BROCADE

STORAGE GROWTH AND ETHERNET

Scott Kipp

September 12, 2011

What is an Exabyte? – 1 Million Terabyte Drives

- Earth created or replicated over 1,000 Exabytes of data in 2010 – that's 143GB for each of 7 Billion people

SI decimal prefixes – short scale			Binary usage	IEC binary prefixes	
Common Name	Name (Symbol)	Value		Name (Symbol)	Value
Thousand	kilobyte (kB)	10^3	2^{10}	kibibyte (KiB)	2^{10}
Million	megabyte (MB)	10^6	2^{20}	mebibyte (MiB)	2^{20}
Billion	gigabyte (GB)	10^9	2^{30}	gibibyte (GiB)	2^{30}
Trillion	terabyte (TB)	10^{12}	2^{40}	tebibyte (TiB)	2^{40}
Quadrillion	petabyte (PB)	10^{15}	2^{50}	pebibyte (PiB)	2^{50}
Quintillion	exabyte (EB)	10^{18}	2^{60}	exbibyte (EiB)	2^{60}
Sextillion	**zettabyte (ZB)**	10^{21}	2^{70}	zebibyte (ZiB)	2^{70}
Septillion	yottabyte (YB)	10^{24}	2^{80}	yobibyte (YiB)	2^{80}
Googol	GoogolByte?	10^{100}			

The world created over a ZB last year!

An Exabyte is not Infinite

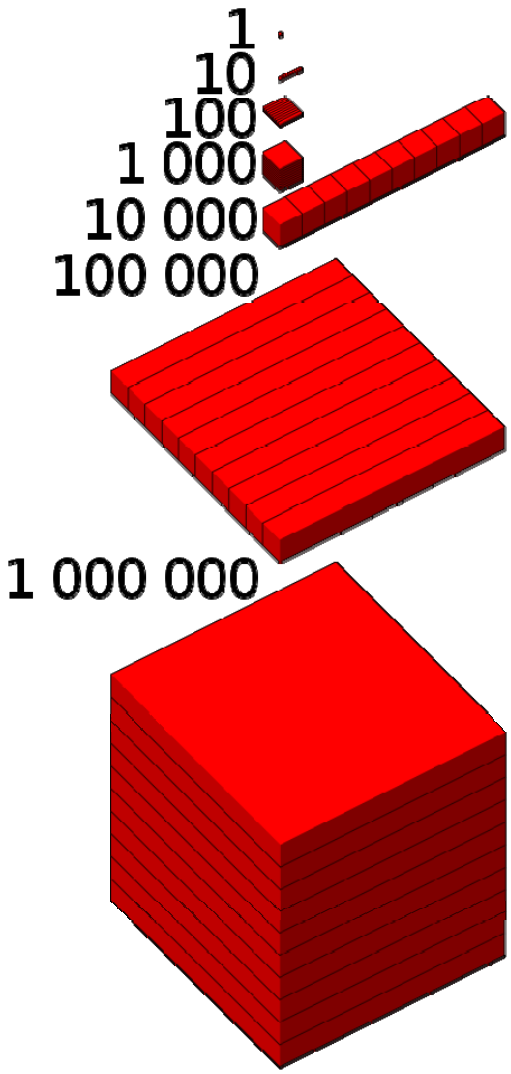

1 TB Hard Disk Drive (HDD)
10 HDD– Just A Bunch of Disks (JBOD)
100 HDD– 1 Storage Subsystem
-Controller +JBODs + Cache + IO Cards
1,000 HDD– 1Row of Storage Subsystems
At 1TB/HDD = 1PB per row
10,000 HDD– 1 Large Data Center
With 10PB of Data
100,000 HDD – 100PB - Storage capacity
of European Grid Infrastructure
1,000,000 TeraByte HDDs
– 1 Exabyte

1 Storage Subsystem
with 72 Disk Drives

The Source of the Data

The Digital Universe Study

- Data is growing 40-50% per year –doubling every two years – compared to IP traffic growth of 30-40%
- 75% of the data is created by individuals, but enterprises have some liability for 80% of it
  - For data creation, think of computer files, music files, Digital Video Recorders, DVDs, backup drives, digital pictures…
  - They don't explain the 80% number well, but I bet an example is that the cable company has liability for the shows on your DVR
- 25% of data is generated by machines and that is growing fast with sensors and remote monitoring
- Over the next decade, the number of servers (physical and virtual) will grow by a factor of 10, storage will grow by a factor of 50 and files will grow by a factor of 75

Source: The Digital Universe Study: http://www.emc.com/leadership/programs/digital-universe.htm

How much will it grow? Into the Zettabytes

1,000 Exabytes is a Zettabyte

- We create more digital data every couple of years than was created in history
- 500,000 Trillion files in 2011

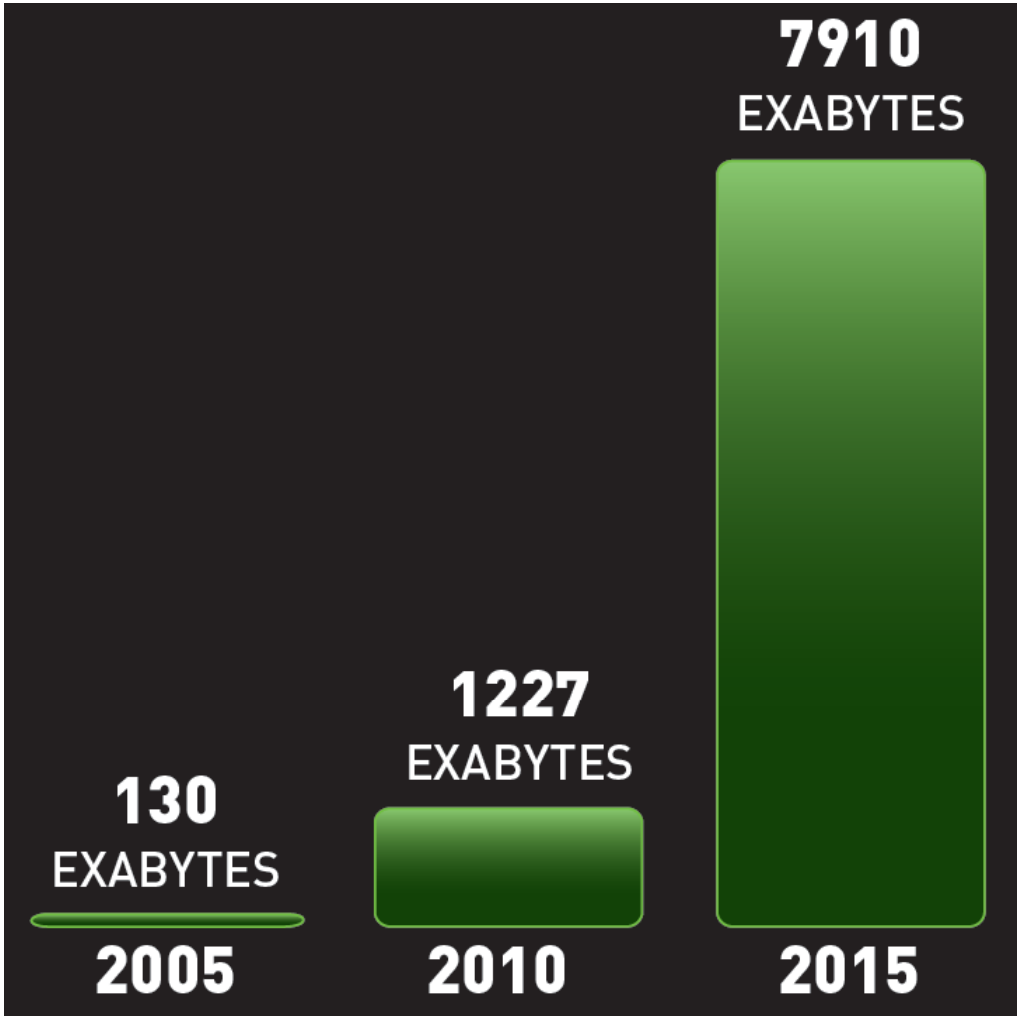

Source: The Digital Universe Study: http://www.emc.com/leadership/programs/digital-universe.htm

Replication is the Great Multiplier

Look back at Andy Bach's NYSE presentation

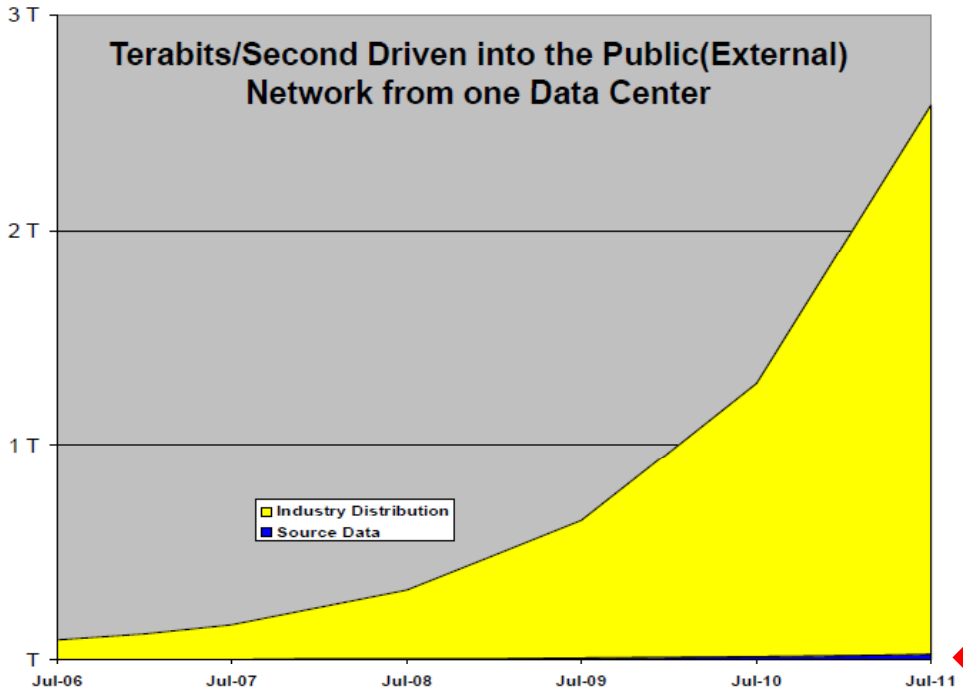

This little blue sliver of data is replicated about a hundred times to different organizations

How many people make their own copy of the data within each organization?

Source:http://www.ieee802.org/3/ad_hoc/bwa/public/jun11/bach_01a_0611.pdf

CERN Case Study

CERN's LHC generates 15 PB of data every year that is distributed over their core network with a 10Tbps capacity

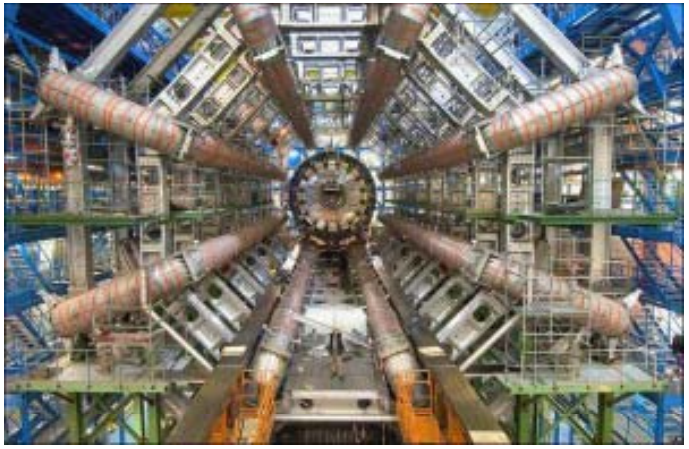

LHC = Large Hadron Collider

Tier 0 storage at CERN – 15 Petabytes

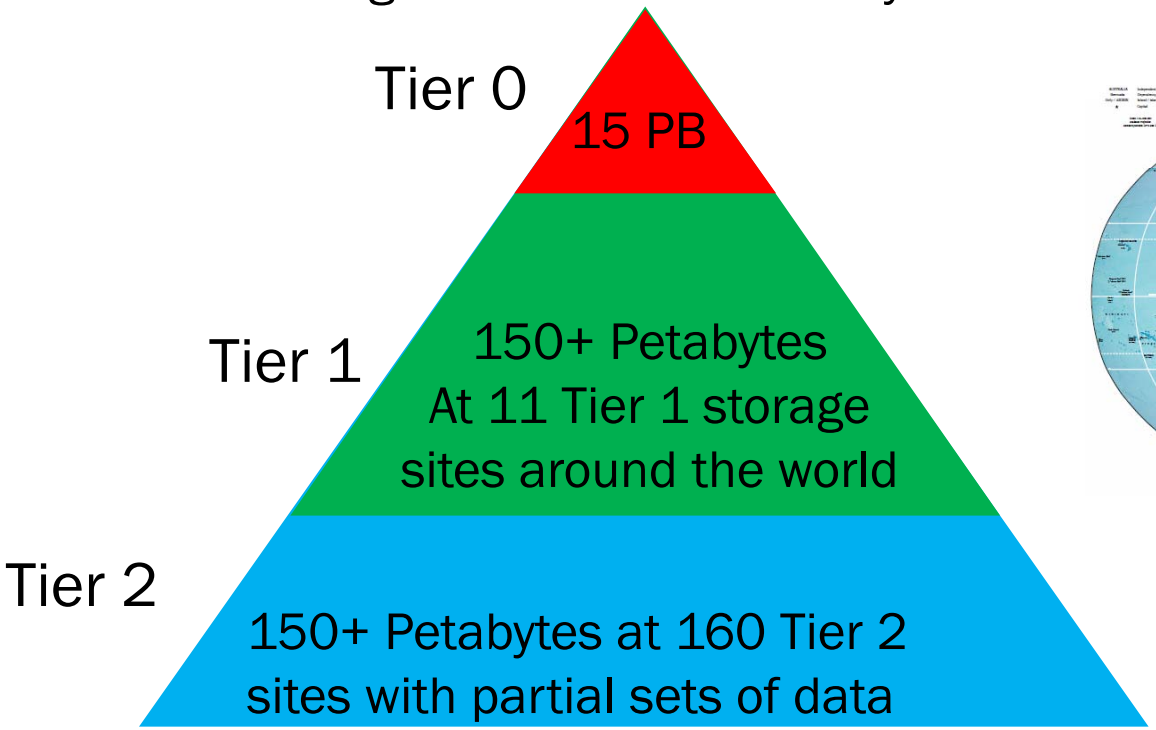

Transferring Large Data Sets – Big Data

- To transfer 15PB would take about:
  - 3.8 Years at 1GbE
  - 137.5 days at 10GbE
  - 13.75 days at 100GbE
  - 33 Hours at TbE

Size of Data to Exchange	Latency of 1GbE	Latency of 10GbE	Latency of 40GbE	Latency of 100GbE	Latency of 400GbE	Latency of 1TbE
1 Gigabyte	8	0.8	0.2	0.08	0.02	0.008
10GB	80	8	2	0.8	0.2	0.08
100GB	800	80	20	8	2	0.8
1 Terabyte	8,000	800	200	80	20	8
1 Petabyte	8M	800000	200,000	80,000	20,000	8,000
10 PB	80M	8M	2M	800,000	200,000	80,000
100 PB	800M	80M	20M	8M	2M	800,000

Video Content – The Growth Component

- Cisco's Visual Network Index (VNI)* predicts ~1 ZB of content will be distributed over Global IP networks in 2015 while there will be almost 8ZB of data produced and replicated that year
- Consumer video streaming is the main bandwidth driver in the future according to VNI*
- 1GB of content can produce 1PB** of data transfers so the storage component is one millionth compared to the networking component in some applications

*http://www.cisco.com/web/solutions/sp/vni/vni_forecast_highlights/index.html
**http://www.emc.com/collateral/demos/microsites/emc-digital-universe2011/index.htm

Breaking IT Down into Information Technology (IT)

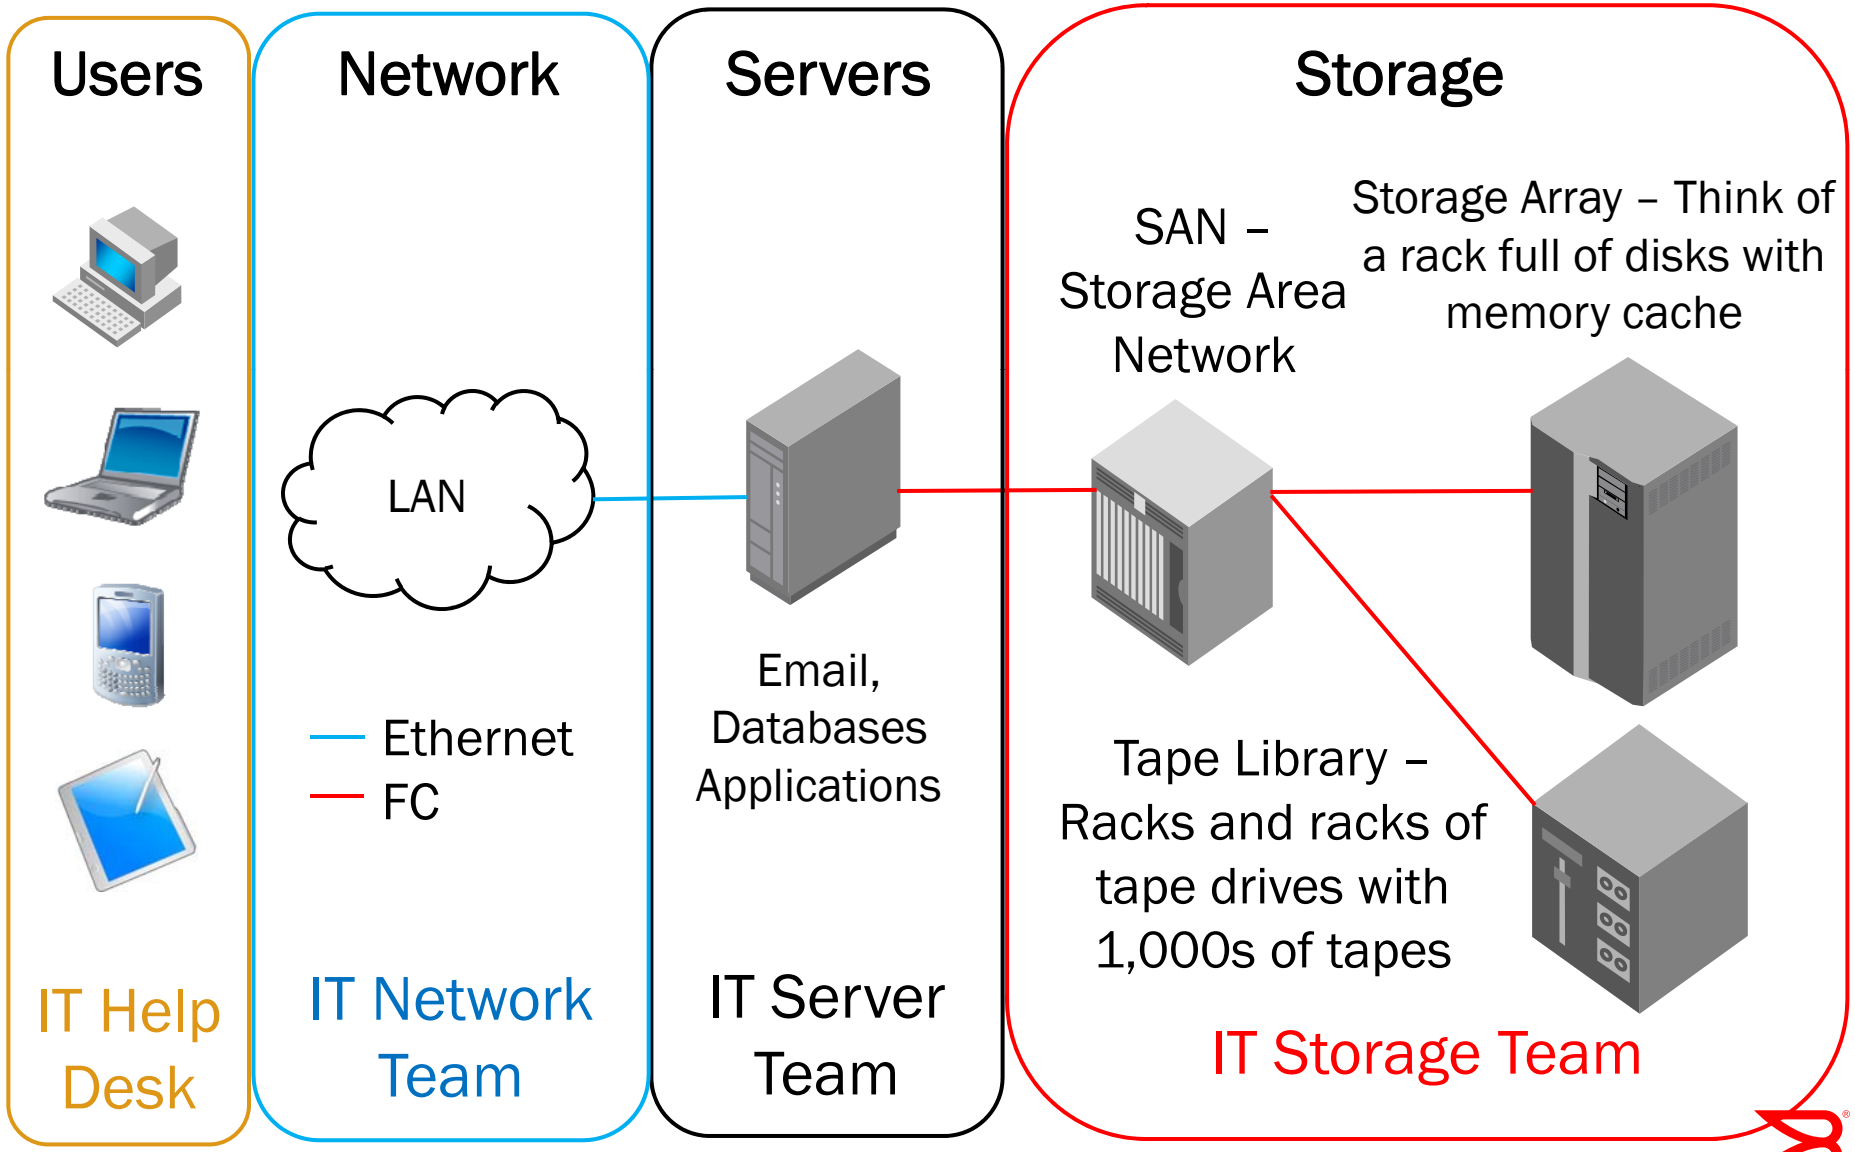

Explanation of Storage Access

DAS = Direct Attached Storage
NAS = Network Attached Storage
iSCSI – Internet Small Computer Systems Interface
LOM = LAN on Motherboard
NIC = Network Interface Card
HBA = Host Bus Adapter
CNA = Converged Network Adapter

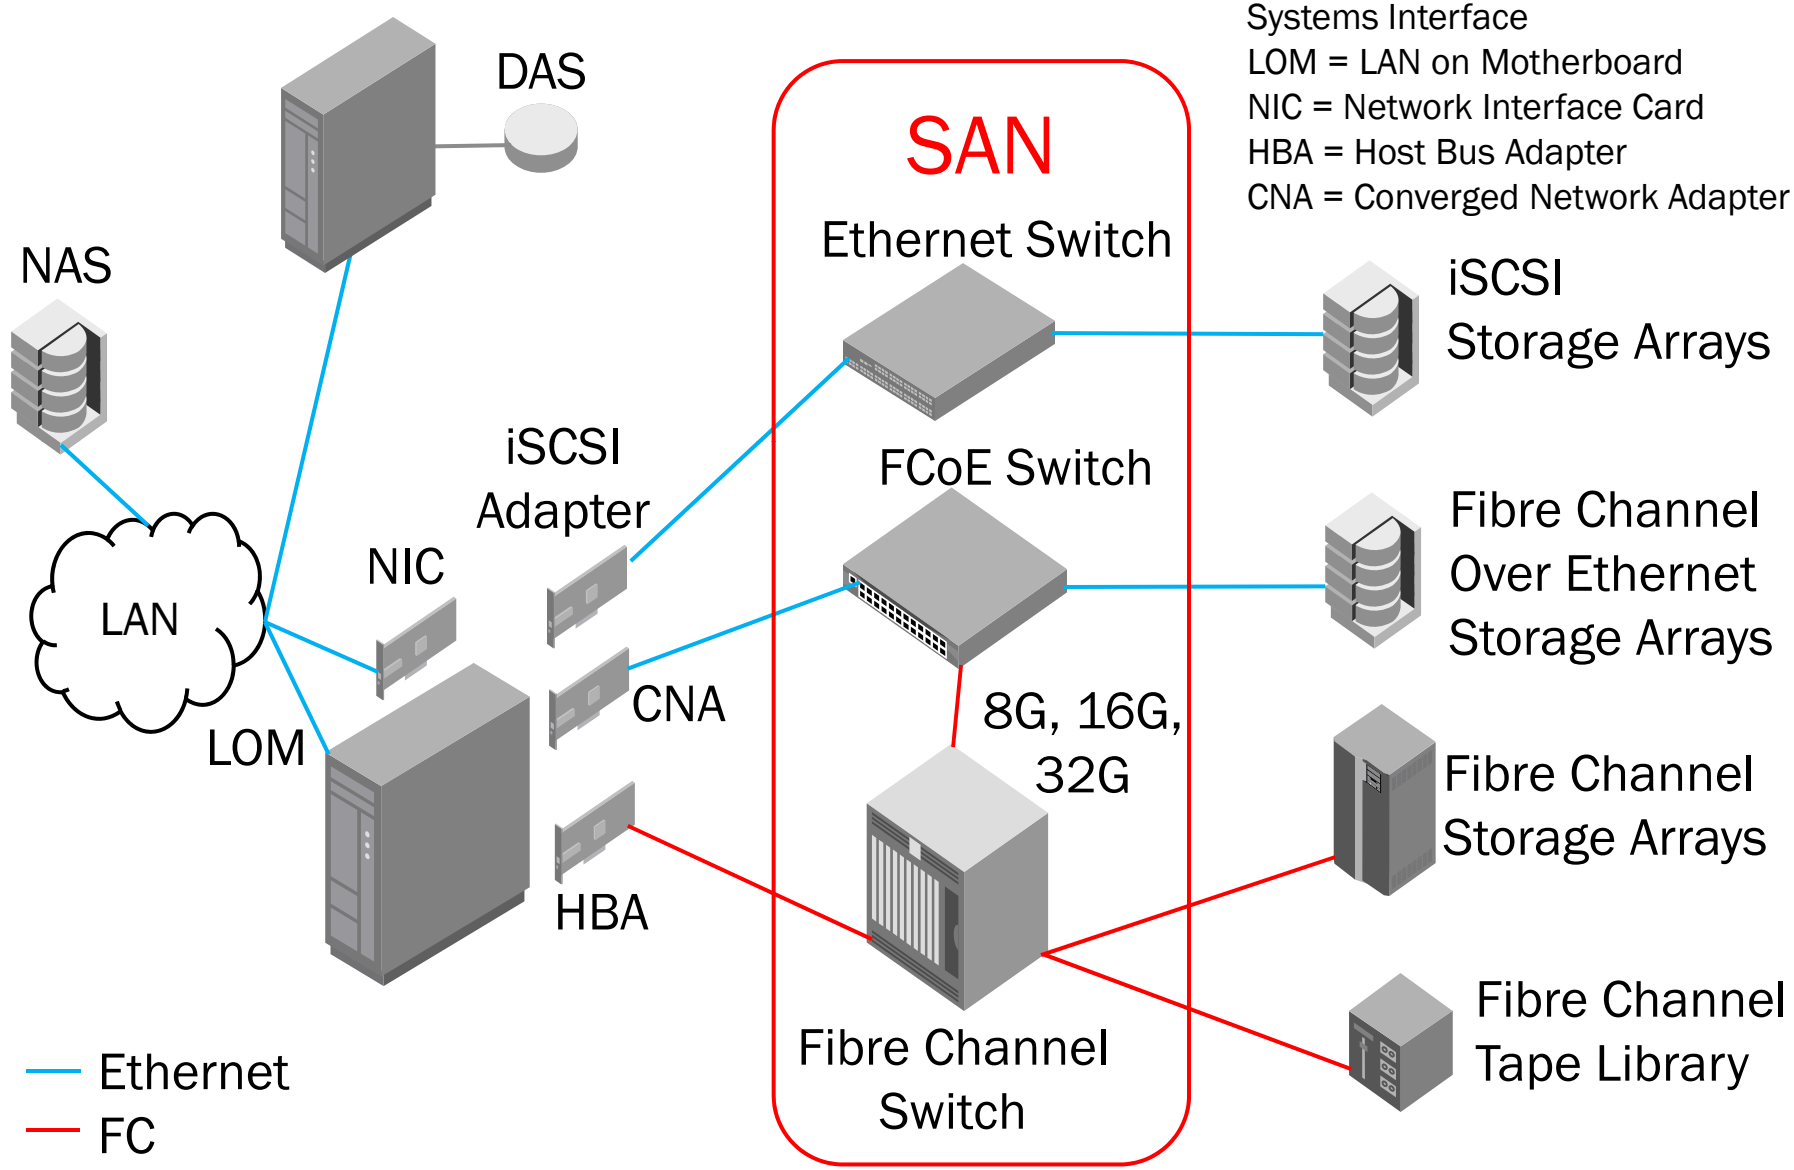

External Storage Sales in Exabytes

- 17EB in 2011 growing to 90 EB in 2015 – About 1% of the digital universe
- Ethernet-based storage expected to grow to over 50% of storage capacity in 2015
- NAS is usually unstructured data, but its supporting more applications

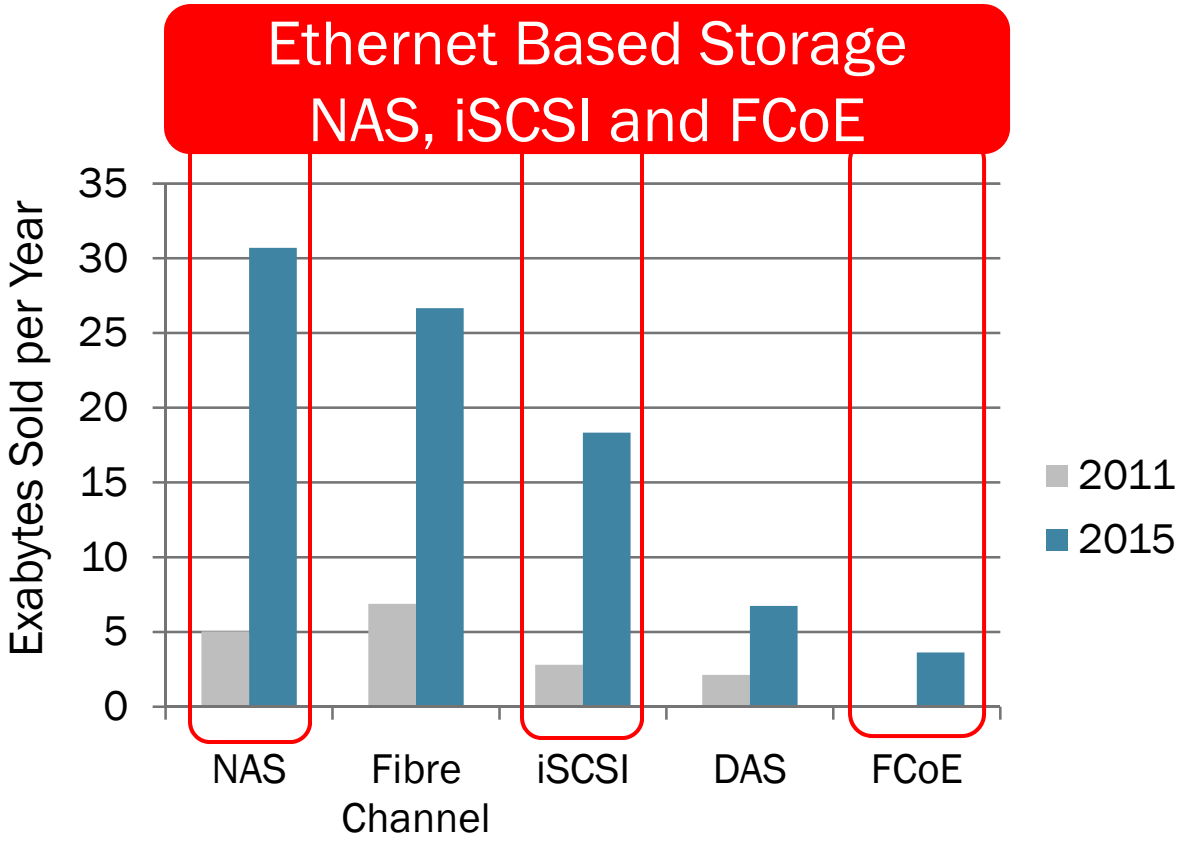

Source: IDC Worldwide Enterprise Storage Systems 2011-2015

NAS – Network Attached Storage

An application server that serves files

Users

Network connections mostly 1GbE, moving to 10GbE and soon 40GbE

LAN

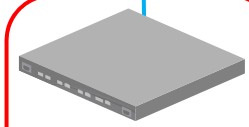

Low End NAS

Low End NAS can be as small as an individual disk drive or tens of disk in a 1U or 2U chassis

High End NAS

NAS – Think of a rack full of disks from with an Ethernet network interface

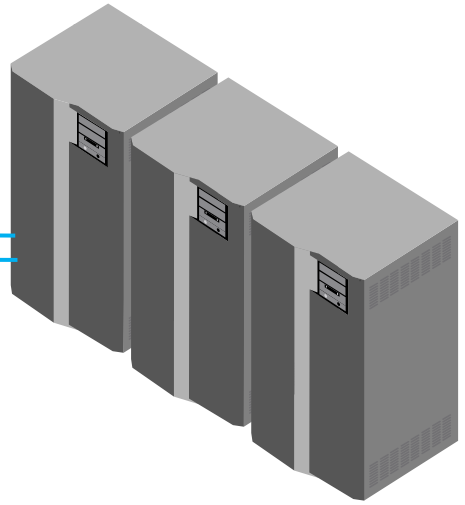

File Server -
NAS Head for unstructured data – serving millions and even billions of files

Ethernet

Fibre Channel Storage Area Network (SAN)

SCSI over Optical Fiber

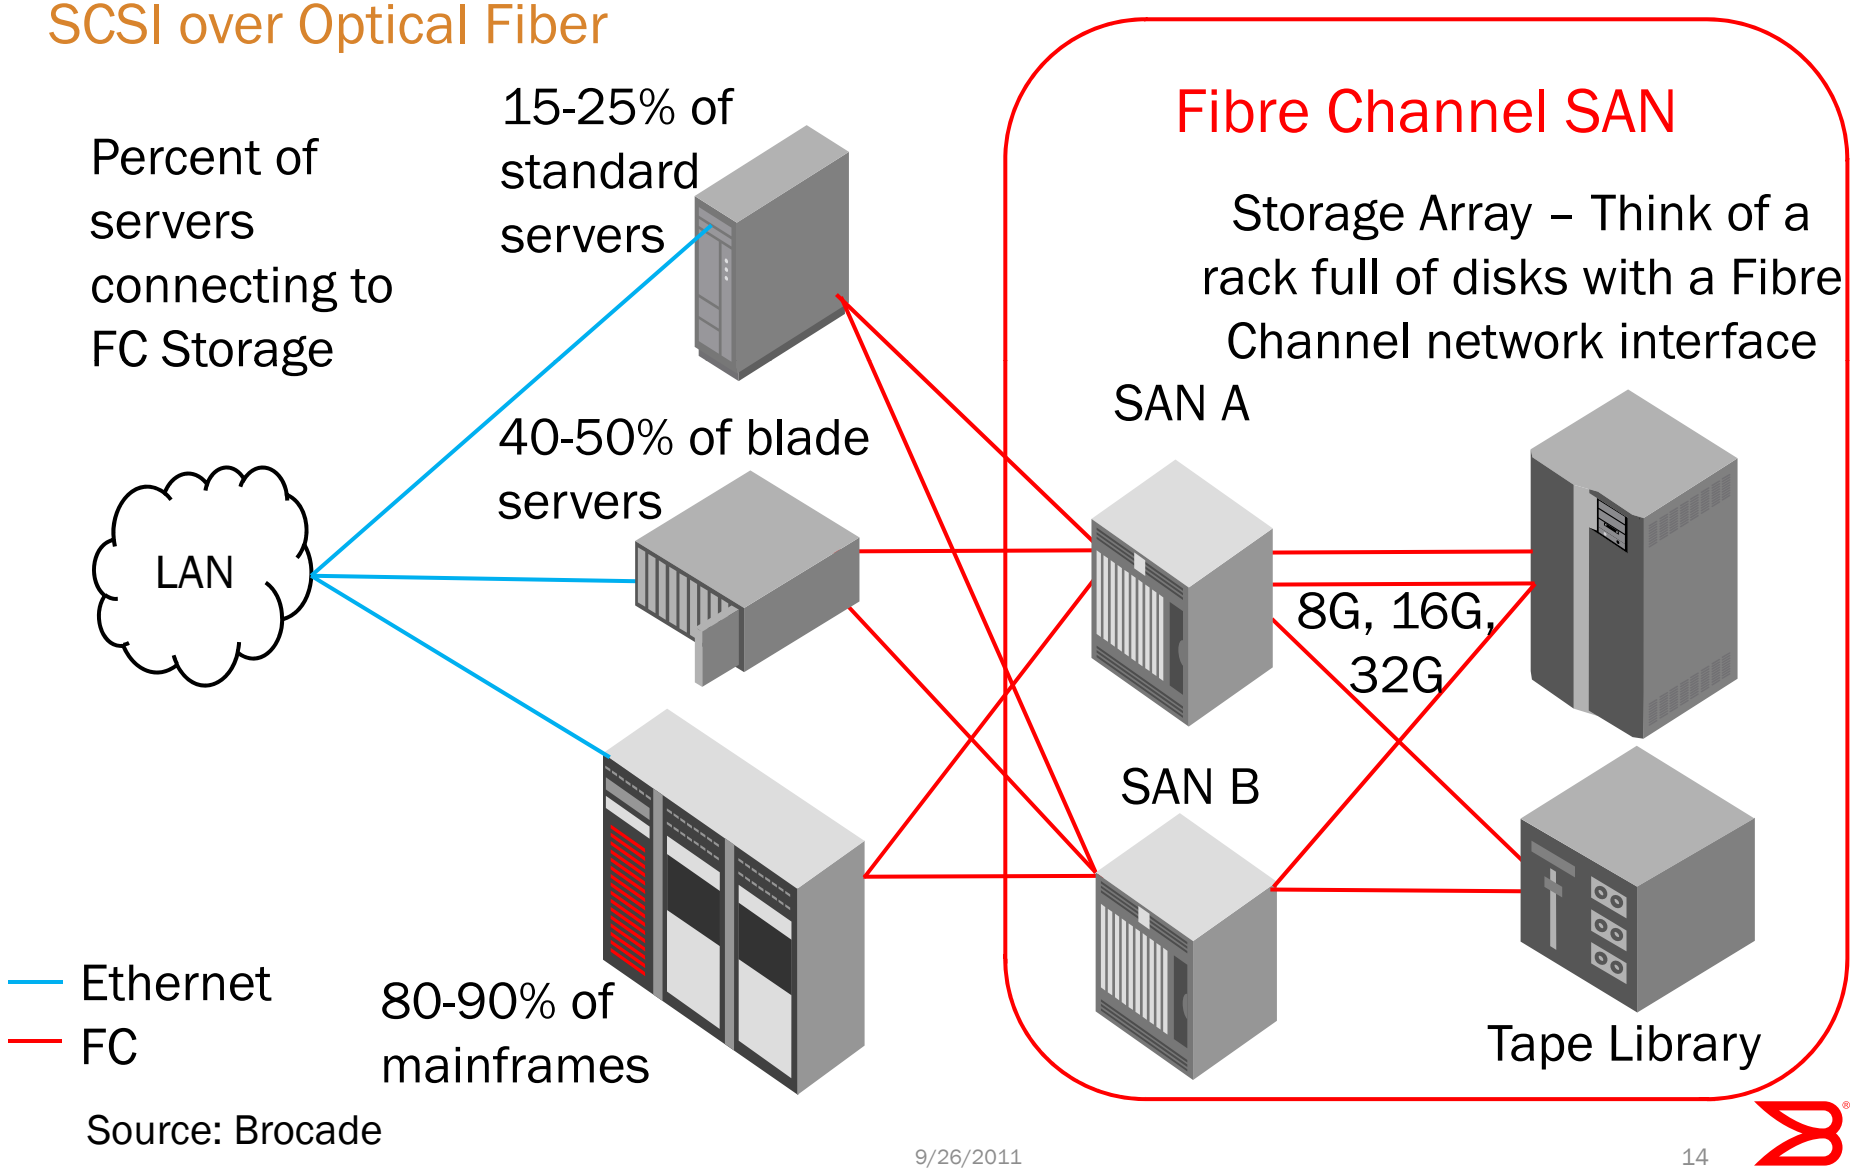

Source: Brocade

iSCSI – SCSI over Ethernet

Still Two Networks

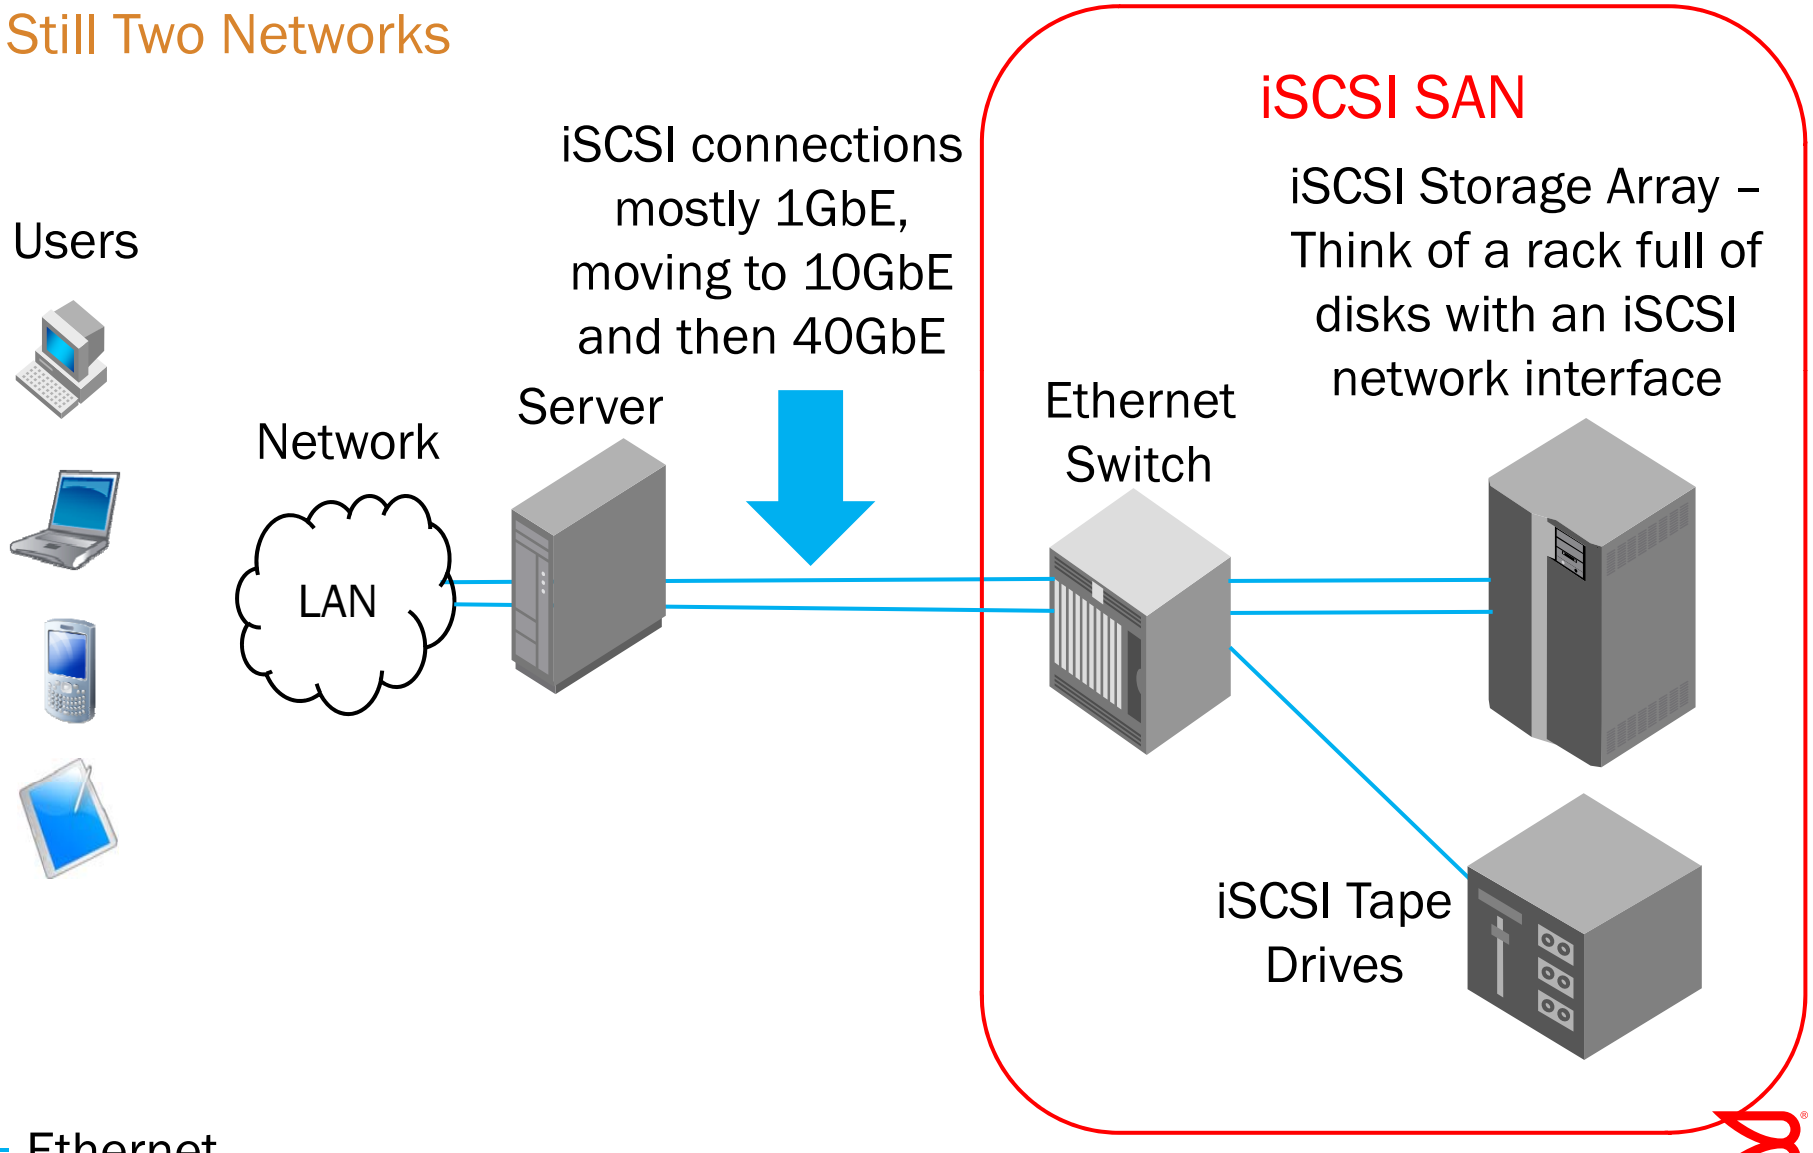

FCoE SAN – Networked Storage Access

Mainly for Top of Rack deployments

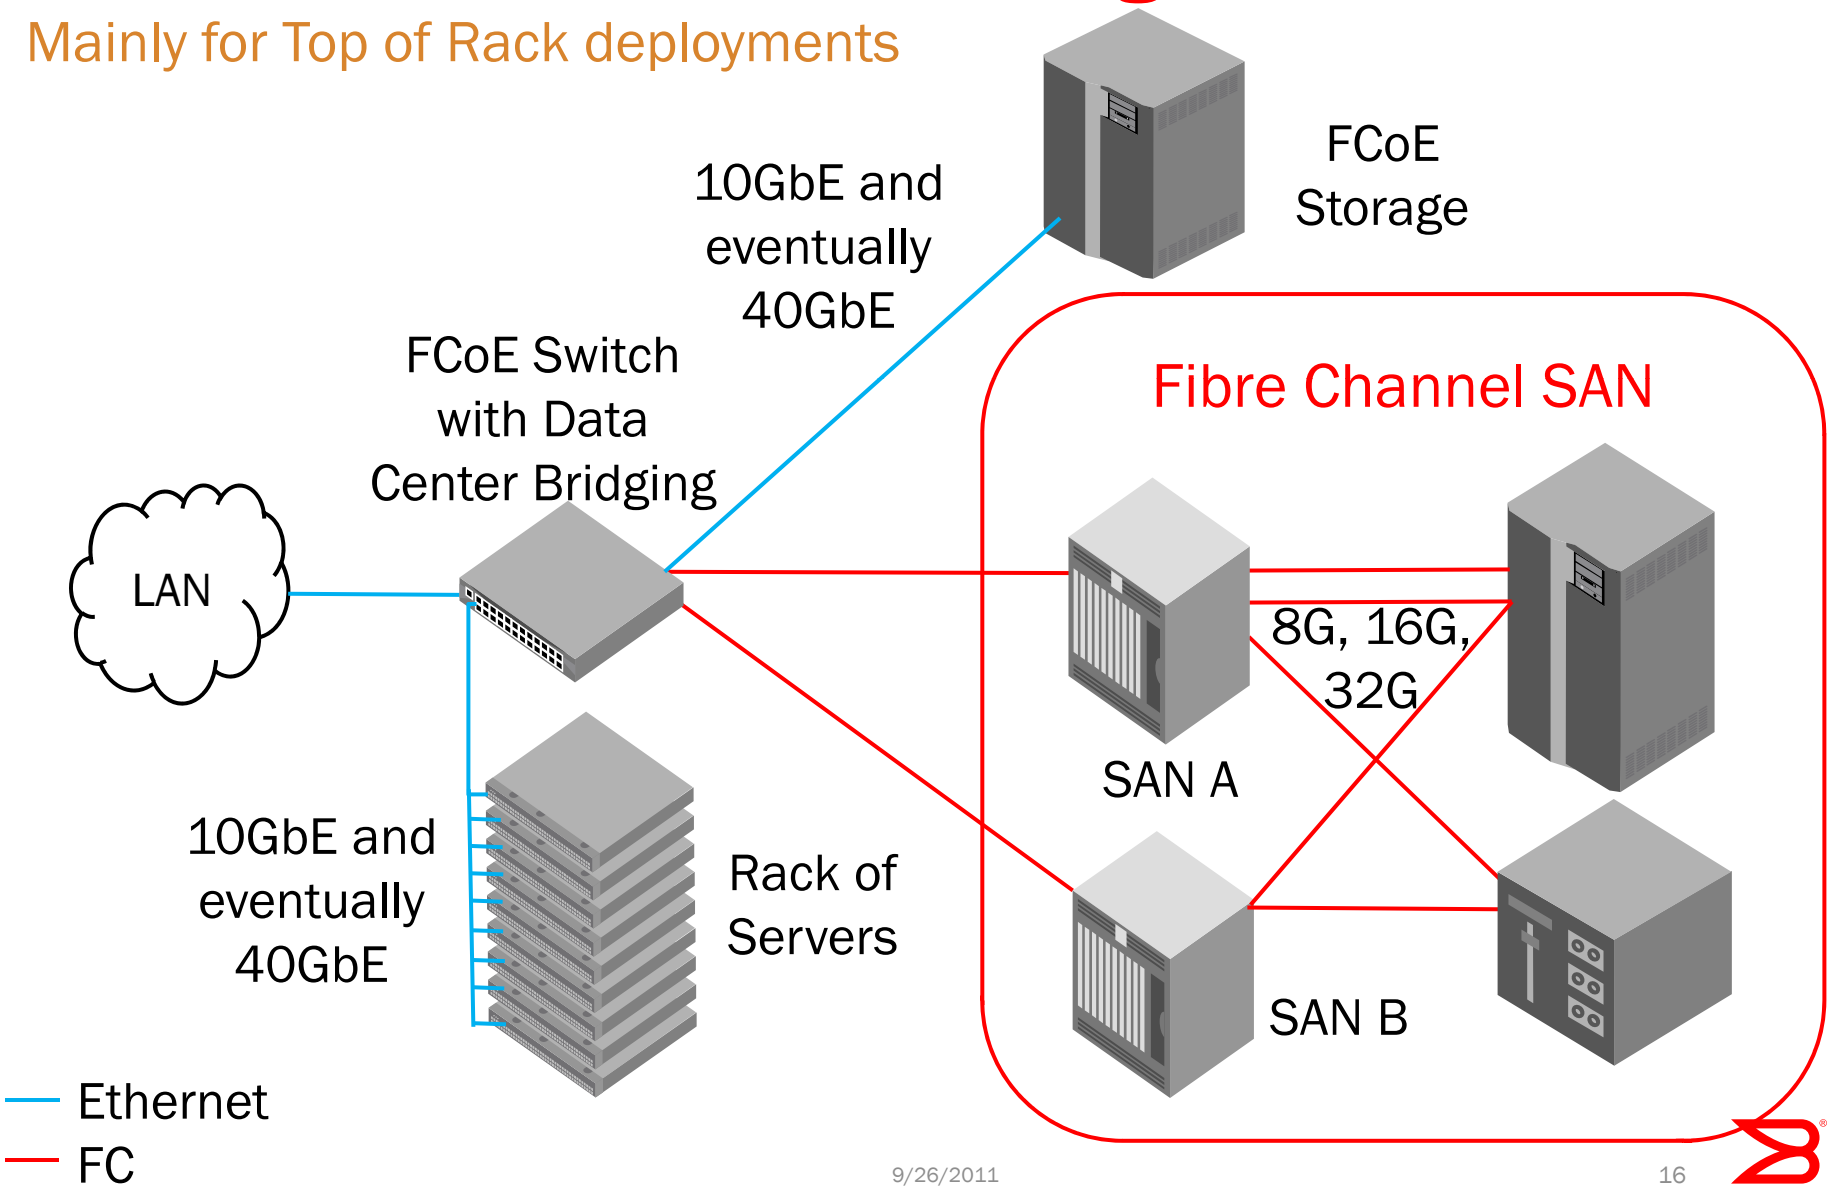

Trends in Storage

- Application migration benefits from networked storage compared to DAS
- Cloud computing requires major data moves
- Virtual Desktop Infrastructure (VDI) leads to centralized storage and increased network traffic
- Solid State Drives (SSDs) or Flash Storage leads to higher bandwidth demands on the network

Server Virtualization and Migration

Data needs to move with the application

For application A to move to a new server with DAS, the data has to be moved over the LAN

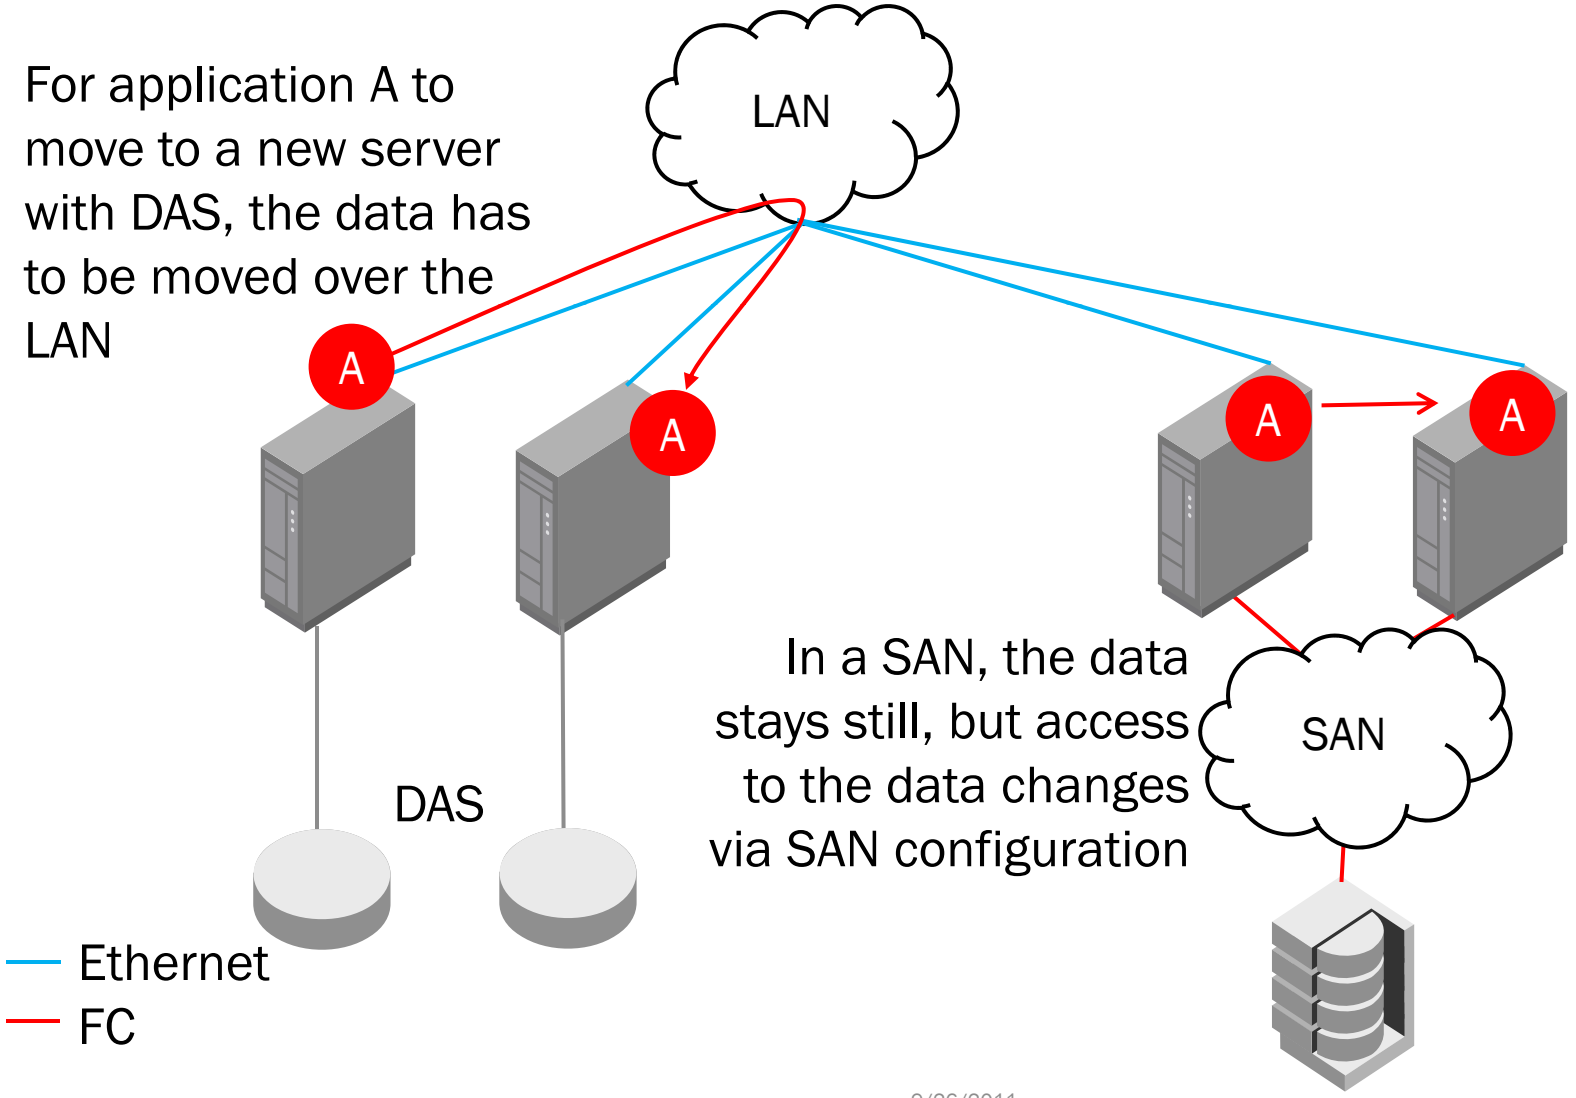

Data in Cloud Computing in 2015

Over 10% of storage could be in the clouds!

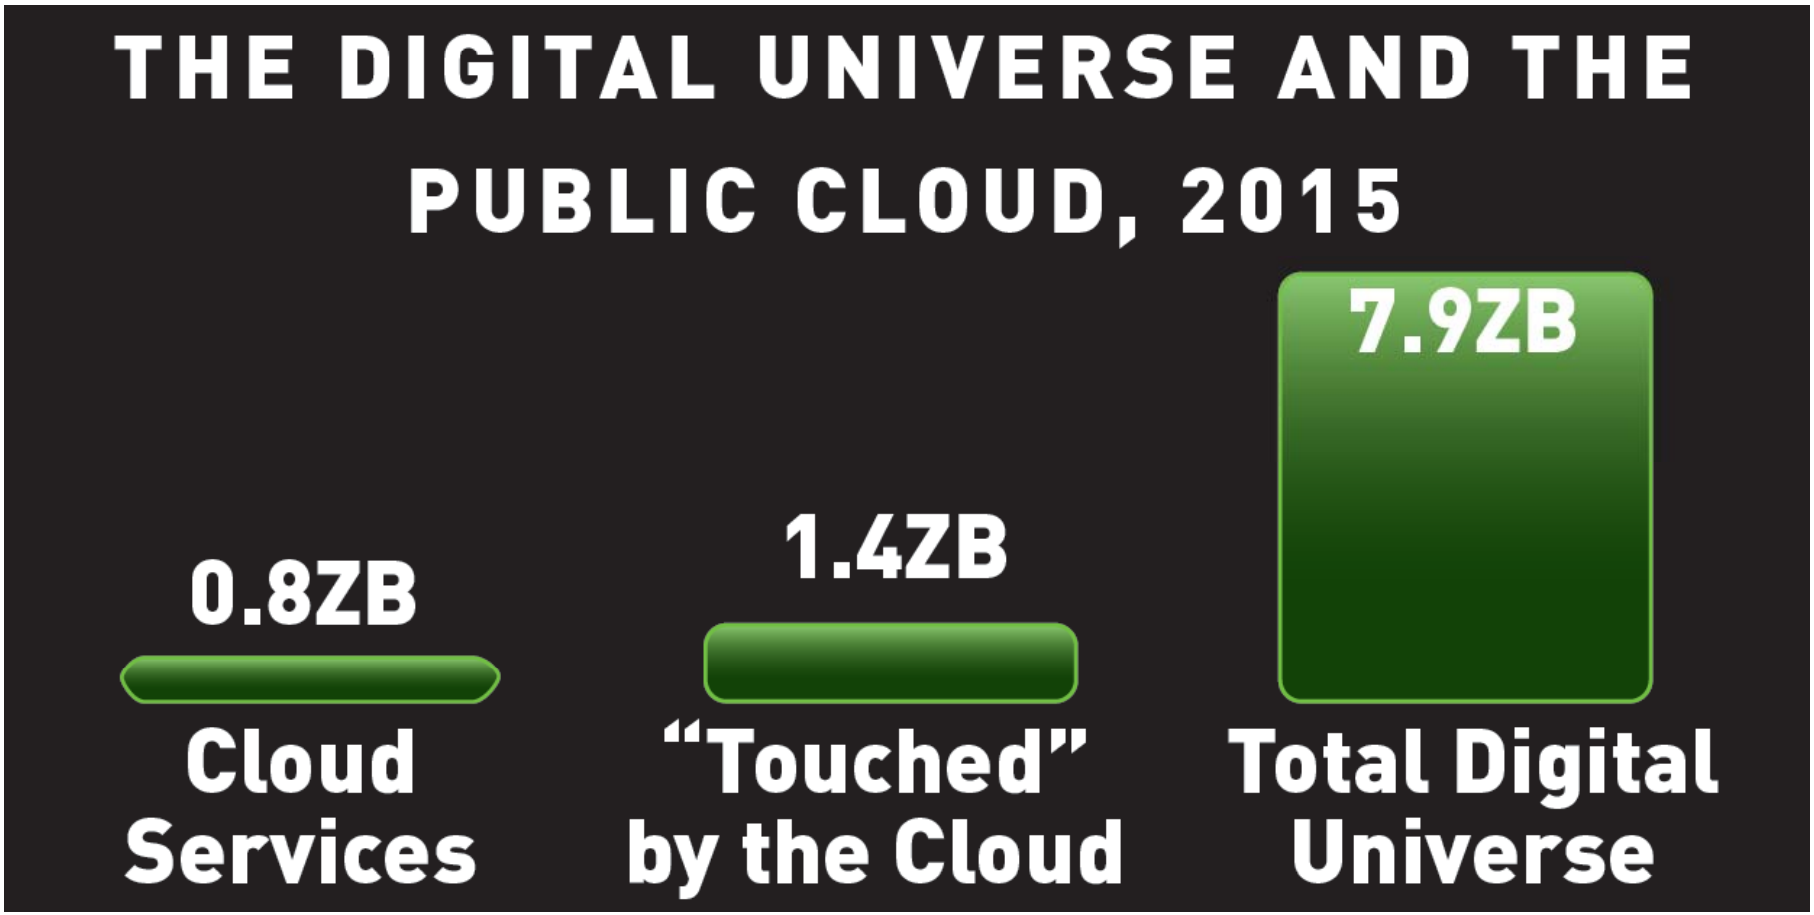

Source: The Digital Universe Study: http://www.emc.com/leadership/programs/digital-universe.htm

Cloud Computing

- Cloud Computing offers the grand vision of hosting and scaling applications from your data center to the cloud provider or another data center on demand
- To enable this transition, the data needs to be exchanged or mirrored first

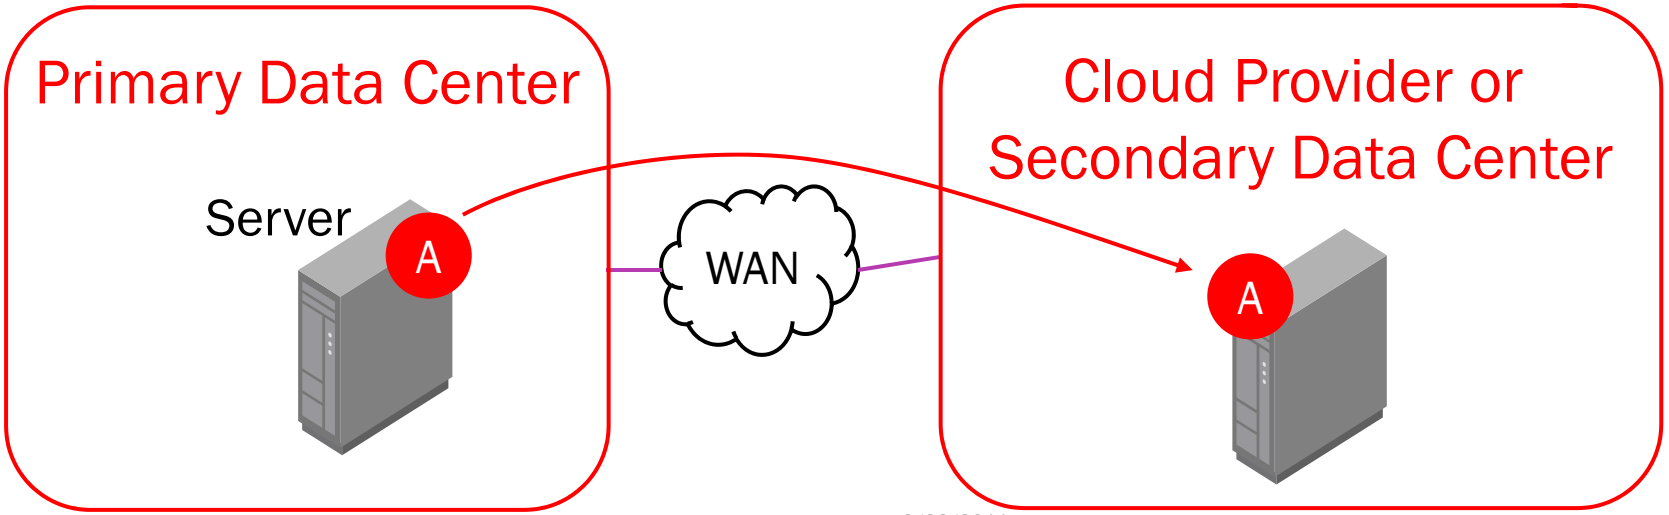

Data Mirroring Between Storage Arrays

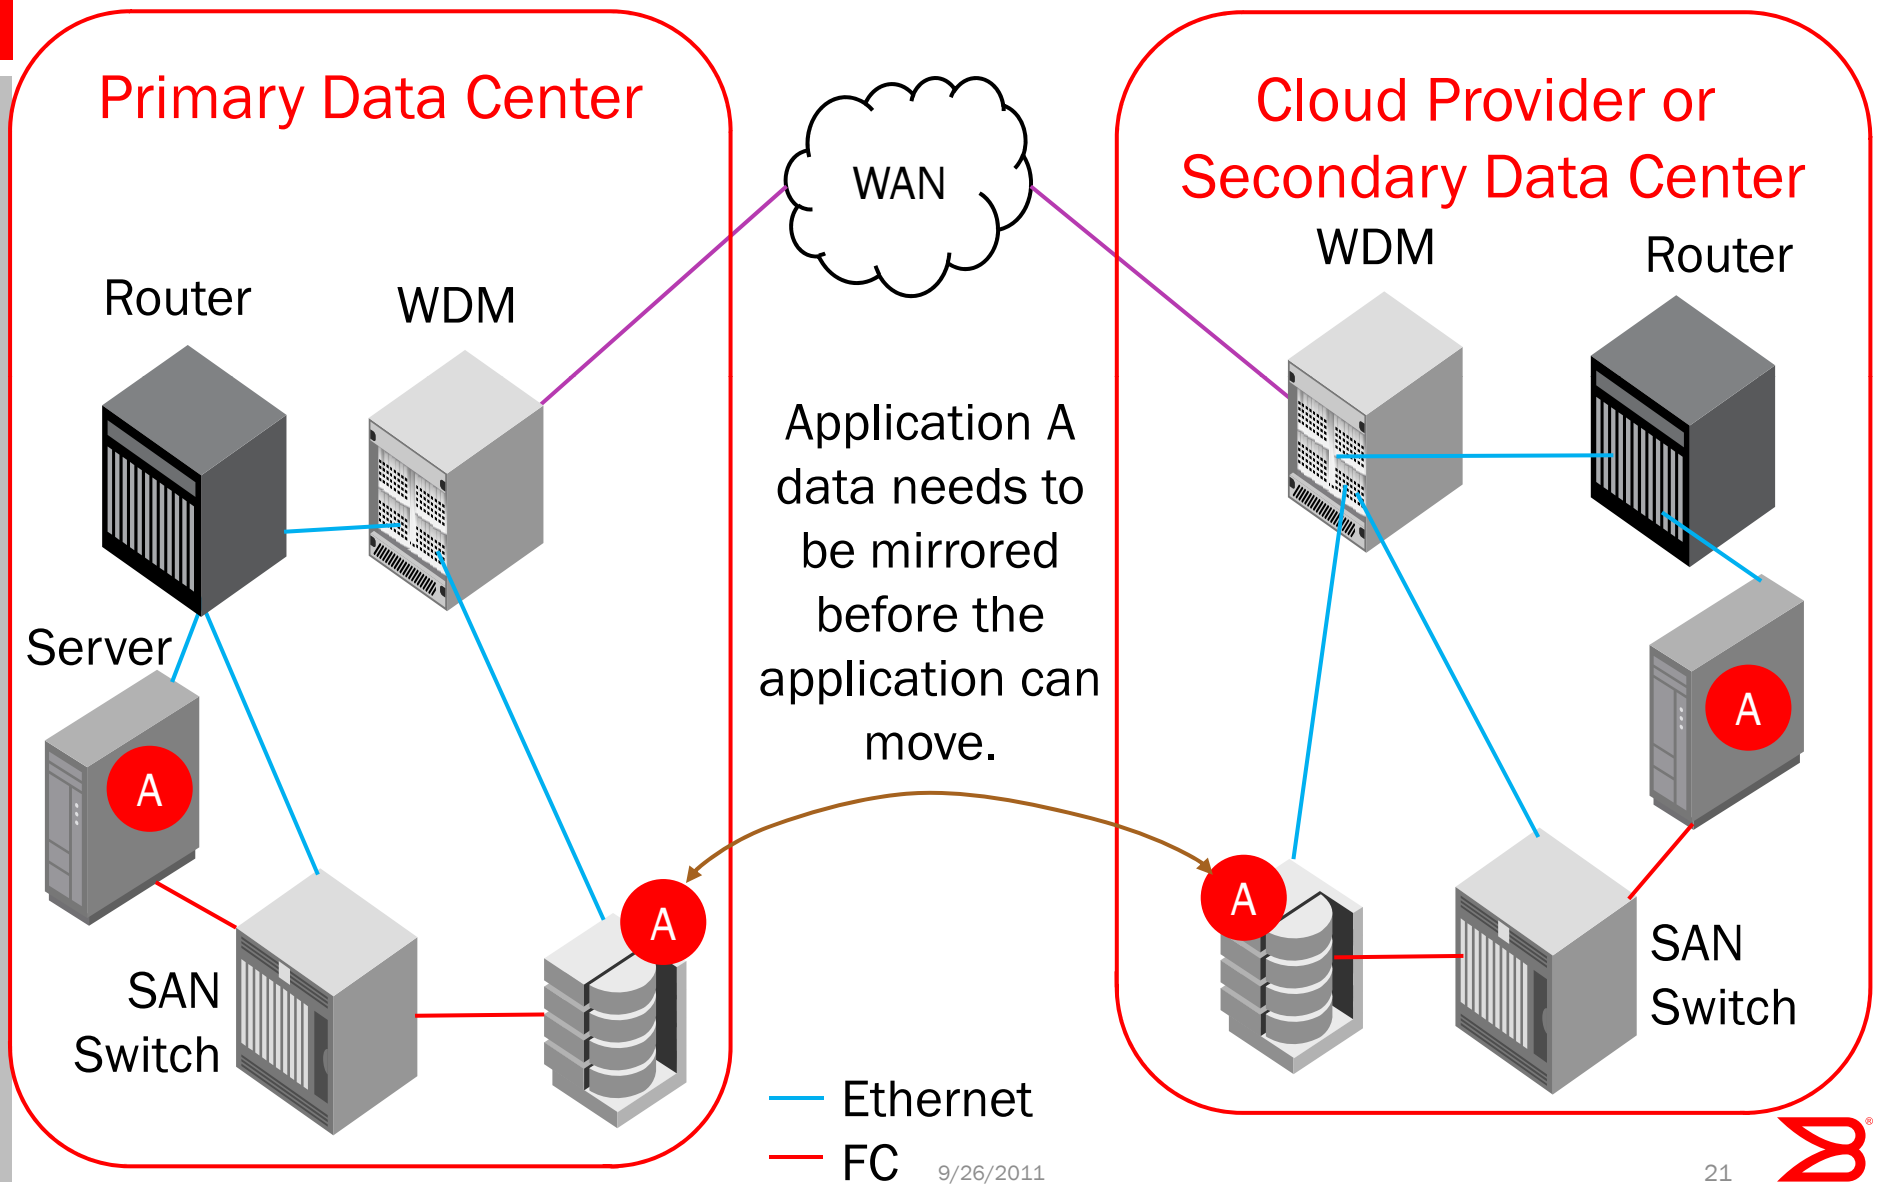

3 Main Ways to Mirror Data over the WAN

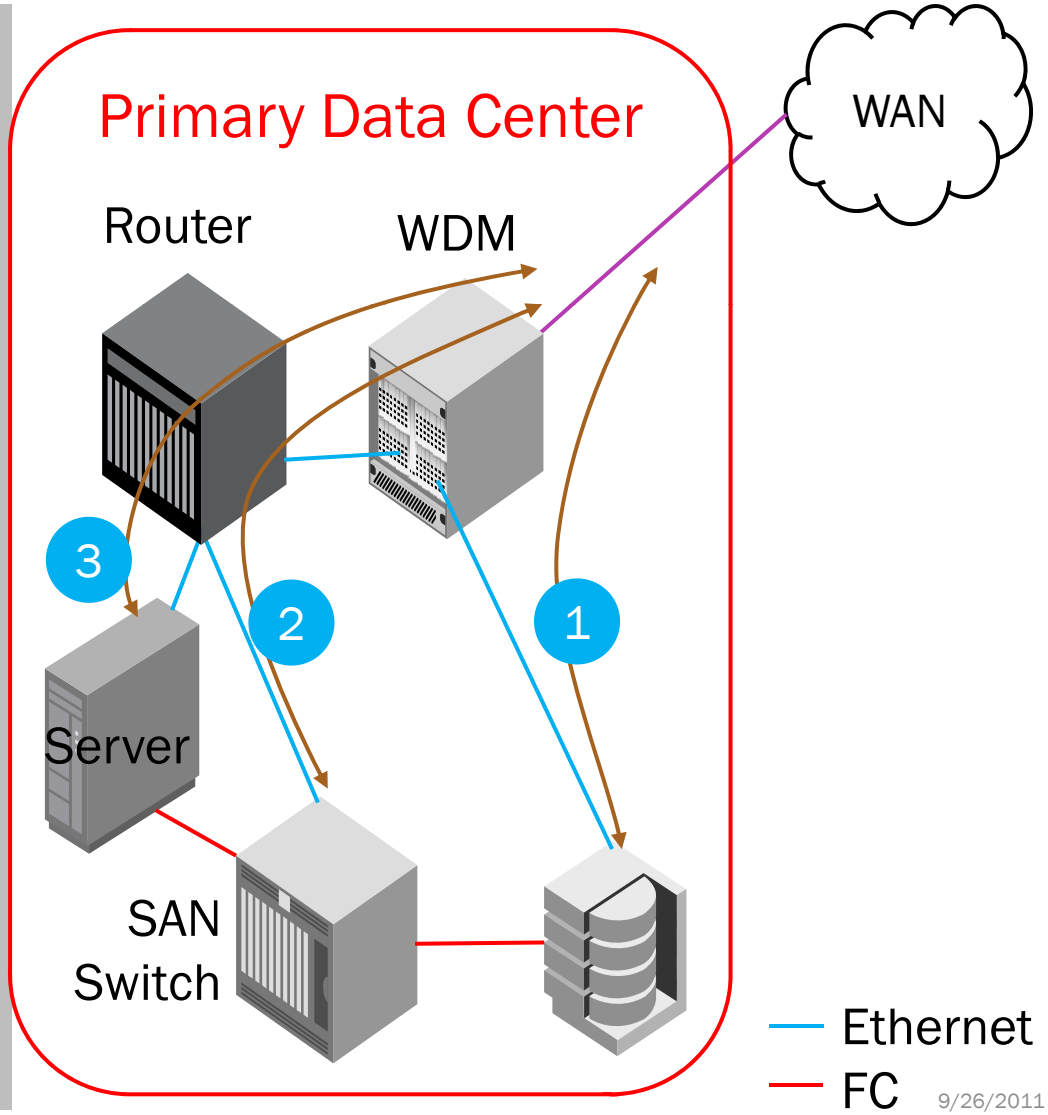

1. Native Fibre Channel, Fibre Channel Internet Protocol (FCIP) or Ethernet Over WDM
2. SAN Switch to SAN Switch via IP and FCIP
3. Server to Server that backs the data up to storage

Virtual Desktop Infrastructure (VDI) Architecture

- VDI enables centralized management and simple upgrades to software and applications and increases LAN traffic

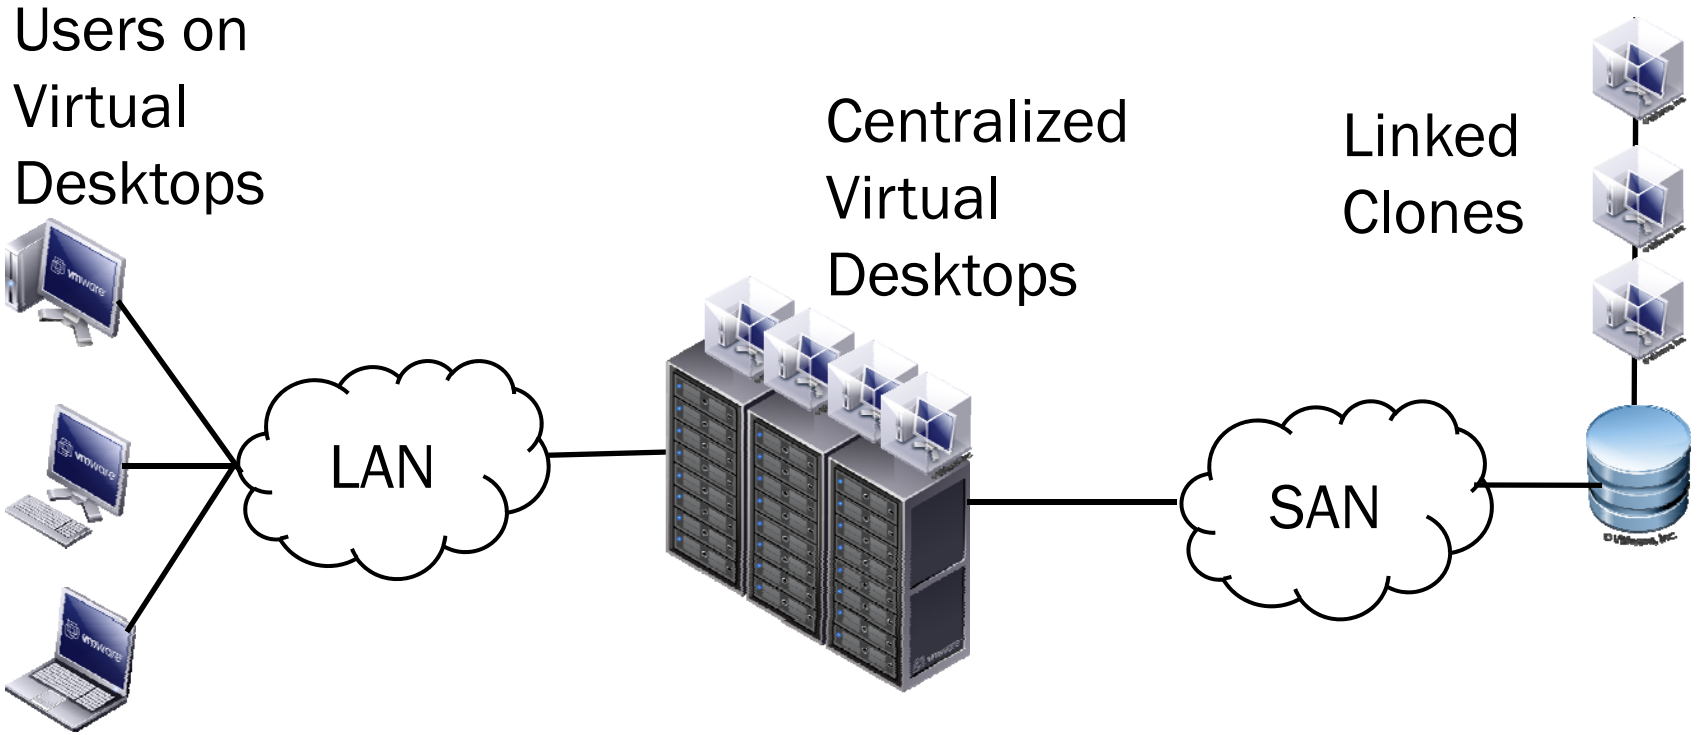

Comparing Server Technologies

	2000	2005	2010	2015
CPU	**1 x** Pentium 4 1.5 GHz	**5 x** Pentium D 2.6 GHz	**15 x** Nehalem Quad 2.6 GHz	**45 x?** Haswell 2.6 GHz?
DRAM	**1 x** DDR1	**4 x** DDR2	**8 x** DDR3	**32 x?** DDR4?
Network	**1 x** 100Mb Ethernet	**10 x** Gigabit Ethernet	**100 x** 10 Gigabit Ethernet	**400 x** 40 Gigabit Ethernet
Bus	**1 x** PCI 32-bit/33 MHz	**15 x** PCIe Gen1 x8	**30 x** PCIe Gen2 x8	**60 x** PCIe Gen3 x8
Fibre Channel	**1 x** 1GFC	**4 x** 4GFC	**8 x** 8GFC	**32 x** 32GFC
Disk	**1 x** 15K rpm hard drive	**1 x** 15K rpm hard drive	**1 x** 15K rpm hard drive	**1 x** 15K rpm hard drive

Source: Nimbus and Brocade

SSDs – Solid State Drives

- Application performance is limited by multiple factors with disk drive latency being one factor
- Order of magnitude improvements in performance
  - While traditional spinning disk drive seek times are in the millisecond range, SSD seek times are in the microsecond range
  - SSDs often referred to as Tier-0 storage while disk drives are Tier-1
  - Capacities in the hundreds of GBs per drive
  - Very energy efficient compared to spinning disks
  - Most SSDs provide over 50,000 IOPs per drive
- One flash storage system supports 500,000 IOPS and 8 GBps (64 Gbps) of throughput

	Latency	Drive IOPs	Array IOPS
HDD	2-10 mS	100-300	400-40,000
SSD	50-250 uS*	40k-150k	50k-500k

* This is based on Flash memory and multiple parallel processing

Conclusion

- We entered the Zettabyte era last year– 1M TB/year of new data
- More data is created every two years than all previous years combined
- Virtualization causes the need for networked storage of all varieties (SAN and NAS)
- All storage technologies are improving except disk drive access times and disk rotational speeds
- New applications and devices are driving more data access and higher bandwidths

BROCADE

Thank You