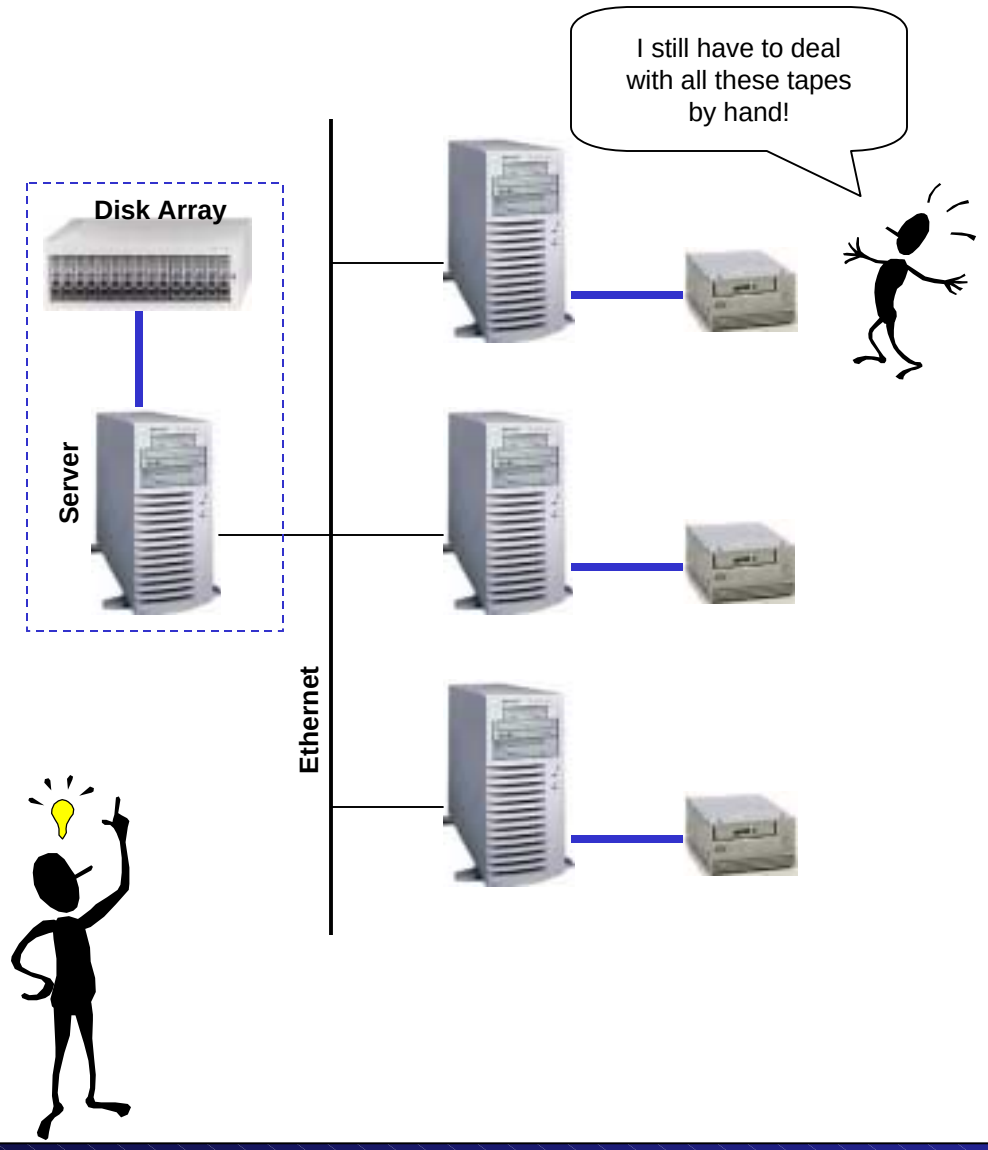

Why NAS?

Now I have an integrated disk array and server (a filer) to save setup time.

I've also centralized my backup instead of using a lot of stand-alone tape drives!

NAS Filer

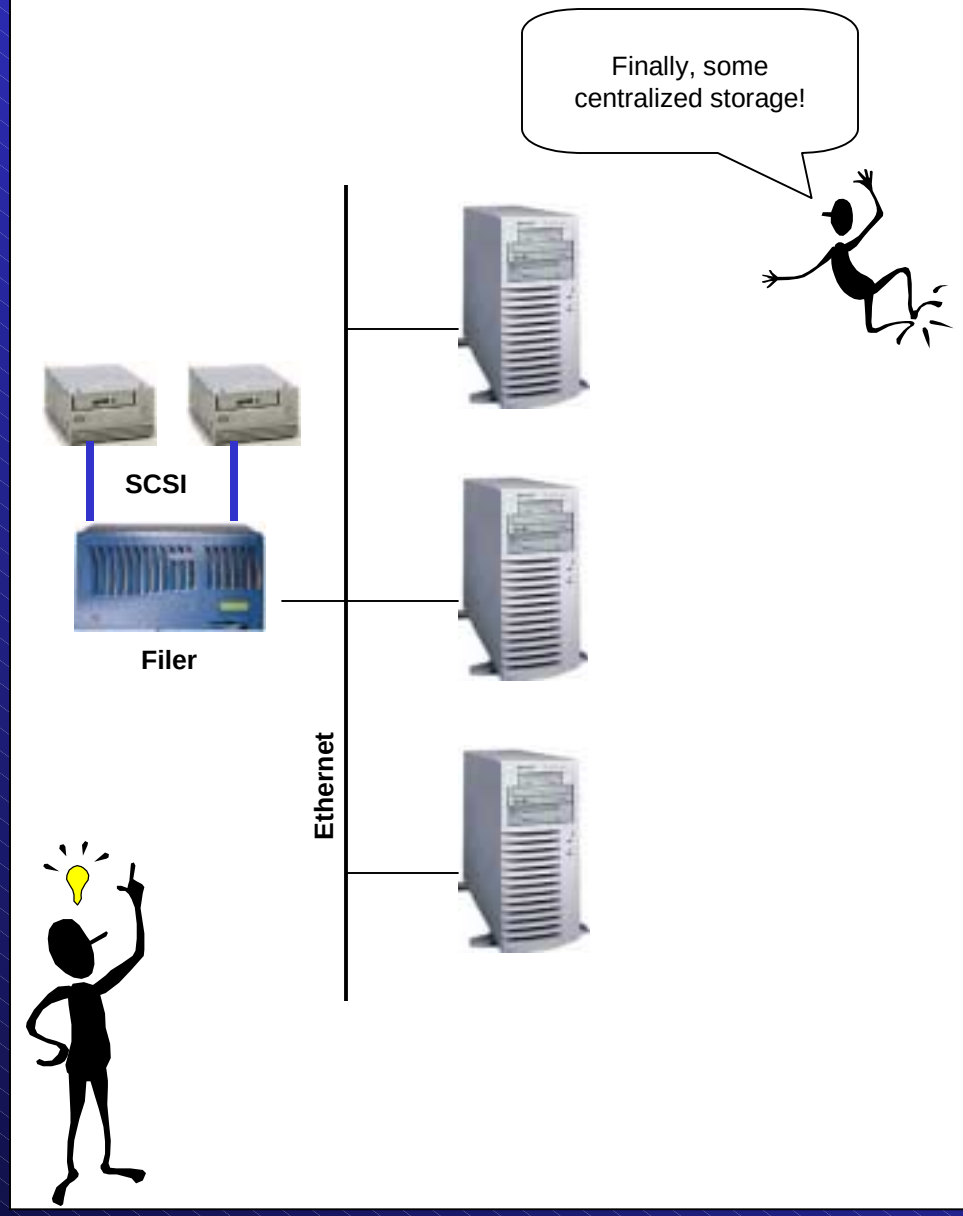

Why NAS?

To really make things easier, I can automate my backups with a tape library…

But what if I need to add more (disk) filers or (tape) backup? How will they connect?

NAS Filer with automated backup

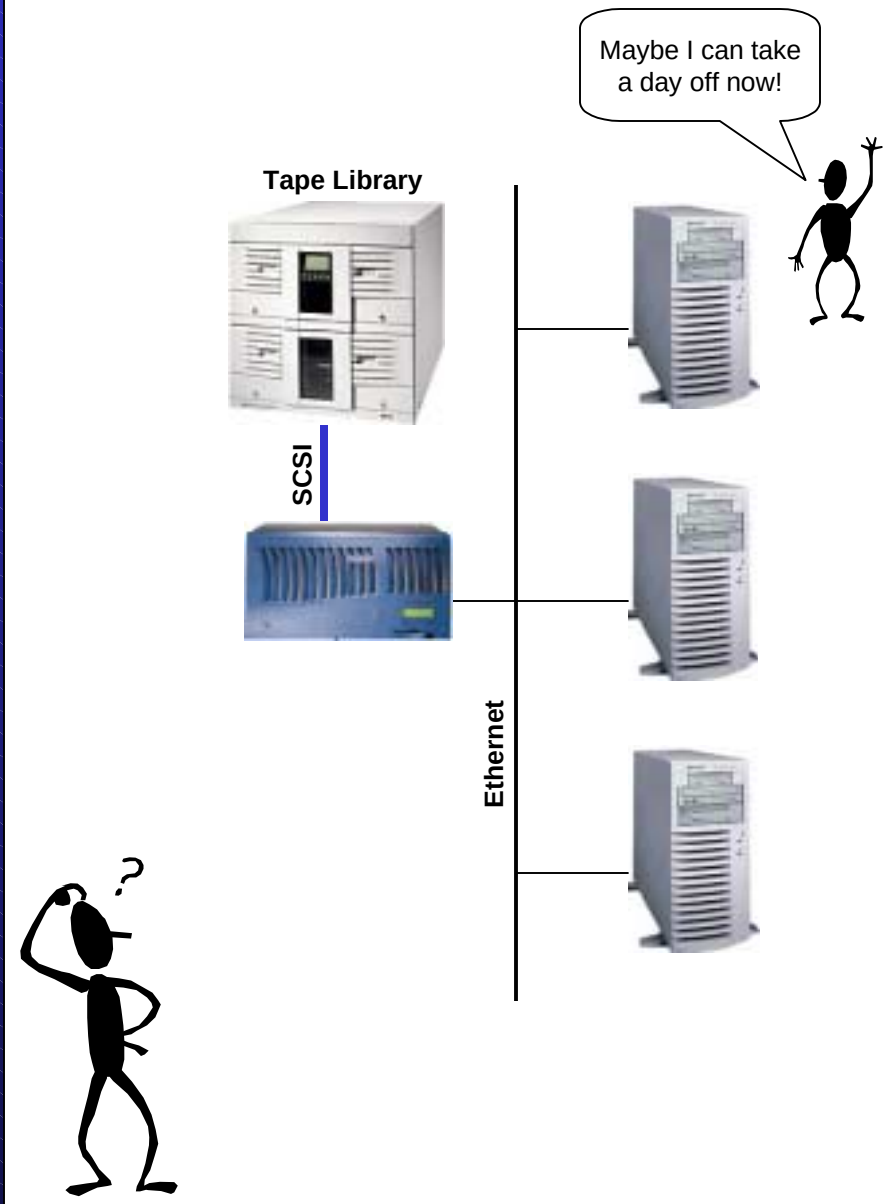

NAS+SAN?

With a SAN,
I can add servers,
filers, or libraries
whenever I want…

But I've created a new
type of network that
my team needs to
learn…and will all these
devices really work
together?

NAS Disks + SAN Tapes and Servers

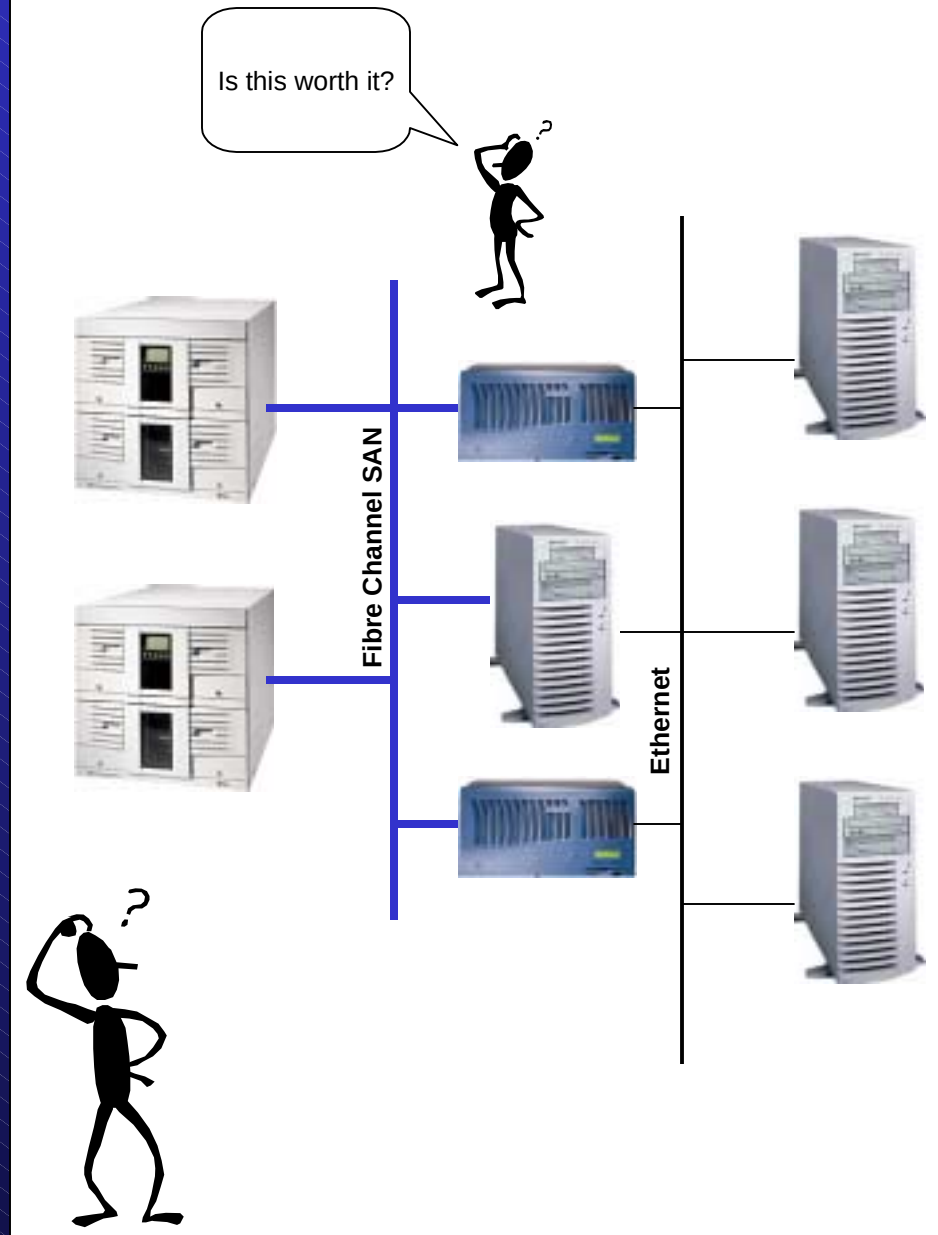

NAS Evolution

This would be nice!

Servers, filers, and libraries working together on the LAN my team already understands!

NAS Disk and NAS Tape

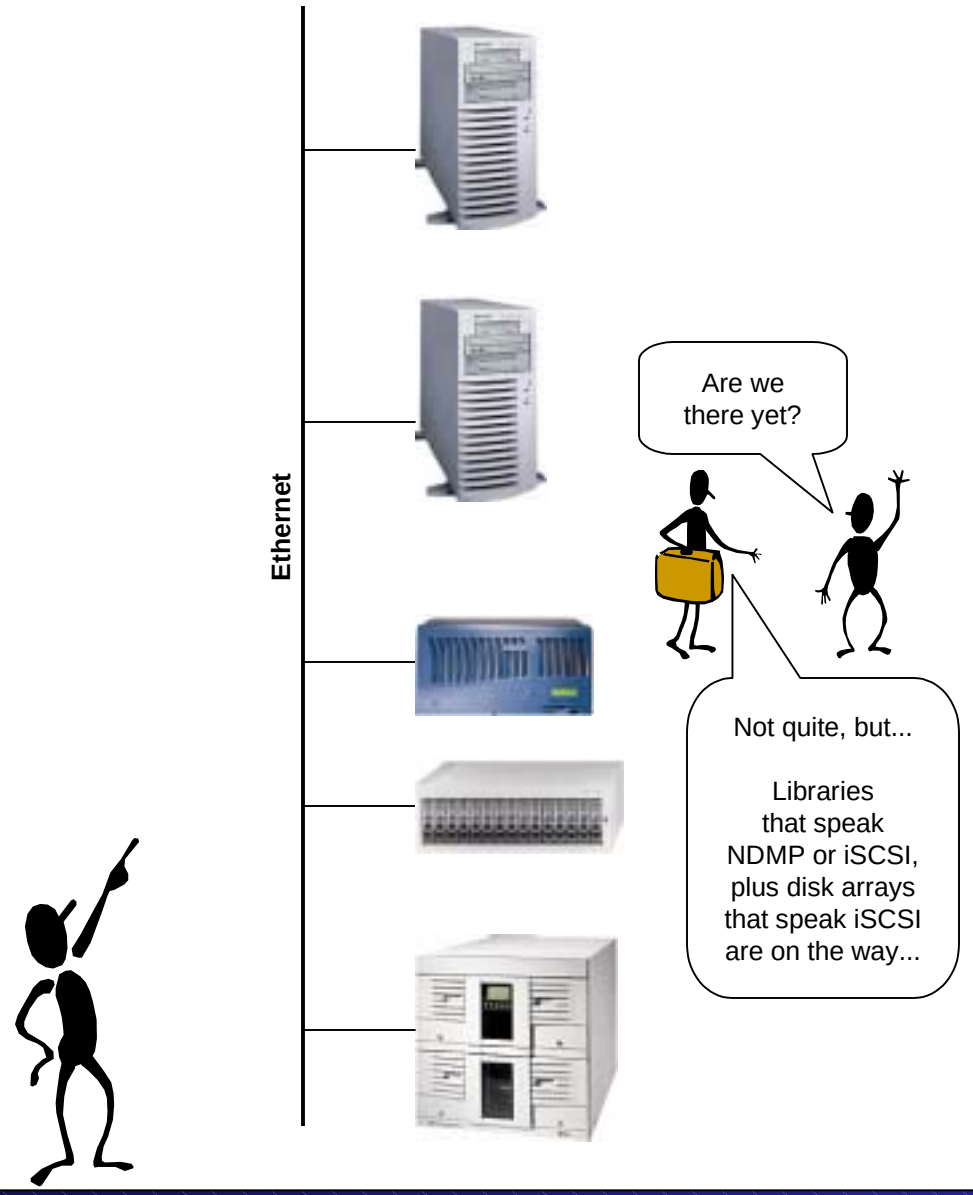

NAS Evolution

Servers, disk arrays, and tape libraries: all talking iSCSI.

Scaleability for disk and tape without a new infrastructure!

Gigabit Ethernet, 10-GbE, and TOE NICs will help!

Look for this in 2002...

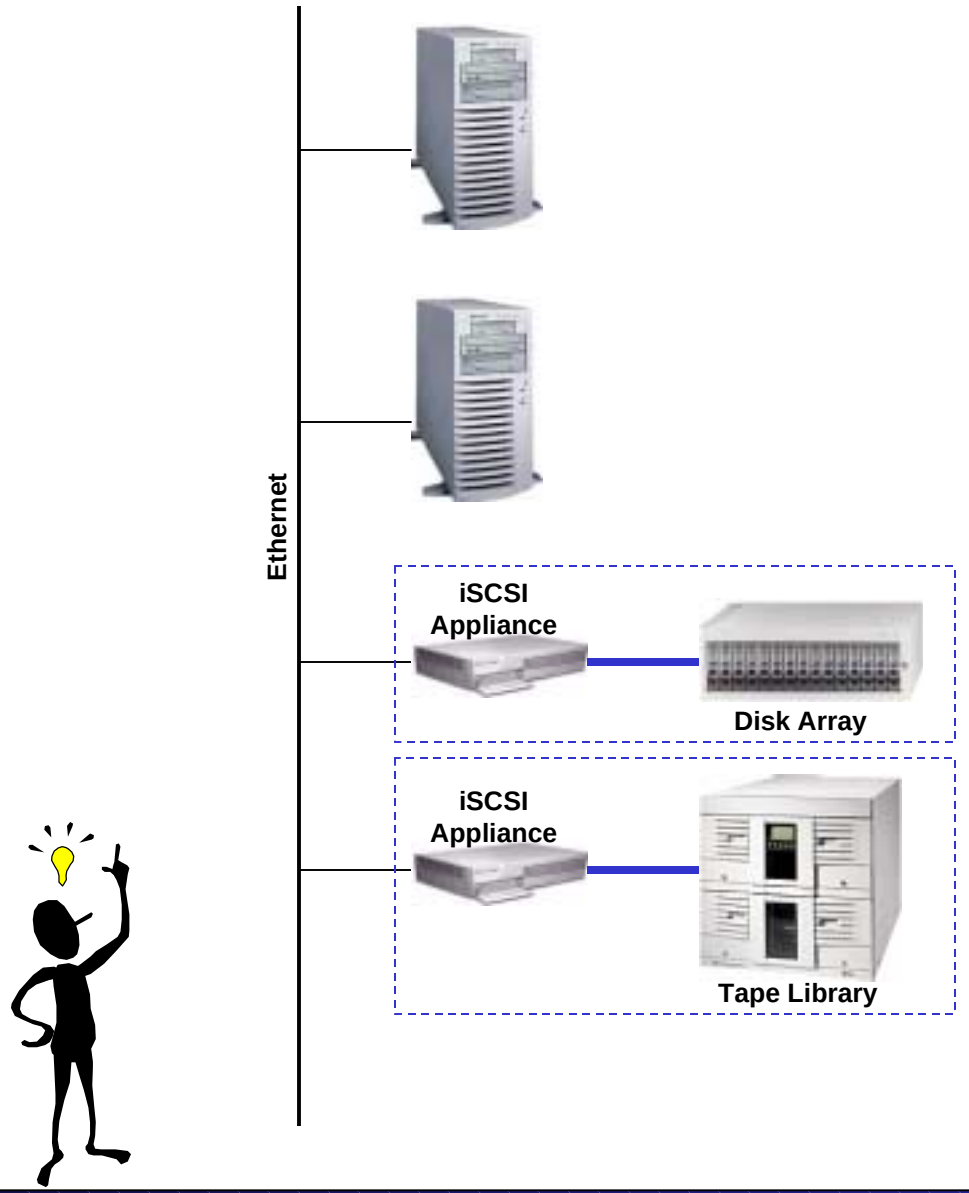

BACKUP SLIDES

- Gigabit Ethernet -

Gigabit Ethernet Backup

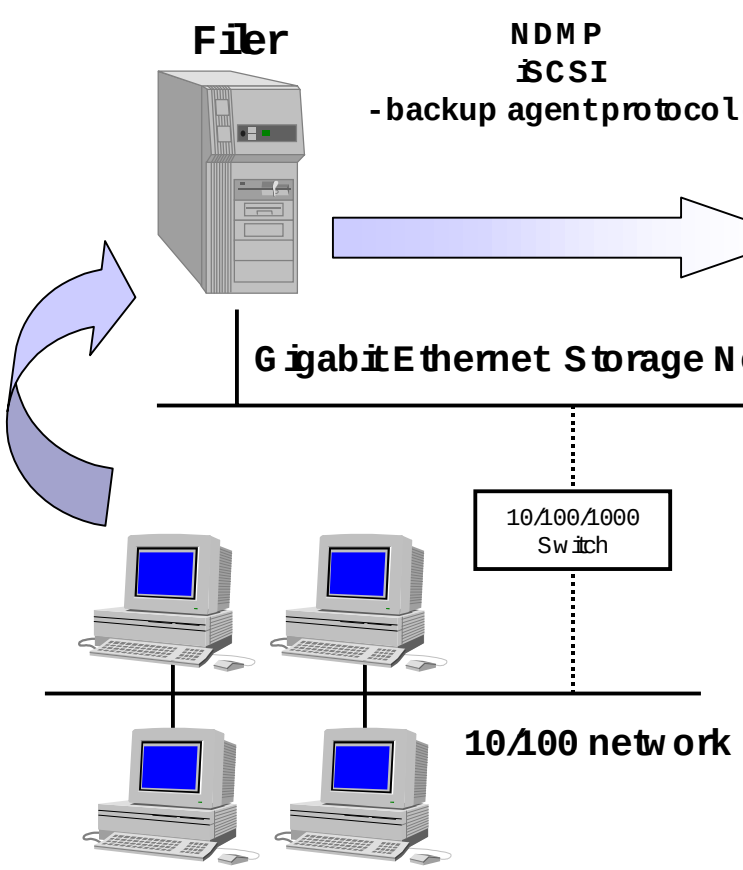

Making full use of gigabit speeds [For Backup]

- Filer consolidates disk storage
- Workstations mount to filer
- Tape library consolidates backup
- Filer backed up (over GbE) using:
  - NDMP
  - Future: iSCSI
  - or backup agent protocl

What kind of speed should I really expect?

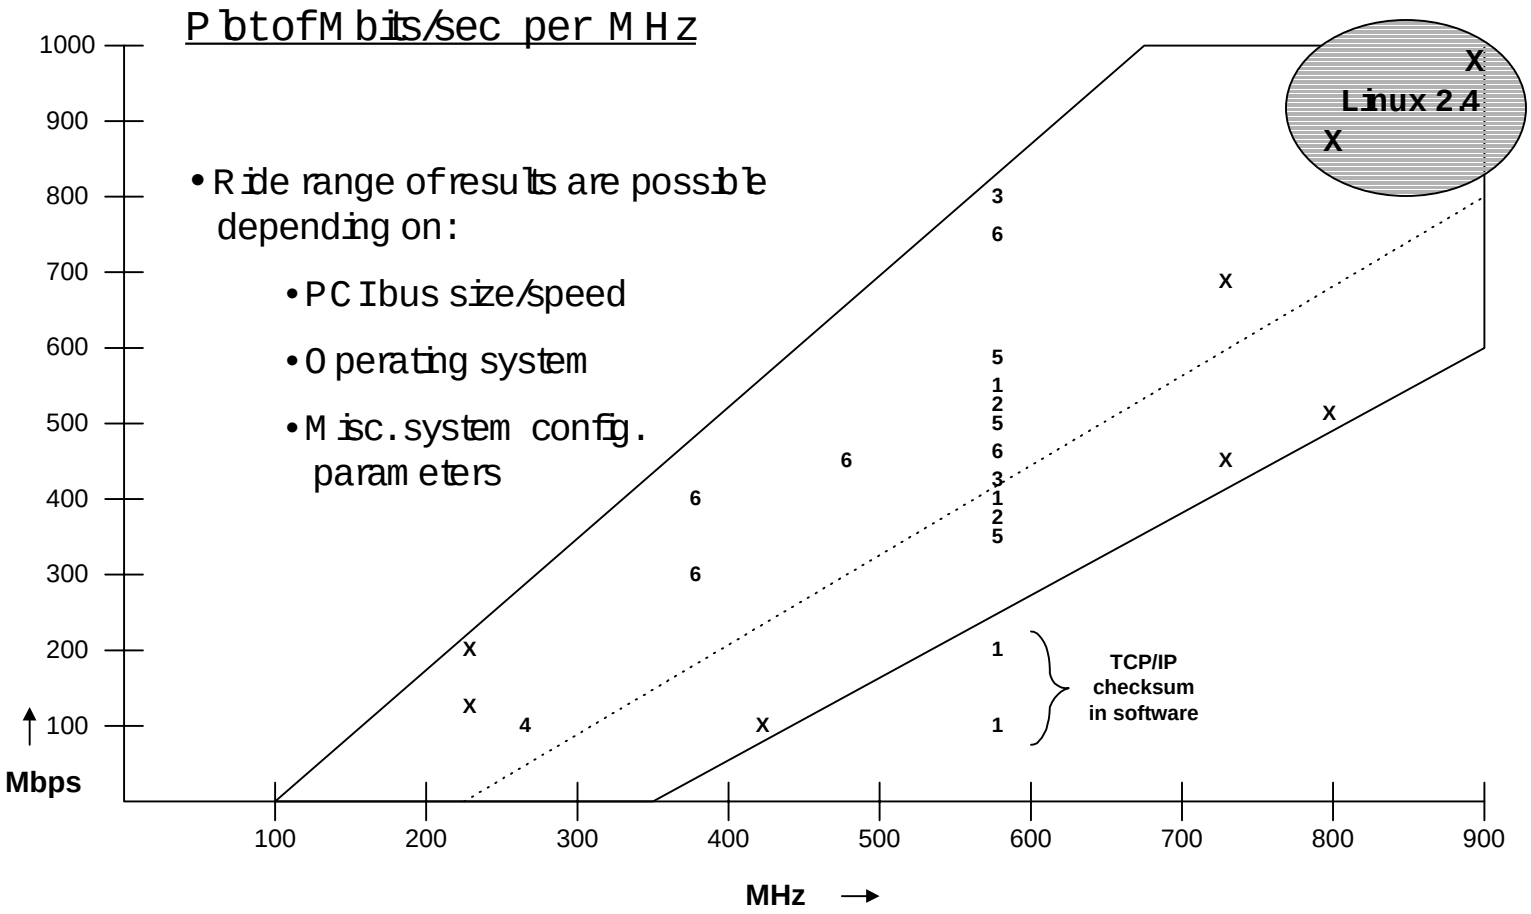

Data comes from: University studies, Trade periodicals, HP research

Plot of Mbits/sec per MHz: Selected Sources

Duke computer science dept. study
http://www.cs.duke.edu/ari/trapeze/slides/freenix/index.htm

"Gigabit choke points"
http://www.nwfusion.com/new/1999/0705feat/html?nf

"Gigabit Ethernet hits second gear"
http://www.nwfusion.com/research/2000/0320revgig.html/?nf

"Gigabit Ethernet Switches Set to Take On The Enterprise"
http://www.nwc.com/916/916rlside4.html

"Boosting Network Server Adapter Performance by 800%"
http://www.tradespeak.com/htmldocs/1897.html

"Gigabit Linux Server Bottlenecks"
http://www.cs.helsinki.fi/linux/linux-kernel/2000-06/0268.html

"Selected Benchmark results obtained on the IBM cluster"
http://www.scl.ameslab.gov/Projects/IBMCluster/Benchmarks.html

"Alteon Networks Configuration"
http://archiv.tu-chemnitz.de/pub/1998/0010/data/vortag/fischer/Perf04.html

One other good source for real-world data has been the Linux-acenic mailing list:
Send mail to the following for info and FAQ for this list:
<linux-acenic-info@sunsite.auc.dk>
<linux-acenic-faq@sunsite.auc.dk>

Other data from HP's Automated Storage Division testing