Storage Virtualization

Peter A. Rigsbee
StorageTek

April 15, 2002

Agenda

- Definitions
  - Separating hype from reality
  - Defining virtualization
- Customer benefits
- Taxonomy
  - Overview
  - Different types of virtualization, examples
- Futures

Virtualization – Definitions

Storage Networking Industry Association, 2002

Virtualization – Definitions

- **"First IT buzzword of the 3rd millennium"**

Storage Networking Industry Association, 2002

Virtualization – Hype

- **New and cool invention**
- **Products and solutions**
- **Will solve all your IT problems**

Virtualization – Hype

- **New and cool invention**
- **Products and solutions**
- **Will solve all your IT problems**

Reality

- **New – in places; simple concept**

Virtualization isn't New

- **Computer systems have virtualized resources for years**
  - **Memory**
  - **Processors**

Virtualization is pretty simple

- **Maps a virtual resource to a physical resource**
- **Complexity comes from:**
  - **Reliability**
  - **Performance**
  - **Utilization**
  - **Monitoring**
  - **Administration**

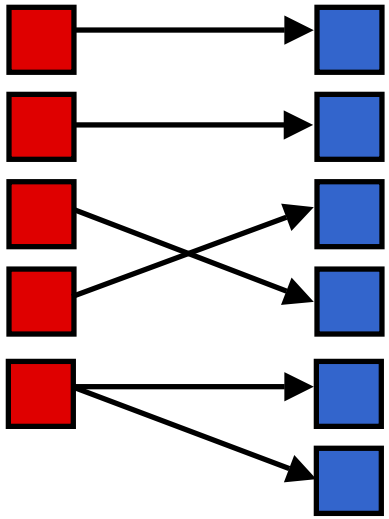

Virtualization – Hype

- New and cool invention
- **Products and solutions**
- Will solve all your IT problems

Reality

- New – in places; simple concept
- **Technology, not a solution**

Virtualization – Hype

- New and cool invention
- Products and solutions
- Will solve all your IT problems

Reality

- New – in places; simple concept
- Technology, not a solution
- May solve some of your IT problems

Virtualization – Definitions

- "First IT buzzword of the 3rd millennium"
- **"An abstraction of storage that separates:**
  - **Host view**
  - **Storage system implementation"**

Storage Networking Industry Association, 2002

Physical Storage has Limitations

- **All storage technologies have physical limitations that cause business inefficiencies**
  - **Performance**
  - **Capacity**
  - **Reliability**
  - **Names**
  - **Price**

Virtualization Abstracts Them Away

- **Virtualization abstracts away those limitations to:**
  - **Present a simpler view of storage to servers**
  - **Separate storage management from server management**
    - **Allow storage administrators to deal with these limitations**
    - **Allow server administrators to ignore these limitations**
  - **Reduce total costs of administration**

Storage Administration Problem

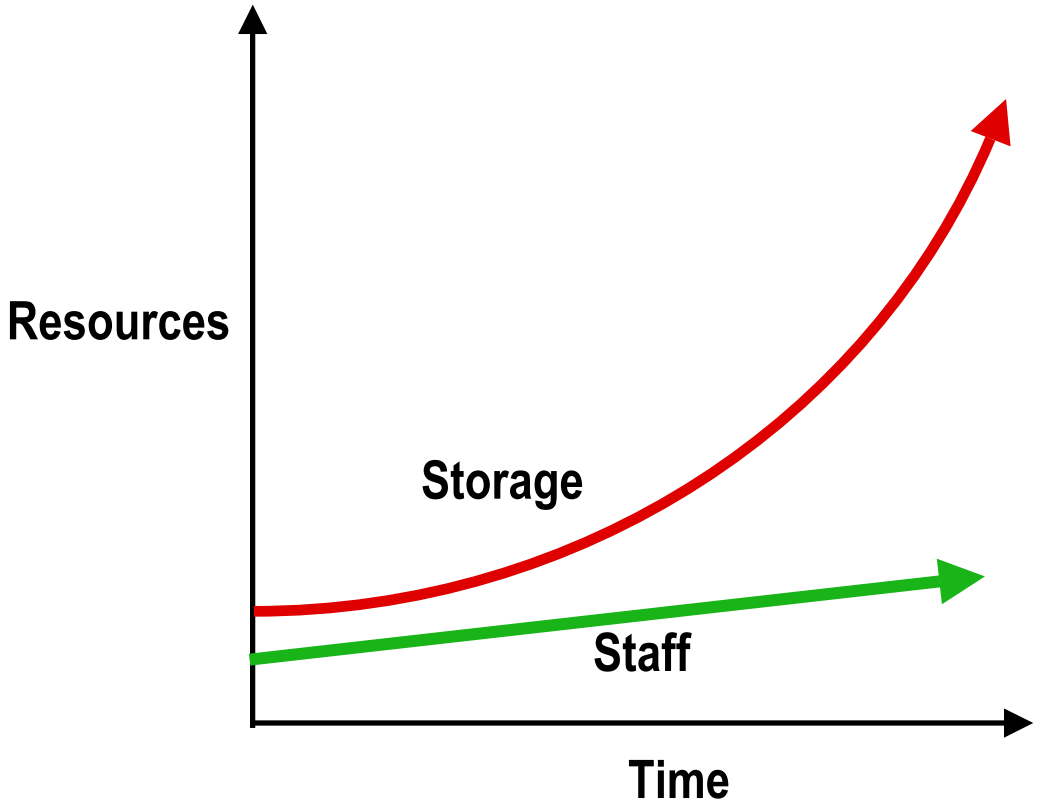

Analyst predictions continue to show dramatic increases in storage growth, while IT administrative staff will grow little (if at all)

Virtualization offers customers the ability to address this gap

Virtualization – Benefits

- **Increased availability**
- **Increased performance**
- **More flexibility**
- **Reduced administration**
- **Reduced cost of ownership**

Virtualization Taxonomy

What?	(Basic) disk virtualization	Disk block virtualization	File/record virtualization	Tape virtualization
Where?	Storage device or server	Storage device, network (in-band, out-of-band), or server	Network (in-band) or server	Storage device or network (in-band)
Examples	Disk firmware, files	RAID firmware, volume managers, virtualization appliances	NAS appliances and heads, virtual filesystems	Virtual tape products

Virtualization – What?

- (Basic) disk virtualization
  - Clean, linear view of disk blocks
- Disk block virtualization
  - Combine physical disks at the block level
- File/record virtualization
  - Combine files and filesystems
- Tape virtualization
  - Better utilization of tape components

Virtualization Vendors

	(Basic) disk virtualization	Disk block virtualization	File/record virtualization	Tape virtualization
What?				
Who?	Most disk and OS vendors	Storage device: RAID: most disk vendors; virtual disk: StorageTek, IBM, Compaq Network (in-band): DataCore, FalconStor, HP (StorageApps), Veritas, Brocade (platform) Network (out-of-band): Compaq, StoreAge, TrueSan Server: Veritas	NAS: EMC, NetApp, IBM, etc. Virtual filesystems: Veritas, Sistina, SGI	Tape media: StorageTek, IBM, Sutmyn Tape drives: StorageTek, NearTek

Virtualization – Where?

- Storage, Network, Server?
- No single right answer, depends on:
  - Granularity of virtualization:
    - Smaller => closer to the device
  - Combining/sharing devices => network or server
  - Sharing across servers => network or storage
- Possible to combine virtualization that resides in different places

In-band or out-of-band?

- **In-band or symmetric**
  - Requires no special server software
  - Reduces switch requirements and headaches

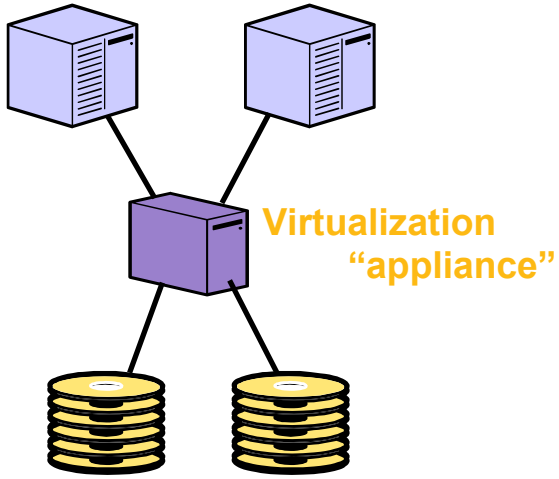

- **Out-of-band or asymmetric**
  - Separates control and data paths
  - Avoids single point of failure

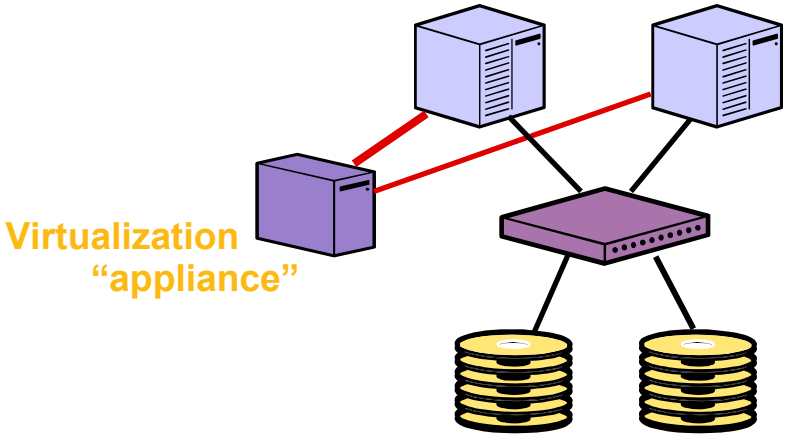

(Basic) Disk Virtualization

- Low-level virtualization
- Has been around for many years
- Disk firmware presents a clean, linear view of disk blocks
  - Don't have to worry about cylinders, heads, sectors
  - Don't have to worry about "holes" due to media defects
- Operating systems present a clean, linear view of disk
  - File (stream of bits)
  - Operating system (filesystem) keeps track of where it is on disk
    - Or in memory in a disk cache
  - Storage applications extend this concept
    - Hierarchical storage managers (HSMs)
    - Backup applications

Disk Block Virtualization

- Higher level virtualization
- Most common form of storage virtualization
  - Storage- and server-based products have been around
  - Network-based products are getting most of today's hype
- Present "virtual disk" that combines physical disks at block level
  - Increased reliability
  - Increased performance
  - Increased capacity
  - Fewer points of administration

Disk Block Virtualization Approaches

- RAID
  - Combine disk for higher reliability and performance
  - Storage-based virtualization
- Virtual Disk
  - Over-allocate virtual volumes, allocate physical space on write
  - Storage-based virtualization
- Volume Managers
  - Virtual volume made up of physical volumes
  - Server-based virtualization
- Virtualization Appliances
  - Features of both Volume Managers and RAID
  - Network-based virtualization
    - Both in-band and out-of-band implementations
    - Some vendors offer appliances, others software-only

Example: StorageTek SVA

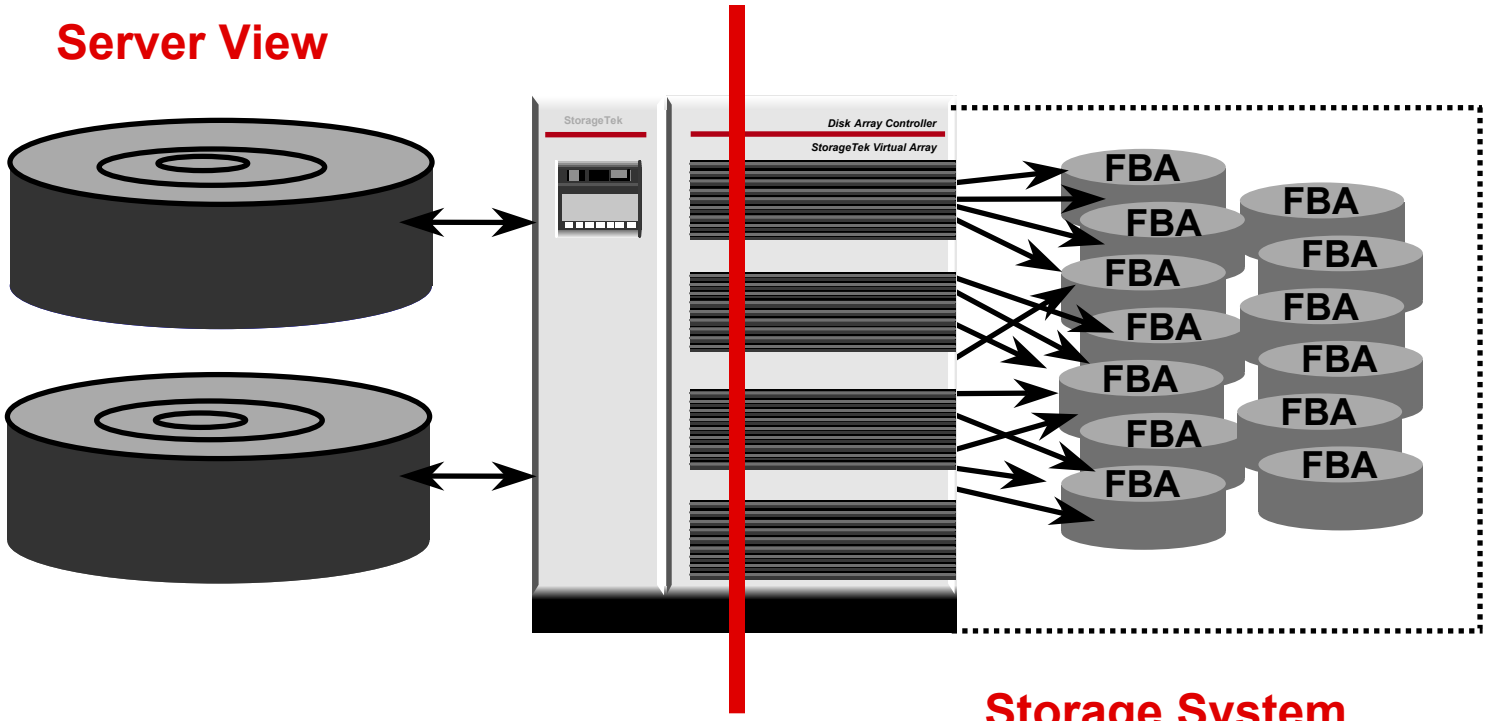

Storage-based, Disk Block virtualization (virtual disk)

Example: DataCore SanSymphony

Server View

Storage System

Network-based, Disk Block virtualization (in-band appliance)

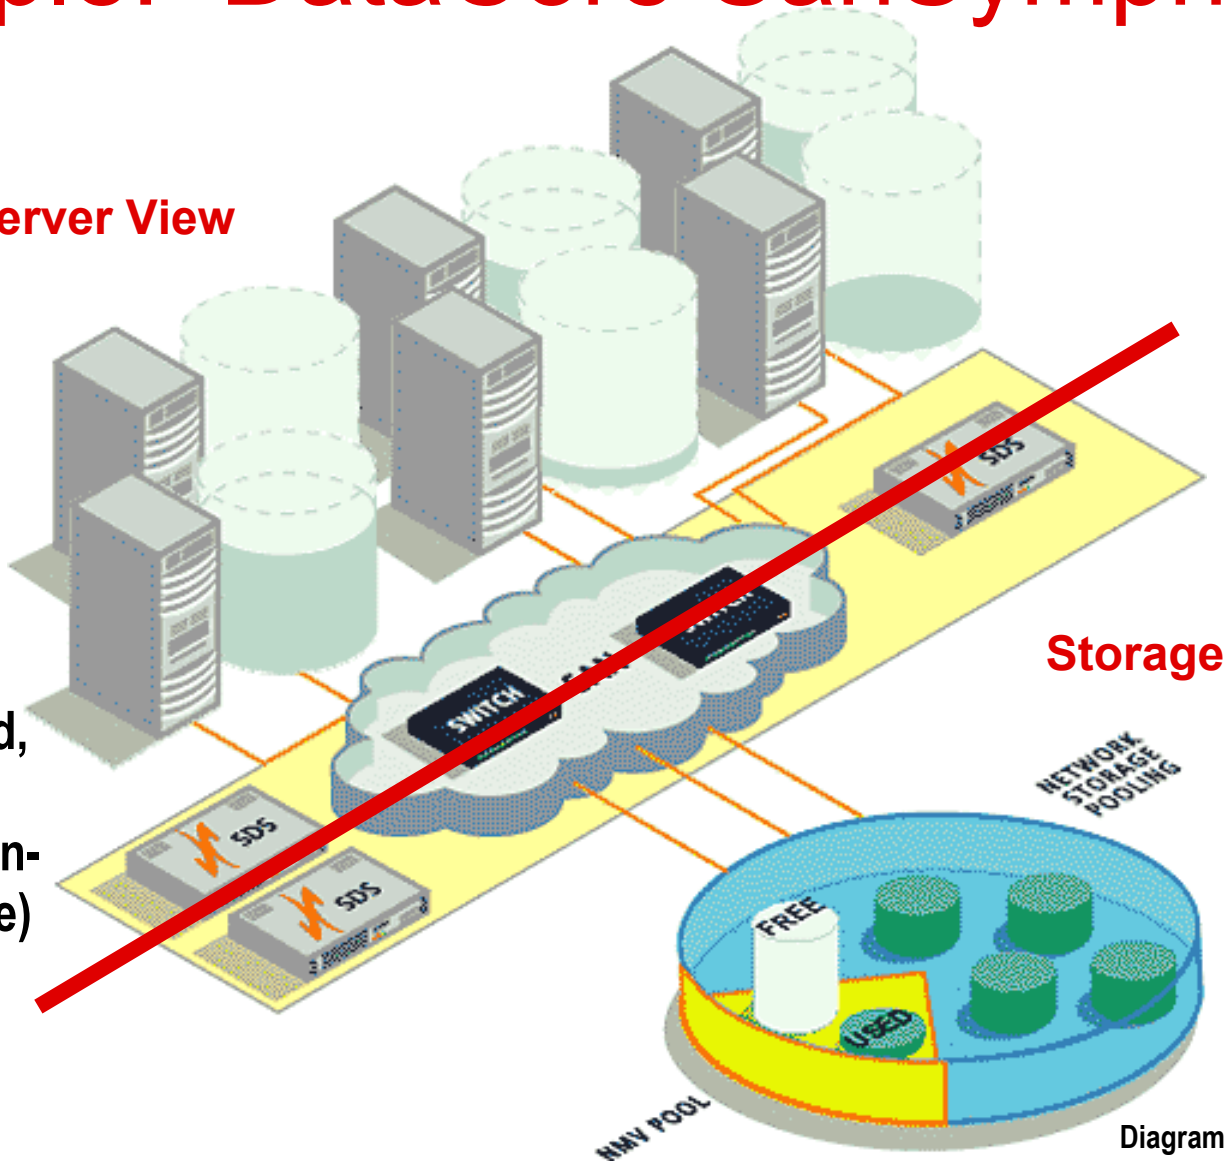

Diagram from DataCore web site

Example: StoreAge SVM

Server View

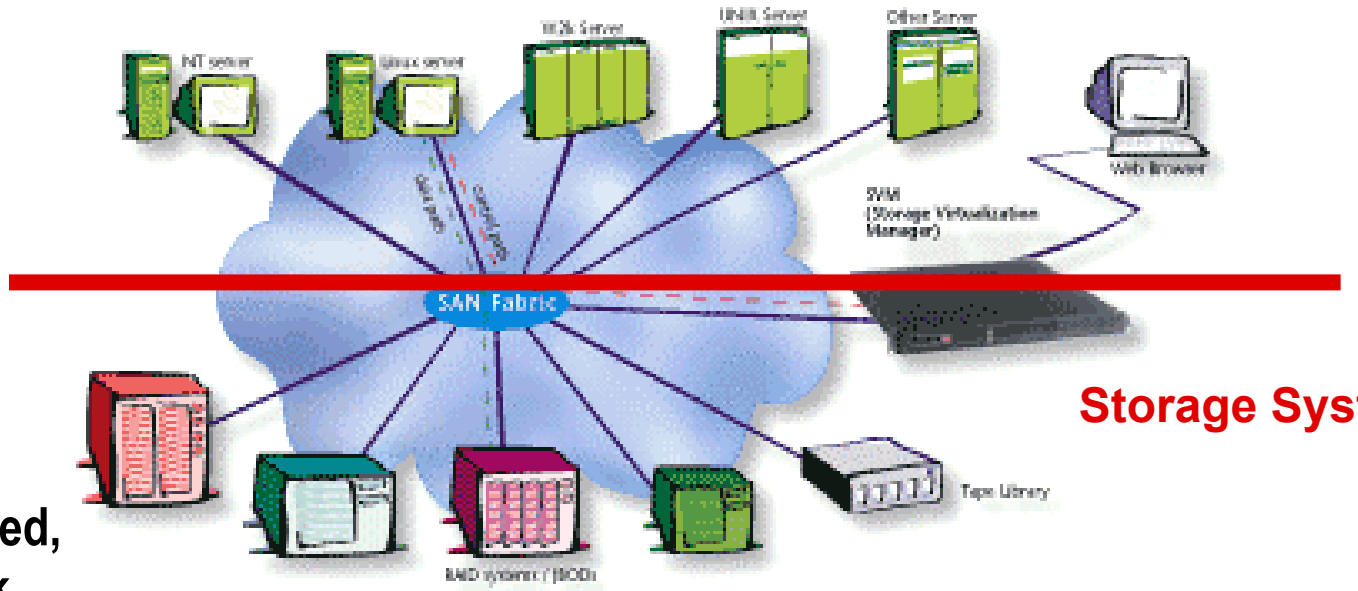

Storage System

Network-based, Disk Block virtualization (out-of-band appliance)

Diagram from StoreAge web site

File/Record Virtualization

- Highest level of virtualization
- Products in the market are often not positioned as virtualization
- File virtualization (NAS)
  - Present a file that may be made up of multiple files
- File system virtualization
  - Present a virtual file system that may be made up of multiple file systems

Example: NetApp filer

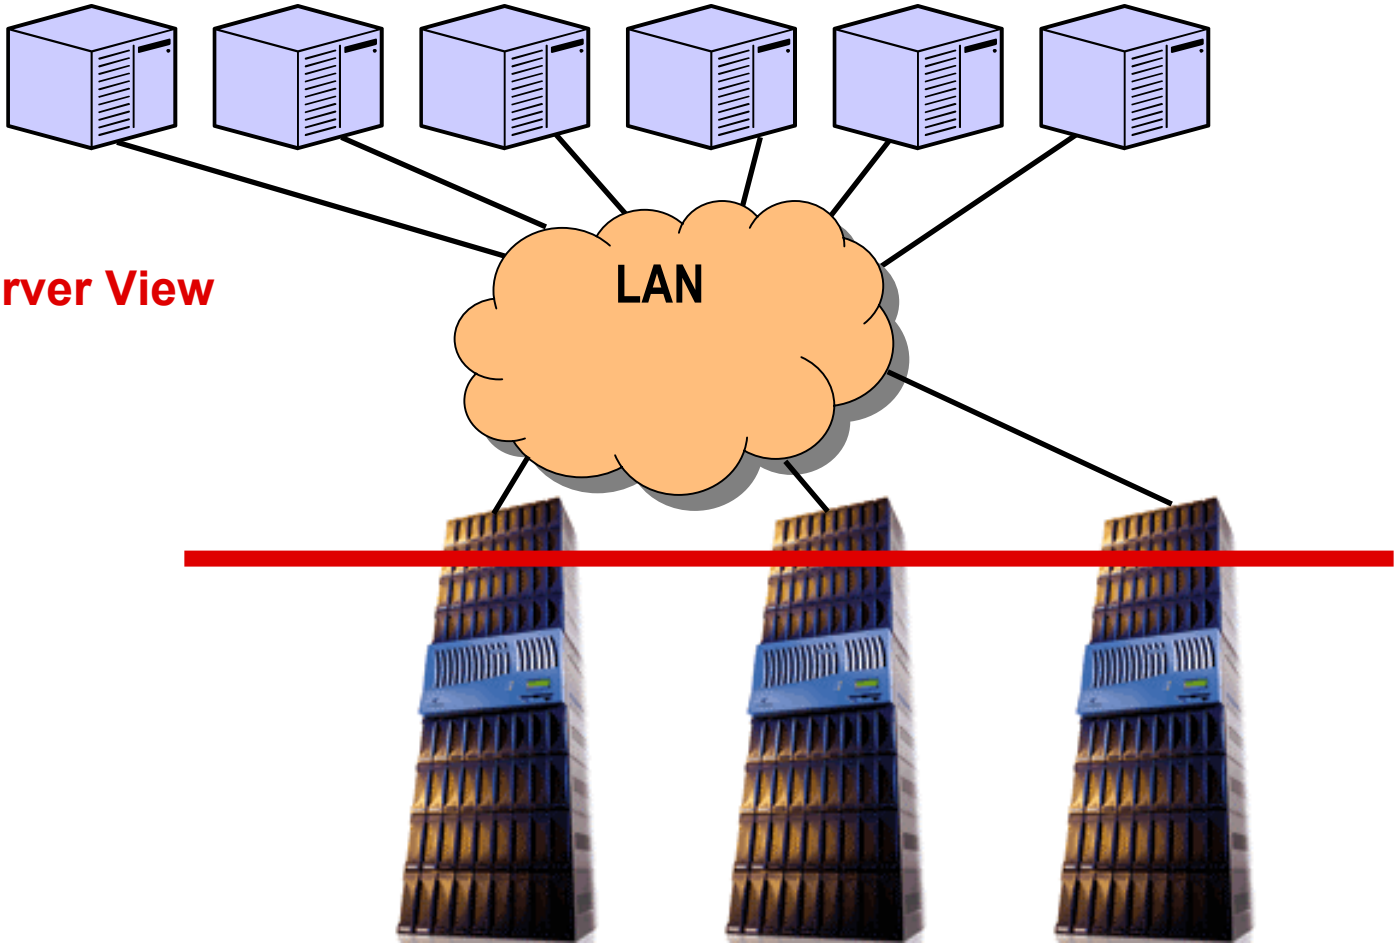

Example: Sistina GFS

Server-based, File virtualization (virtual filesystem)

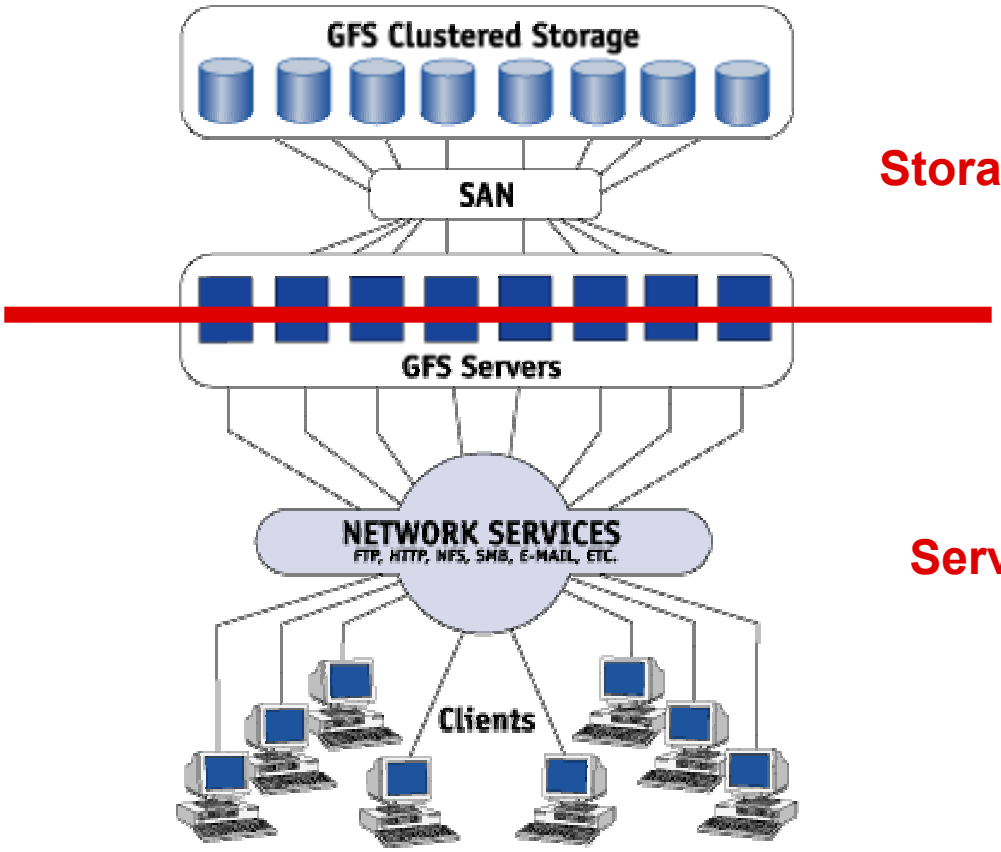

Storage System

Server View

Diagram from Sistina web site

Tape Virtualization

- Higher level virtualization
- Not as visible or common as disk block virtualization
  - Not as many vendors offer these products
- Tape media virtualization
  - Present a tape cartridge that is as fast as disk
  - Write data to disk; then combine and stage to tape
- Tape drive virtualization
  - Present a tape drive that's always there and never changes

Example: StorageTek VSM

Storage-based, Tape Media virtualization

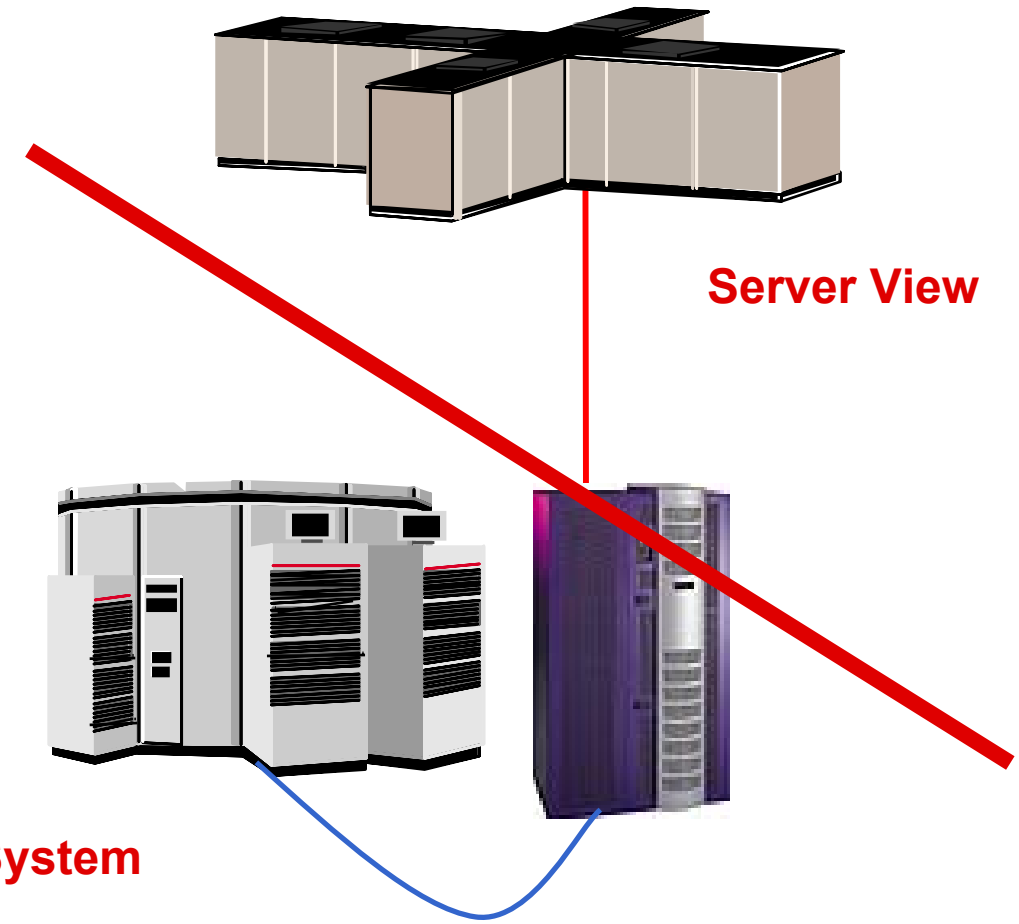

Example: StorageTek SN6000

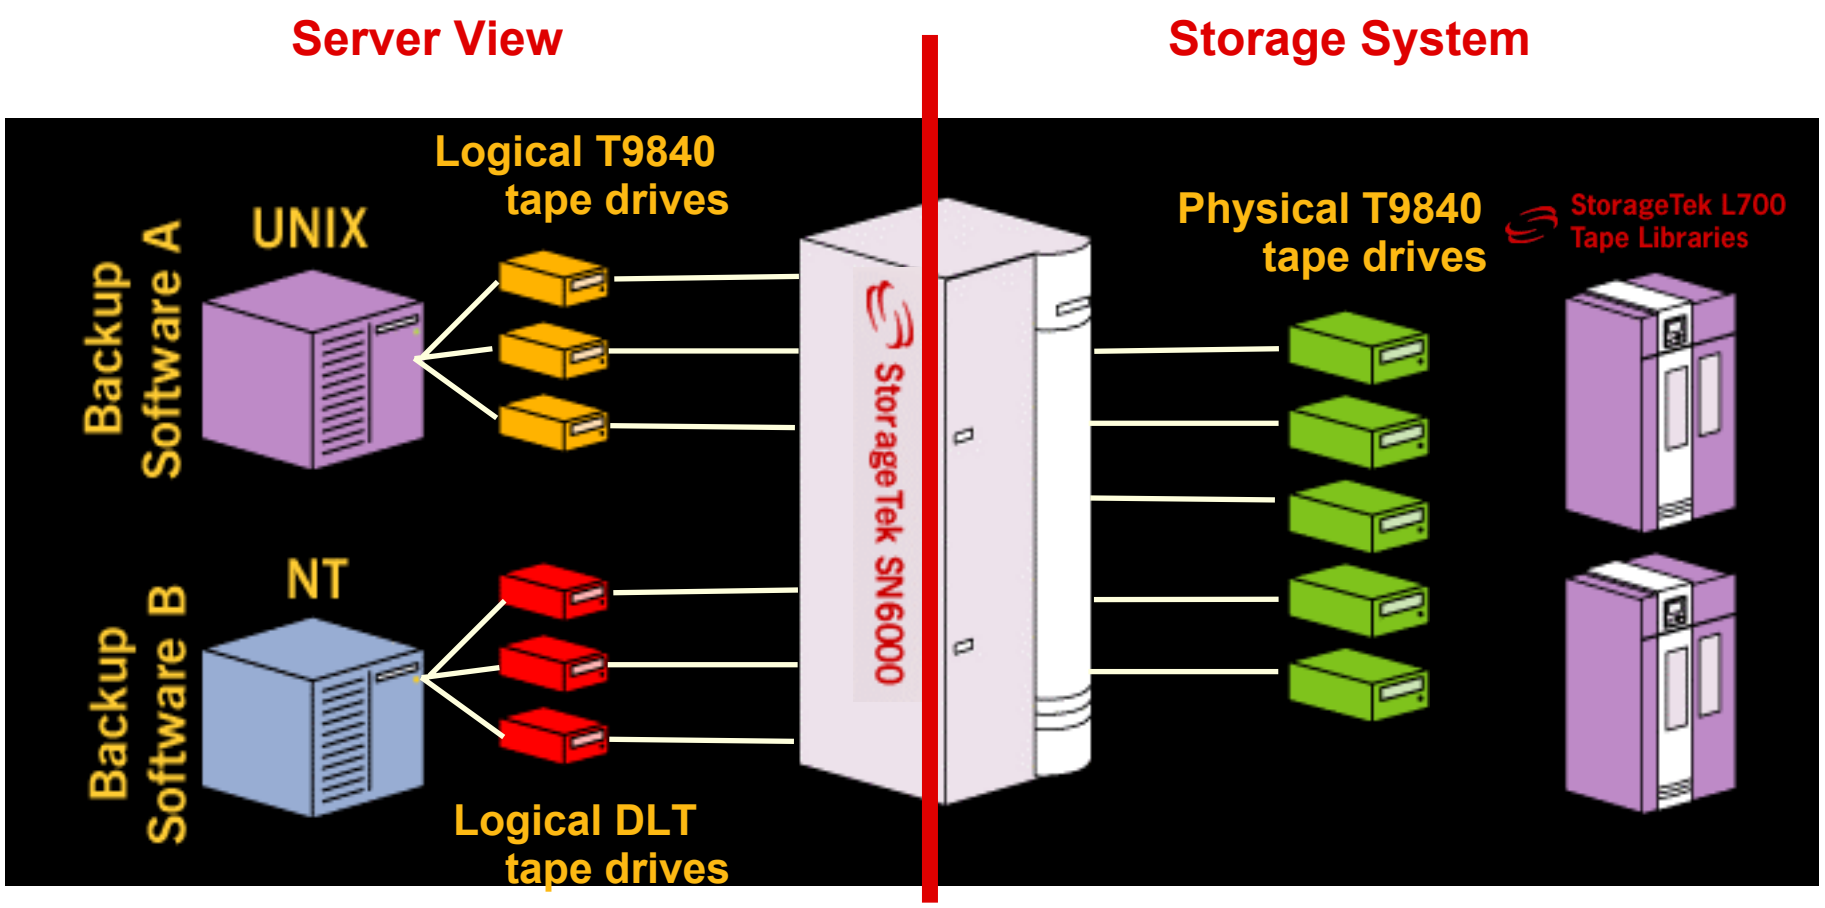

Network-based, Tape Drive virtualization

Virtualization Taxonomy

What?	(Basic) disk virtualization	Disk block virtualization	File/record virtualization	Tape virtualization
Where?	Storage device	Storage device, network (in-band, out-of-band), or server	Network (in-band) or server	Storage device, network (in-band), or server
Examples	Disk firmware	RAID firmware, volume managers, virtualization appliances	NAS appliances and heads, virtual filesystems	Virtual tape products

Futures?

- Storage virtualization will continue to thrive and evolve
  - Storage administration problem (shown earlier) is only going to get worse
- Network-based virtualization will continue to move forward
  - Expect to see consolidation as market selects products and vendors
  - Storage- and server-based virtualization products will also remain viable

Storage Virtualization

- A technology that is delivered as part of a solution
- Not a panacea
- An abstraction of storage that separates host view from storage subsystem
- Can reduce administration and total cost of ownership
- Can be found in servers, network, and storage devices

Questions?