The Evolution of Storage Area Networks (SAN's)

Winston W. Hait
Senior Product Manager
Syncsort, Inc.
50 Tice Boulevard
Woodcliff Lake, NJ 07677
Phone: (201) 573-5179, Fax: (201) 930-8290
E-Mail: whait@syncsort.com

Overview - From where to where?

- The Movement behind SAN's
- Examine Network/Hardware-the push begins
- 3 Stages of SAN's Evolution
- SCSI vs. Fibre Channel
- LAN-less Backup
- 3rd Party Copy
- InfiniBand
- SSP's
- NDMP in the SAN-7 different configurations

What is a SAN?
Storage Area Network

- A SAN is a secondary network whose primary focus is off loading the traffic associated with data storage and movement from a primary network

The Movement Behind SAN's: Explosion of Data and Technology

- Data Growing 50-400% per year
- "Internet storage estimates say that capacity needs will double every three months"
  - Network World Fusion, March 15, 2000
- Overall forecasted installed storage capacity worldwide for 2003 is 300 times the capacity installed in 1993
  - IDC, Mass Storage News, March 2, 2000

The Movement Behind SAN's: Explosion of Data and Technology

- By 2004, the 100 largest companies in the world will have an average storage capacity that exceeds 150 TBs.
- - Forrester Research, 12/4/00- Computerworld magazine

The Movement Behind SAN's: Explosion of Data and Technology

- How handle the explosive growth in an easily managed, highly dynamic environment?

SAN's - Storage Area Networks

The Movement Behind SAN's: Explosion of Data and Tape Tech

- Move from Centralized to Distributed...and back
- Data storage and backup - backing up to tape has been done for years starting with reel to reel
  - DLT came along and helped revolutionize tape because of its speed, capacity and reliability
  - 4mm and 8mm were good, and now are overcoming quality and ruggedness issues

From Centralized to Distributed
The Evolution Begins

- Large data centers centralized administrative tasks including data storage and backups
  - Not much data moving at first over network
  - Backups not a big impact to the corporate enterprise since they are done locally
- Distributed computing becomes the rage
  - Reduced TCO (total cost of ownership)
  - Simpler (supposedly)
  - Lower loads on the over all network(s)
  - Administrative nightmare

From Centralized to Distributed
The Evolution Begins

- Data growth begins, the network impact grows
  - More data traversing the network
  - Link and node utilization's rise
  - Off-hours backups necessary
  - Must add additional storage, but how manage?
  - Distributed backups to isolated tape & libraries
  - Which systems should get backed up vs. do
- Data continues to grow, on a geometrical basis
  - Networks have a hard time keeping up

SAN's Emerge
Goal: Off-load increasing traffic

- Goal/purpose: a Storage Area Network (SAN) is a dedicated, storage-only secondary network that off-loads storage traffic from the primary or enterprise network
- Helps reduce the impact of data movement
- Remember, server backups are only 20% of the equation, the other 80% is the restore and how quickly you can do it
  - 10% and 90% ???

Review:

Let's review
networking capacities
and
backup media/hardware

or....

Where is the bottleneck?

Step 2 - Review:
Examine network speeds

In a perfect world...

- 10Base-T: **3.6GB/hr**
- 100Base-T: **36 GB/hour**
- 1000Base-T: **360 GB/hour**
- 10 GBit Ethernet
  - IEEE 10GBit task force-spec not ready until March'02 says NWW-6/00
  - Ten-Gigabit Ethernet is full duplex, just like full-duplex Fast Ethernet and Gigabit Ethernet, and thus has no inherent distance limitations

Step 2 - Review:
Examine network speeds

In a perfect world...

- FDDI: **36 GB/hour**
- ATM: **270 GB/hour**
- Fibre Channel: **360 GB/hour**
  - 200MB/Sec now shipping
- Enterprise Management Associates' research consistently finds that SANs today are made up of combinations of Ethernet, Fibre Channel and SCSI technology.

Review:
Examine network speeds

Let's be practical...

Ethernet - Divide by 2

Fibre - Think 90%

Review:
Hardware Specifications: 1/2”

3480	1.5 MB/sec	.2GB cap	5.4 GB/hour
comp*	3 MB/sec	.4GB	10.8 GB/hour
3490	3 MB/sec	.4GB cap	10.8 GB/hour
comp*	6 MB/sec	.8GB	21.6 GB/hour
3490E	3 MB/sec	.8GB cap	10.8 GB/hour
comp*	6 MB/sec	1.6GB	21.6 GB/hour
3590 Magstar	9 MB/sec	10GB cap	32.4 GB/hour
comp*	18 MB/sec	20GB	64.8 GB/hour

* based upon 2:1 compression

Review:
Hardware Specifications: 4mm

DDS2	336KB/sec	4GB cap	1.2 GB/hour
comp*	772KB/sec	8GB	2.4 GB/hour
DDS3	1.2MB/sec	12GB cap	4.32 GB/hour
comp*	2.4MB/sec	24GB	8.64 GB/hour
DDS4	3MB/sec	20GB cap	10.8 GB/hour
comp*	6MB/sec	40GB cap	21.6 GB/hour

* based upon 2:1 compression

Review: Hardware Specifications: 8mm

8mm Mammoth	3MB/sec	20 GB	10.8 GB/hour
comp*	6MB/sec	40 GB	21.6 GB/hour
8mm Mammoth-2	12MB/sec	60 GB	43.2 GB/hour
comp**	30MB/sec	150 GB	108 GB/hour
8mm Mammoth-3	20MB/sec	120 GB	72 GB/hour
comp**	50MB/sec	300 GB	180 GB/hour

- due Q2 or Q3/2001

* based upon 2:1 compression ** based upon 2.5:1 compression

Review:
Hardware Specifications: 8mm

8mm AIT	3MB/sec	25 GB cap	10.8 GB/hour
comp*	6MB/sec	50 GB	21.6 GB/hour
8mm AIT-2	6MB/sec	50 GB cap	21.6 GB/hour
comp*	12MB/sec	100GB	43.2 GB/hour
8mm AIT-3	11MB/sec	100 GB cap	39.6 GB/hour
comp*	22MB/sec	200GB	79.2 GB/hour

- due mid-2001

* based upon 2:1 compression

Review:
Hardware Specifications: DLT

DLT 4000	1.5 MB/sec	20 GB cap	5.4 GB/hour
comp*	3.0 MB/sec	40 GB	10.8 GB/hour
DLT 7000	5 MB/sec	35 GB cap	18 GB/hour
comp*	10 MB/sec	70 GB	36 GB/hour
DLT 8000	6 MB/sec	40 GB cap	21.6 GB/hour
comp*	12 MB/sec	80 GB	43.2 GB/hour

* based upon 2:1 compression

Review:
Hardware Specifications: Latest

Ecrix (8mm)	3 MB/sec	33GB	10.8GB/hour
comp*	6 MB/sec	66GB	21.6GB/hour
DLT1	3 MB/sec	40GB	10.8GB/hour
comp*	6 MB/sec	80GB	21.6GB/hour
9840 (STK)	20 MB/sec	20GB	54GB/hour
comp***	80 MB/sec	80GB	288GB/hour(7-to-1?)
SuperDLT	15 MB/sec	100-500GB	54GB/hour
comp*	30 MB/sec	200GB-1TB	108GB/hour

* based upon 2:1 compression *** based upon 4:1 compression

Review:
Hardware Specifications: Latest

LTO (Ultrium)	15 MB/sec	100-500GB	54GB/hour
comp*	30 MB/sec	200GB-1TB	108GB/hour
1/2” tape			

LTO (Accelis)	15 MB/sec	25GB	54GB/hour
comp*	30 MB/sec	50GB	108GB/hour

- 8mm tape - mid-point mount (10 second access)

* based upon 2:1 compression

Three Stages of SANs Evolution:

- Multi-Hosting
- Arbitrated Loop
- Switched Fabric

First Stage in Evolution: Multi-Hosting

- Multi-hosting – direct SCSI connect server to tape library
  - Cheap – cost of SCSI cables
  - Fast – 40MB/sec –most tape drives can only write at 5MB/sec
  - Off-loads traffic and maximizes tape investment – speeds are fast enough to keep tape spinning all the time
  - Only downside – distance - 25 meter limit of SCSI

First Stage in Evolution:
Multi-Hosting

Network Speed = 10 MB/sec

Library Speed = 20 MB/sec

-10 MB/sec
Lost Throughput

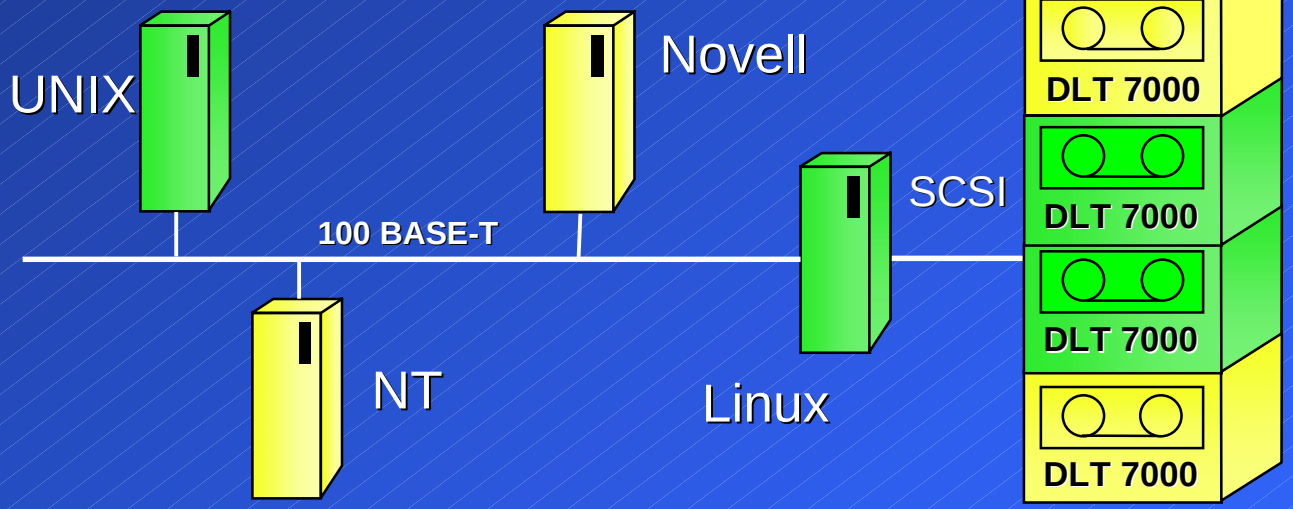

First Stage in Evolution: Multi-Hosting

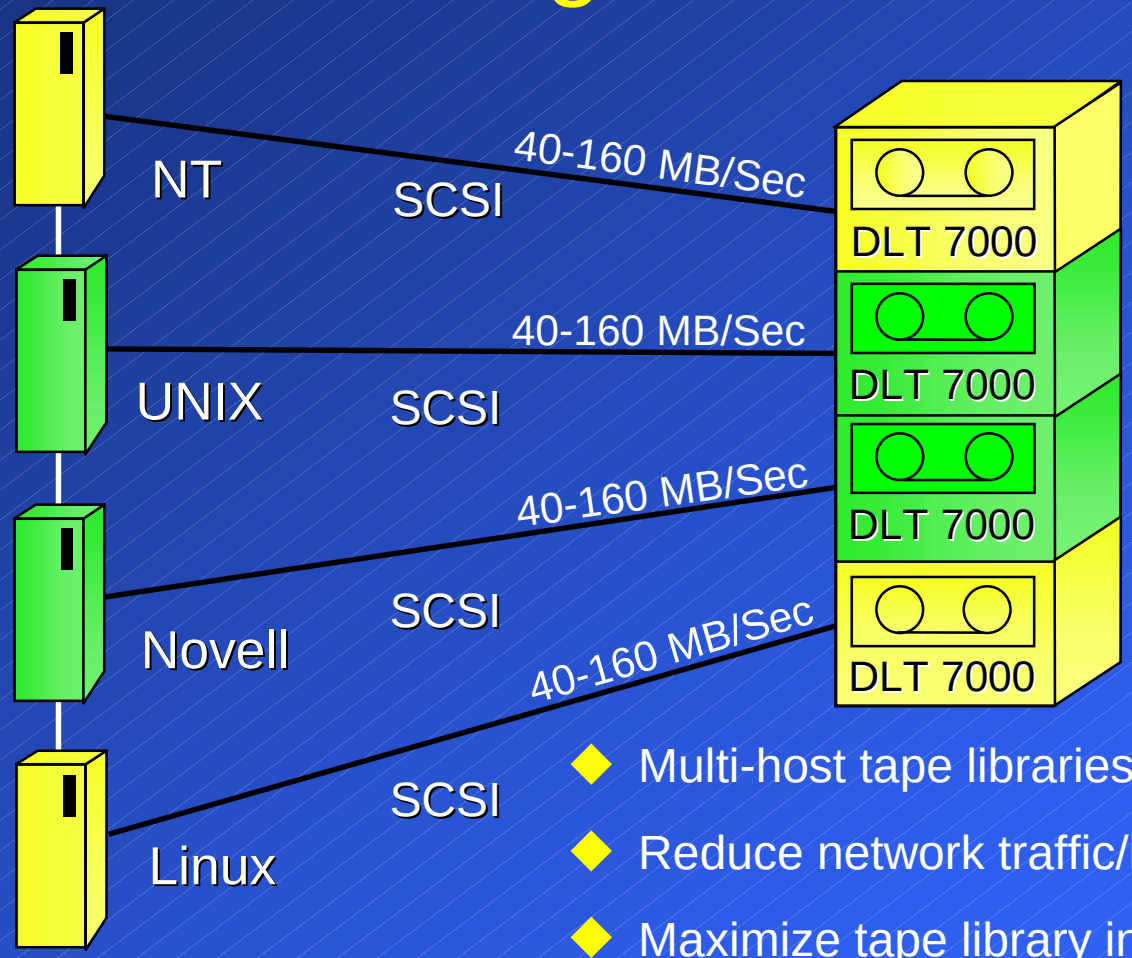

After

TAPE LIBRARY

- Multi-host tape libraries with multiple servers
- Reduce network traffic/improve performance
- Maximize tape library investment

Second Stage in Evolution: Fibre Channel - Arbitrated Loop

- 100 MB/sec speed
- Arbitrated - circuit based - only one conversation at a time without redundant loops, hubs or using a switch
- Not a bandwidth consideration, but latency of arbitrating connections
- 126 devices on one loop vs. 16 for SCSI
  - Current suggested/recommended is 3-6 servers
  - Must bring down entire loop to add or remove a device without node-bypass circuitry or a switch

Second Stage in Evolution:
Fibre Channel - Arbitrated Loop

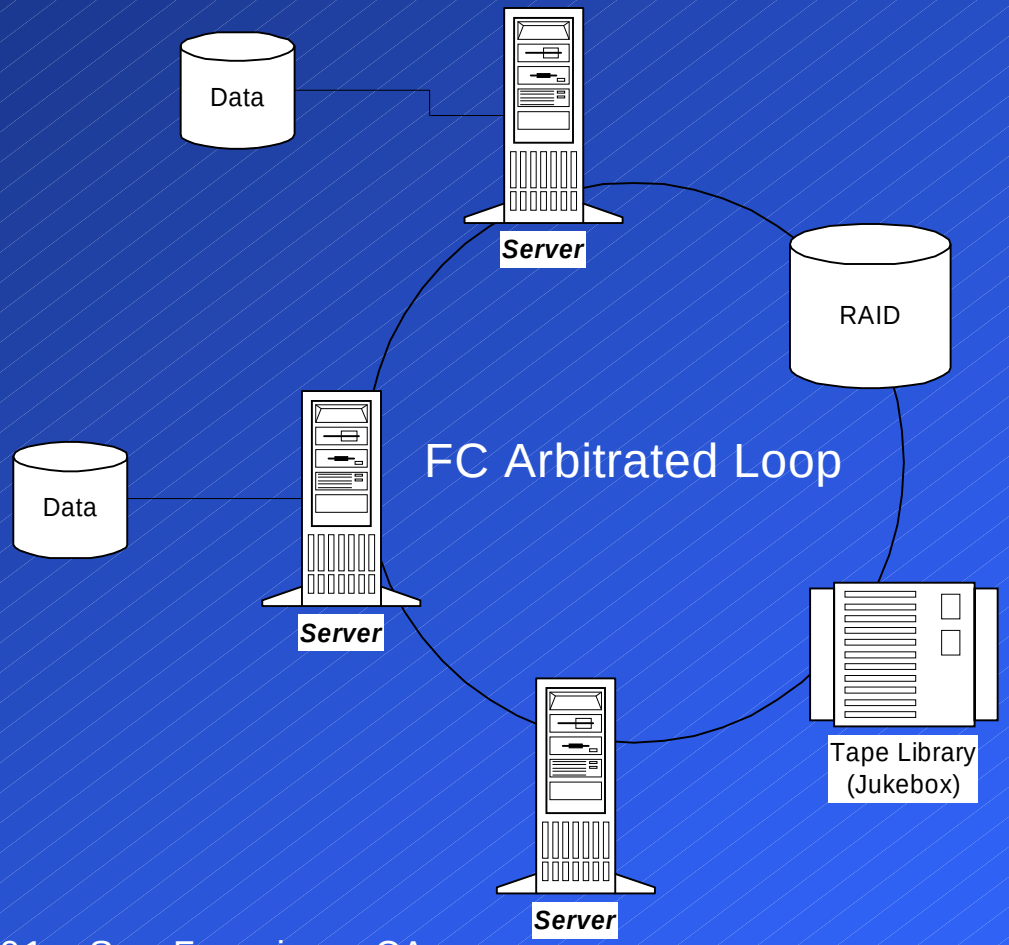

Second Stage in Evolution:
Fibre Channel - Arbitrated Loop

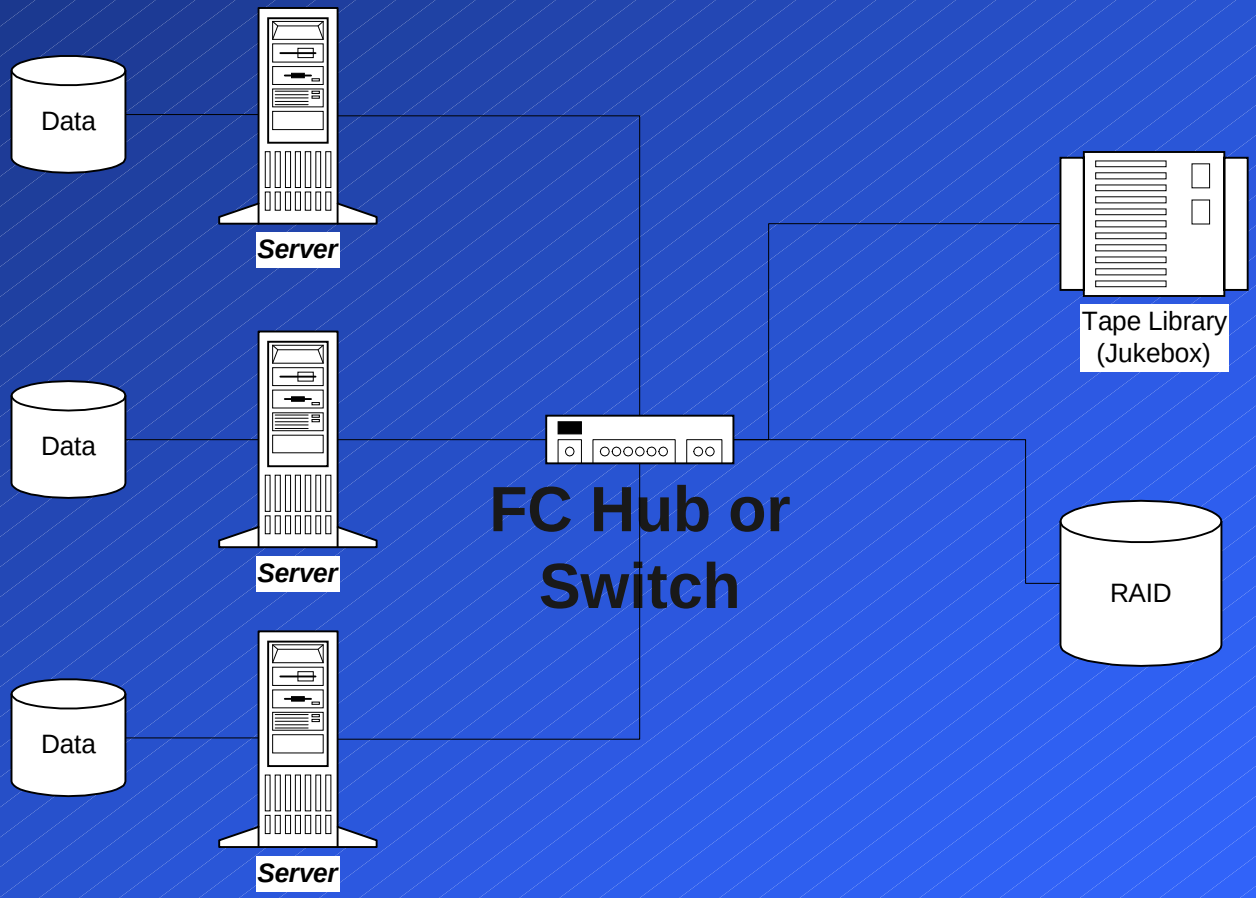

Third Stage in Evolution: Fibre Channel - Switched Fabric

- More expensive
  - Multiple switches, hubs and routers
  - Can help some with SCSI to FC converters at the end points, but the fabric in the middle is where the costs climb quickly
- Fast speeds
  - Multiple paths allow multiple channels of gigabit speed in system
    - 8 simultaneous 100MB/sec pipes through a 16 port switch
- Redundancy/failover
- Zoning

Third Stage in Evolution:
Fibre Channel - Switched Fabric

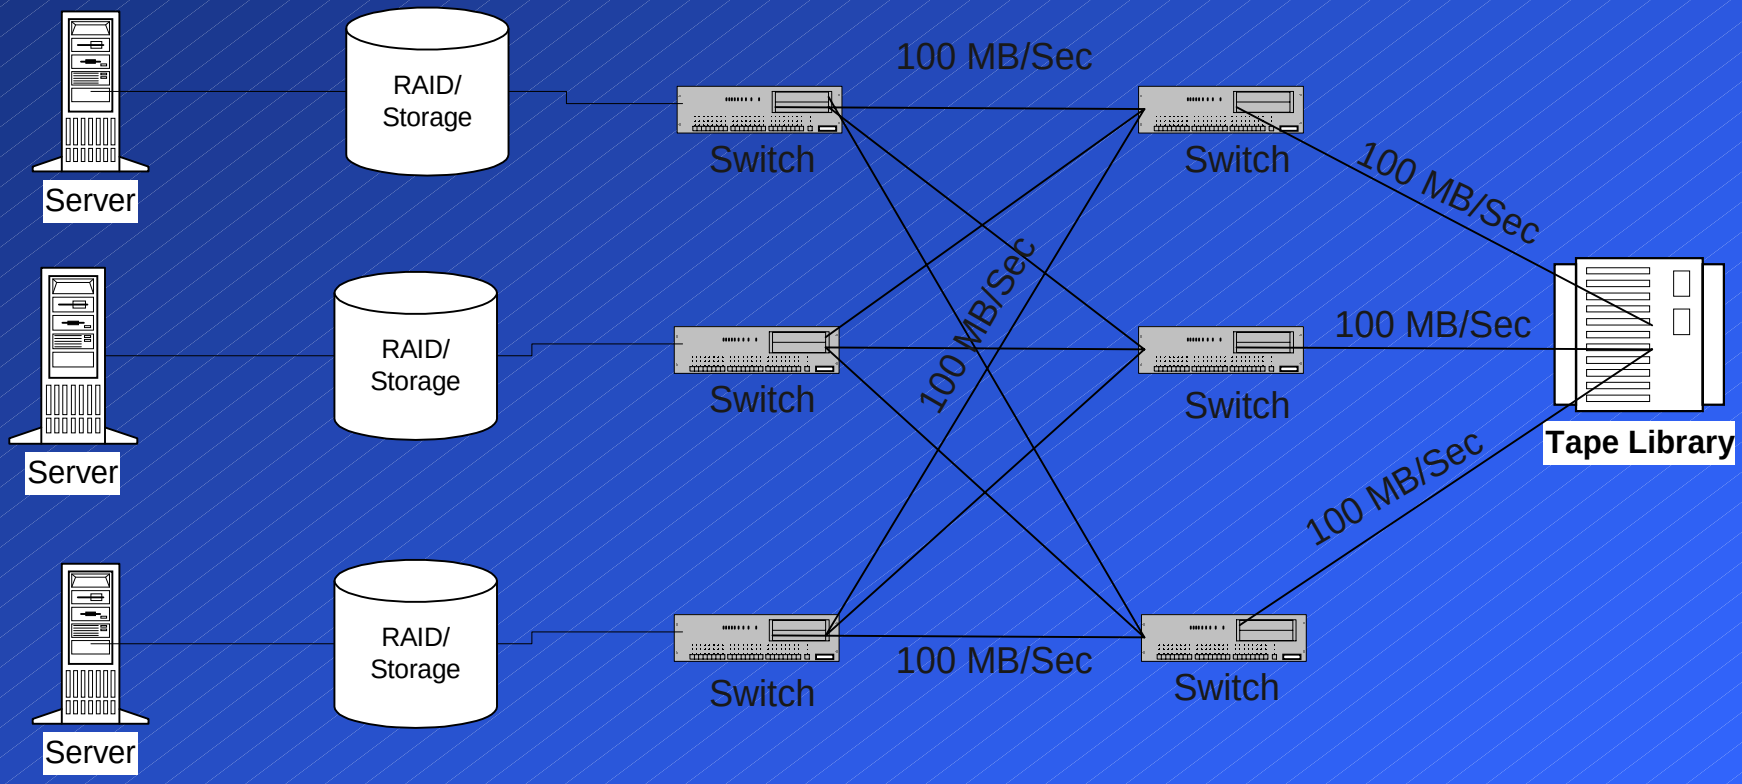

Comparison: SCSI vs. Fibre Channel

- SCSI
  - Risk is low
    - Time tested
    - Widely installed
    - Simple to use
    - Inexpensive - mostly just cost of cables
- Limited to maximum distance of 25 meters from server
- Most versions of SCSI slower than FC
  - UltraSCSI III - 160 MB/sec vs. FC - 100MB/sec
  - FC - 200MB/sec, non-aggregated standard

Comparison: SCSI vs. Fibre Channel

- Fibre Channel - Risk is higher, but is dropping
  - No standard…yet: FibreAlliance & SNIA
- Greater distance – up to 10km on a single run
- Heterogeneous systems are still 1-2 years away
- Cost is higher, but dropping
  - Managed hub - $500-700/port, Switches - $625/port
  - Fully redundant switches (failover) - $4,000/port
- Zoning -
  - Security
  - Fail-over capabilities
  - Management / Load balancing

What is the best for you? Determine your need

- First question - how important is availability
  - - i.e. how much is it worth?
- ERP systems with verbose, highly redundant subsystems
  - worth it
    - Easy to justify
    - Link aggregation can help scale up as needs grows
    - Helps handle data transfer bursts (i.e. mirroring/backups)
- Sales Force Data - lead tracking, etc.
  - May only need simple backup depending on data
  - High volume, try multi-hosting
  - Has the largest backup window

What is the best for you?
Server Farms

- Greatest advantage is distance, or lack there of
  - SCSI
    - Distance isn't a problem
    - Simple – SCSI is an accepted standard
    - Inexpensive – just the cost of cables
    - Multiple servers can be connected to one tape library
    - Configure a server with more than one SCSI adapter
    - Multi-hosting - Multiple channels from one source
  - Fibre Channel
    - More expensive, but could be first step to implementing a SAN
    - Faster speeds - not that important if use multiple SCSI connects

What is the best for you?
Campus or MAN

- FC or SCSI
  - Distance is solved with FC connects or SCSI routers
    - Simplifies management
- Arbitrated Loop (AL)
  - May work if only have a few servers - dedicated loops
- Switched Fabric (SF)
  - If you have great value to the data on a 24/7 basis
  - Expensive, but reliable, flexible and dependable

What is the best for you? Enterprise

- Combination based on distance and needed availability
- Can look to multiple “sites” within the organization
  - One main location, but each building it’s own SAN with aggregated links to main data center
  - Will have to use the same vendor for everything to ensure interoperability … for now
- Depending on amount of data to move, a combination of SCSI, FC-AL and FC-Switched may be the best
  - Beware the upgrade path

What is the best for you?
LAN-Less backup

- Designed to not use your main LAN
- How does it work ?
  - What's Involved ?
  - Special Equipment needed? (if any)
  - How is it controlled ?
- Is it practical ?

What is the best for you?
3rd Party Copy - Is the hype worth it?

- Designed to provide low server utilization for backups
- How does it work?
  - Need equipment that supports 3rd Party Copy
  - Build block list
  - Transfer info
  - Who controls?
  - Problem with Database backups

What is the best for you?
InfiniBand

- What is it?
- How does it work?
- What are its advantages?
- What are the ideal applications to use it with?

What is the best for you?
Storage Service Providers (SSP's)

- Highly Scaleable
- Current Players include: StorageNetworks, Exodus Communications, StorageWay, Managed Storage International (MSI) - STK spin off, Storage Access, WorldStor, CreekPath Systems - Exabyte spin off, Centripetal, Electronic Data Systems, IBM Global Services and StorageNetworks. Partners (acting as market makers or intermediaries with end users) include Enron Broadband Services and Xdrive Technologies.
- Performance based on:
  - 1) Your internet access speed - need 20 kbps per user
    - 20 users directly accessing, need T1 - for 200 - T3
    - DSL? What are your speed req's - upload and download?
  - 2) SLA - at least 99% uptime and low latency (<200ms)

What is the best for you?
Storage Service Providers (SSP's)

- Performance based on: (con't)
  - 3) Eliminate routing delays-stay on one backbone
    - SSP using Quest fibre, get ISP connection from Quest
    - Point - to - point connection?
- Industry - $10 million in 1999, $5.6 billion by 2004 (IDC)
- What happens to your data if an SSP goes out of business?

What is the best for you?
Network Attached Storage (NAS)

- NAS devices - how do they work
- What is NDMP
  - Network Data Management Protocol
- Can they work in a SAN?
  - How?
- How big is the industry?
  - .6MM ‘99 -> 7.2B in ‘04 - Gartner Group/Dataquest

What is the best for you?
Early, SAN-like Multi-Hosting

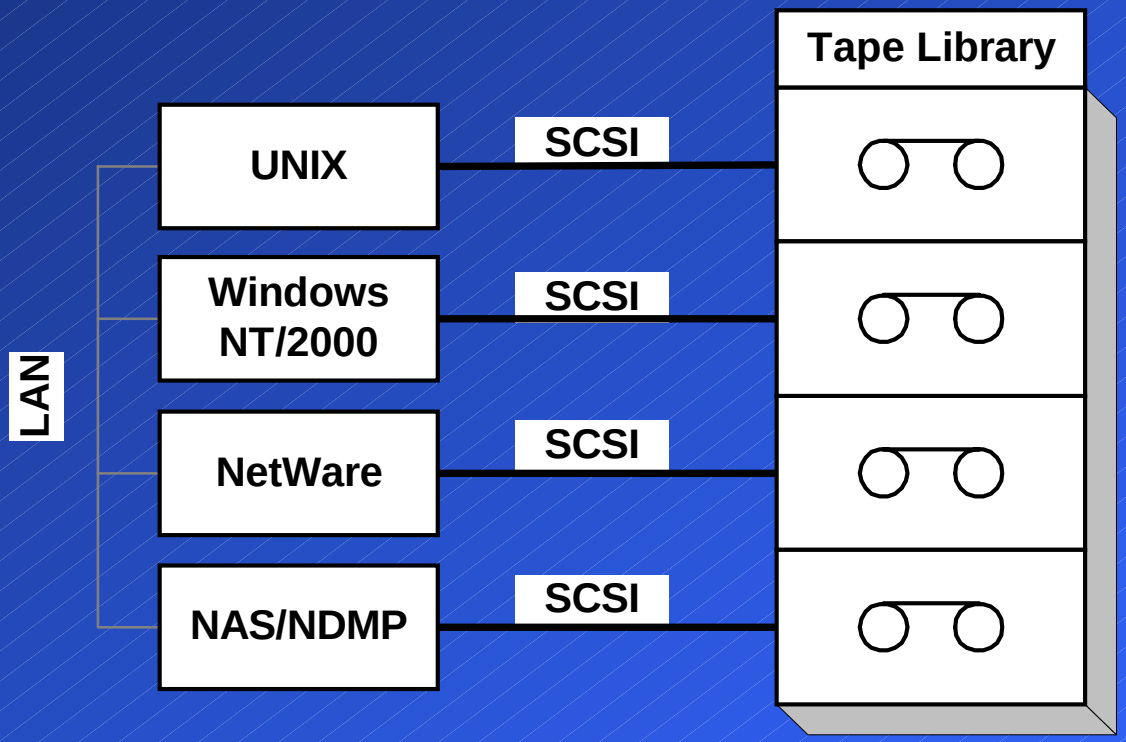

What is the best for you?
Multi-Hosting with SAN & NAS

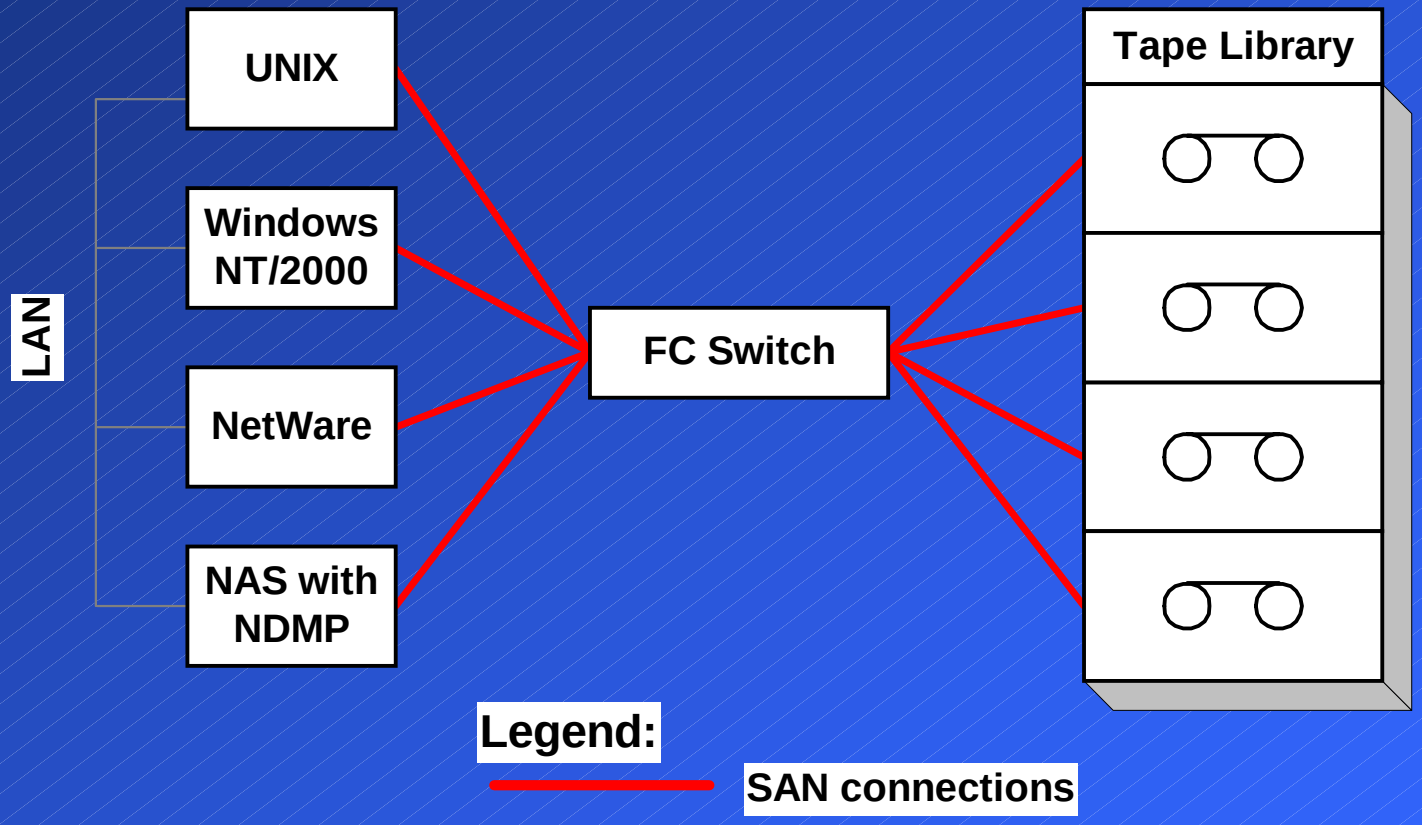

What is the best for you?
Local Backup of NAS via NDMP

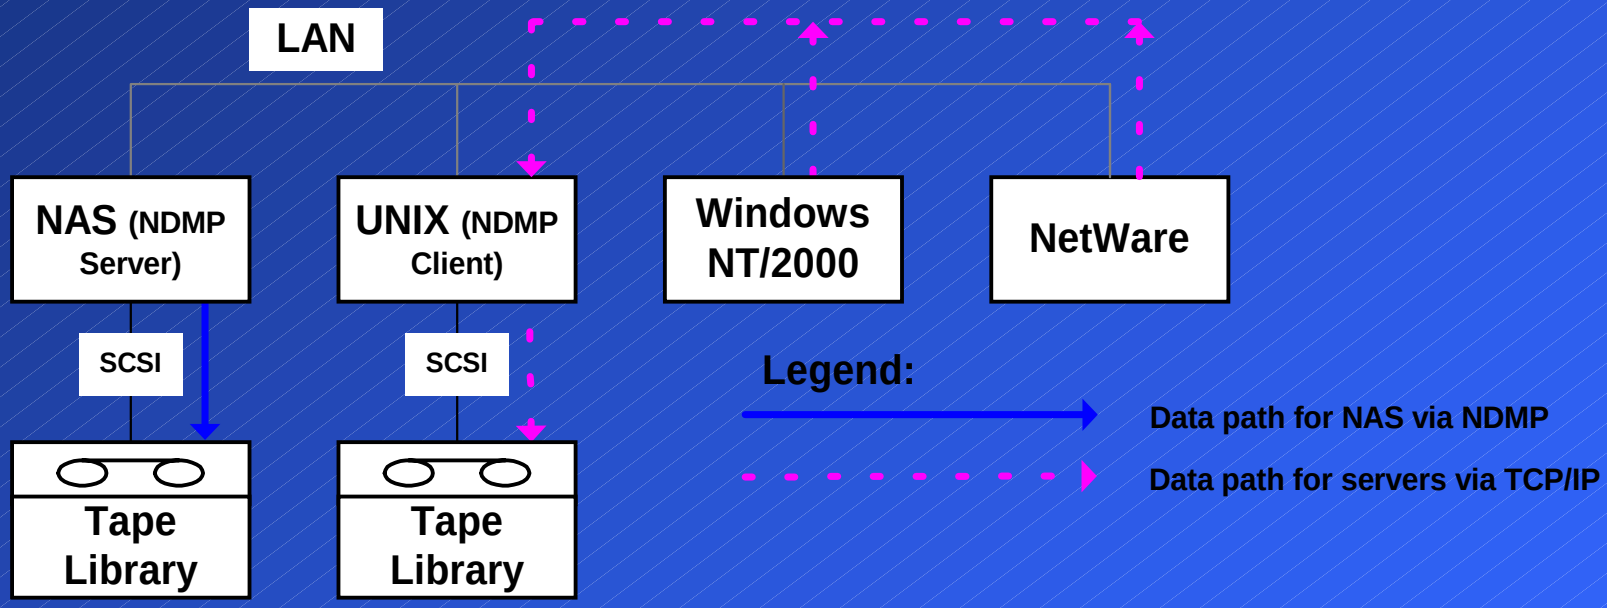

What is the best for you?
Local Backup of NAS with Tape Library Multi-Hosting

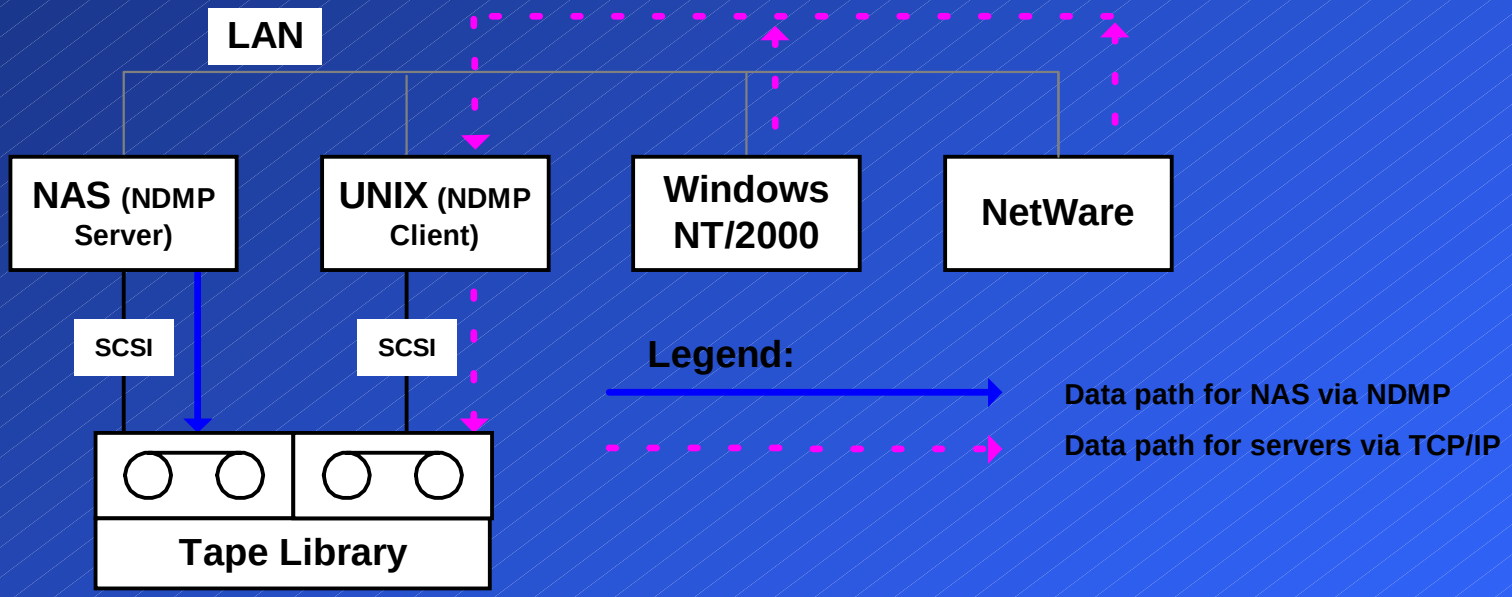

What is the best for you?
NAS-to-Server Backup via NDMP

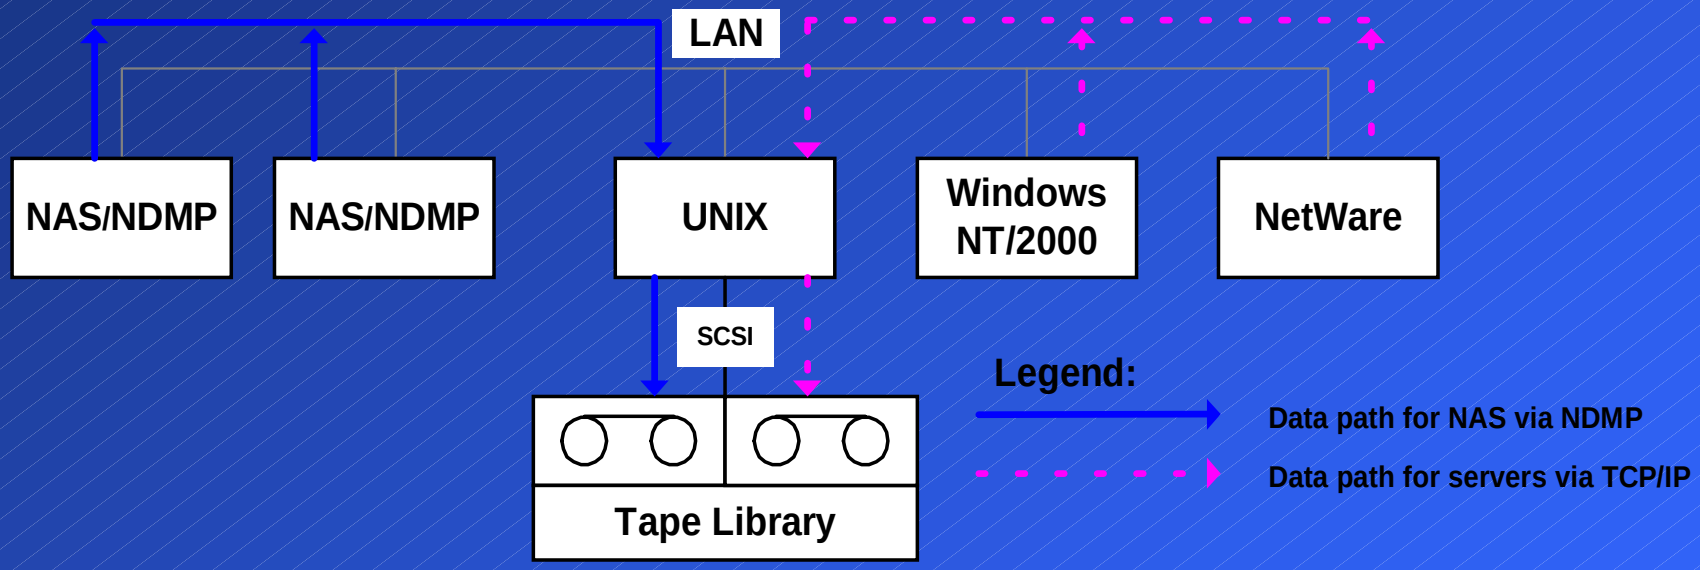

What is the best for you?
NAS backup on a SAN via NDMP

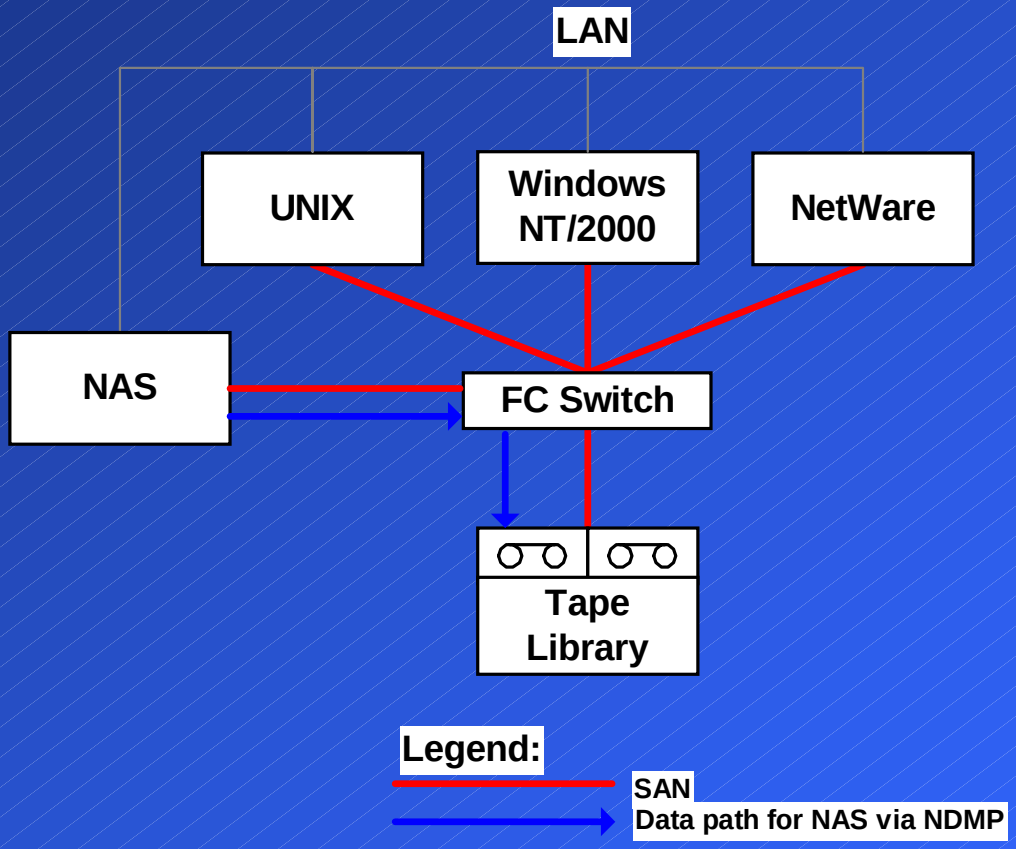

What is the best for you?
NAS-to-NAS Backup via NDMP

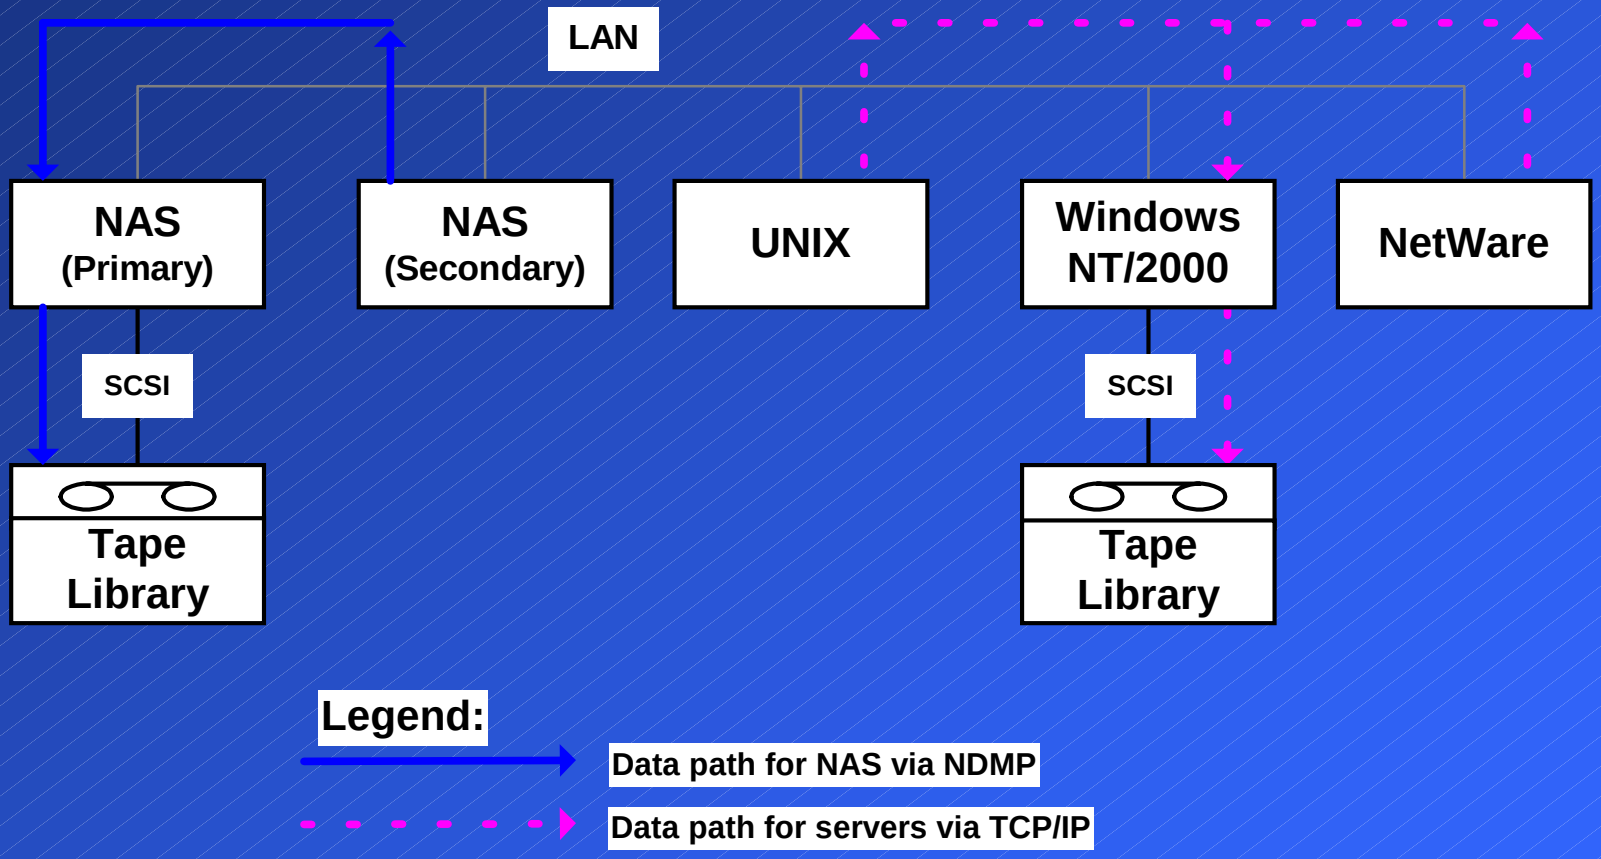

What to do?
How Do I Start?

- Check with vendors and who they have certified with
- If need be, visit their lab and see a demonstration of their system
- References – who's using what, where and how
- Know what your requirements are now AND in the future
  - Evolve your SAN as your enterprise does
- Test in your environment - one piece/segment at a time
- Feedback is the key
  - What do you learn, what's working or not
  - Roll back into the system and continue to fine tune

Summary - From where to where?

- The Movement behind SAN's
- Examine Network/Hardware-the push begins
- 3 Stages of SAN's Evolution
- SCSI vs. Fibre Channel
- LAN-less Backup
- 3rd Party Copy
- InfiniBand
- SSP's
- NDMP in the SAN-7 different configurations

Questions?