

# SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

## Identifying and Eliminating Backup System Bottlenecks

Jacob Farmer, CTO  
Cambridge Computer

# SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.

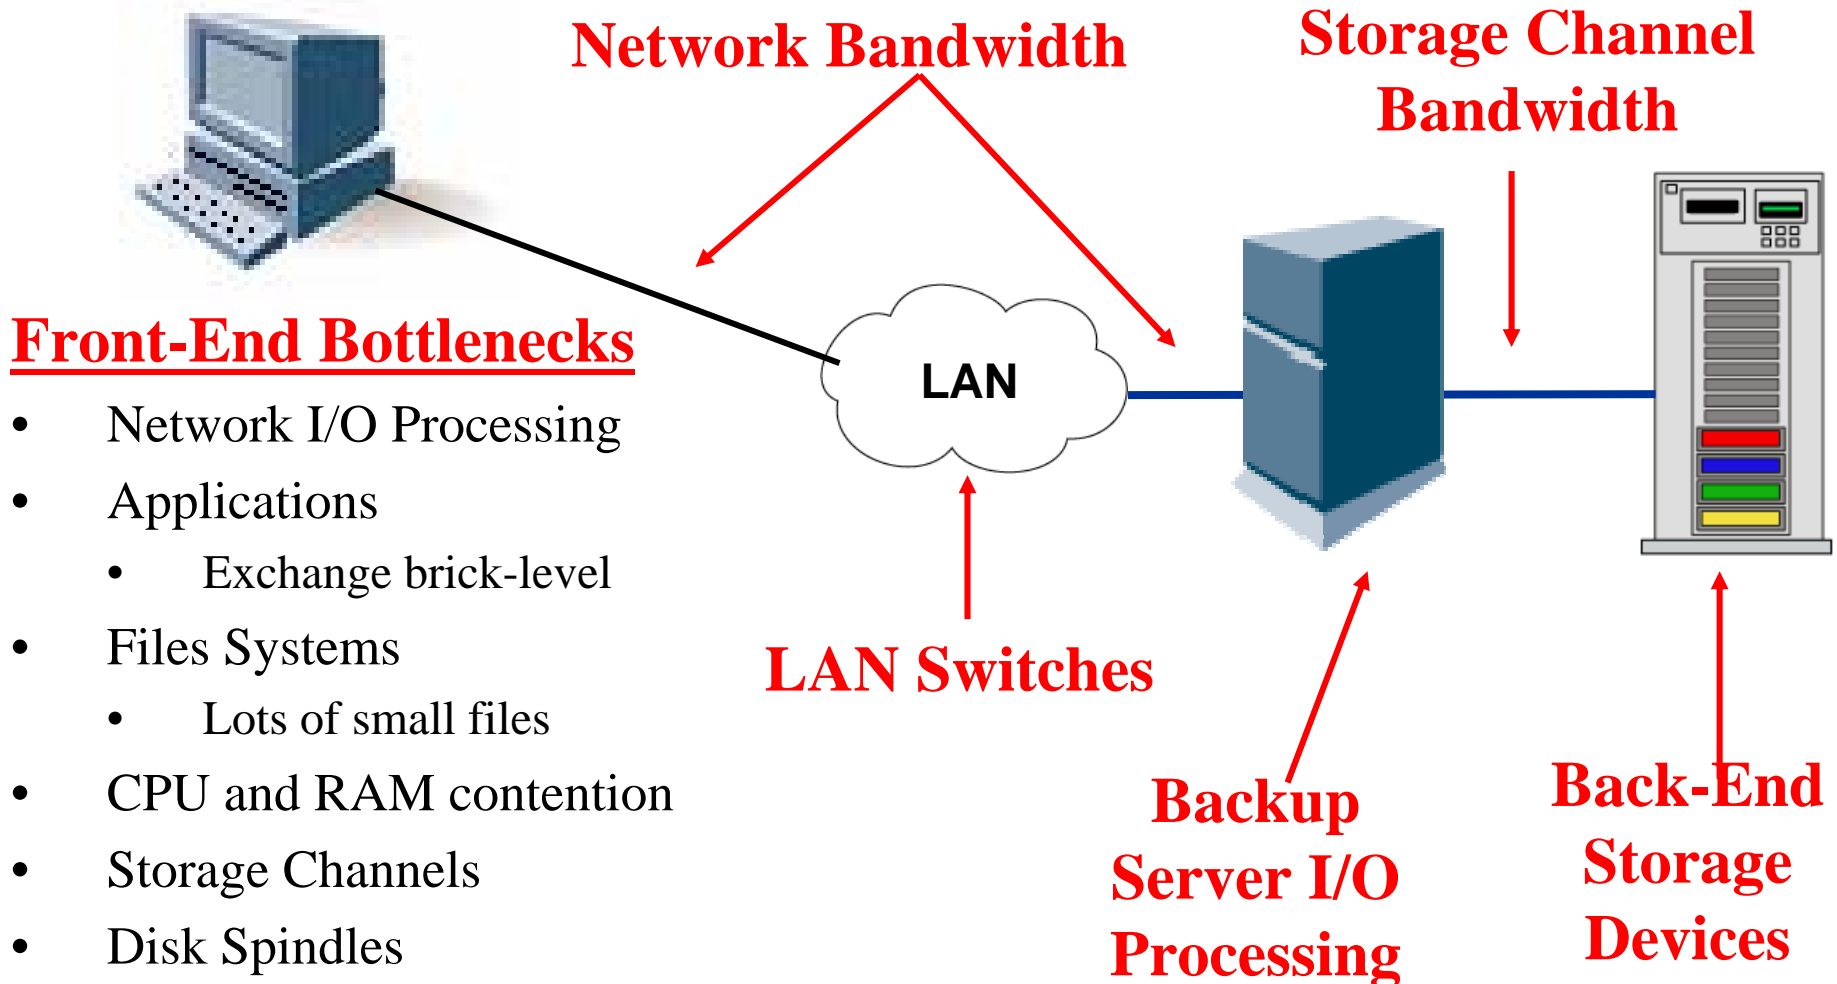
# Session Abstract

- **Identifying and Eliminating Backup System Bottlenecks**
  - This tutorial reveals the obvious and not-so-obvious bottlenecks in enterprise backup systems and offers examples of how one would apply the technologies described in the Data Protection tutorials to achieve one's performance objectives. We start with the assumption that the ultimate goal is to get data on tape for off-site removal. We then explain how to identify backup system bottlenecks and review a few case studies that illustrate how to eliminate them.
  - Technologies covered include: SAN backup paradigms, tape library sharing with ordinary SCSI, tape drive performance, disk staging with ordinary disk and virtual tape, snapshots and replication.
  - You cannot simply buy your way out of backup system headaches, you must design your way out.

# Why are Backups Still Such a Pain?

- The infrastructure for backup does not grow and evolve at the same rate as primary storage.
  - The solutions will be found in smarter software and thoughtful planning.
- Few backup systems are thoughtfully designed.
  - The backup system is often a tack-on to a server or storage purchase, rather than a whole project unto itself.
  - You don't back up storage. You do back up applications.
- Backup systems are typically not re-designed when other major changes take place
  - All moves, add, and changes must be evaluated in terms of their impact on the backup system.

# Where are the Backup System Bottlenecks?



# Pinpointing Bottlenecks

- It's not a perfect science
  - The bottleneck varies from job to job
  - Bottlenecks shift over the course of the night or even within a single backup job.
- Tools for spotting bottlenecks
  - Backup system logs and reporting software
  - Network and SAN monitoring tools
  - Host monitoring tools
  - Cross domain correlation
    - Tools that correlate performance data from tape devices, backup software, SAN, LAN, and host.

# Where to Begin?

## Break Down the Problem

- **Design from Front to Back**
  - The most common mistake is to start with the back-end storage devices without regard to bottlenecks in front of the storage devices.
- **Front-end (host side)**
  - Set RTO and RPO goals for each host, or at least for the important ones.
  - Identify major challenges, e.g. shortcomings in hardware, high volume of small files, etc.
- **Center**
  - Eliminate I/O processing bottlenecks.
- **Back-end**
  - Can your backup devices (tape, disk, or optical) handle the various feeds or do they introduce a new bottleneck?

# Front-End Challenges

- Large volumes
  - How large is “large”? It’s all relative to your environment and the equipment at your disposal.
- Large volumes with zillions of small files
  - Causes a lot of disk thrashing when retrieving files
  - Disproportionate amount of metadata compared to actual data.
    - Secretly increases the capacity of the backup
    - Taxes the backup server, which has to create log entries for each of these files.
- Slower hosts: legacy equipment, slow network connections, etc.

# Eliminating Host-Side Bottlenecks

- Efficiency
  - Incremental Forever Strategies
  - *Synthetic Fulls a/k/a Save Set Consolidation*
  - Block-Level Incrementals
  - Hierarchical Storage and ILM Solutions
- Brute Force
  - Volume image backups
  - Proxy backups
    - Snapshot or split mirror and mount on the backup server or a dedicated backup client
    - Replication to a dedicated backup client
  - Multiplexing data streams with “push” agents

# Efficiency: Why would you ever back something up more than once?

- If you want multiple backup copies, back it up once and then make copies of the backup.
- Incremental Forever
  - After the initial full backup only changed files are backed up.
  - Back end processes ensure that a full restore could be performed efficiently.
- Synthetic Full / Save Set Consolidation
  - The most recent full and its subsequent incrementals are consolidated into a current full backup.
  - Software must compensate for deliberate deletions.

# More Granular Incrementals


## *Go beneath the file*

- Examples of sub-file-level incremental backups
  - Block-level incremental
  - Transactional changes
  - Object-level backups
  - Commonality Factoring
- Instead of moving an entire changed file, only take action on the portions of the file that have changed.
  - For instance, change the word “dog” to “cat” in your 100 page word document. Do you back up the full 100 pages, or just “cat”.
- Sub-file-level tracking of changes is the key to a variety of data protection technologies:
  - Snapshots
  - Data replication
  - Continuous Backup

# Efficiency: Hierarchical Storage and ILM

- Files, database records, and/or email that are not in active use are either deleted or migrated to a secondary storage system.
  - Eliminates redundant backups.
  - Allows for speedier recovery.
  - Frees up backup system resources for other jobs.
  - Saves on tape media and reduces capacity needs of back-end storage devices.
- **What's the catch?**
  - Stub files and symbolic links can confuse the restore process.
  - Backup performance might not improve as much as you might think.
    - Watch out for meta-data processing bottlenecks.


# Interleaved Tape from Multiplexing



## Results in:



Roughly one in four blocks is relevant to the restore. With higher levels of parallelism, the problem is that much worse. Tape copy and restore operations suffer.



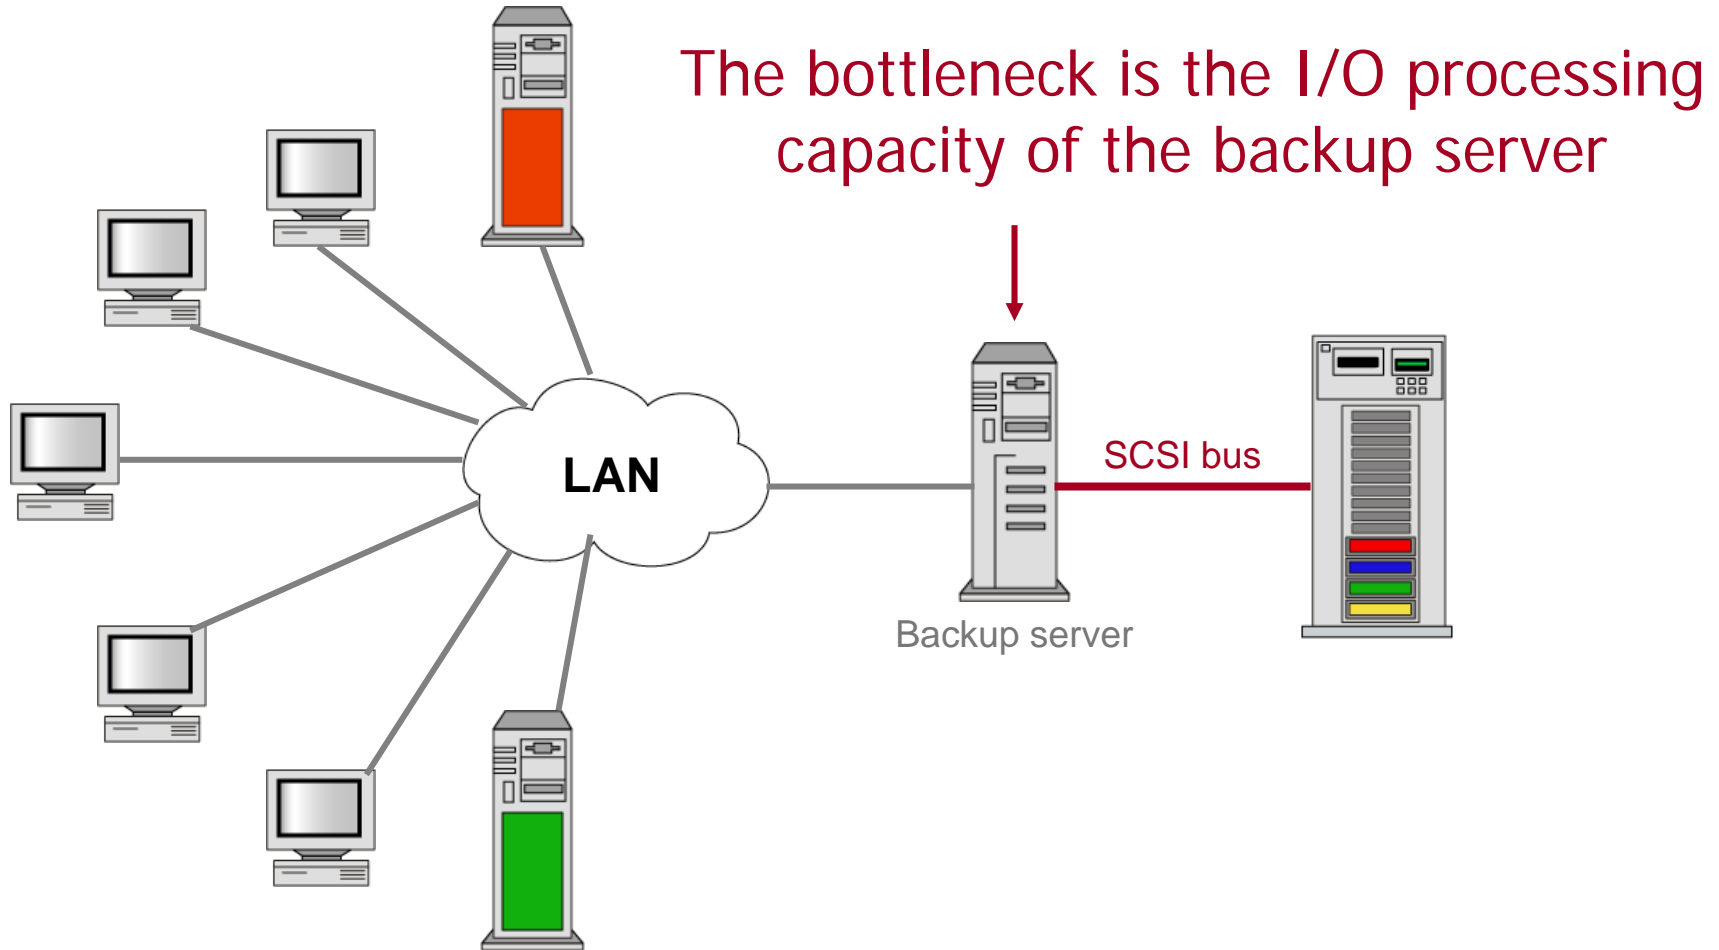
# SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION


## Eliminating Central Bottlenecks

# Traditional Centralized Backup System



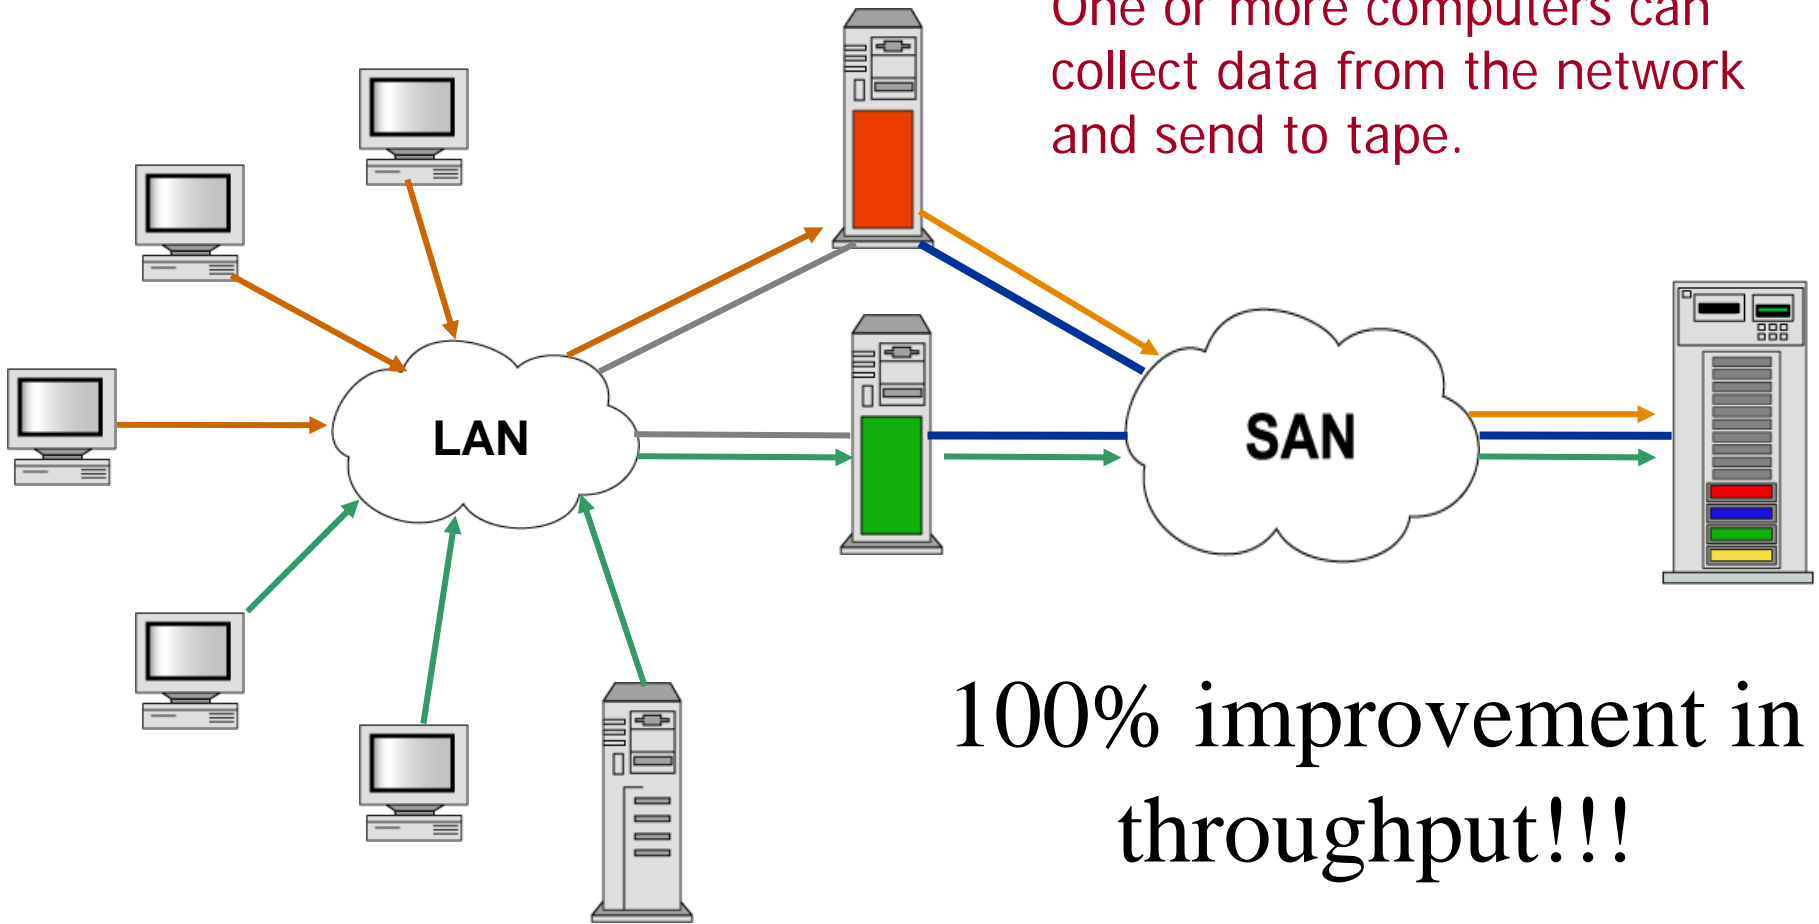
# Dedicated High Speed Ethernet Backbone

Problem: Still have I/O bottleneck on backup server.



# SANs Enable the Backup Server to Seek Help from a “Slave Server”

One or more computers can collect data from the network and send to tape.



100% improvement in  
throughput!!!


# Common Names for Backup System “Slave Servers”

- There is no industry standard term for “slave server”. Common names include:
  - Storage Node
  - Media Server
  - Device Server
  - Tape Server
  - Smart Clients
  - Media Agents
  - Others??

# SAN Backup Without SAN Disk

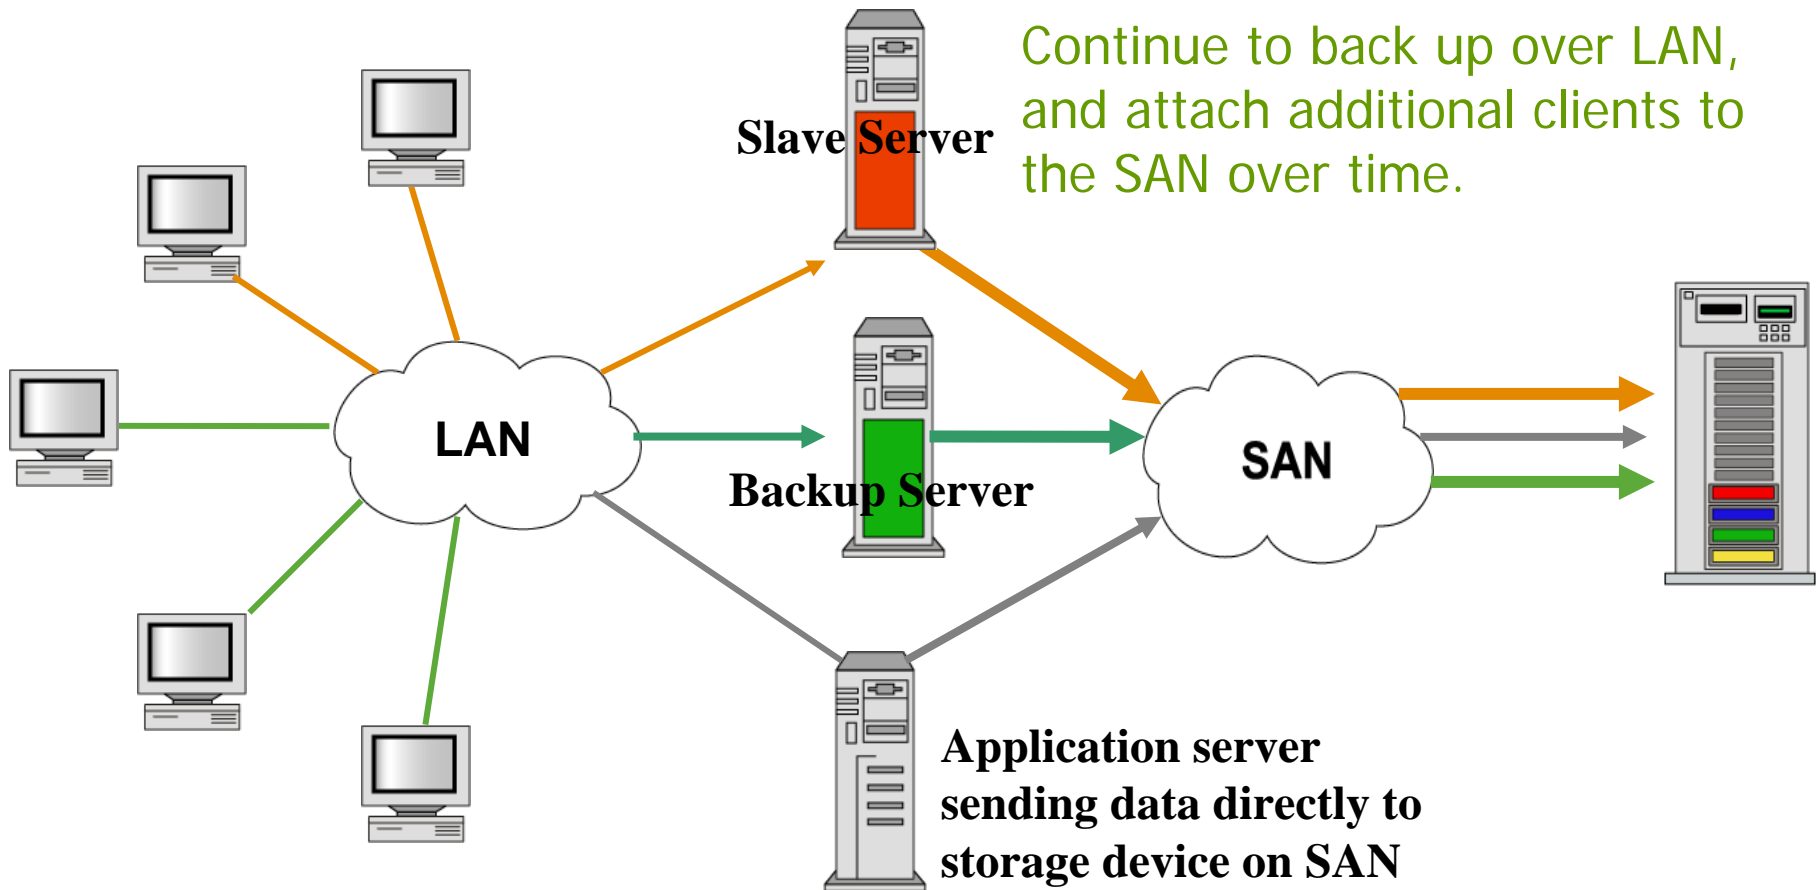
- You don't need to have SAN disk to have SAN backup.
  - For that matter, it often makes sense to do LAN backup, even if you have disk on a SAN.
- “LAN-Free” backup does not require disk to be on the SAN.
- SANs for backup only do not require fancy SAN infrastructure.
  - FC-AL loop switches are okay.
  - Non-redundant fabrics are okay .
  - Older model switches are okay.
  - Multi-port SAN routers can form entire SAN.

# SAN Backup Without a SAN !!!




- You don't need a SAN to have SAN backup!!!
- Use dedicated drives and dedicated SCSI channels.
- Add Fibre Channel or iSCSI functionality and dynamic drive sharing as a future upgrade.

# SANs Enable Circumventing the Backup Servers




# Hybrid of LAN and SAN Backup



# SAN v. LAN Client Connection: How to Choose

- Just because is not a valid reason
  - Don't use the SAN to back up just because you happen to have a SAN connection.
- Select SAN Backup When
  - You have a host that can generate substantially more I/O than your network backup servers can handle.
  - You have an aggressive RTO that cannot be met over the LAN.
  - You have back end storage devices that can keep pace.
- Select LAN when LAN offers suitable performance.
  - Focus your efforts on improving LAN backup performance.
    - Consider staging backups to disk
    - Consider multiplexing



# SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

## Eliminating Back End Bottlenecks

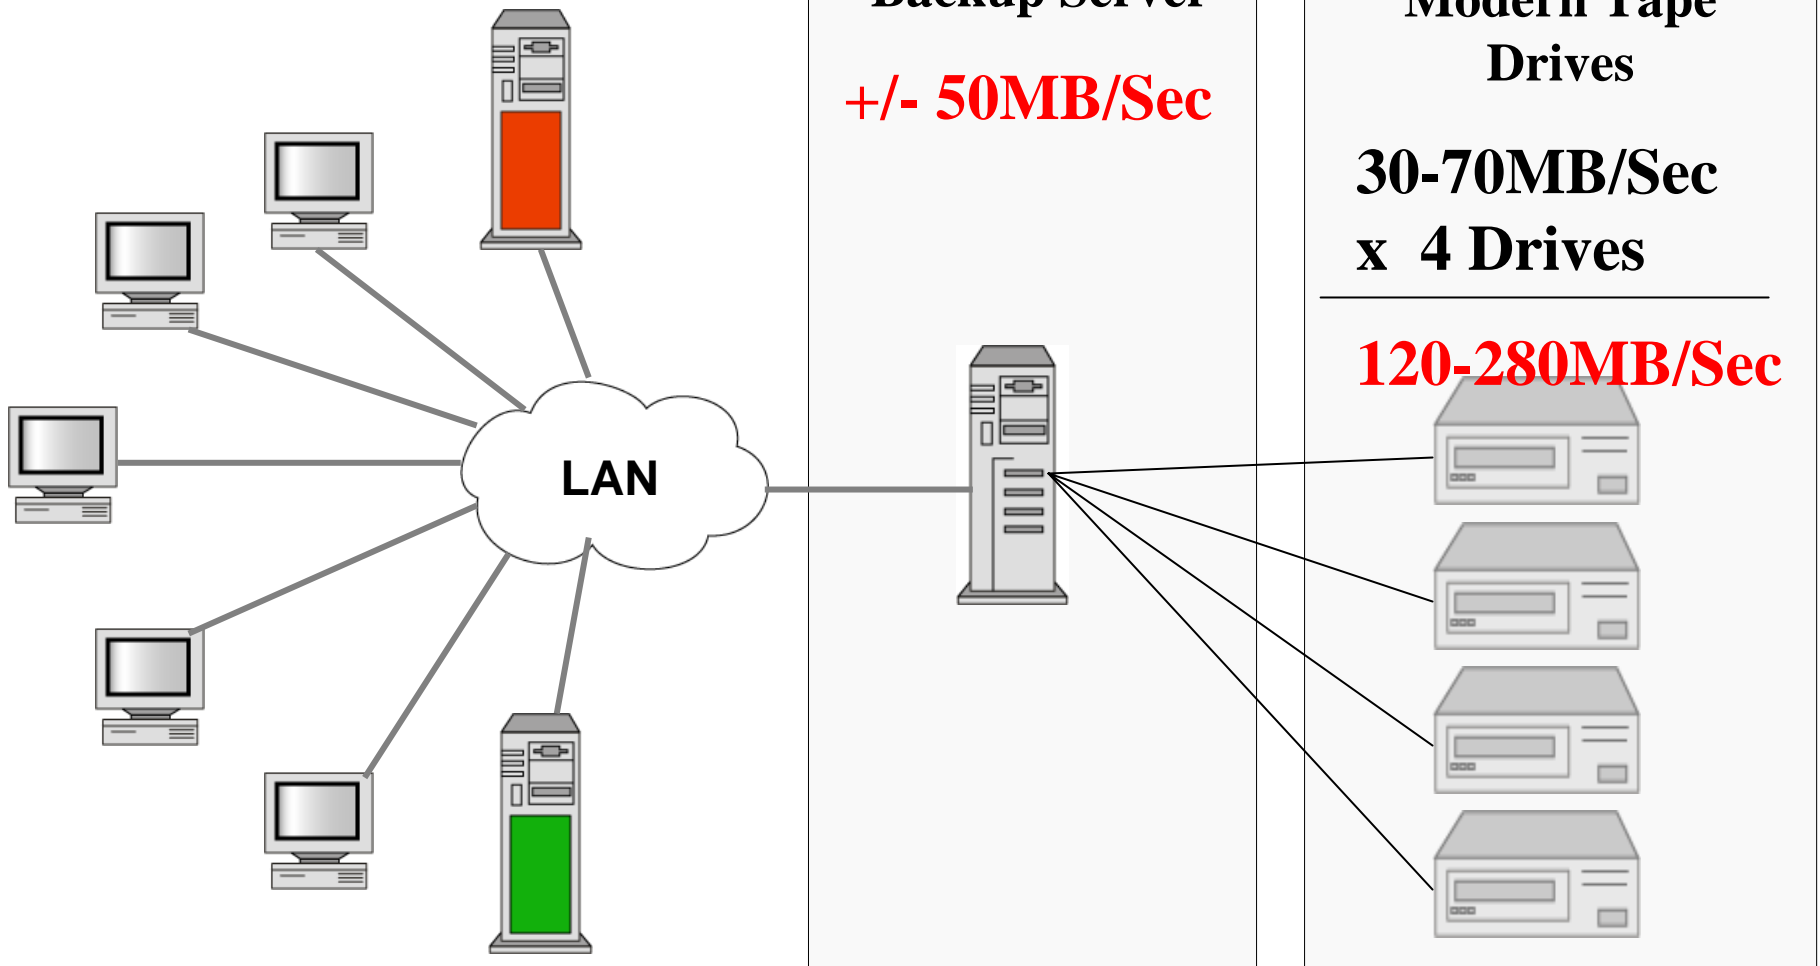
# “Shoe Shining” or Back-Hitching

- Performance problem caused by the mechanical operations of linear tape drives.
- The tape media is moving too fast to stop. When the data buffer empties. It keeps on going, then rewinds, then starts again.
- With erratic and slow data patterns it's like two steps forward, one step back.
- Larger drive buffers and throttling of the drive mechanism help compensate for media repositioning, but problem still persists in many situations.


# Consequences of Shoe Shining

- Backup performance suffers.
- If you use multiplexing to stream the drives, then restore performance suffers.
- Tape copies, reclamation, and job consolidation take much longer to complete.
- Wear and tear on the drives and media.

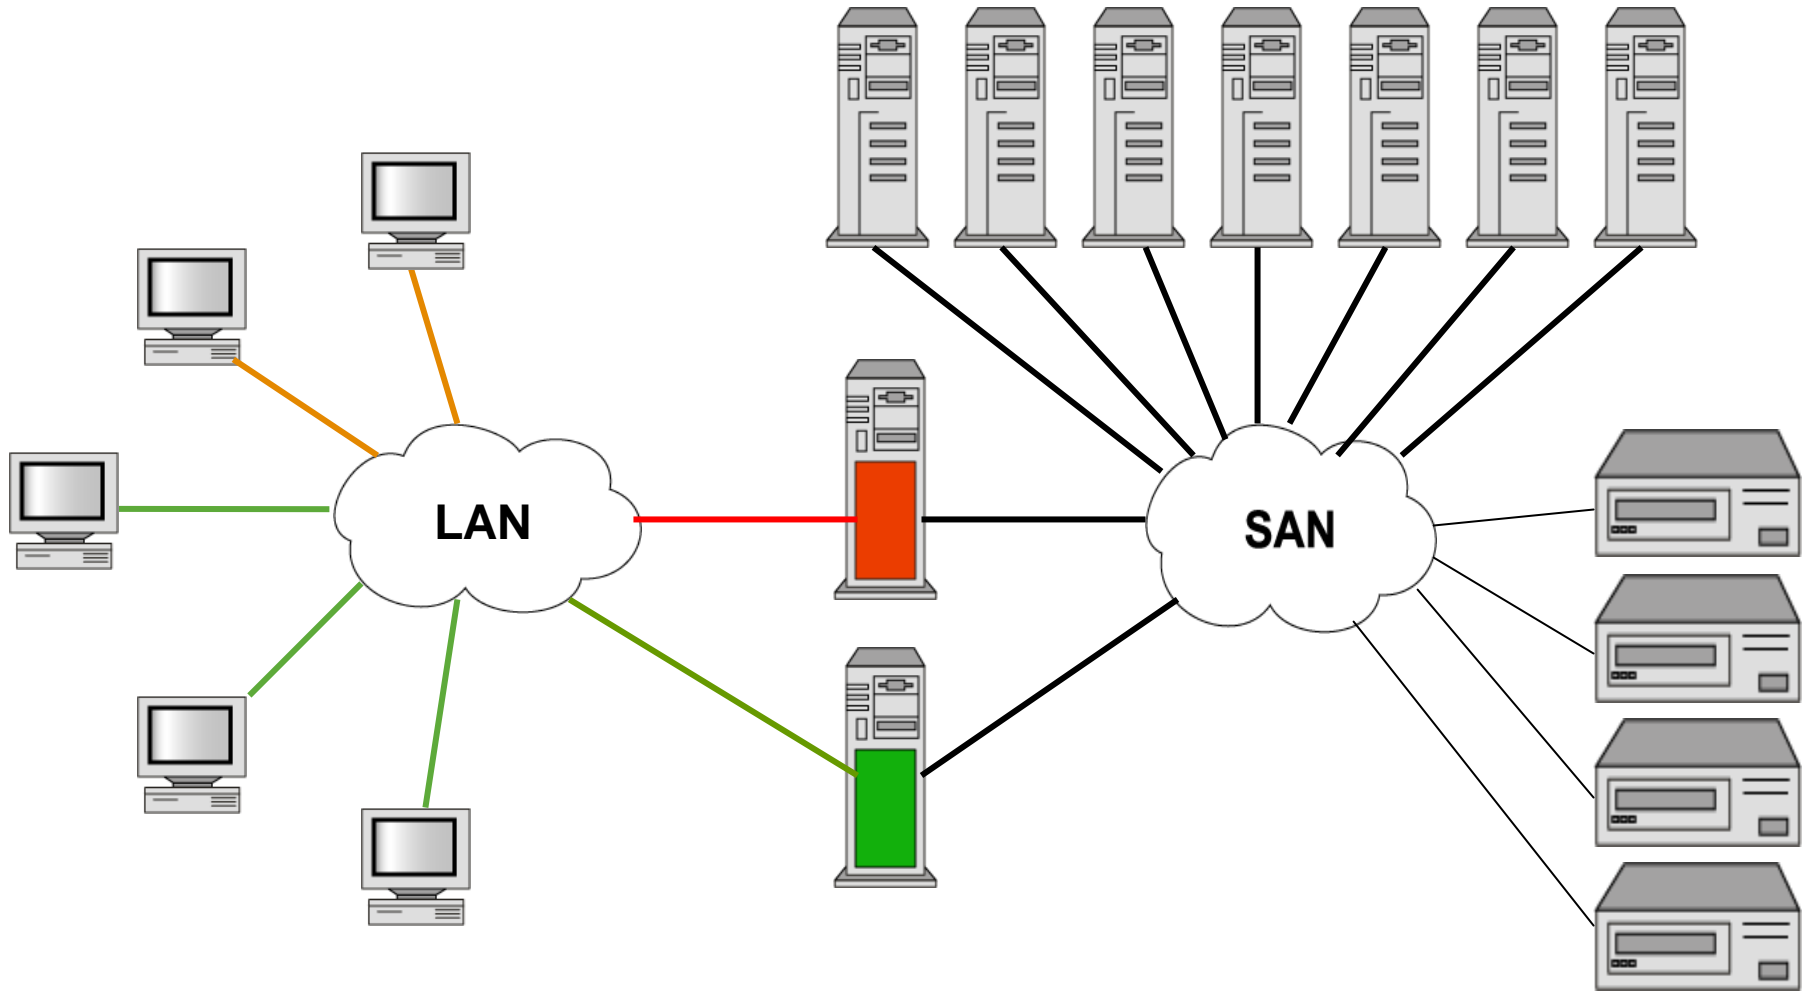
# You Can Have Too Many Drives



# Distribute Tape Drives Across Multiple Hosts



# Use the SAN with High-Speed Hosts to Get the Drives to Stream



# Tape Backup on the SAN: Some Things to Watch out For

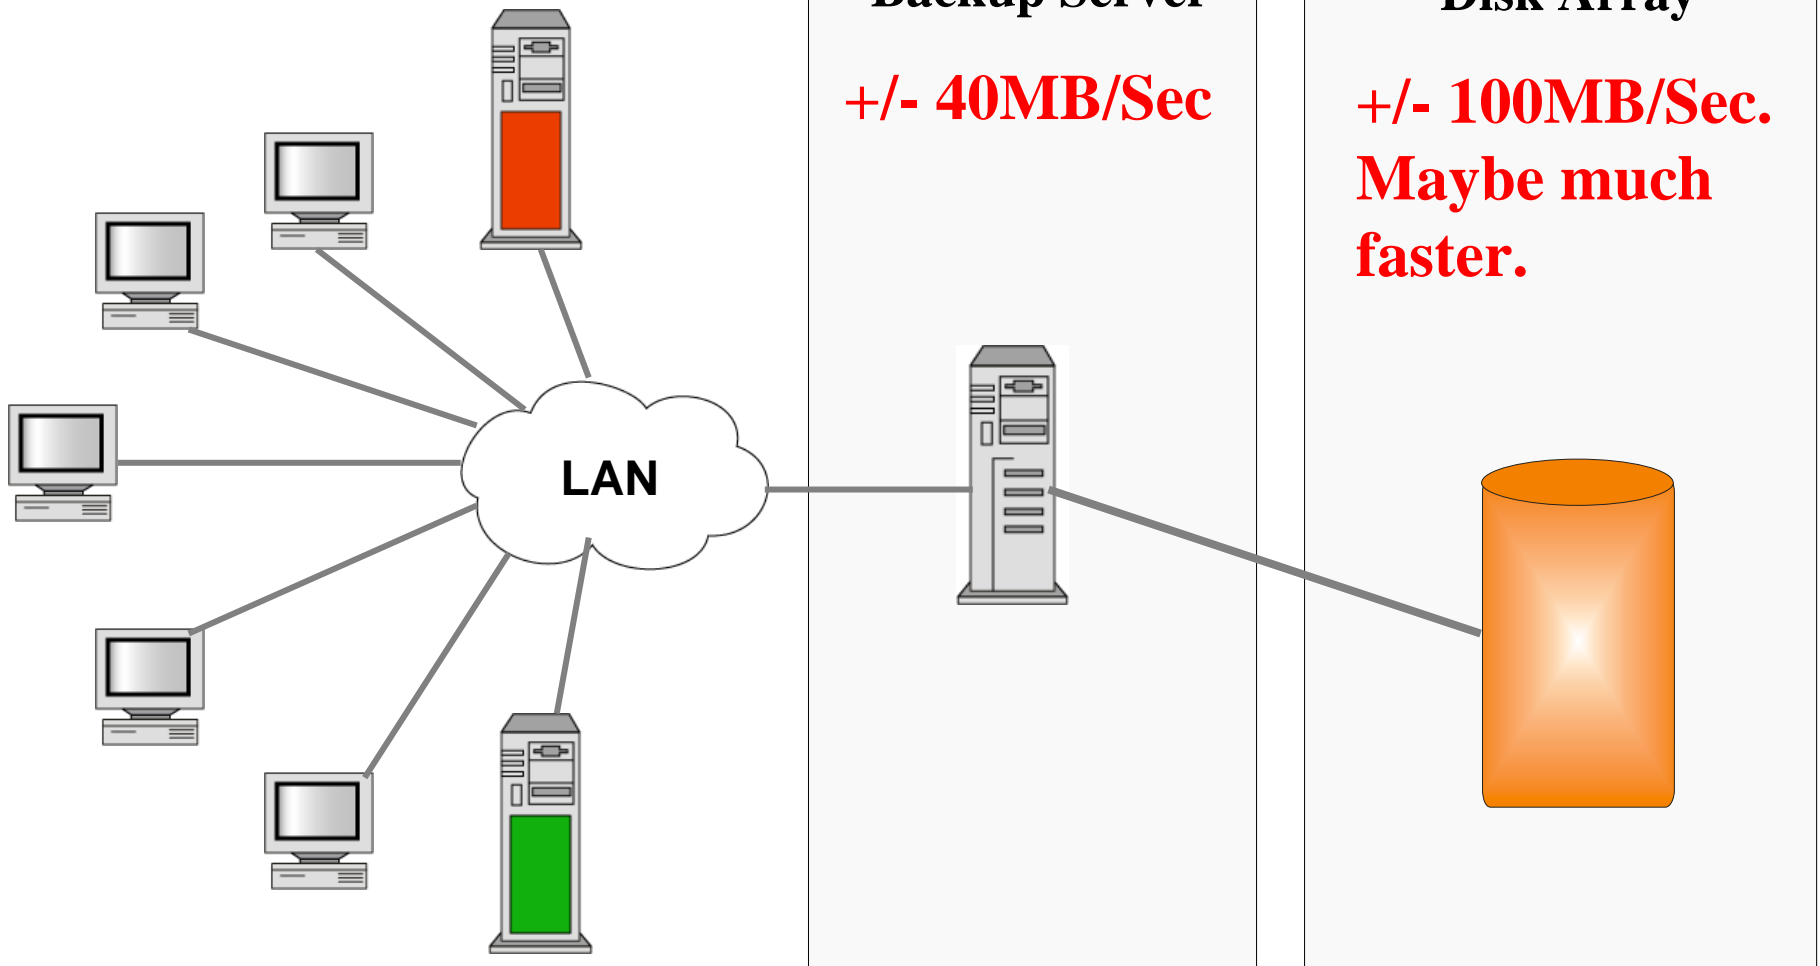
- Not all hosts can push enough data to stream the drives directly.
  - Slower hosts can tie up an expensive tape drive. These hosts should back up over the LAN.
  - When you deploy new models of drives with higher shoe-shine thresholds, you might need to move some clients from the SAN *back* to the LAN.
- Some hosts might be able to go faster than the tape drives.
  - It's a shame to let the tape drive slow them down.
- Slower hosts can tie up an expensive drive
  - It's a shame to waste a drive on these hosts.

# Staging Backups to Disk

- Disk arrays don't shoe-shine!!!
- As linear drives progress in throughput, it will be harder and harder to keep them streaming without first staging backups to disk.
  - Similarly, fewer hosts will be able to deliver the throughput to justify direct communication with the tape device. Staging to disk will get the most out of your tape investment.

# Disk is NOT the Panacea!

## Remember the I/O Bottleneck



# Provisioning Secondary Disk: Try Not to Create a Headache

- Conventional disk staging either involves direct attached storage or SAN storage.
- Provisioning disk for backups can be a pain.
  - Low cost disk devices might not have the same provisioning capabilities as high end disk.
  - It is difficult to accurately predict needs.
  - Utilization is a big challenge since disproportionately more disk is needed for full backups than incrementals.

# Disk Staging on the LAN v. Direct to Tape on the SAN

You must always re-evaluate whether to back up any given host over the LAN or SAN. As things change, you might need to re-design your backup system.

- **Disk Staging Breaths New Life into LAN backup**
  - Disks eliminate the uncertainty of tape performance.
  - I/O Processing on the LAN is getting better all the time.
  - Smarter software that takes advantage of disk minimizes the amount of data that has to be transferred in the backup process.

- **Tape on the SAN still has some advantages.**
  - No fighting the I/O bottleneck.
  - High end hosts can stream the drives.
  - Some drives don't shoeshine or compensate for varying data rates.
  - Tape drives offer compression with a performance boost instead of a performance hit.



# Using Disk Staging and SAN Backup Together

- Great idea, but there are some catches
  - Provisioning is a pain. You have just duplicated the provisioning problem you had for primary storage.
    - Every time you provision primary storage, you have to provision secondary storage for disk staging.
  - Sure ATA disk is cheap, but not if you waste it.
    - Secondary disk needs are heavy during full backups and light for incremental backups.
  - Okay for slave servers. Not great for SAN-enabled backup clients.
    - SAN client has to backup to disk and then copy to tape.

# SAN-Enabled Backup-To-Disk Requires a Sharing Mechanism


- Sophisticated volume manager integrated with backups.
  - But few vendors offer this and it might be too much to deploy volume management just for this purpose.
- SAN file system that allows multiple hosts to write to the same volume.
  - It might be overkill to deploy a SAN file system just for backups.
- CIFS/NFS Share
- Virtual tape library

# The Role of Shared Storage in Disk Staging




# Virtual Tape Library On a SAN

## Physical and Virtual Library



# Virtual Tape Library on a SAN

## Virtual Front-Ends the Physical



# Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: [track-datamgmt@snia.org](mailto:track-datamgmt@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

*SNIA Education Committee*

**Jacob Farmer**

# SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.

# Refer to Other Tutorials

- Please use this icon for referral to other tutorials where appropriate



- Please use this icon for referral to demonstrations of the technology in the SNW Interop Lab when appropriate



# Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: (speaker please note: choose your tutorial reflector from below and insert reflector address here) [track-datamgmt@snia.org](mailto:track-datamgmt@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

*SNIA Education Committee*

Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here

Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here  
Name of contributor here