BROCADE

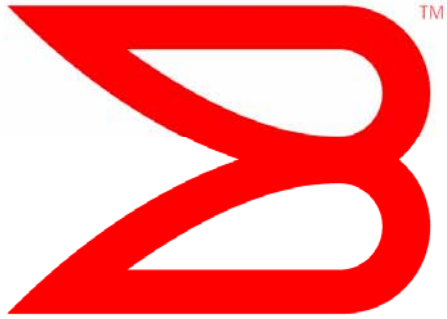

SAN and NAS Bandwidth Requirements

Exploring Networked Storage

Scott Kipp
Office of the CTO
Brocade Inc.

July 2007

Categorizing Storage - DAS – SAN - NAS

Directly Attached Storage – DAS – Storage that is connected point-to-point with a server

- Originally on the SCSI Bus but also serial attached devices

Internal

External

Storage Area Network – A network that enables block level (raw) storage to be accessed from multiple servers

- Mainly based on Fibre Channel but iSCSI has about 5% revenue

Networked Attached Storage – Shared files (cooked storage) that are accessed over the LAN

- Originally a server that delivers shared files
- ~25% of the revenue of External Storage

Over $13B spent on external storage in 2006 to surpass internal for the first time.

Fibre Channel Replaces the SCSI Bus

SCSI Bus lead to captive storage with only 1 server (initiator) per bus

Fibre Channel liberated storage so that multiple servers could access it at the same time

Fibre Channel extended the distance and increased the speed of Storage I/O to Gigabit per second speeds (Gbps)

Internet SCSI (iSCSI) has mapped the SCSI protocol on top of Transmission Control Protocol (TCP/IP)

Fibre Channel Over Ethernet (FCoE) is being designed to encapsulate FC frames and transfer them over Ethernet networks

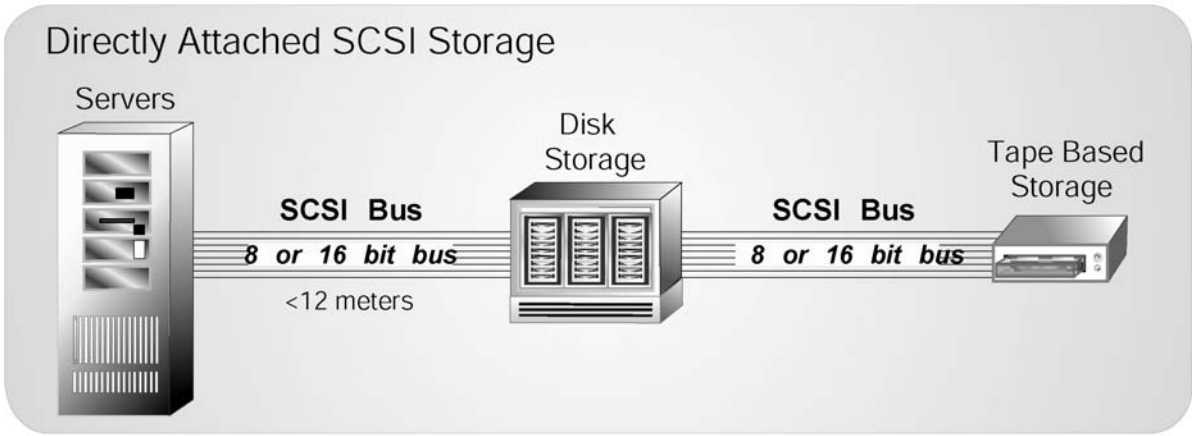

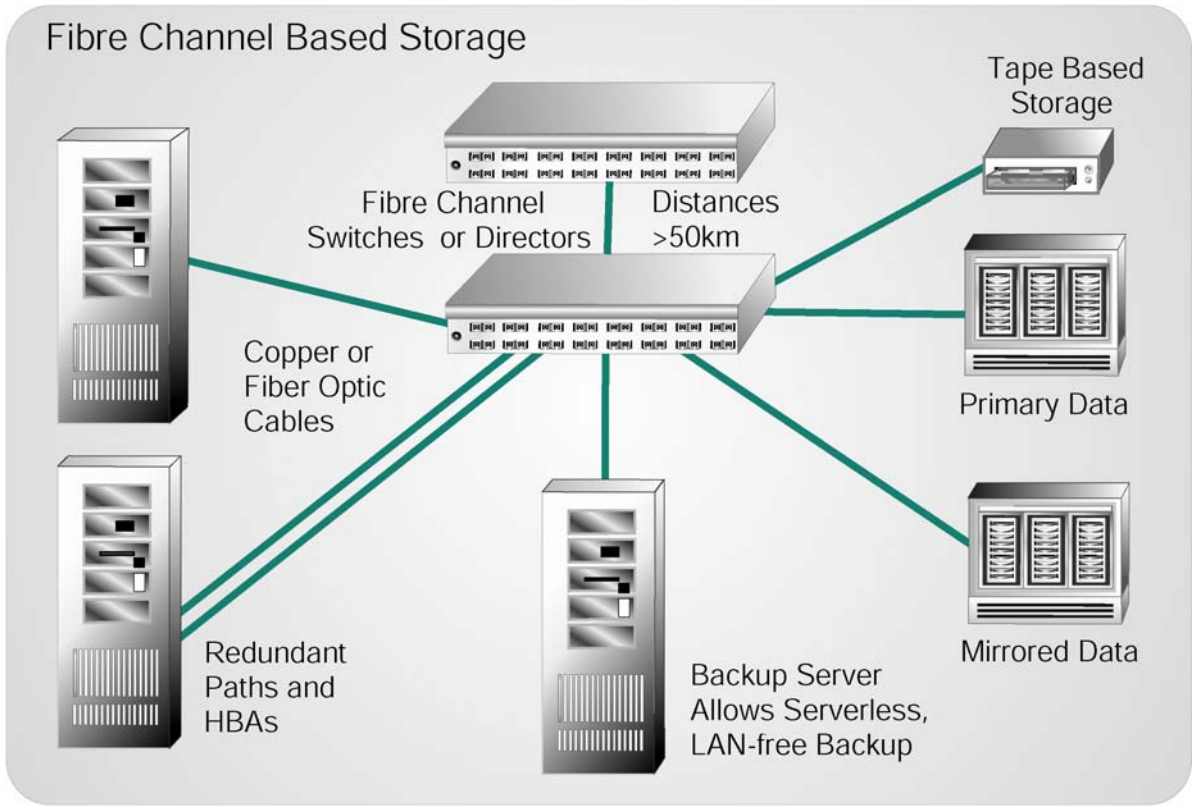

High Level Overview

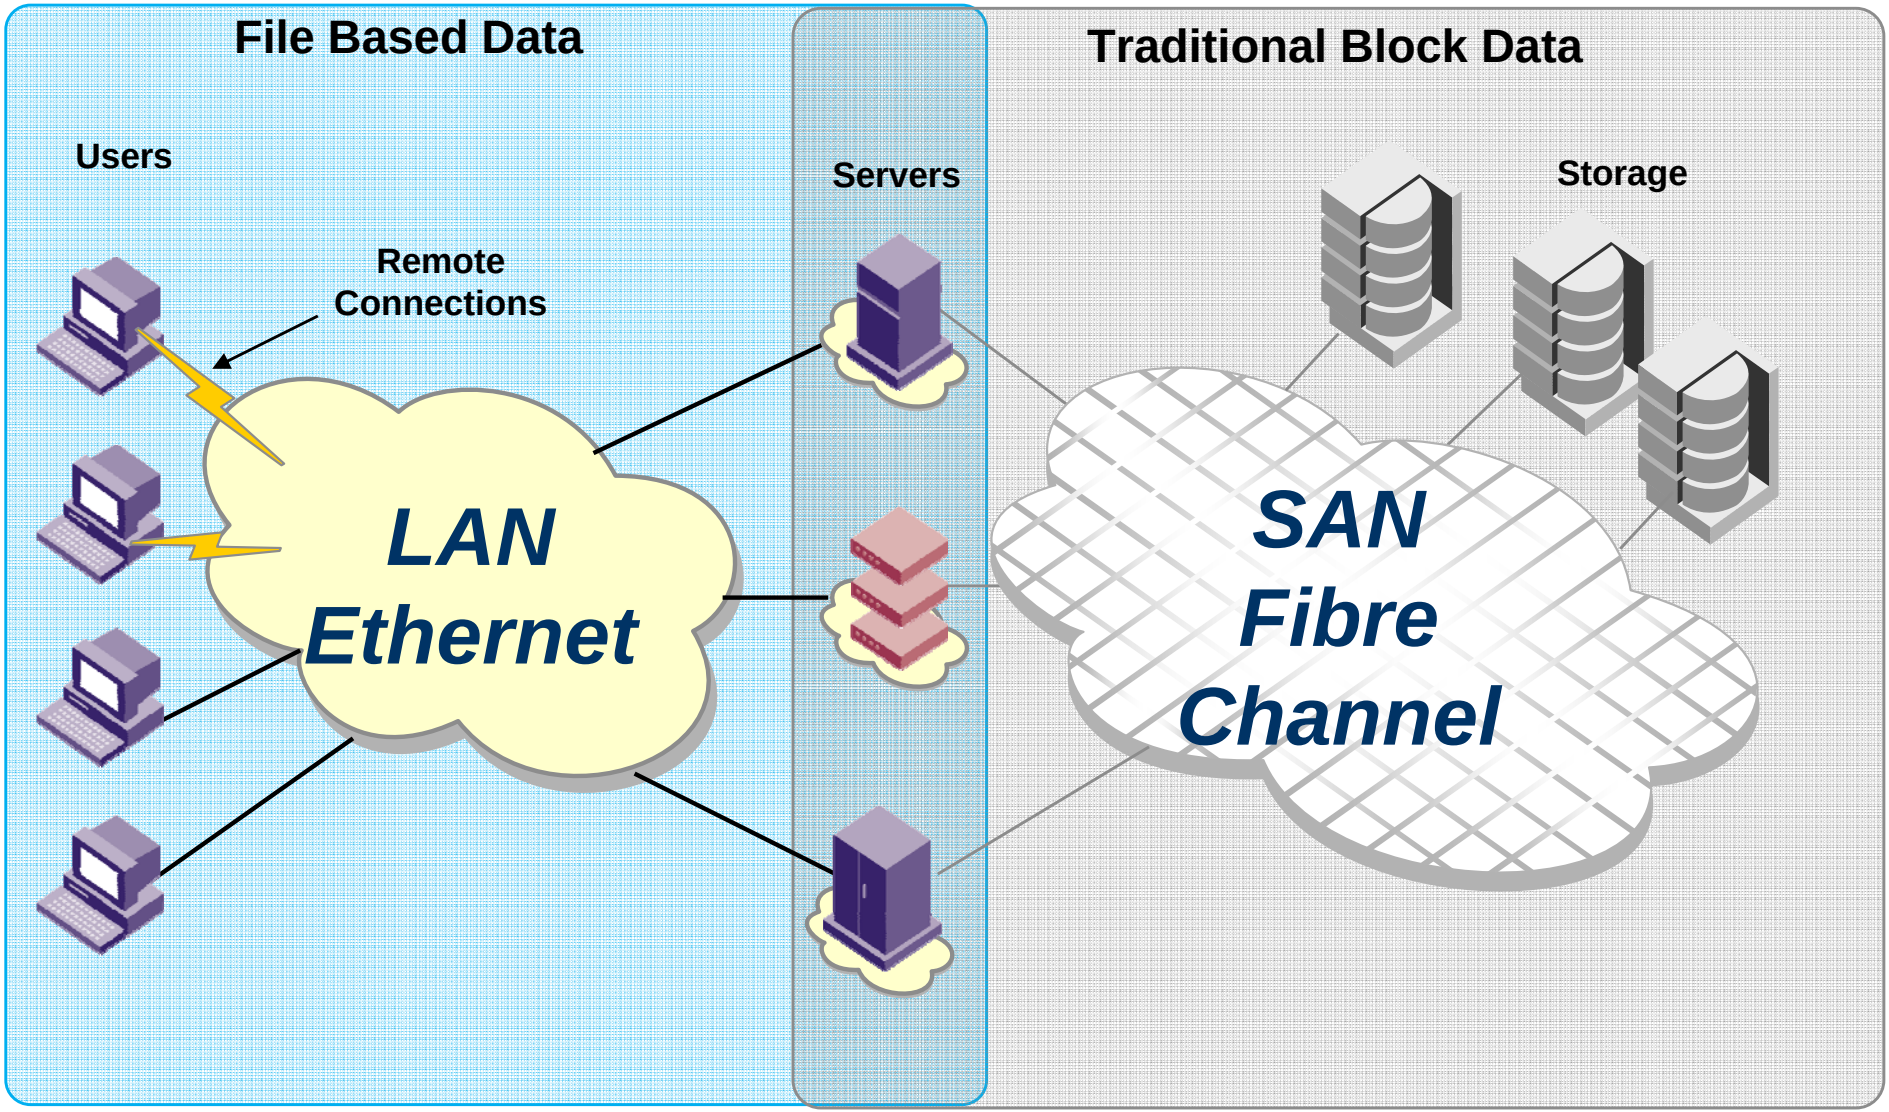

Fibre Channel Roadmap

Base2 99% of volume with all End Devices

Product Naming	Throughput (MBps)	Line Rate (GBaud)†	T11 Spec Technically Completed (Year)‡	Market Availability (Year)‡
1GFC	200	1.0625	1996	1997
2GFC	400	2.125	2000	2001
4GFC	800	4.25	2003	2005
8GFC	1600	8.5	2007	2008
16GFC	3200	17	2009	2011
32GFC	6400	34	2012	Market Demand
64GFC	12800	68	2016	Market Demand
128GFC	25600	136	2020	Market Demand

Base10 Only Inter-Switch Links

10GFC	2400	10.52	2003	2004
20GFC	4800	21.04	2008	2008
40GFC	9600	42.08	TBD	Market Demand
80GFC	19200	84.16	TBD	Market Demand
160GFC	38400	168.32	TBD	Market Demand

SAN Generations

SAN Generation	Release Date	# of FC Ports in a Corporate SAN	New Speeds
1st	1998	10s	1GFC
2nd	2002	100s	2GFC
3rd	2005	1,000s	4/10GFC, GE
4th	2008	10,000s with some NPIV	8/10GFC, 10GE
5th	2011	100,000s with NPIV	16/20GFC

NPIV = N_Port_ID Virtualization is a technique to acquire multiple N_Port_IDs or Fibre Channel Addresses for Virtual Servers
GFC = Gigabit/second Fibre Channel, GE = Gigabit Ethernet

Volume to Mid-Range Storage

Volume Storage Arrays contain 10s of disks for a few Terabytes of capacity

- 2 to 8 2GFC Links/ storage Array
  - 2-10+ Gbps of throughput

Midrange Storage Arrays contains 100s of disks for 10s of Terabytes of capacity

- 4 to 24 2/4GFC Links/ storage Array
  - 10s of Gbps of throughput

GFC = Gigabit/second Fibre Channel

High End Storage

Over 1,000 disk drives for 100s of TBs in a storage subsystem

- 64+ lanes of 2-4GFC
  - exceeding 100 Gbps of throughput

Server Types

IDC reported 7.5 million servers shipped in 2006 and breaks out three types of servers:

- Volume Servers (<$25k and 97.5% of volume)
  - Mostly use Directly Attached Storage
  - Blade servers are fastest growing segment with a 60% SAN attach rate
    - Each blade is counted as a server
    - Blade servers use more virtualization and multi-core processes to drive I/O
- Midrange Enterprise Servers ($25k to < $500k and 2.4% of volume)
  - Largely SAN attached
- High-End Enterprise Servers ($500k or more and 0.1% of volume)
  - Mostly SAN attached
  - Mainframes and UNIX

Current Blade Server Architecture

Each Blade Server has an FC lane to an embedded FC Switch

- Usually 10-16 server blades with 4GFC connections to switch
  - Up to 64 Gbps of throughput from servers
- 4-8 external 4GFC ports for up to 32Gbps of throughput / switch

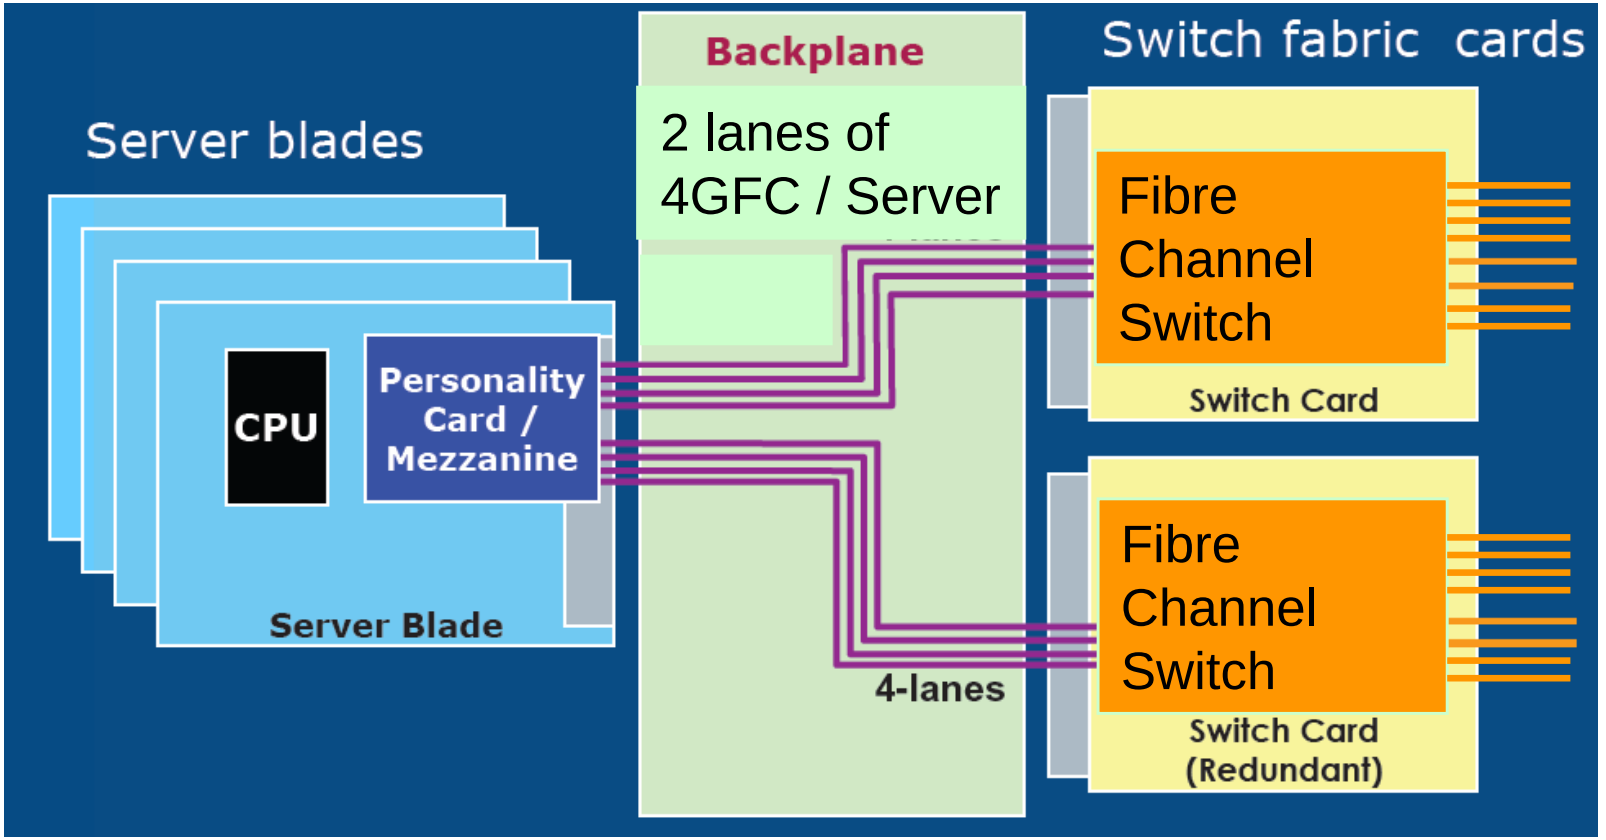

32 Gbps to the FC Fabric

32 Gbps to the FC Fabric

Virtualization Drives Bandwidth

Each Blade Server has multiple virtual machines (VMs) running different applications that increase the utilization of the server and drive bandwidth

Think Grid or adaptive computing

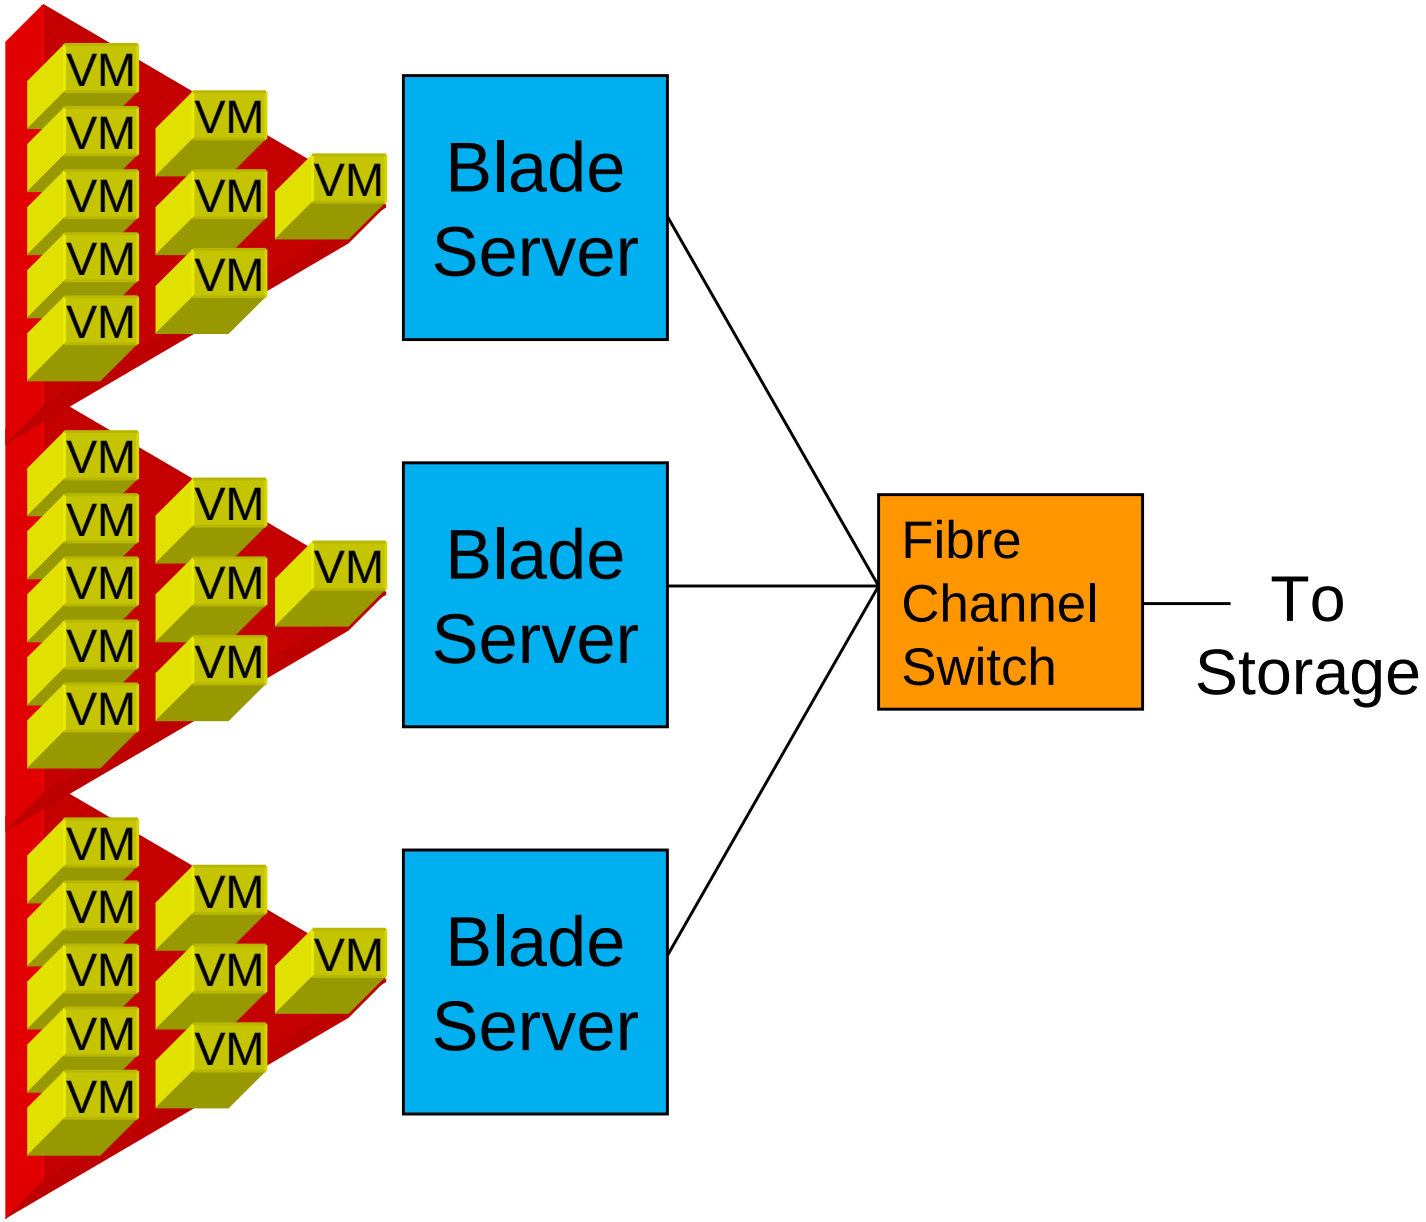

Midrange and High End Servers

Midrange Servers have a few processors with a few cores to a couple dozen processors

- A few FC Host Bus Adapters (HBAs) with 4-25 physical ports running at 4GFC
  - 10s of Gbps* to storage

High End Servers (mainly mainframes and UNIX) have tens of processors

- Hundreds of FC ports at 4GFC
  - 100s of Gbps* to storage

*These bandwidths are current architectural limits and not the bandwidth that is being driven through the servers.

Fibre Channel Switches and Directors

10s to 100s of 4GFC Ports

- Soon will support 8GFC
- Already supports 10GFC for Inter-Switch Links
- Supports Ethernet ports for iSCSI servers and Fibre Channel over Internet Protocol (FCIP)
- Supports distance extension over SONET/ATM/GFP/PW

Architectural Details

- http://www.brocade.com/san/pdf/whitepapers/Achieving_Ent_SAN_Perf_48000_Dir_WP_01.pdf

GFC = Gigabit Fibre Channel, GFP = Generic Framing Protocol, PW = Pseudo Wire

Trunking of InterSwitch Links

8 lanes of 4GFC are trunked to effectively create a 32GFC link

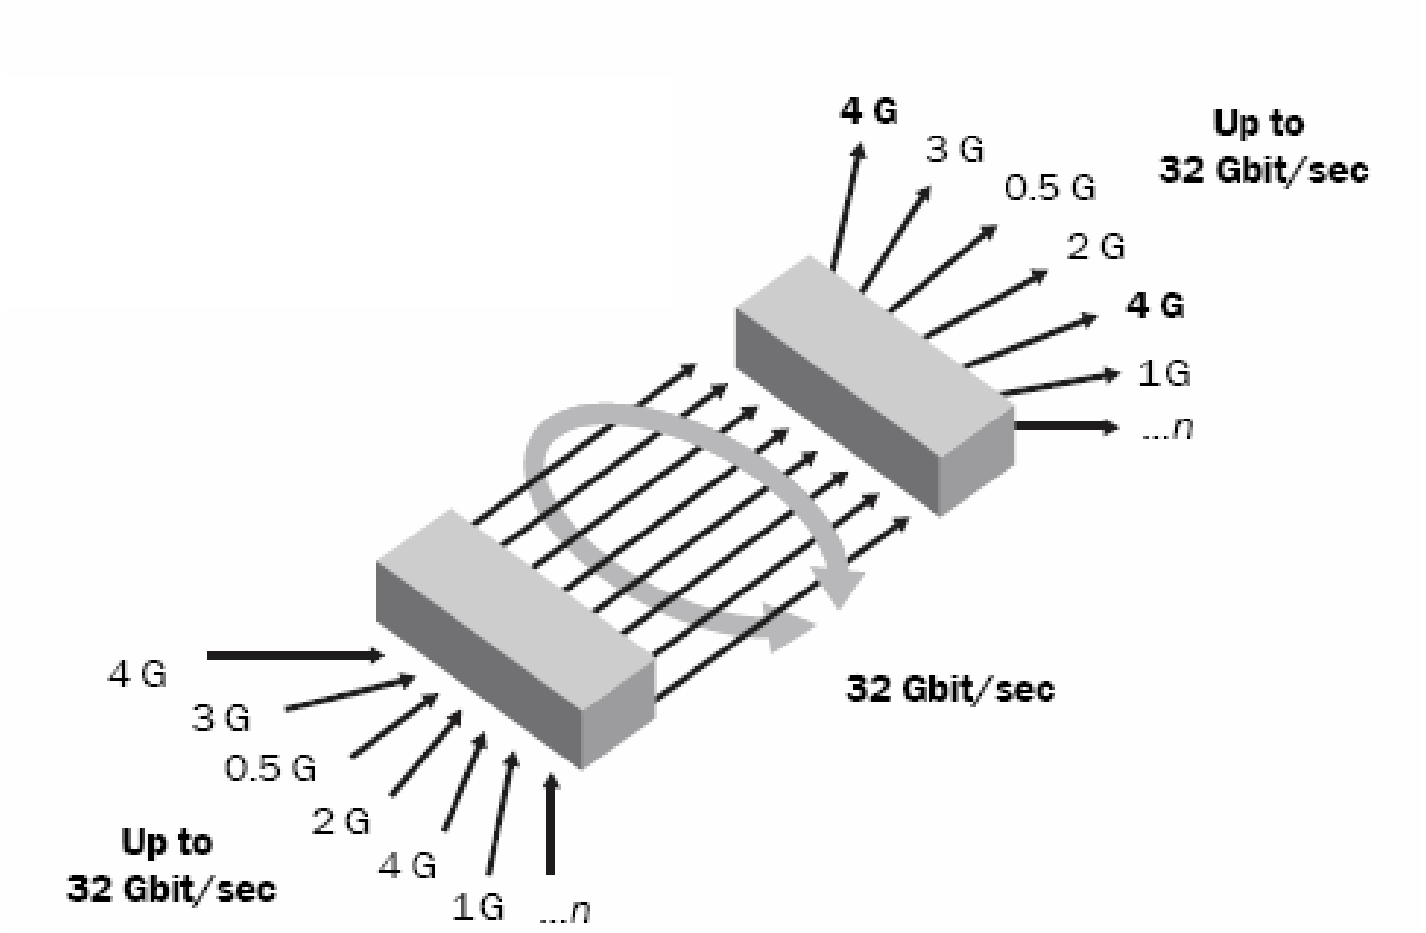

Network Attached Storage (NAS)

NAS is storage on the LAN that share files (file level storage) while SANs share Block level storage behind servers

- NAS products host shared drives (think m: or v: drives)
- End users quickly began storing and sharing large amounts of computer files
- Unstructured file (end user) data is the fastest growing data type – think powerpoint, excel, photos and MP3s

Large NAS heads or filers have Ethernet front ports and Fibre Channel back ports

NAS Products

NAS file-based products follow the same progression of SAN's block-based storage products

Volume systems may have a few TeraBytes of Directly Attached Storage with a few Gigabit Ethernet front end ports

High end systems hold hundreds or TBs of SAN attached storage and tens of FC and GE ports and even 10GE ports

High Level Overview with NAS

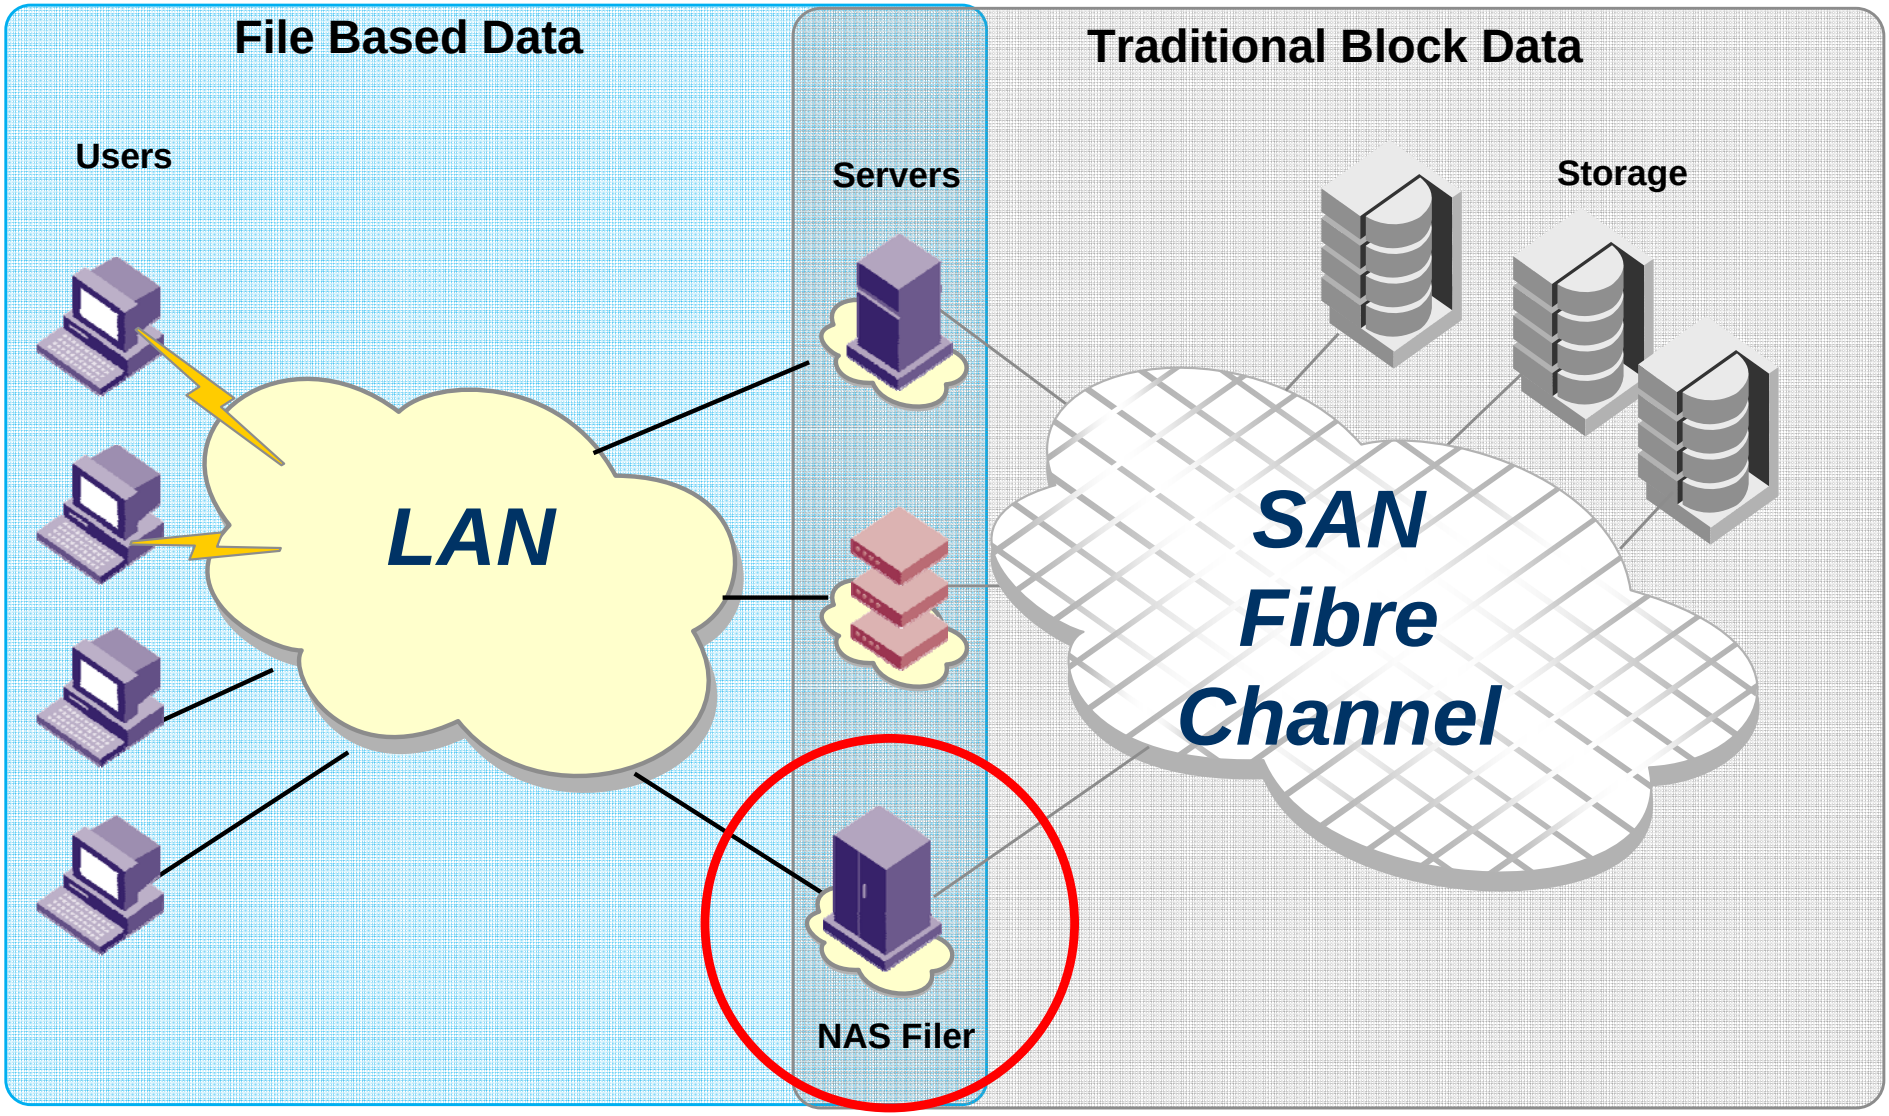

Data Convergence

Data Centers of the future will merge multiple protocols (FC, Ethernet, FCoE) and technologies

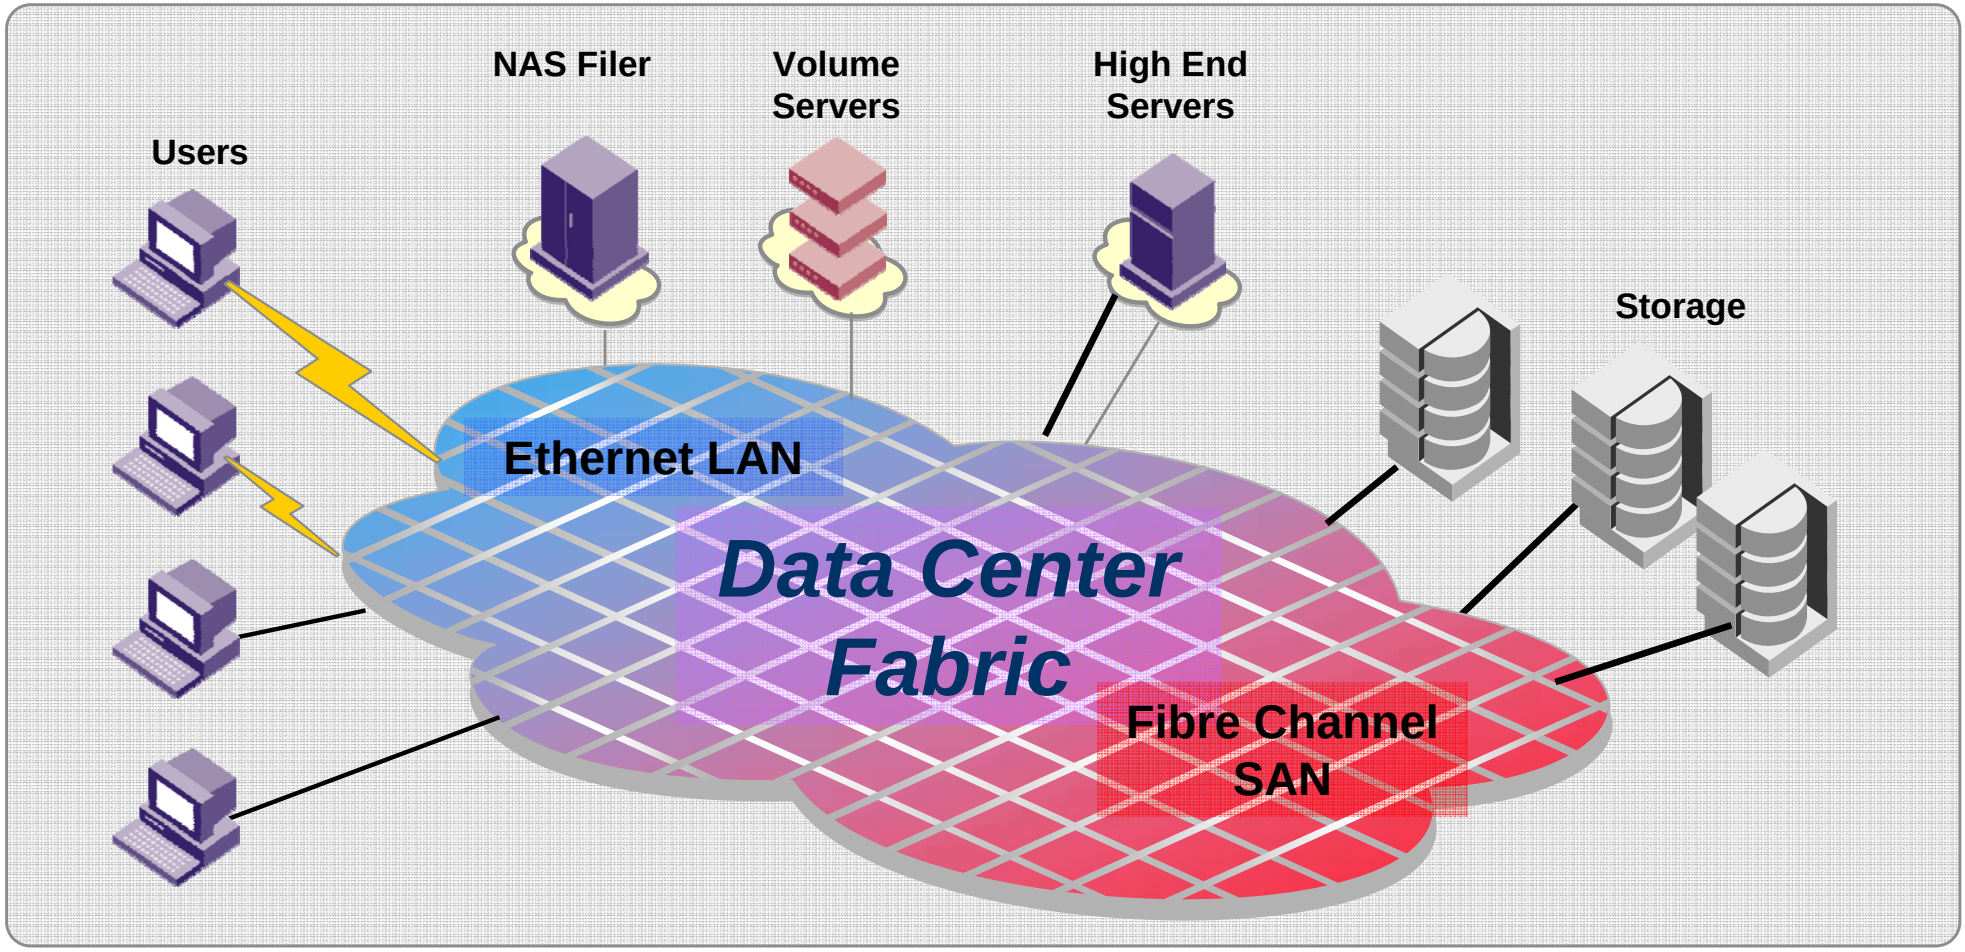

Conclusion

SANs are the basis for storage and server consolidation that drives 10s to 100s of Gbps of throughput

NAS is an application that consumes a large amount of storage and serves a large amount of files in the few to tens of Gbps range

Virtualization of Servers, switches and storage increases utilization of physical assets and bandwidth in the SAN space

Corporate networks are converging with Fibre Channel over Ethernet (FCoE) being defined as the latest “Unified Wire” protocol in T11 and IEEE

Backup Slides

SANs May be a Collection of Fabrics

Corporations may have multiple Fibre Channel Fabrics and use new virtualization techniques

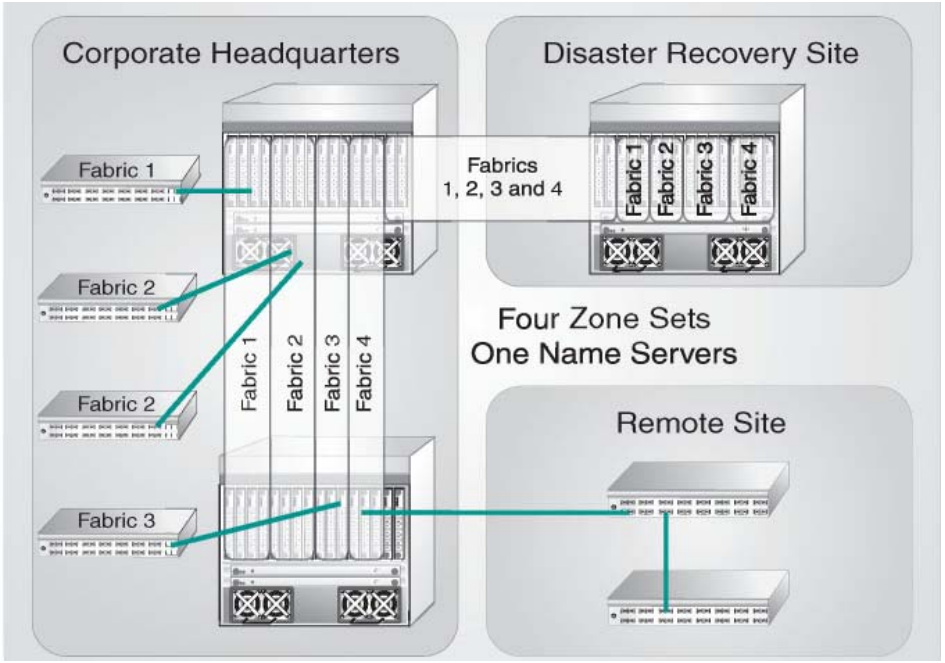

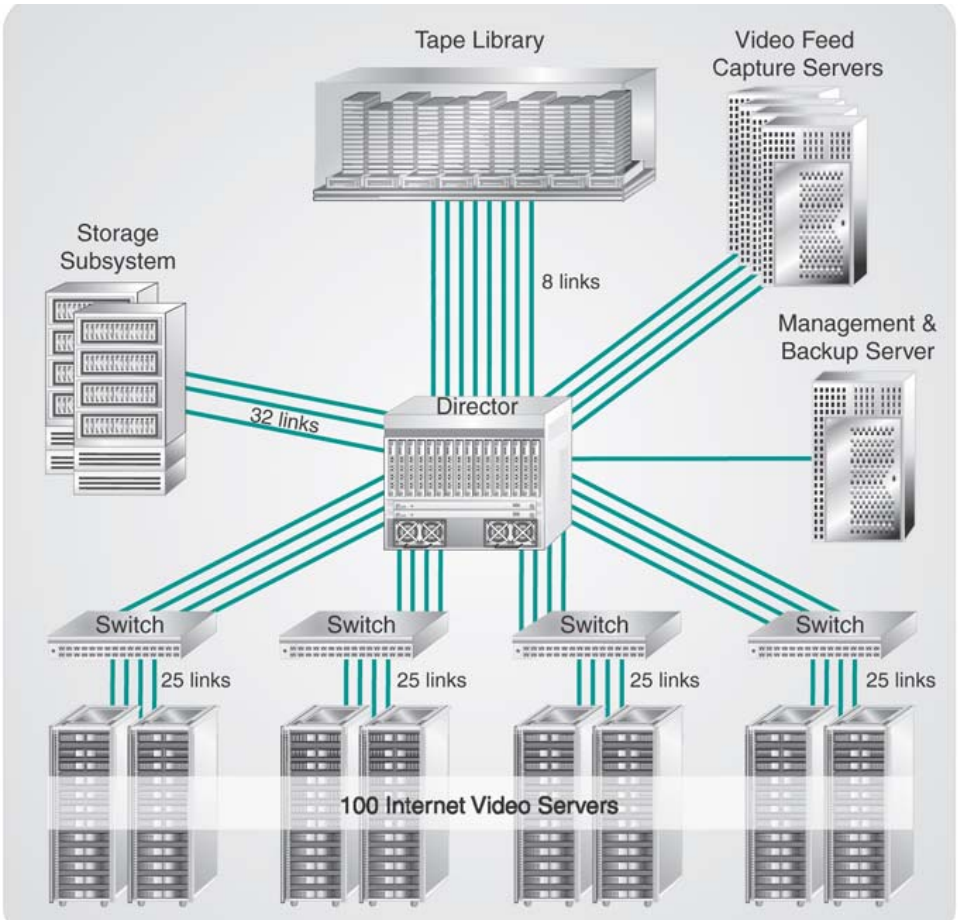

Server and storage Independence

Many servers can access networked storage

One server can access many types of shared storage

Typically there are 4-8 server ports for every storage port