#SHAREorg

SHARE

Buzz Fibrechannel To 16G and Beyond!

Howard Johnson – howard.johnson@brocade.com

Monday, March 10, 2014: 9:30 AM-10:30 AM
Platinum Ballroom Salon 4 (Anaheim Marriott Hotel)

Session 14382

QR Code

SHARE in Anaheim

Abstract

- In this jointly presented session, the major players in storage networking will discuss:
  - FICON speed roadmap - per the standards.
  - Current customer trends in bandwidth utilization.
  - Do you need 8G, do you need 16G, do you need 32G ?
  - How does zHPF play into FICON speeds ?
  - What about FCoE - how does this play into FICON ?
- At the end, there will be time for Q&A.

Agenda

- **Trends and Drivers**
  - Bandwidth Drivers
  - Fibre Channel Speed Evolution
- **FICON Influences**
  - Channel Speed Evolution
  - zHPF
- **16G and Beyond**
  - Fibre Channel Roadmap
  - FCoE
- **Let's Talk about Light**
  - Modal Dispersion
  - Light in Flight
  - Measuring Light Signals

Bandwidth Drivers

Fibre Channel Speed Evolution

TRENDS AND ROADMAPS

What is driving bandwidth demand?

- Applications increasing in scale and number
- Server virtualization
- Multi-core processors
- Large Memory increases
- Solid State Disks (SSD)
- Faster PCIe rates

The Internet Minute demands greater bandwidth and faster deployment from telecommunication manufacturers, operators and service providers. (Courtesy of Intel)

Prolific Applications Server Virtualization

- Applications keep growing in number and breadth
- Multiple servers need to access data from shared storage
- Database applications drive bandwidth
- Server virtualization creates multiple virtual machines (VMs) for each application, so each physical server is producing more Input/Output (I/O)

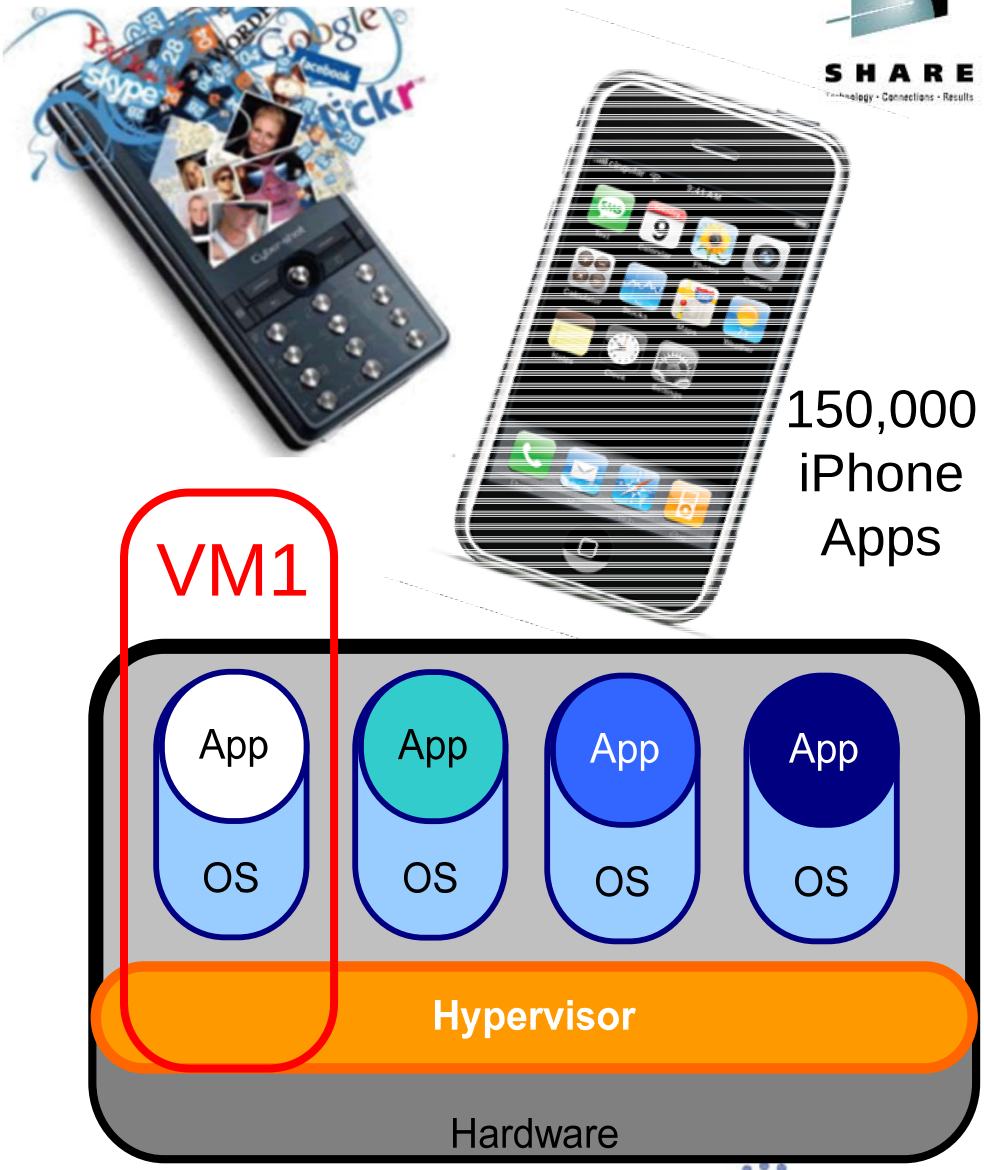

Faster Processors

- IBM has the Power7 that has 8 cores and supports 50 GBps of peak IO and directly interconnects 32 of these processors on a server
- NehalemEX has 8 cores and 16 threads and uses Intel QuickPath Interconnect at 6.4 Gigatransfers per second delivering up to 25 GigaBytes/second (GBps)
- AMD has 8-core and 16 core processors that support 32 threads and HyperTransport 3.0 to support 4.8 gigaTransfers/second
- Sun's UltrasparcT3 chip has 16 cores and supports up to 128 threads
- A single, multi-processor server supports 10s or 100s of cores

Nehalem-EX in two-chip configuration

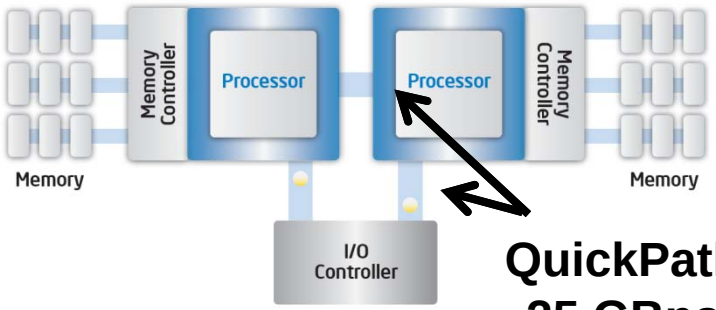

IBM's Power7

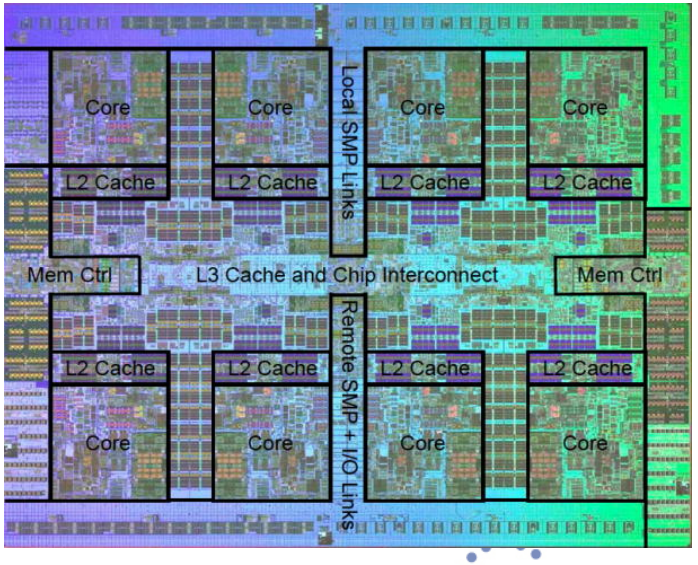

Increased Memory in Servers

- Memory has limited virtual servers in the past
- Server performance and number of VMs is dependent on memory capacity in servers
  - Gartner: Midrange servers averaged 32GB of memory in 2009 and were expected to triple to 96GB in 2012
  - Registered Dual-Inline Memory Modules (LRDIMM) already come in 32GB packaging
  - Dell's 2U PowerEdge R710 supports 144GB of memory
  - Sun SPARC M9000-64 offers 4TB memory capacity
  - VMWARE supports 1TB/server and 255GB/VM
- Memory drives more applications that drive more storage I/O traffic

32GB RDIMM

SSDs – Solid State Drives

- Performance of applications 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 I/Os per second per drive
- Texas Memory Systems RamSan-630 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

PCIe Continues Ramping

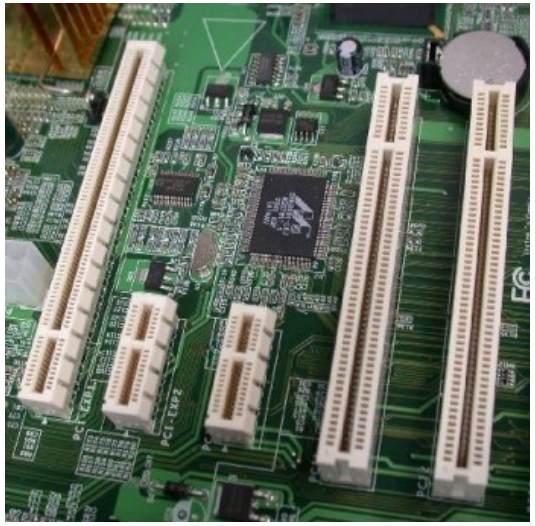

- PCIe 2.0 increases in speed to support dual ported 16G FC HBAs
- PCIe 3.0 will support quad ported 16G FC HBAs
- But they use multiple lanes (wire links) to do it

	Number of Lanes	Speed per Lane (MBps)	Directional Bandwidth (Gbps)	Ports Supported
PCIe -1.0	4	250	8	1 – 8GFC
PCIe -1.0	8	250	16	2 – 8GFC
PCIe -2.0	4	500	16	1 – 16GFC
PCIe -2.0	8	500	32	2 – 16GFC
PCIe -3.0	4	1000	32	2 – 16GFC
PCIe -3.0	8	1000	64	4 – 16GFC

More Applications Drive more Bandwidth

- 16G FC was designed for servers over the next few years that will use these technologies

More Applications ⟵ **More Bandwidth** ⟶ **More Data**

More Virtual Machines

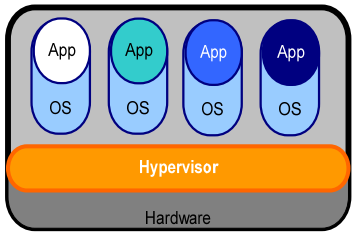

Multi-core Processors

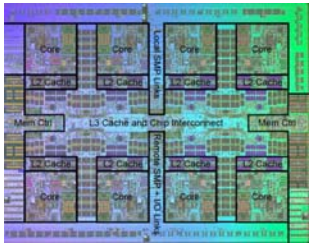

More Memory

Faster PCIe

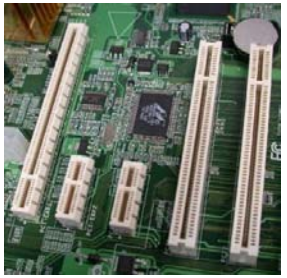

Faster SSDs

The Evolution of Fibre Channel Speeds

- Five generations of Fibre Channel have been delivered to the market
- Speed doubling about every 3-years
- Fibre Channel dominates the storage market

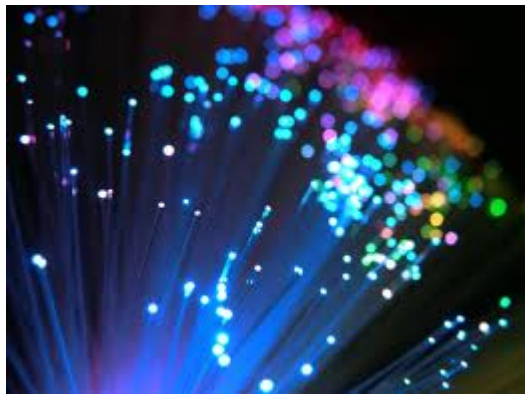

Generations of Fibre Channel

The newest speed in Fibre Channel - Keep it Serial Stupid

Generation	1st Gen	2nd Gen	3rd Gen	4th Gen	5th Gen	6th Gen
Electrical / Optical Module	1GFC / GBIC/ SFP	2GFC / SFP	4GFC / SFP	8GFC / SFP+	16GFC / SFP+	32GFC / SFP+
Electrical Speeds(Gbps)	1 lane at 1.0625	1 lane at 2.125	1 lane at 4.25	1 lane at 8.5	1 lane at 14.025	1 lane at 28.05
Encoding	8b/10b	8b/10b	8b/10b	8b/10b	64b/66b	64b/66b
Availability	1997	2001	2006	2008	2011	2014

GBIC

← SFP / SFP+ →

FC Dominates the Backbone Storage Network

What is the predominant storage network backbone you use?

- By asking for the predominant storage network, we find where the heart of storage networking lies. With 82% selecting FC, the answer is clear.
- From our latest technology roadmap, non-FC storage network technologies are used in greater percentages than appear in this chart. FCoE is *in use* by 8% of respondents, and 10Gbps Ethernet (used for storage) is *in use* by 31%.

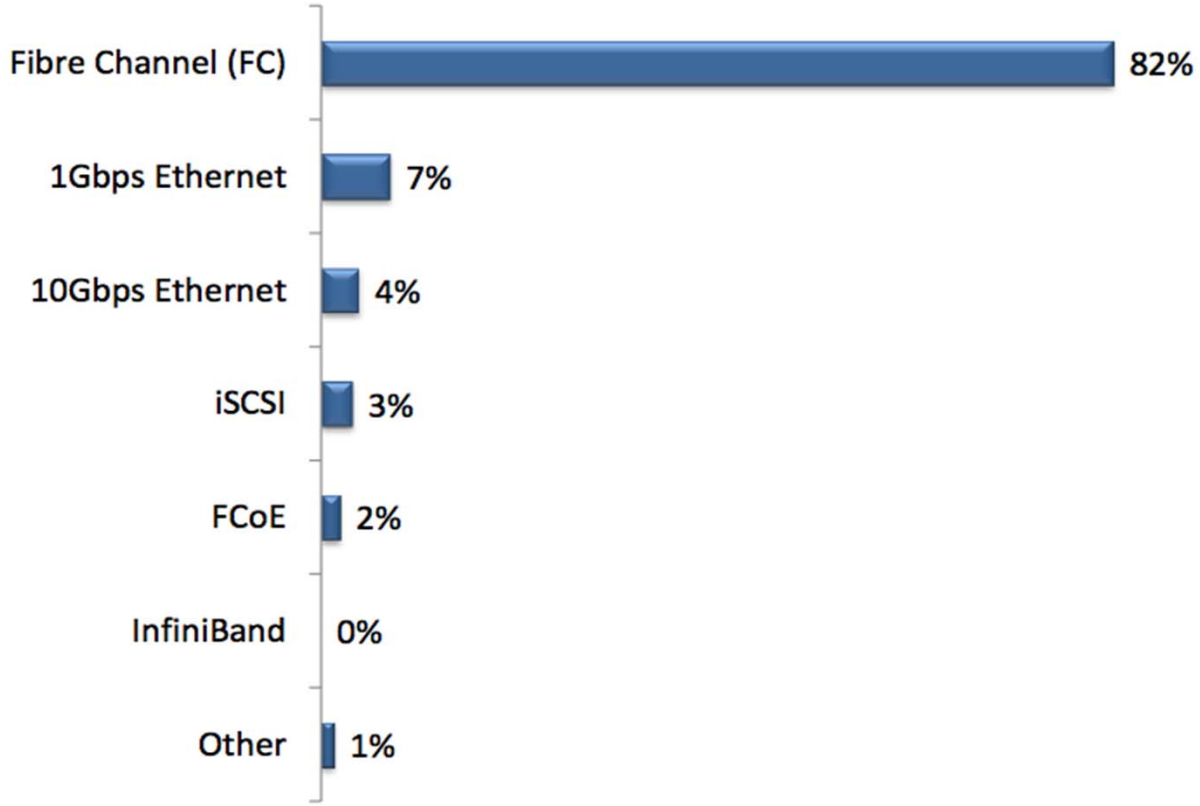

Source: The InfoPro, 1H CY 2012

Channel Speed Evolution
zHPF

FICON INFLUENCES

Mainframe and Open Systems Time Lines

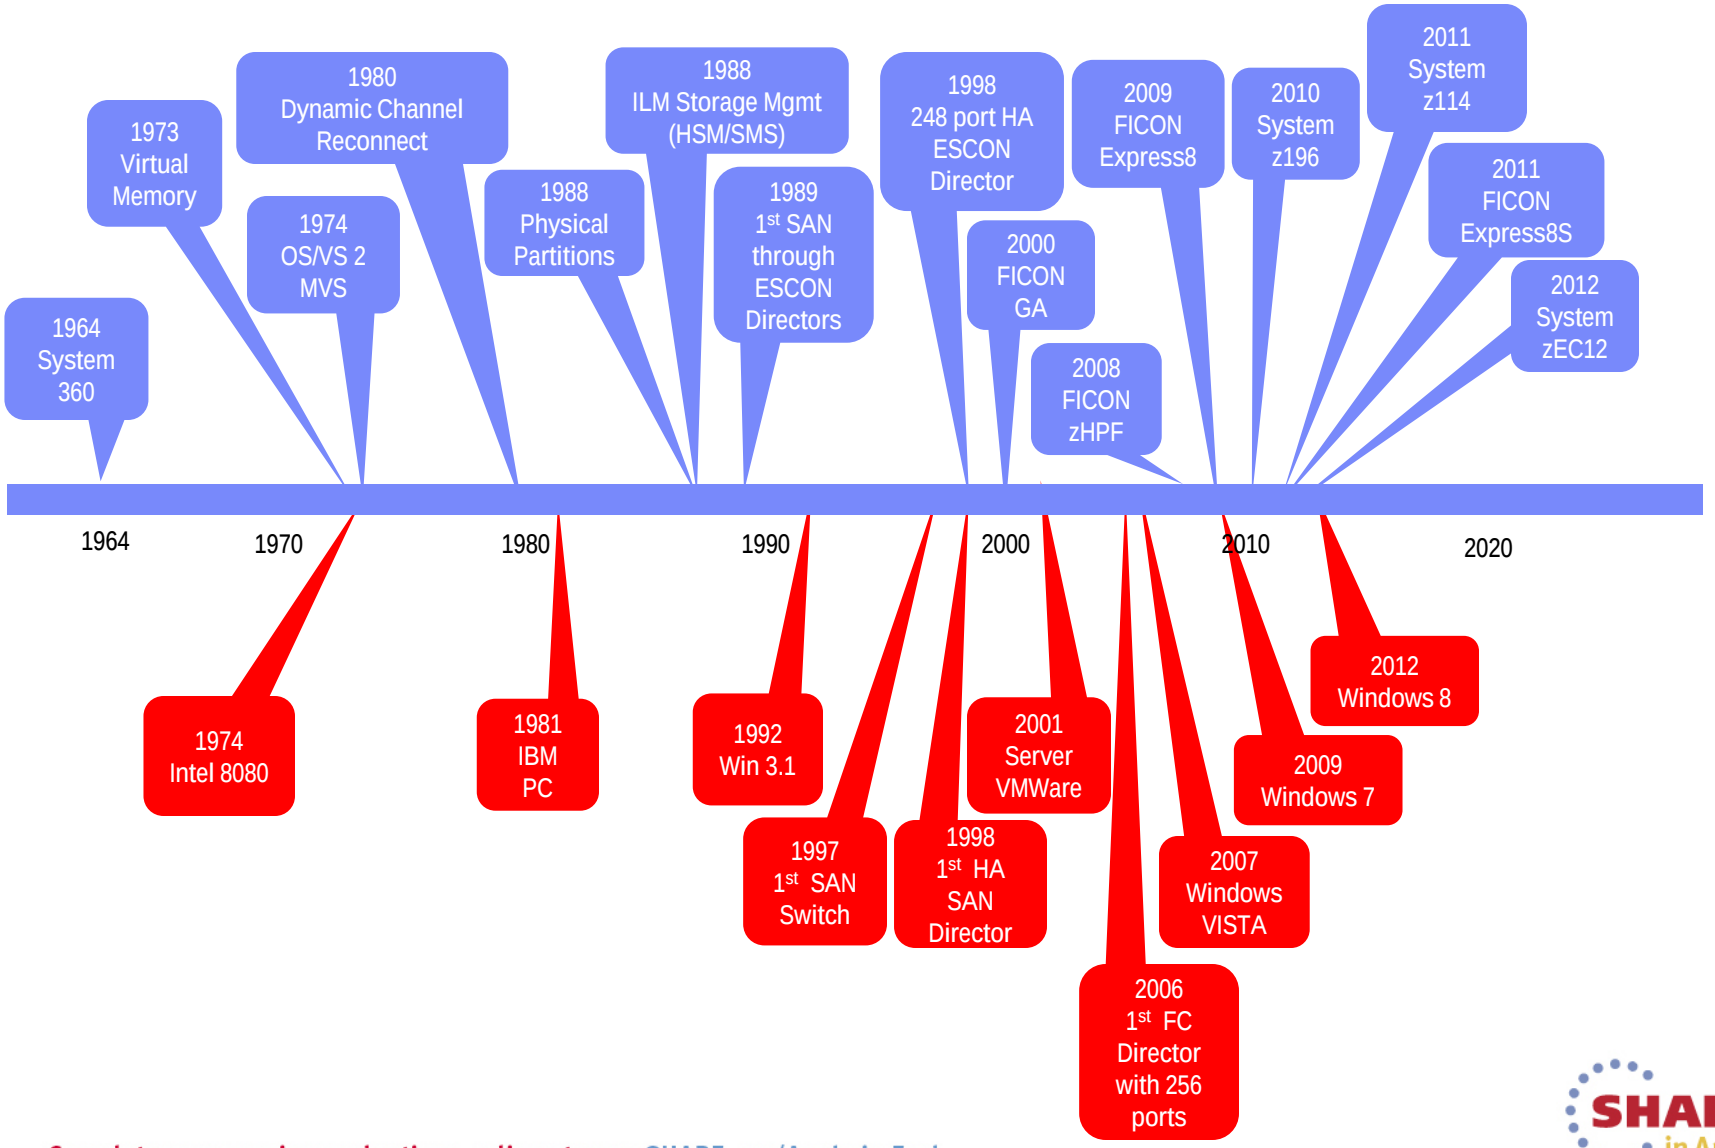

FICON performance – Start I/Os Historical Actuals

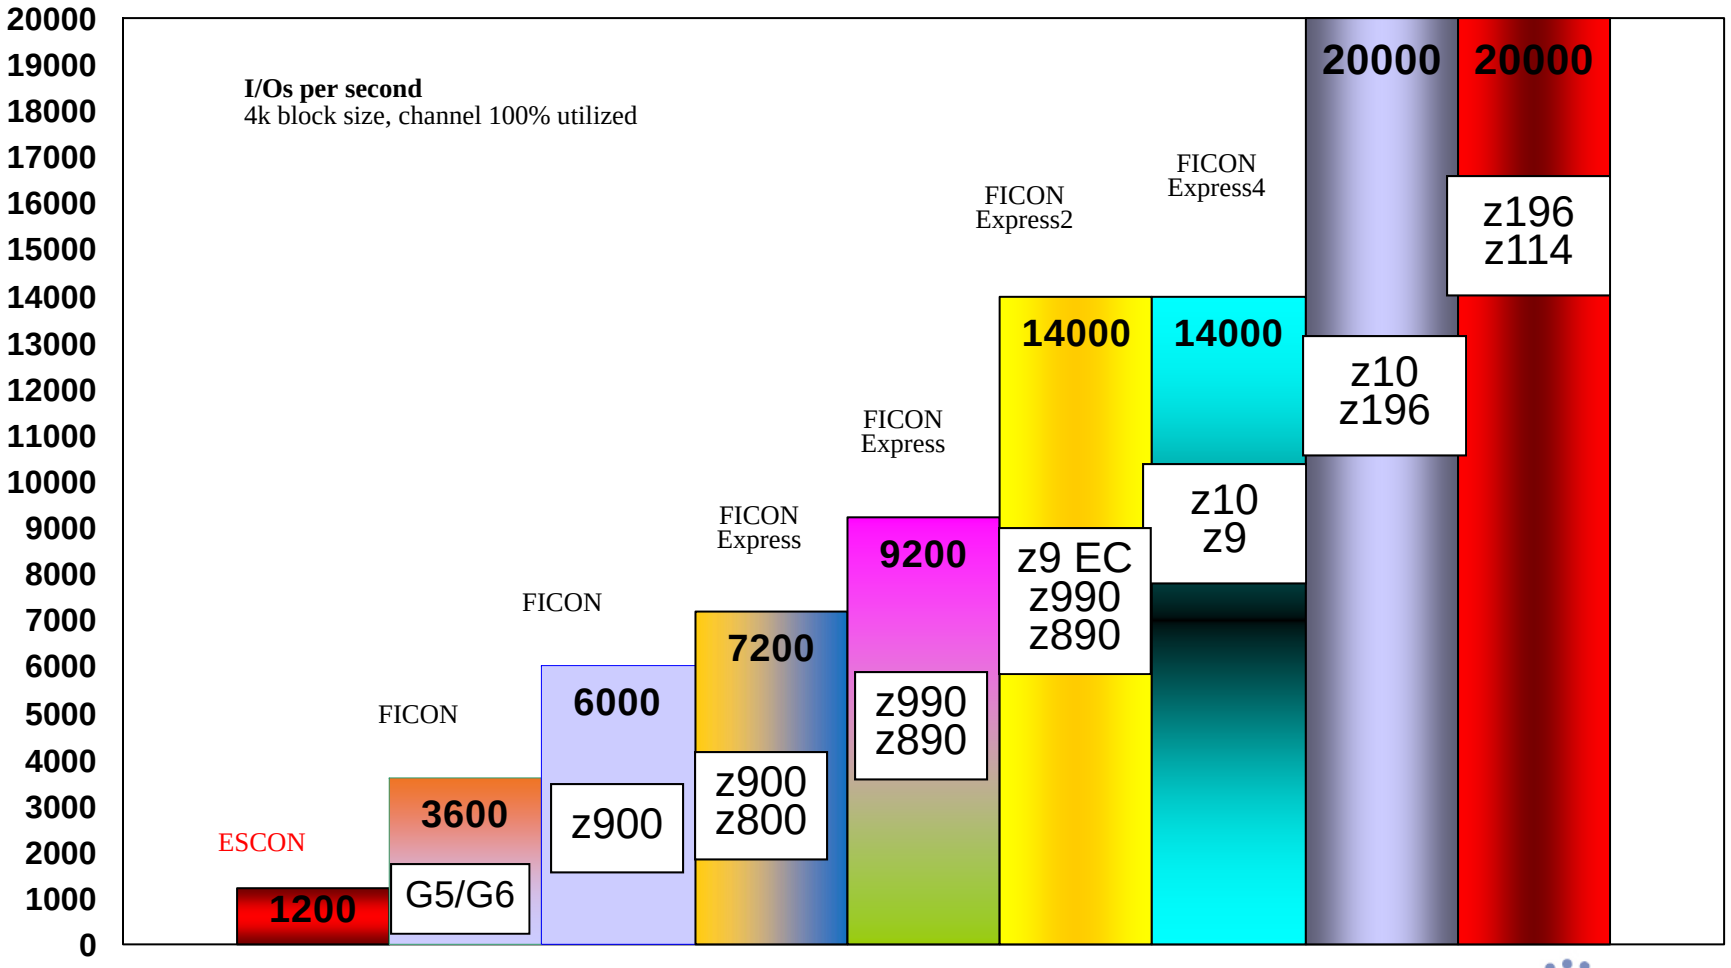

FICON performance on System z

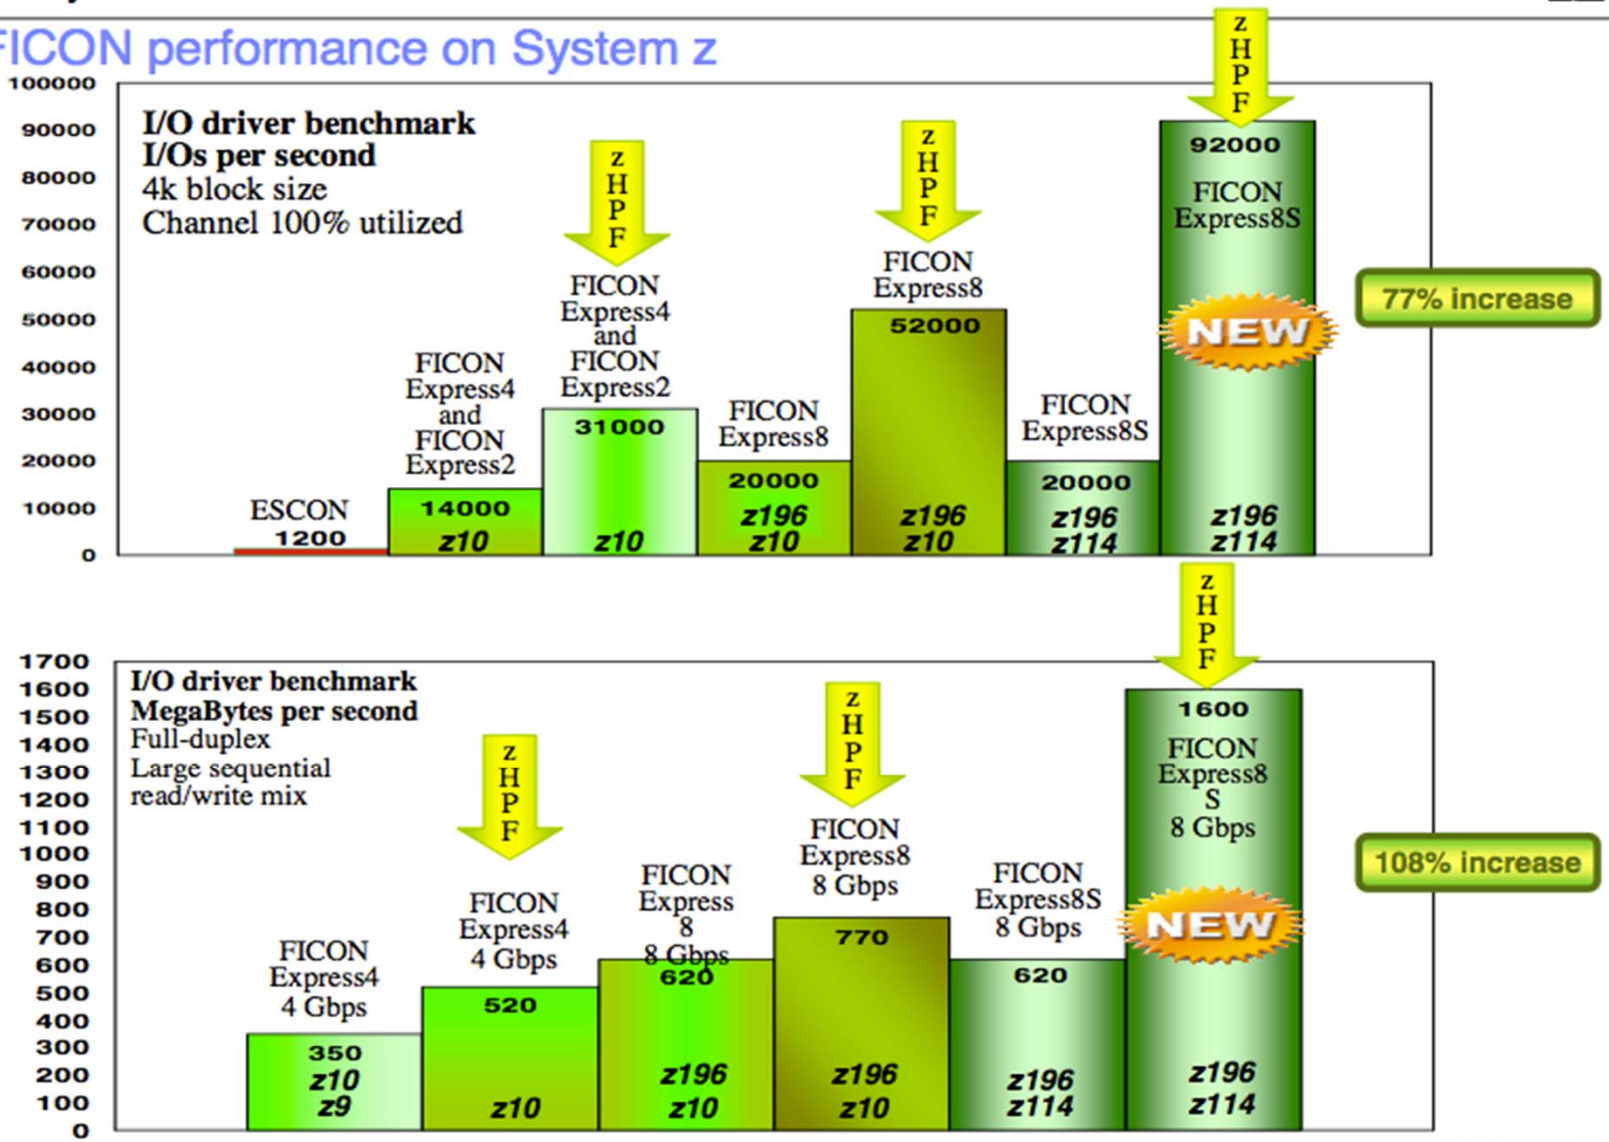

FCP performance on System z

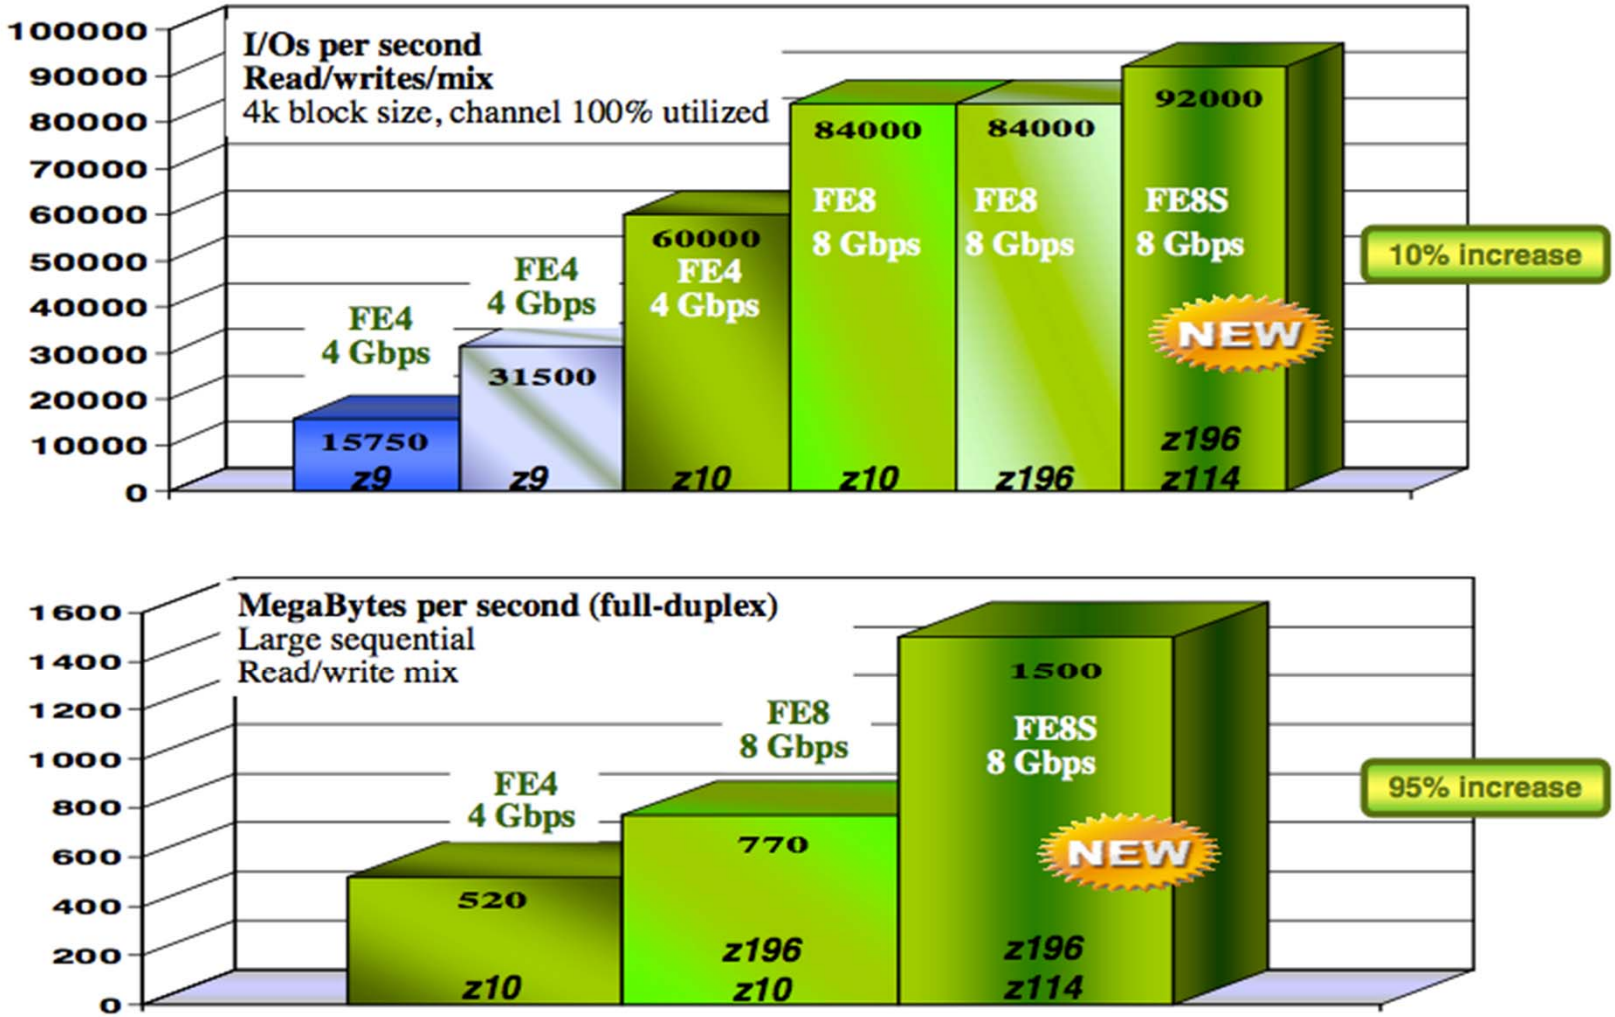

What's Ahead?

Fibre Channel Roadmap
FCoE

16G AND BEYOND

Storage Networking: Technology Roadmap

What is your status of implementation for this technology?

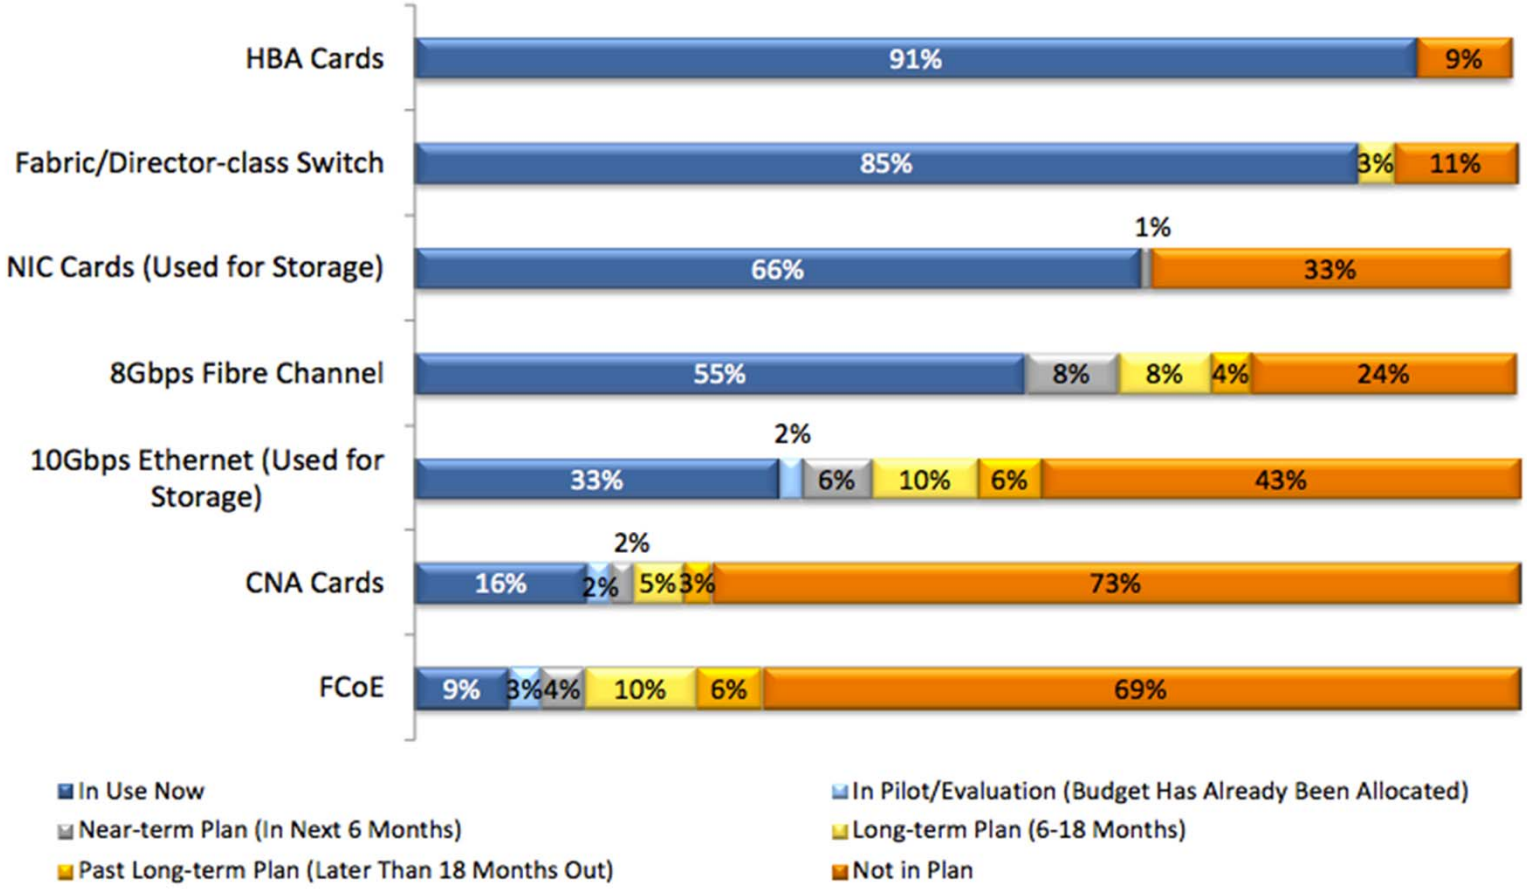

Source:
The InfoPro,
1H CY 2012

The Benefits of 16GFC

- 16GFC is 100% faster than 8GFC and 40% faster than 10GE and leads to these benefits
  - Higher performance lets servers process more data
  - Fewer links to do the same job
  - Easier cable and device management
  - Less power consumption per bit

Innovation

Performance

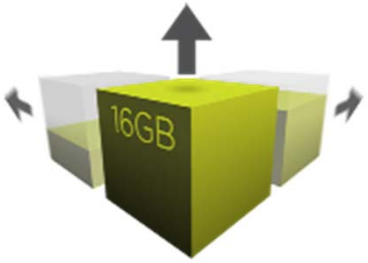

Scalability

Environment Friendly

Characteristics of 16GFC

- Double the throughput over backplanes, 100 meters and 10 kilometers

Fibre Channel Physical Interfaces 5 (FC-PI-5) standardized 16GFC

Speed Name	Throughput (MB/sec)	Line Rate (Gbps)	Encoding	Retimers in the module	Transmitter Training	OM1/2/3/4 Link Distance (meters)
1GFC	100	1.0625	8b/10b	No	No	300/500/860/*
2GFC	200	2.125	8b/10b	No	No	150/300/500/*
4GFC	400	4.25	8b/10b	No	No	50/150/380/400
8GFC	800	8.5	8b/10b	No	No	21/50/150/190
10GFC	1200	10.53	64b/66b	Yes	No	33/82/300/*
16GFC	1600	14.025	64b/66b	Yes	Yes	15/35/100/125

* FC-PI-5 didn't standardize distances for OM4 fiber for 1/2/10GFC

The Benefits of 32GFC
Keep It Serial Stupid

- Highlights of FCIA 32GFC MRD for T11 standard:
  - Stay serial and single-lane; 2 * 16Gb FC = 28.05Gbaud
  - 32GFC = 1 * 32G lane (serial)
  - T11.2 FC-PI-6: Was stable by August ‘11, and released Feb ’12
  - 70 to 100 meters on OM3 optics, 7 meters on copper
  - <=50% Watts/Port compared to 40GE and <=50% $/port compared to 40GE
- 2014 products:
  - Leverage work from multitude of technologies
    - “Perfect Storm” flocking towards 25Gbaud range
    - Expect feasibility for FC core markets around 2014-2015
      - *Ethernet 100G mandates a 25G/lane technology 2015*
        - *100GE = 4 * 25G lanes*
      - *IB will have 25G per lane option in 2014/2015*
- Any 8GFC a user buys today will work with a 32GFC infrastructure
  - Now THAT is a Safe Investment!

Speeds Double Through 2020

- 32G FC Standard Stabilizing
- 64G FC Standard Starting

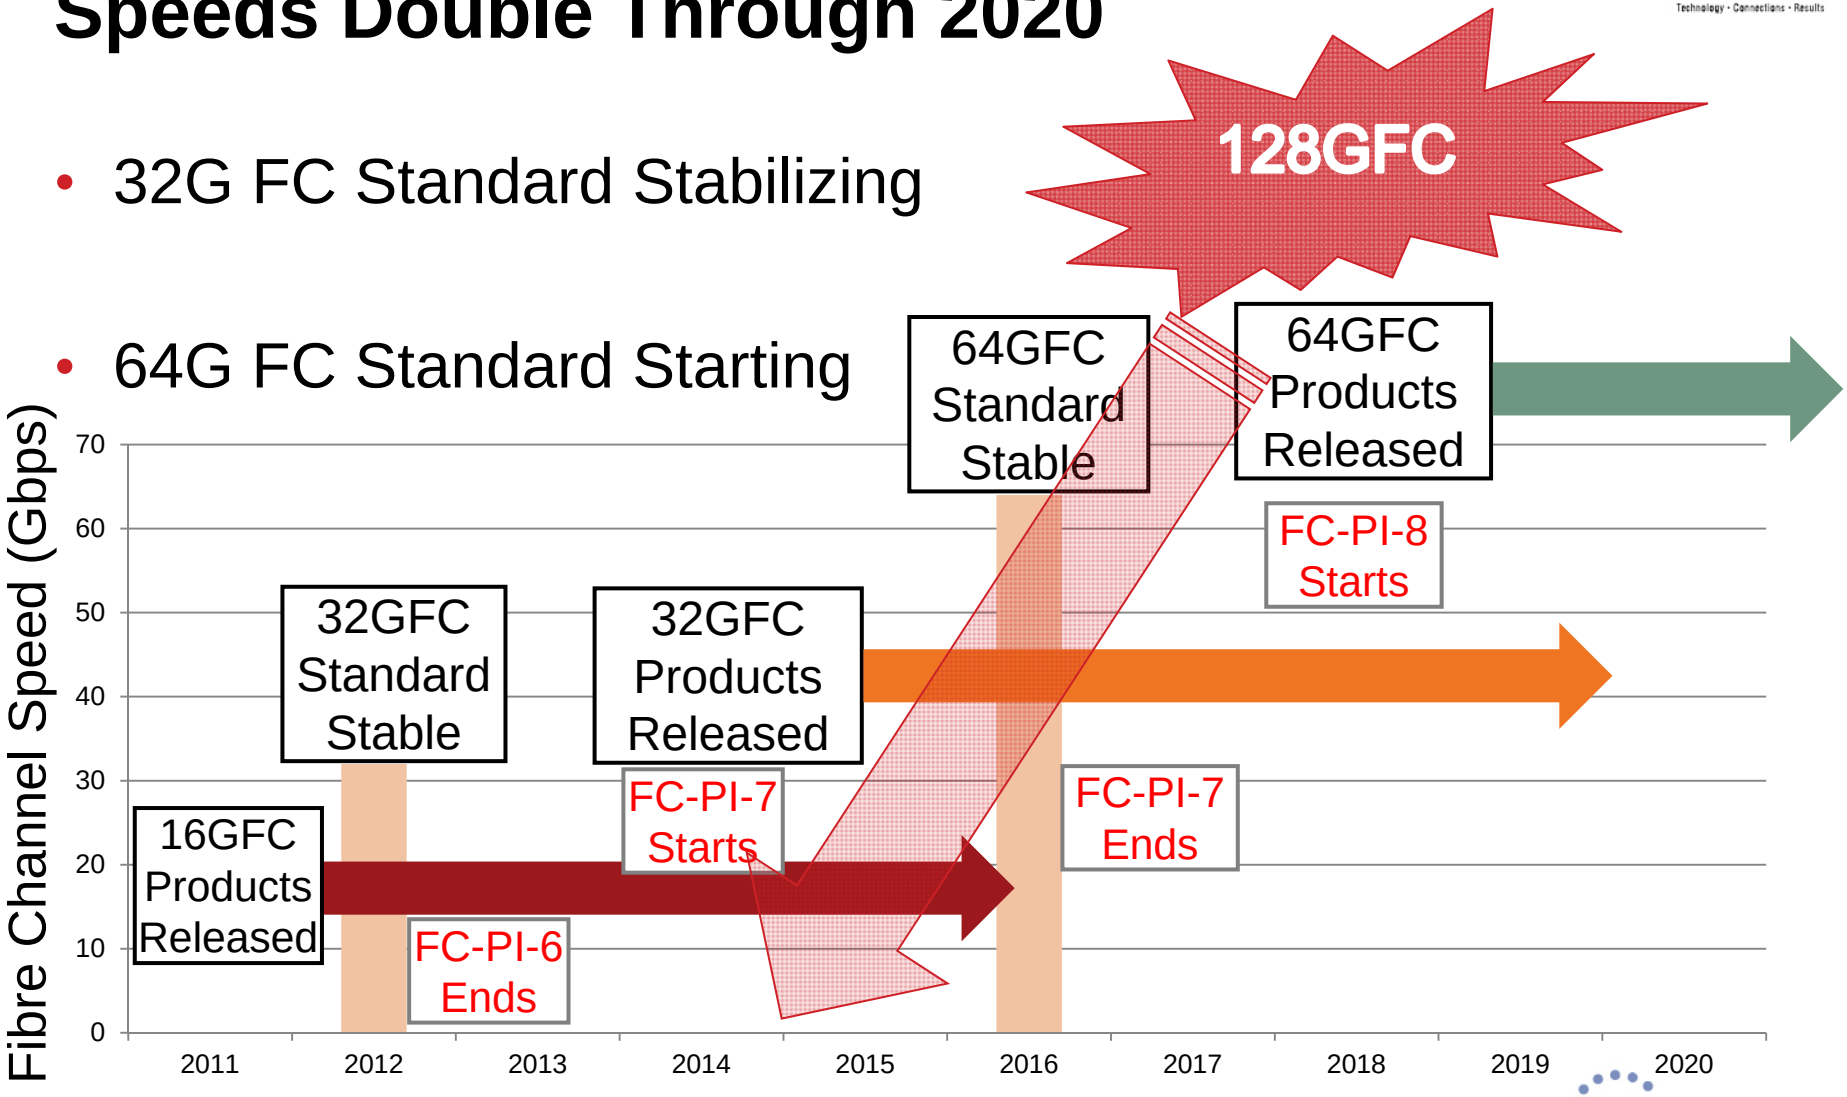

Consolidating the Data Center Fabric

Many networks, One Infrastructure

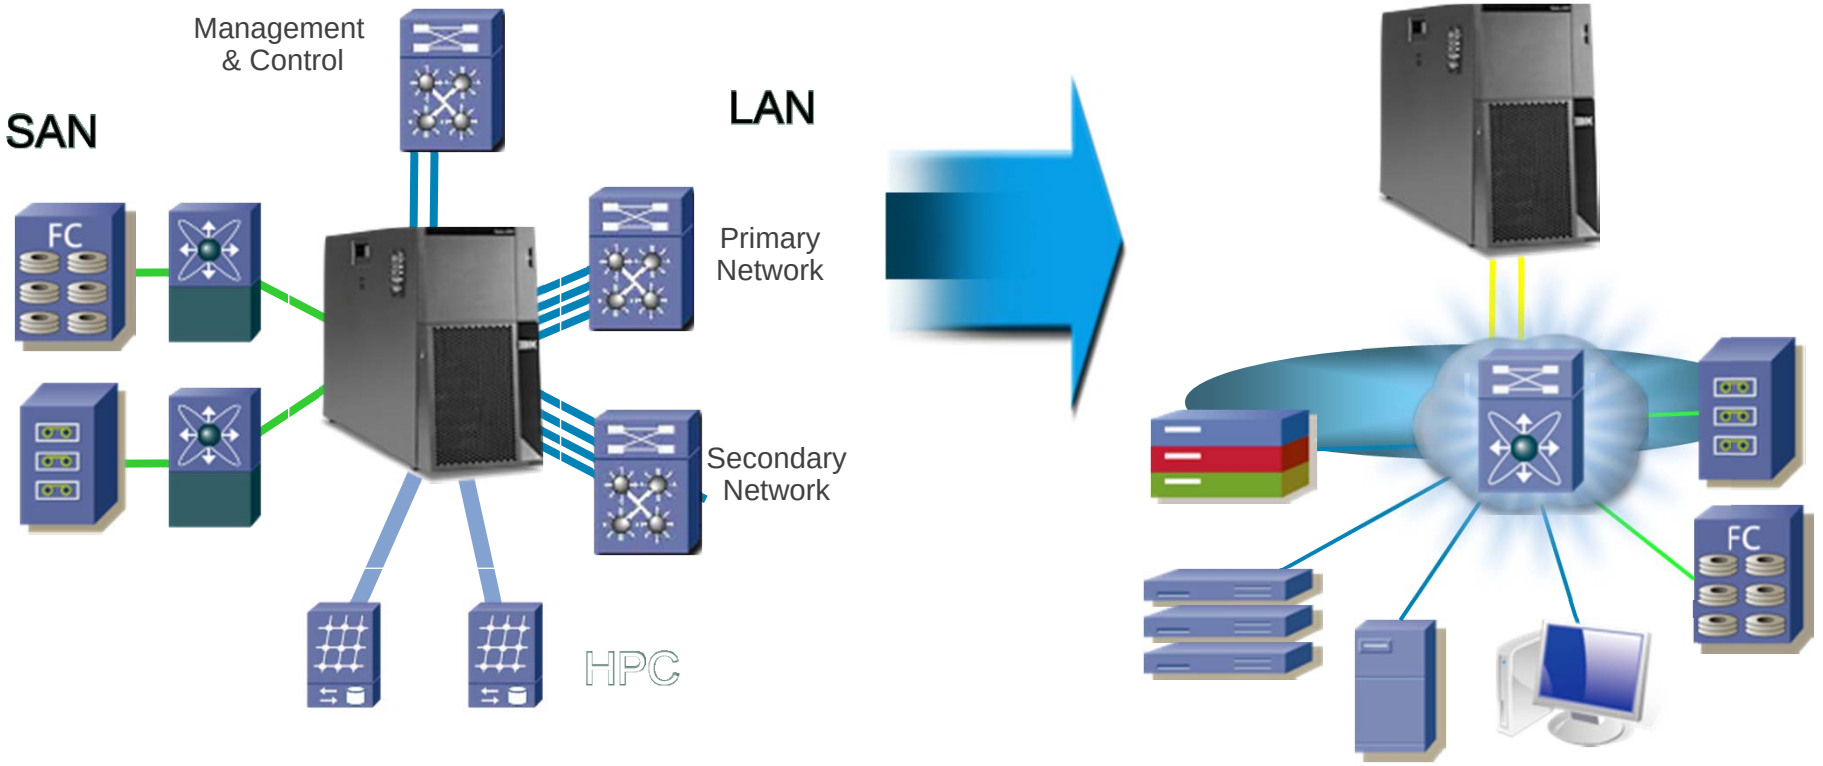

Complexity,
Cost, Power

Increased Efficiency,
Simpler Operations

Primary FCoE Use Case

Before Unified I/O

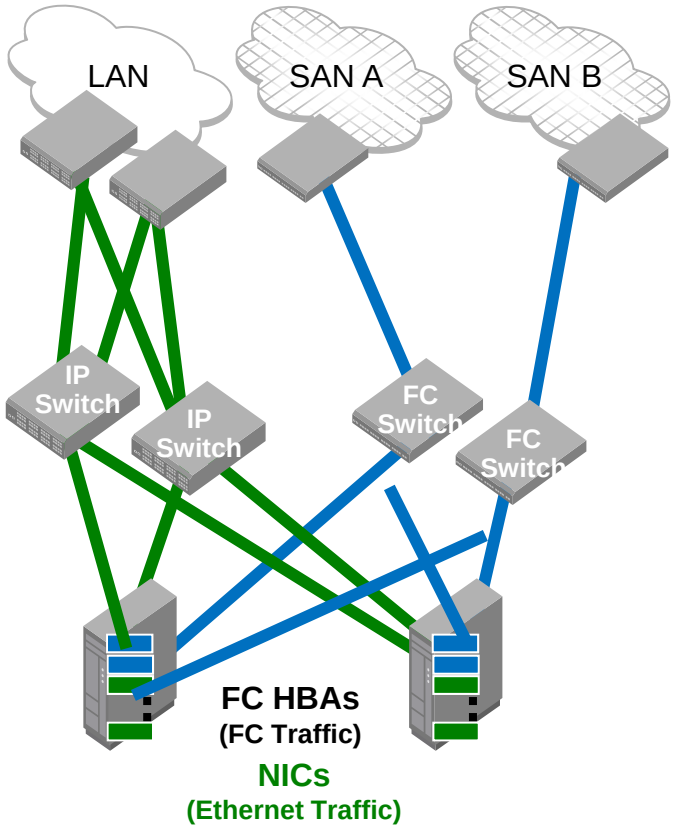

After Unified I/O

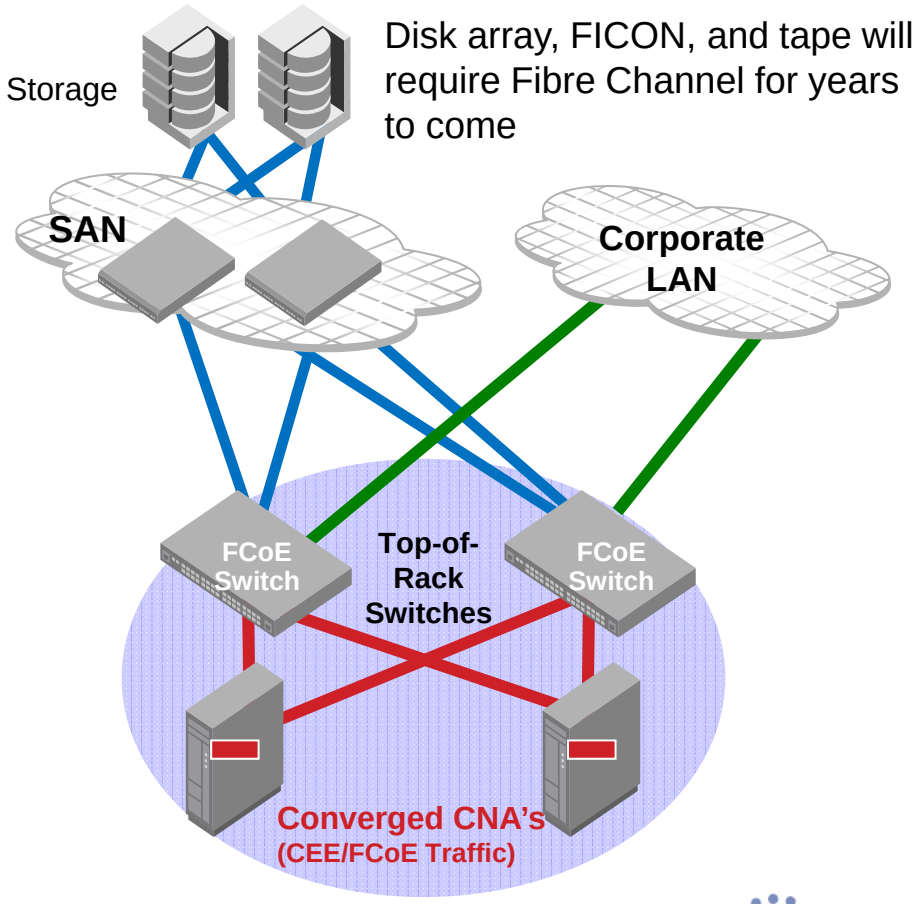

FC Ethernet FCoE and CEE

10G on LAN On Motherboard (LOM) A Game Changer

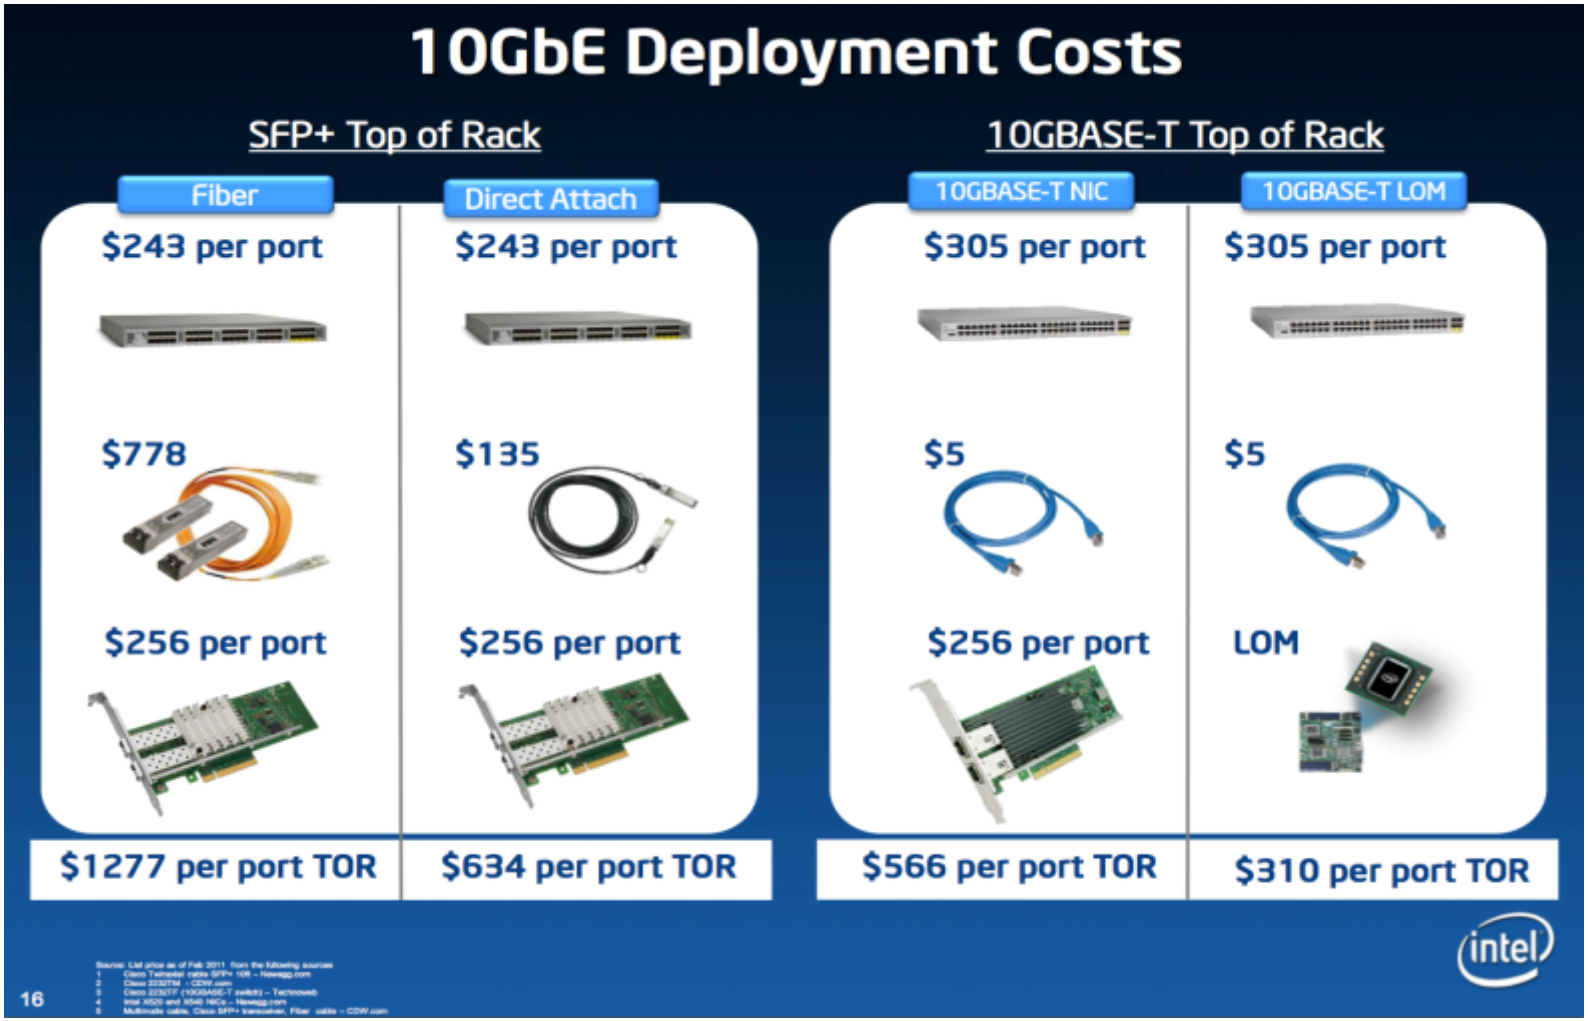

Protocol Roadmaps

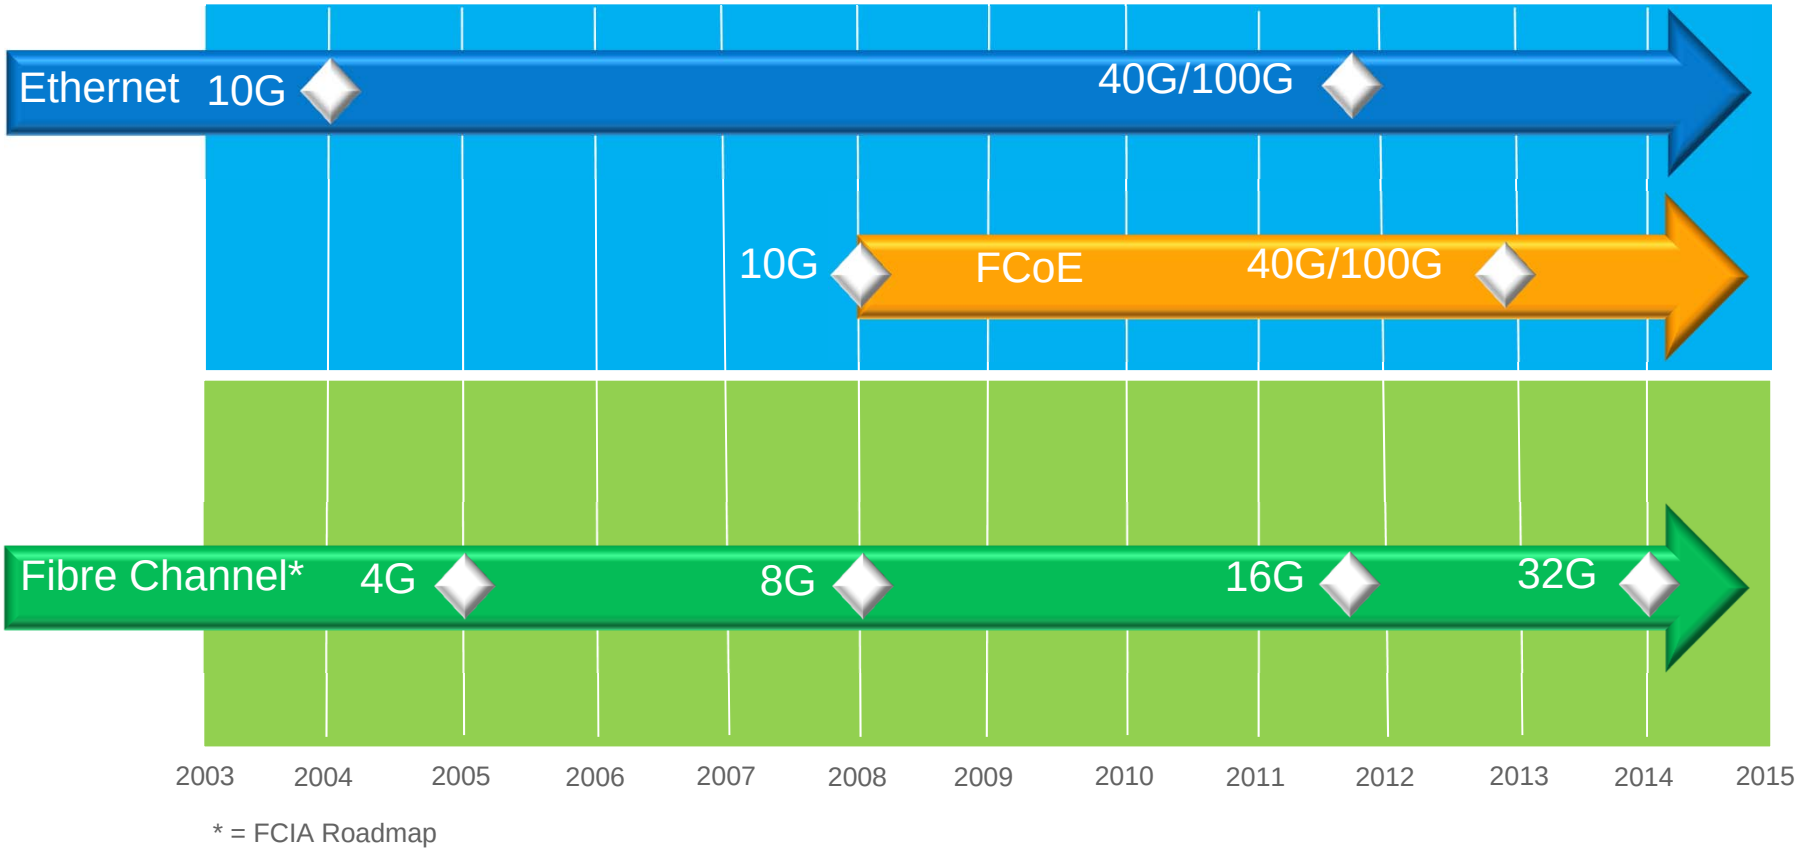

* = FCIA Roadmap

Ethernet is set to surpass Fibre Channel on throughput

Fiber Channel Bandwidth Roadmaps

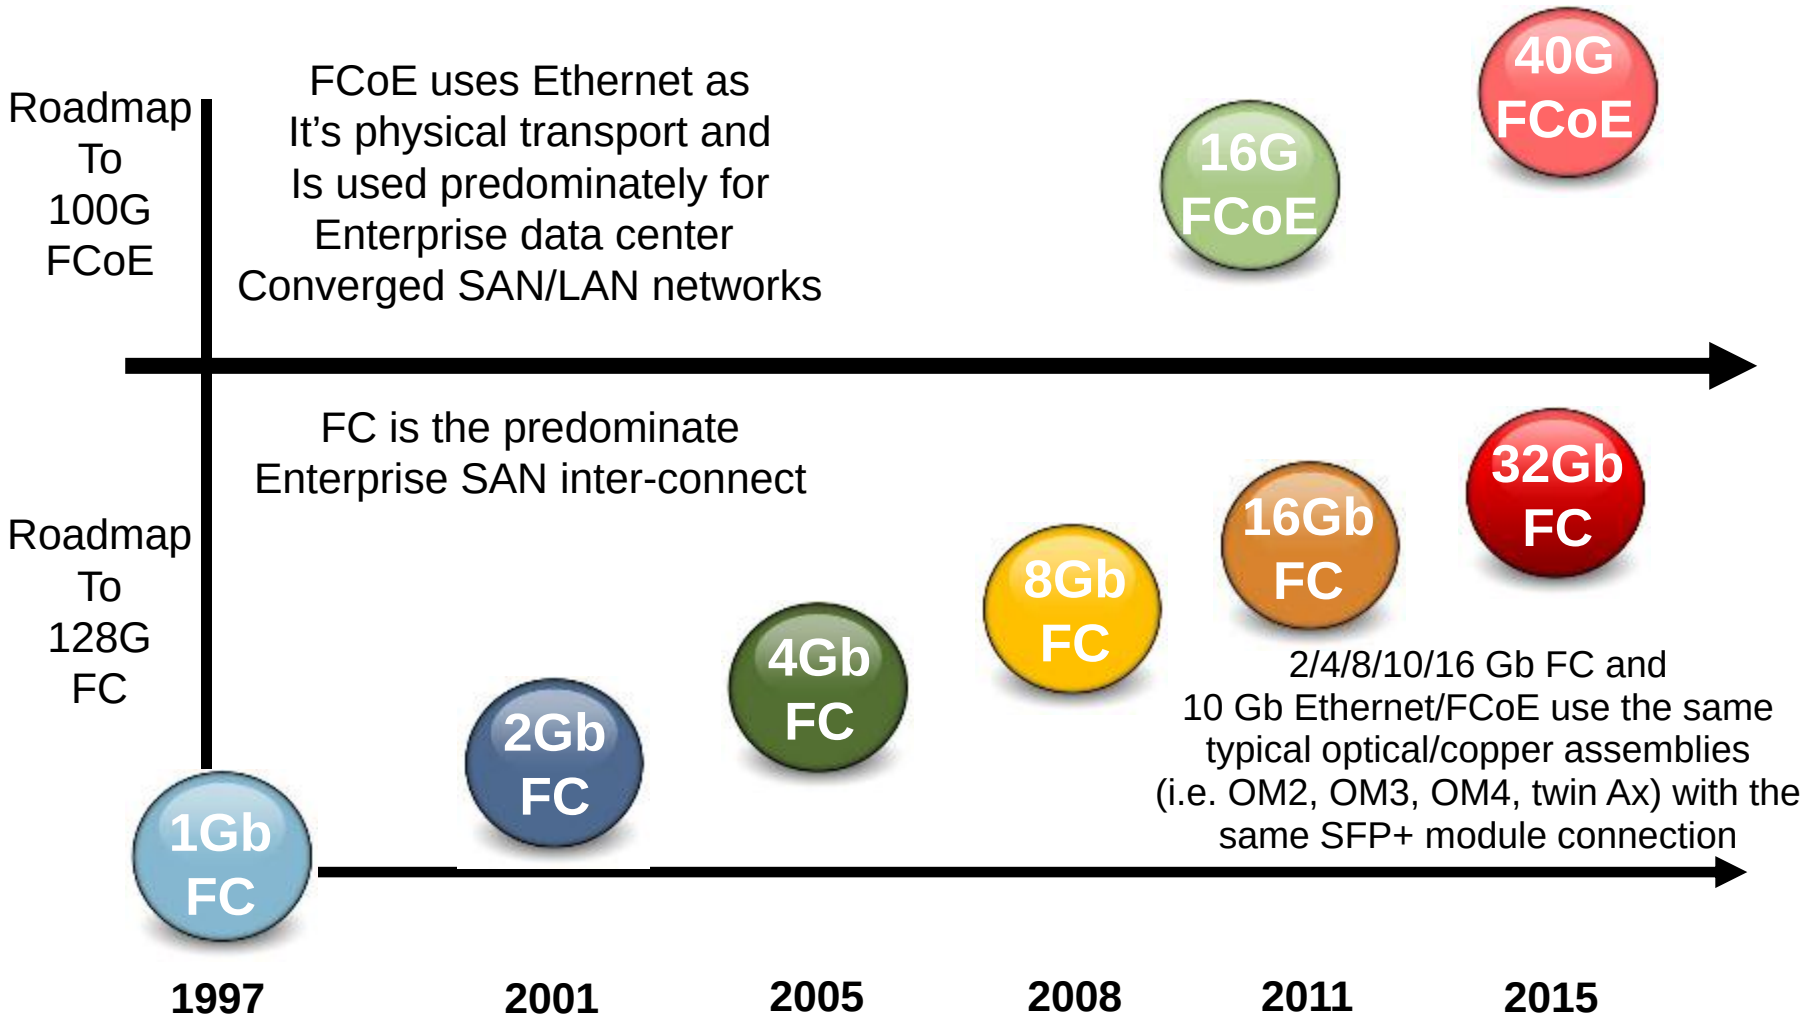

TOTAL Investment Protection!

Storage Technology Hype Cycle Curve

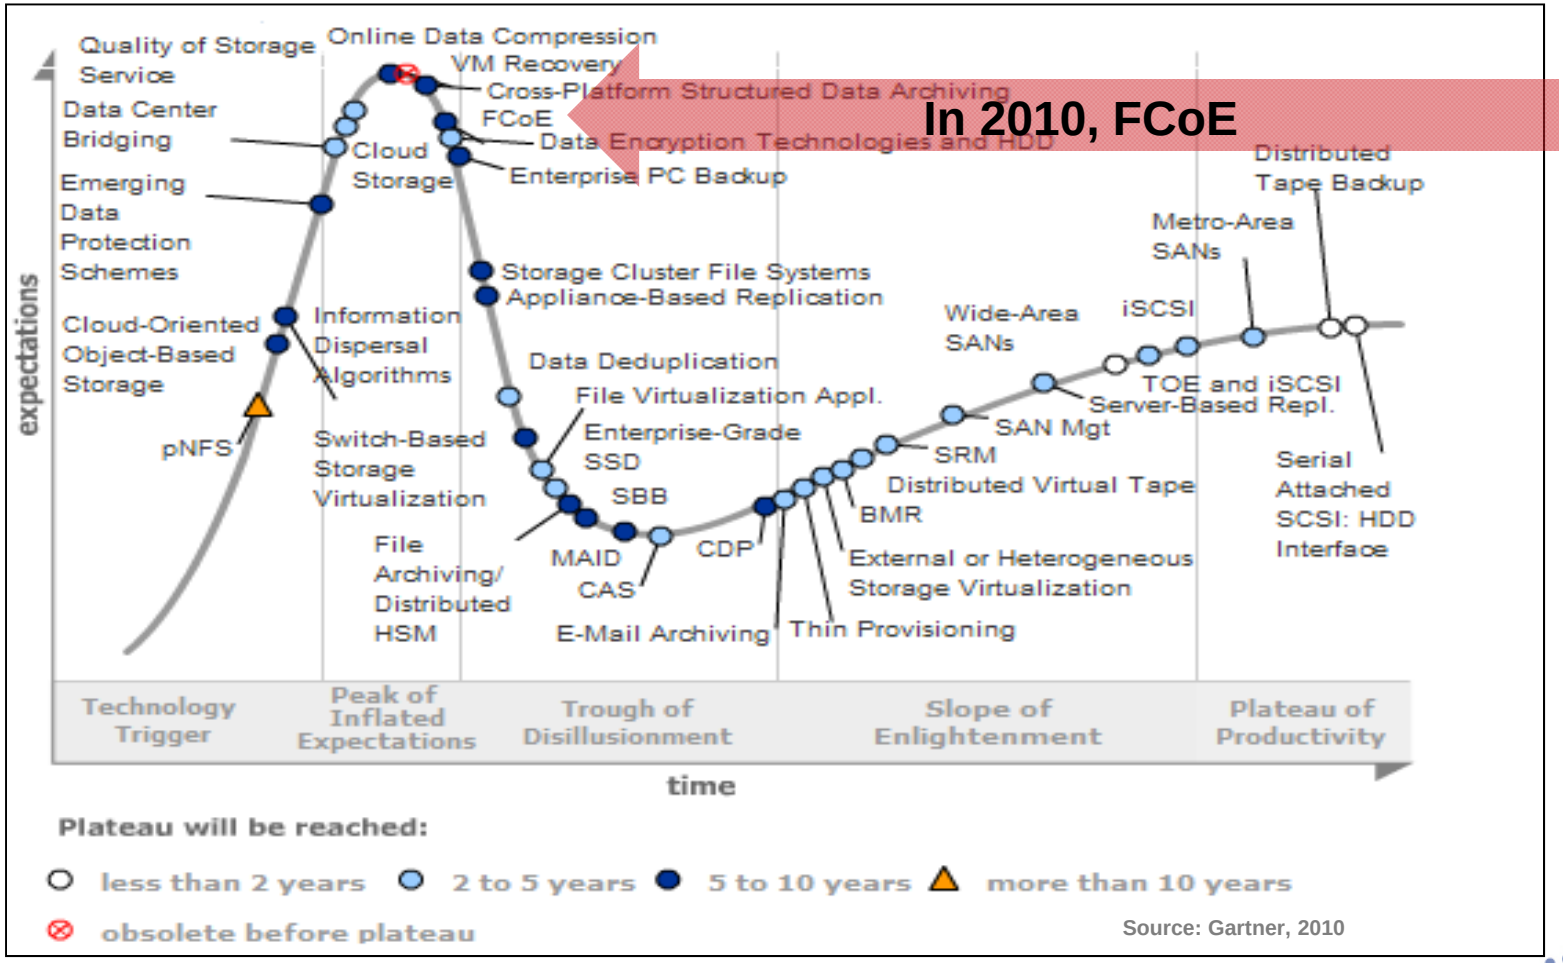

Source: Gartner, 2010

Storage Technology Hype Cycle Curve

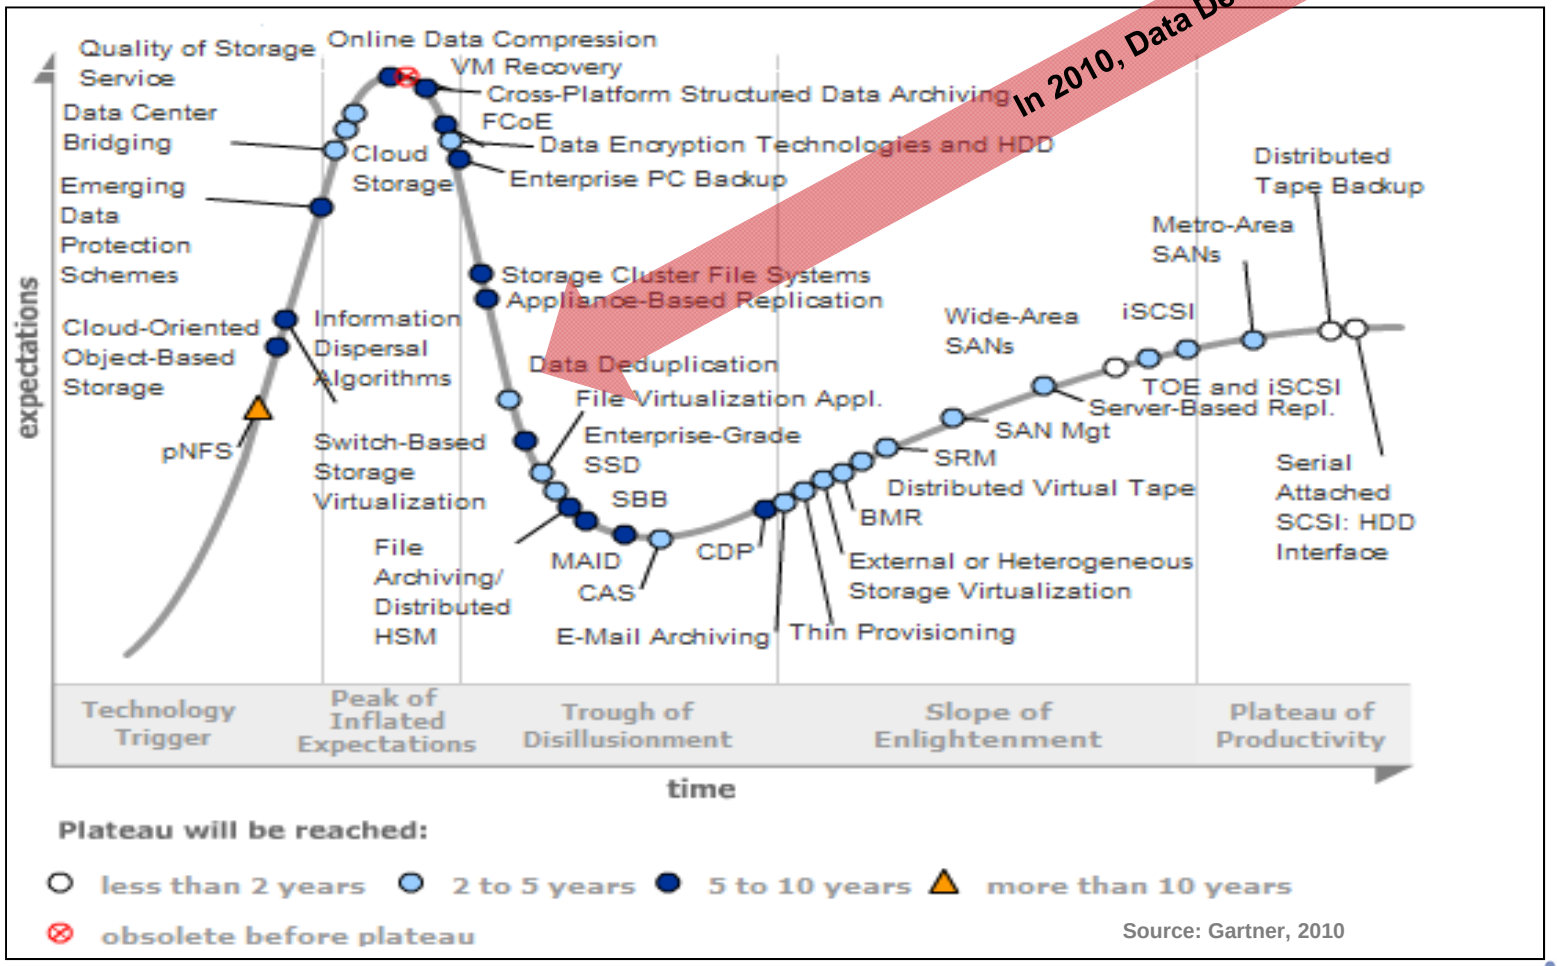

Storage Technology Hype Cycle Curve

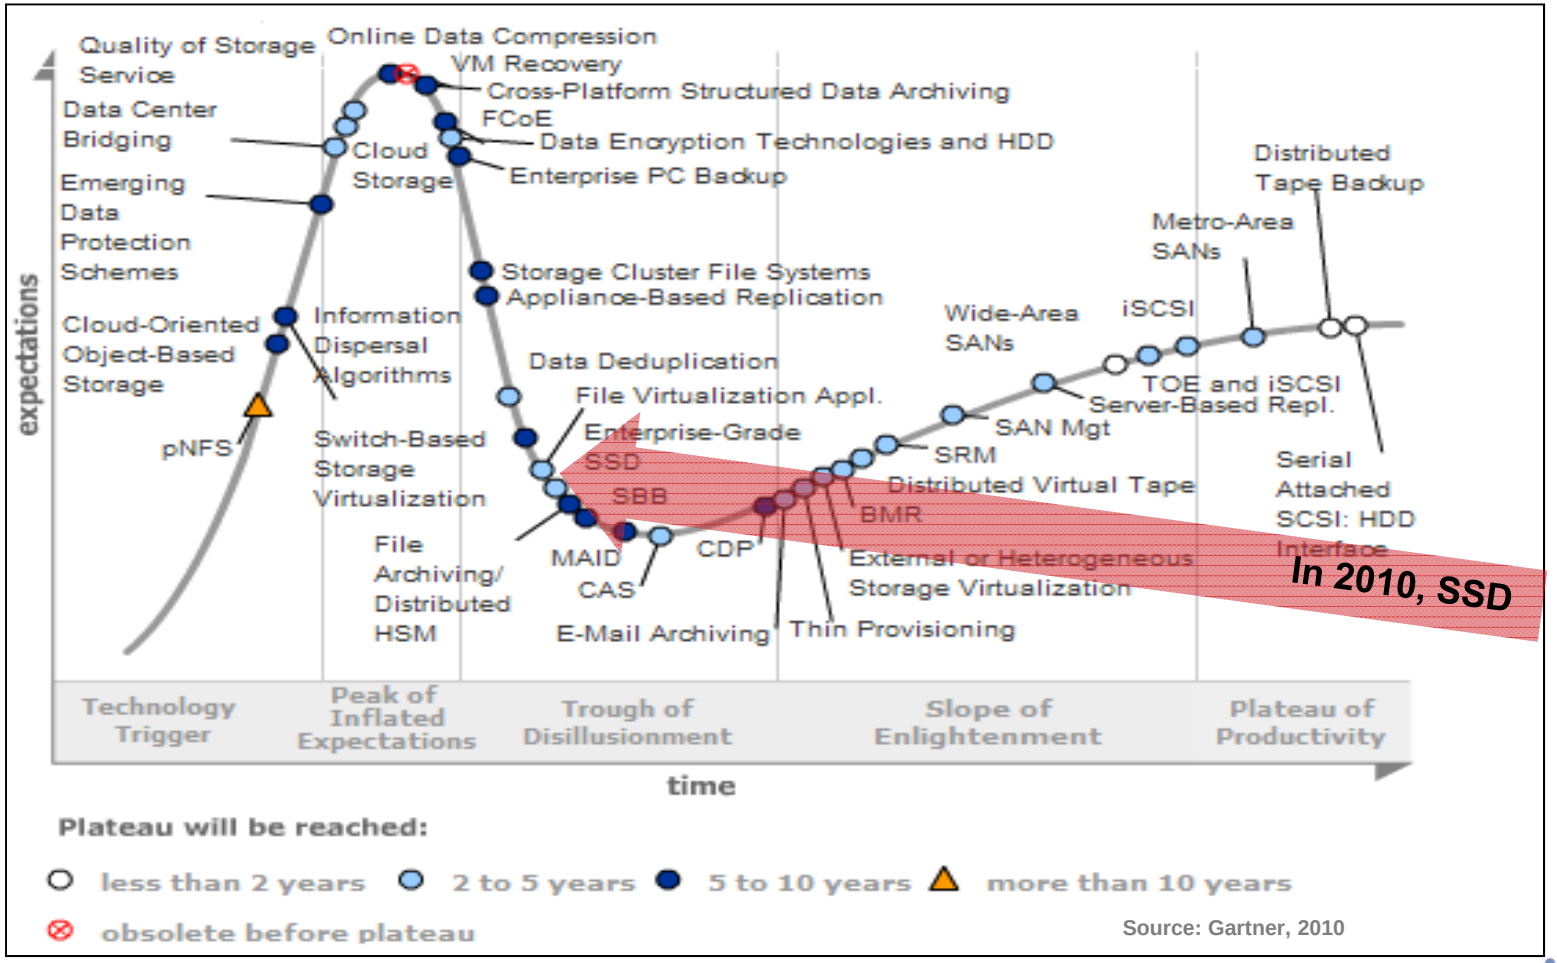

Network Technology Hype Cycle Curve

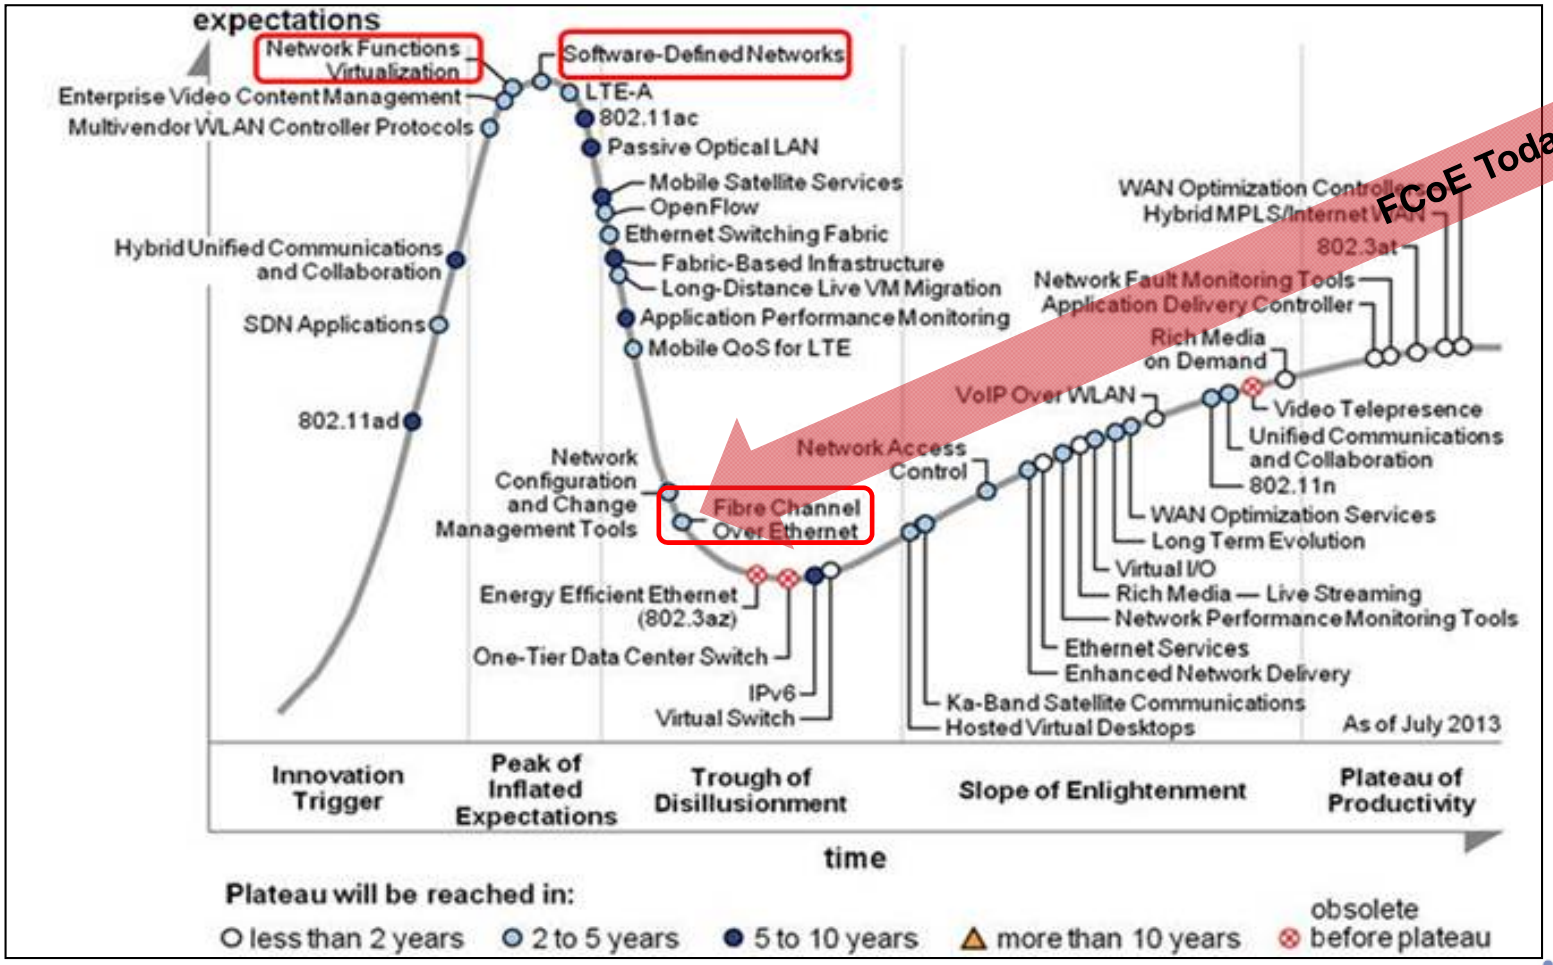

Storage Growth Has No End In Sight

- Fibre Channel is forecasted by IDC to remain on top for SAN connectivity

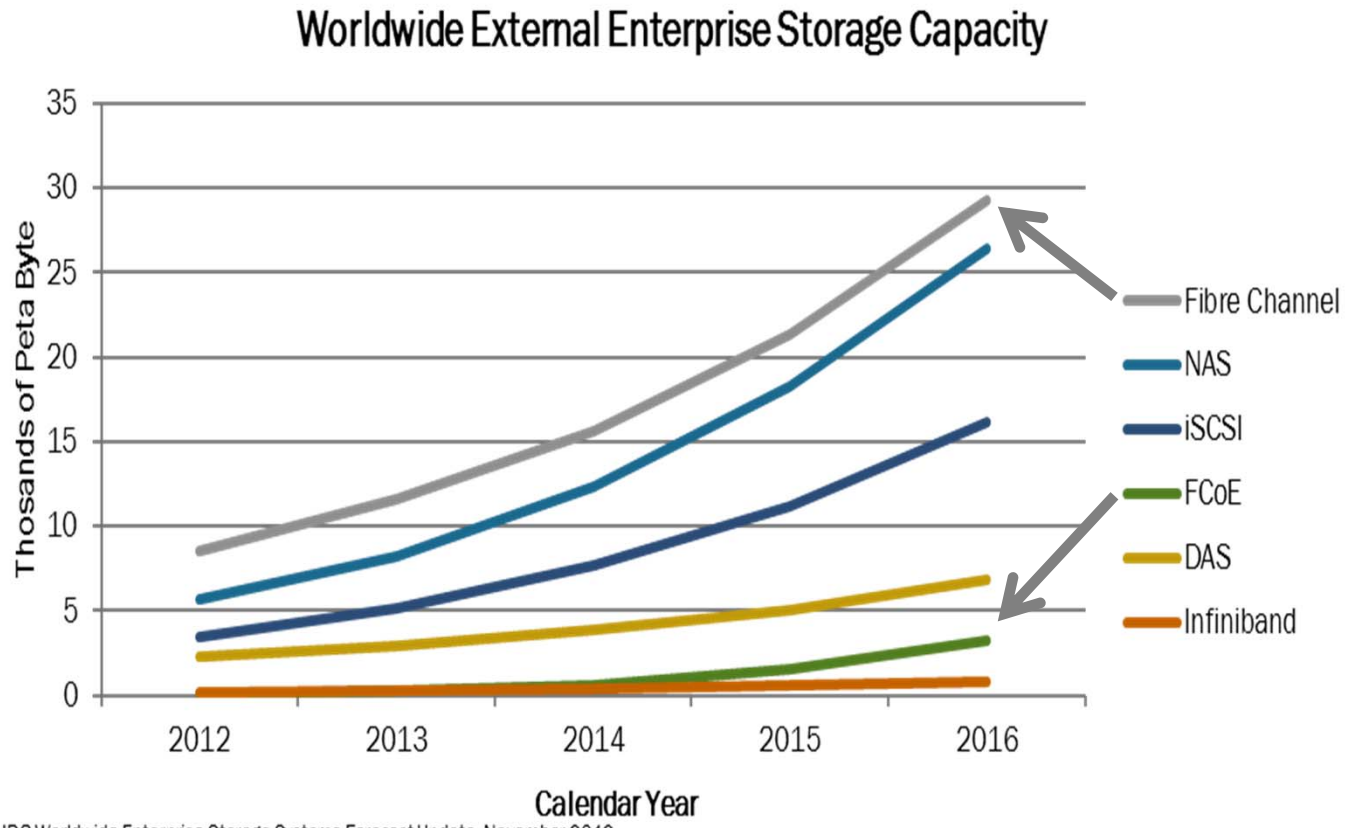

IDC Worldwide Enterprise Storage Systems Forecast Update, November 2012

FC Storage is predicted to have ~36% CAGR (2012/2016)

Fibre Channel is “Purpose Built”

- **Simplicity**
  - FCP was designed with a singular purpose in mind, and does not have to contend with a complex protocol stack.
- **Performance**
  - A native 16Gbps FC port is 40% faster than a 10GbE network, and it too can be trunked to provide aggregate ISL bandwidth up to 128 Gbps.
- **Low Latency**
  - FC fabric is not penalized by the additional 2-hop latency imposed by routing data packets through a NAS server before it’s written to disk.
- **Parity of Cost**
  - The dramatic reduction in expense promised by FCoE has failed to materialize. The complexity and cost of pushing data at NN_Ghz is fairly consistent, regardless of what protocol it used.
- **Efficiency**
  - Having a Fibre Channel back-end network supports such capabilities as LAN-less backup technology, high speed data migration, block-level storage virtualization, and in-fabric encryption.

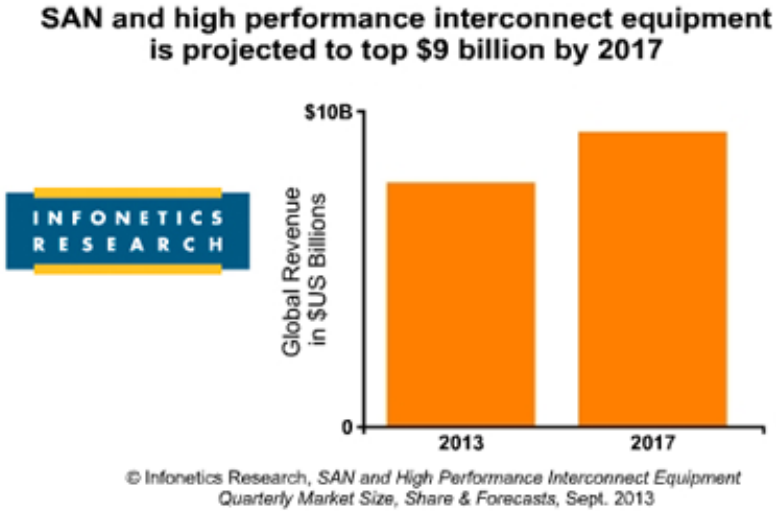

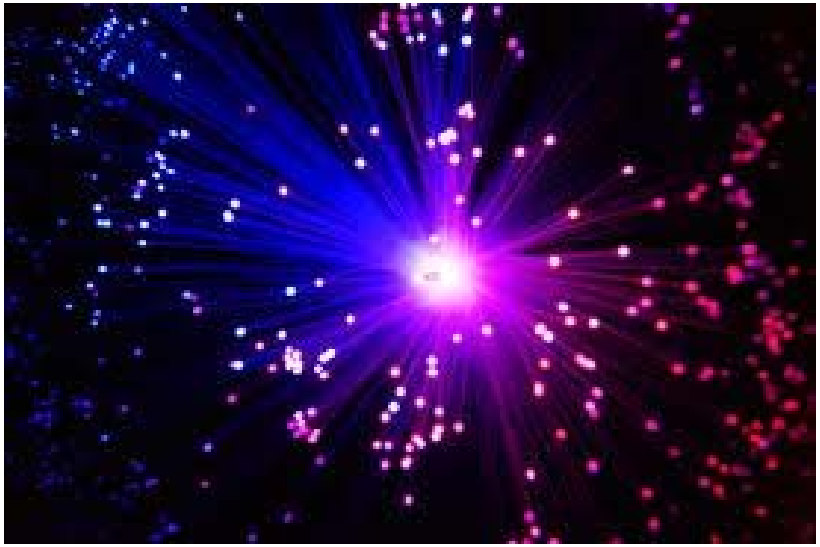

Modal Dispersion

Light in Flight

Measuring Light Signals

LET'S TALK ABOUT LIGHT

FC Storage Networking Terminology

Fiber Channel Links

- **Modal dispersion** is a distortion mechanism occurring in multimode fibers in which the signal is spread in time because the propagation velocity of the optical signal is not the same for all modes.
- Modal dispersion limits the bandwidth and distance of multimode fibers.

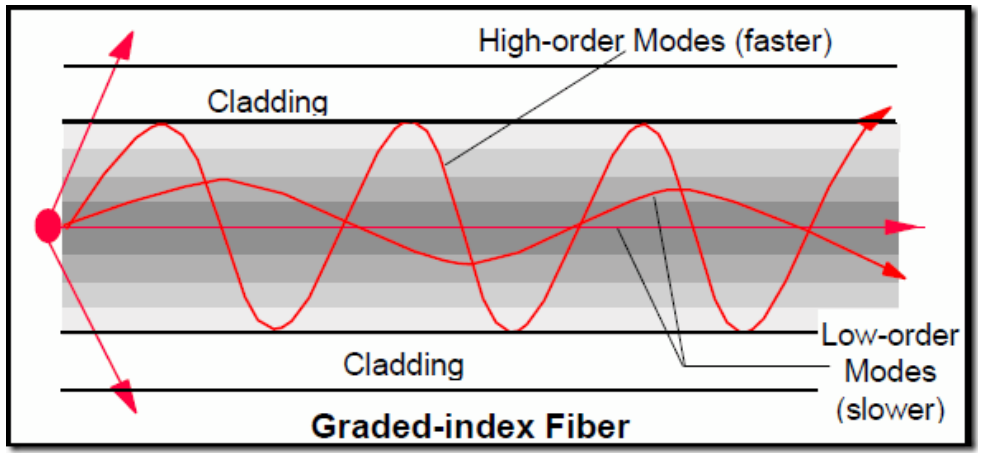

OM3 and OM4

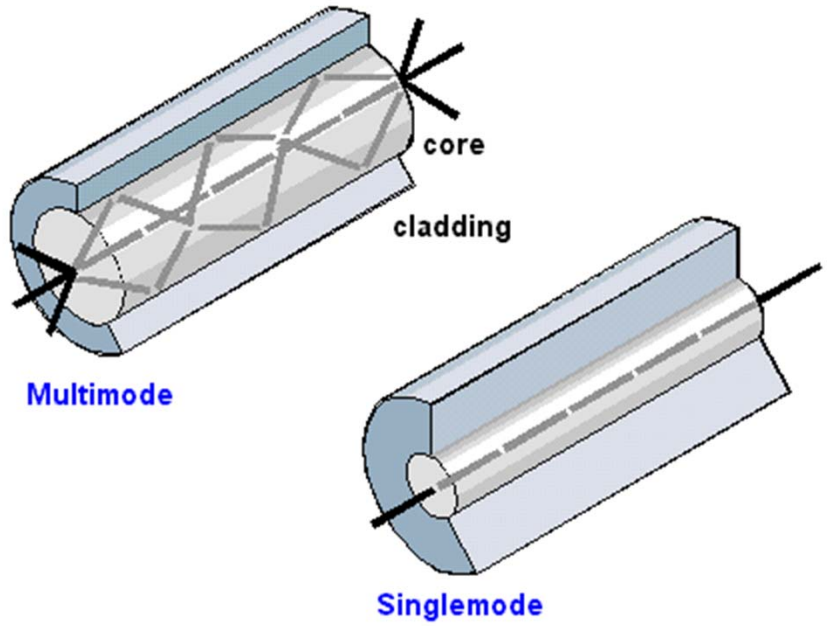

FC Storage Networking Terminology
Fiber Channel Links

- **Photos of Modal dispersion**
  - As you can see, a beam of light travels from side to side as it travels from one end of the cable to the other. This is how fibre optics can transmit data across long distances while not confined to being straight line of sight paths.

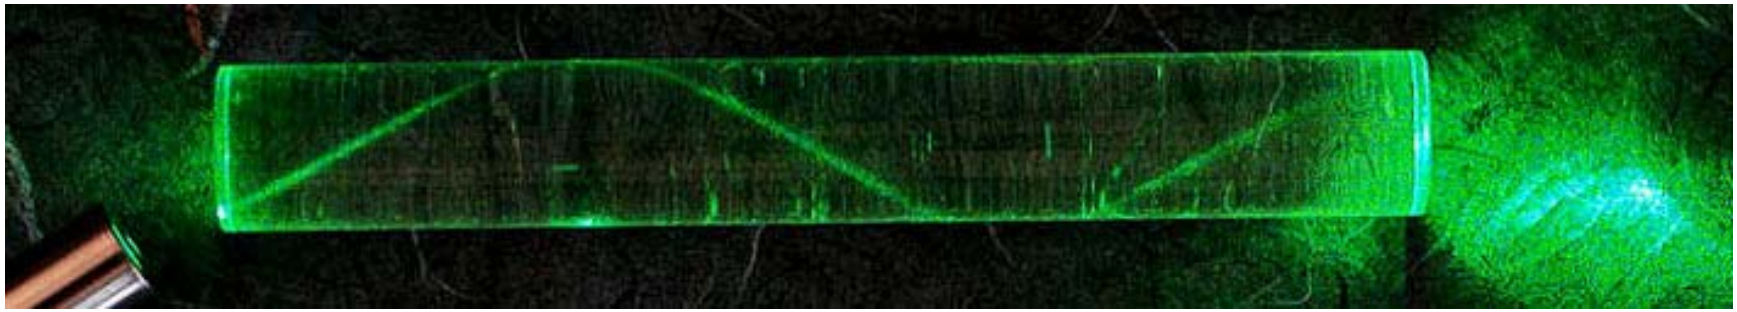

Light enters the cable

Light carries through the cable with a little dispersion

Without the cable light dispersion happens quickly

We send Data using Light

http://www.ted.com/talks/ramesh_raskar_a_camera_that_takes_one_trillion_frames_per_second.html

- **Light in Flight**
  - There is now a camera that can take a trillion frames per second. Below is a photo of light in flight from a laser pointer. The distance of the light shown below is the total distance that light travels in atmosphere in a Femtosecond.
  - A femtosecond (10^{-15} seconds) is one quadrillionth, or one millionth of one billionth of a second. Put another way: a femtosecond compares to a second, as a second compares to 30 million years.

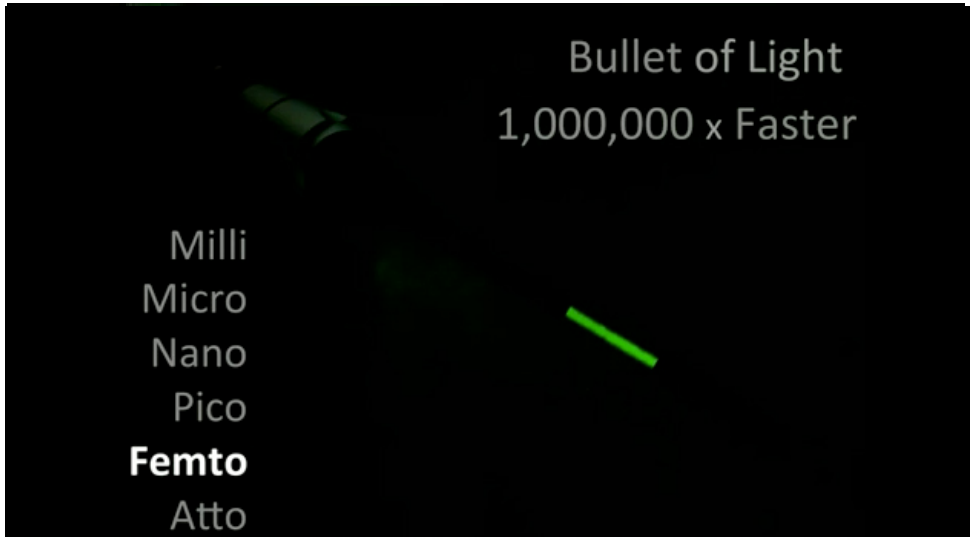

Unit	Size	Notes
attosecond	10^{-18} s	shortest time now measurable by scientists
femtosecond	10^{-15} s	pulse width on world's fastest lasers
picosecond	10^{-12} s	switching time of the world's fastest transistor
nanosecond	10^{-9} s	time for molecules to fluoresce
microsecond	10^{-6} s	length of time of a high-speed, strobe light flash
millisecond	0.001 s	time for a housefly's wing flap

Measuring Light Signals

- Technology is pushing our capabilities to measure the data in a light signal
- 20 picoseconds is about our technological capability to be able to measure the rising and falling of light in pulse in order to determine the information that the light pulse is carrying – but a femtosecond of time can carry a lot of data

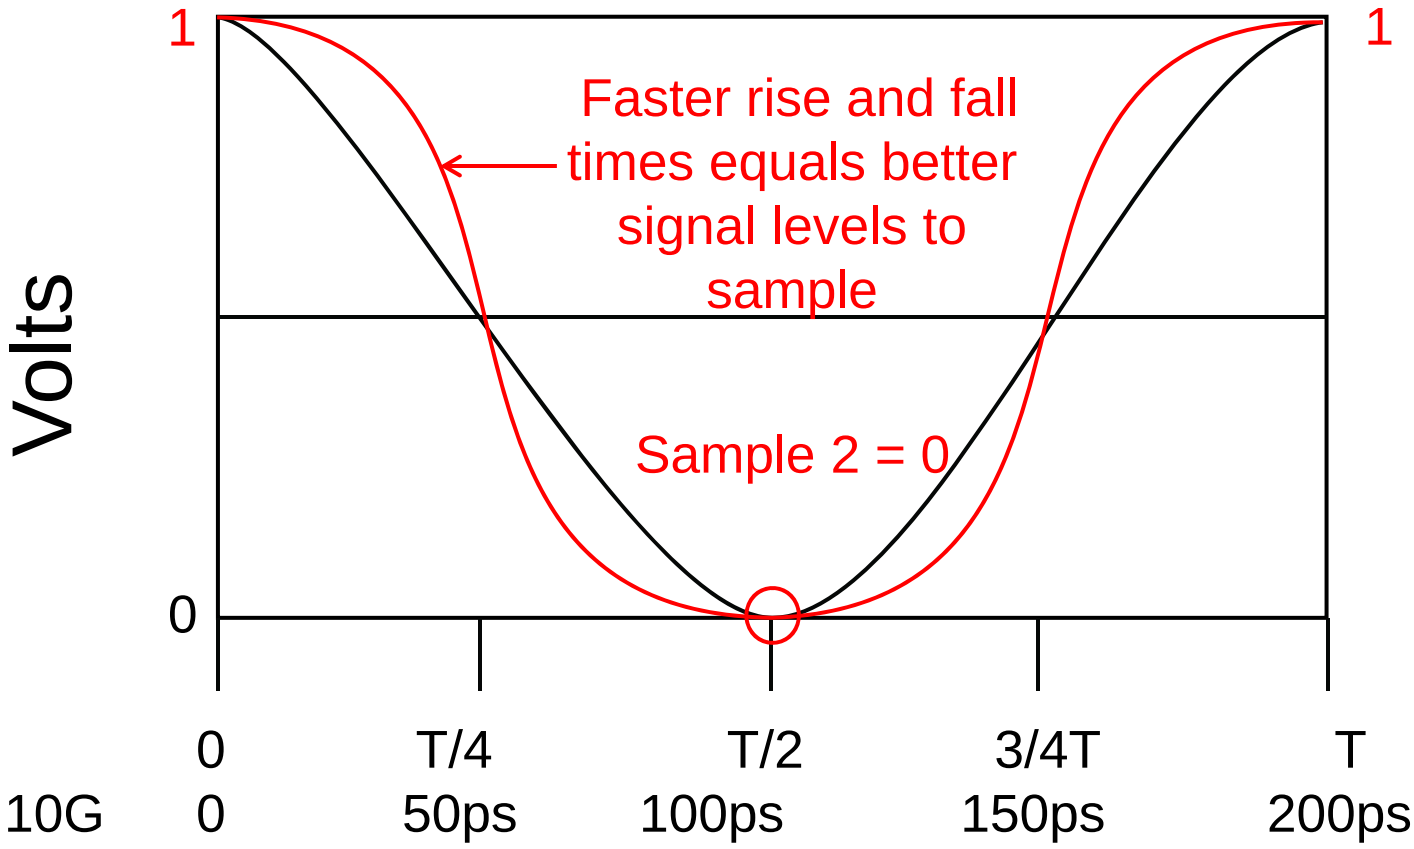

Reflection / Discussion

- This is all interesting data-points but –
  - When do your applications need greater than 8G ?
  - When will the servers have higher speed availability ?
  - When will the Disks / Tapes / VTLs have higher speeds ?
  - What and When are the Technology inflection points ?
  - Which Technology(s) will have solutions sooner ?

References

- Fibre Channel Standard
  - www.t11.org
- Fibre Channel Industry Association
  - www.fcia.com
- Storage Networking Industry Association
  - www.snia.org
- Ethernet Alliance
  - www.ethernetalliance.org

Sessions of Interest

- Steve Guendert
  - Session 14882 (Wednesday, March 12, 2014: 8:00 AM-9:00 AM)
    - Replication for Disaster Recovery – Don't Forget the Tape!
  - Session 14374 (Wednesday, March 12, 2014: 11:00 AM-12:00 PM)
    - 'Why' Ultra High Availability and Integrated EMC-Brocade Data Replication Networks
  - Session 14986 (Wednesday, March 12, 2014: 1:30 PM-2:30 PM)
    - System z Long Distance Extension and FCIPNetwork Primer
- Dave Lytle
  - Session 14375 (Thursday, March 13, 2014: 9:30 AM-10:30 AM)
    - A Deeper Look Into the Inner Workings and Hidden Mechanisms of FICON Performance
  - Session 14769 (Thursday, March 13, 2014: 4:30 PM-5:30 PM)
    - Build FC I/O Fabric Super Highways Using ISL Virtualization

SHARE, Anaheim, March 2014

Buzz Fibrechannel - To 16G and Beyond!

Session 14382

THANK YOU!

QR Code