



Buzz-Fibre-Channel to 16g and beyond

17 April 2013

David Lytle, BCAF

Global Solutions Specialist

System z Solutions & Technologies

dlytle@brocade.com



Fly the cursor over the comments box to see the Speaker notes!

Presenter

This Acrobat PDF file contains notes on some slides. These notes will appear as a cartoon speaker box in the upper left hand corner of the slide being viewed. Just fly your cursor over the cartoon box and you will be able to see the comments. Or you can open the comments by clicking on the lower left hand corner double-boxes.

Also, if you right-click on a slide you will be able to add your own comments to the slide - "Sticky Notes". When you save the file these notes will remain with it.

Cursor to any place on the page and then right click



Click on... "Add Sticky Note"

Add Sticky Note

Add Bookmark Ctrl+B

Select Tool

Hand Tool

Marquee Zoom

Previous View Alt+Left Arrow

Rotate Clockwise Shift+Ctrl+Plus

Print... Ctrl+P

Search Shift+Ctrl+F

Document Properties...

Page Display Preferences...

To open up all of the comments click here!



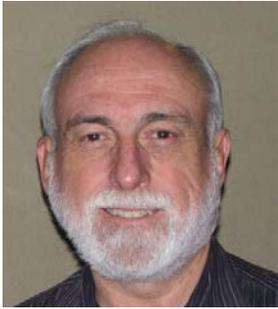
Sticky Notes allows you to add your own comments as you listen to the speaker and want to record some information for later use. Just remember to "Save" the file when you are finished!

ducts detailed in this presentation may still contain specifications, including but not limited to product features, may change. The products may vary and a production version of the products may be released. When a production version is released, it may be materially different from the release version discussed in this presentation.

• NOTHING IN THIS PRESENTATION SHALL BE DEEMED TO CREATE A WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS WITH RESPECT TO ANY PRODUCTS AND SERVICES REFERENCED HEREIN.

• Brocade, Fabric OS, File Li are registered trademarks and are trademarks of Brocade and its subsidiaries, in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Who am I?



David Lytle, BCAF
Principal Engineer / Global Solutions Specialist
System z Technologies and Solutions
44 years in Mainframe Computing
dlytle@brocade.com



- A Subject Matter Expert on Mainframe I/O, I live near Atlanta, Georgia (allowing me to keep my boss 2,000 miles away)



- Published papers which have appeared in z/Journal and in The Journal of Computer Resource Management.



- Held positions such as District SE manager, FICON Solutions Architect, Consulting Systems Engineer, SAN Sales Specialist, Storage Sales Specialist, Post Sales SE, Senior MVS Systems Programmer and Senior Programmer.



- Participate regularly as a SHARE user conference speaker. Have been a speaker at EMC, HDS, HP and IBM events around the world. Happy to have an opportunity to speak at CMG!



- Love to play at Golf (a hacker), SCUBA dive and Travel

Abstract

- In this session, I will discuss:
 - A history of the growth of storage and its bandwidth
 - Current customer trends in bandwidth utilization.
 - Fibre channel speed roadmap - per the standards.
 - Requirements for 8G (Gen4), 16G (Gen5), and 32G (Gen6)
 - What about FCoE - how does this play in user environments?
 - Then, if there is time, a fun little topic about Light!

Data Ubiquity In the Digital Age

What's Fueling The Storage And Optical Network "Big Bang"?

- Enterprise-wide transactional computing
- 7 x 24 x forever global operations
- Private/Public cloud computing
- Huge email, text, twitter growth
- Explosive Internet growth
- Big Data (unstructured)
- Digital Content
 - Especially Medical and Public Entertainment

The Past Gives Us A Roadmap to the Future!



Language Big Bang – 100,000BC +



Artistic Expression Big Bang



Art

40,000 to 50,000
Years Ago +

Language

100,000 Years Ago +

Archiving Information Big Bang



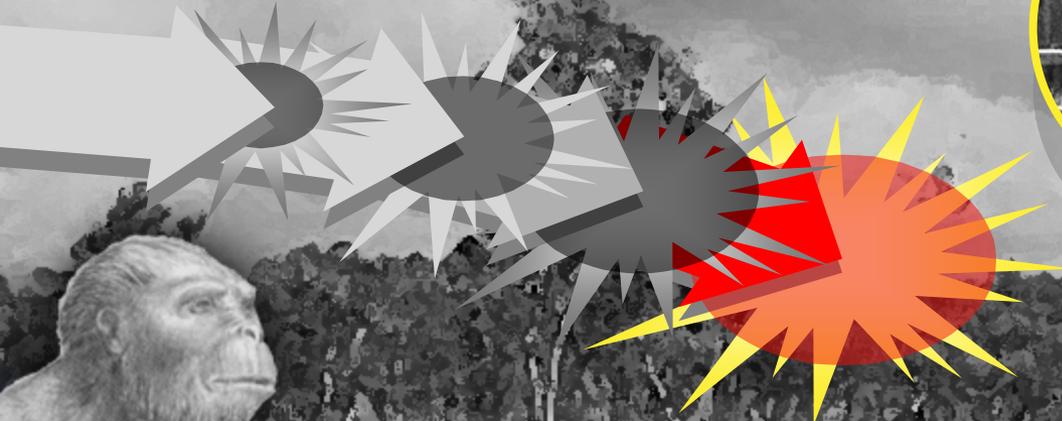
Language
100,000 Years Ago +

Writing
3,500 BC





Sharing Ideas Big Bang



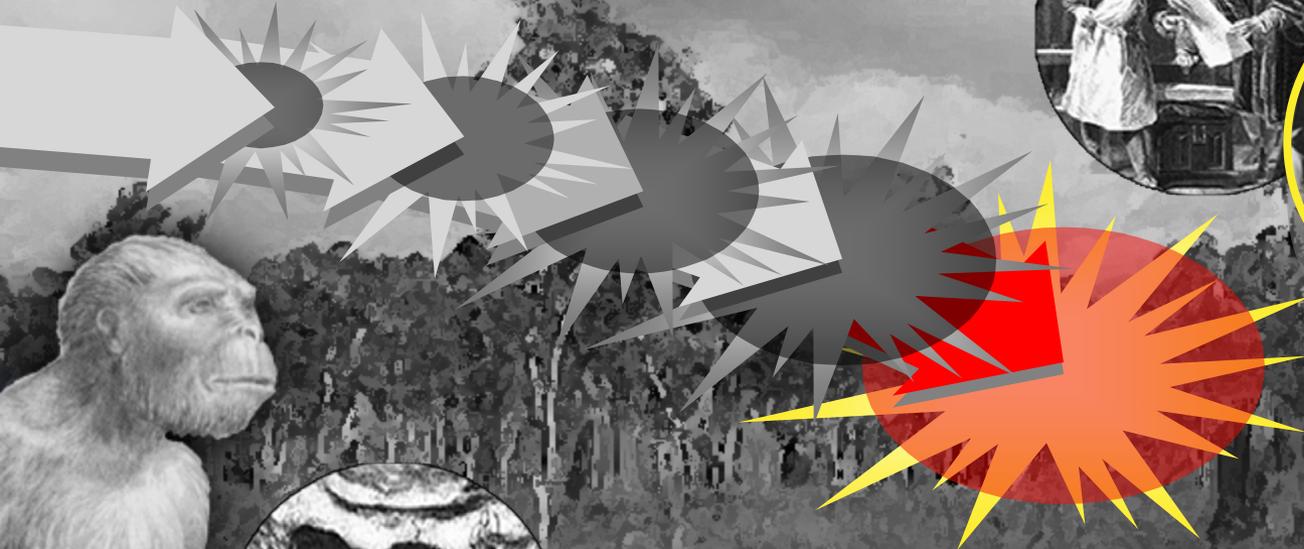
Printing
1450 AD



Language
100,000 Years Ago +



Information Processing Big Bang



Computing
1947 AD



Language
100,000 Years Ago +

Each Big Bang Is
Getting Larger and Larger



Instantaneous Communication Big Bang



Language
100,000 Years Ago +

**Internet
&
email**
1968 AD

Each Big Bang Is
Getting Larger and Larger

Data Storage and Retrieval Big Bang



Language

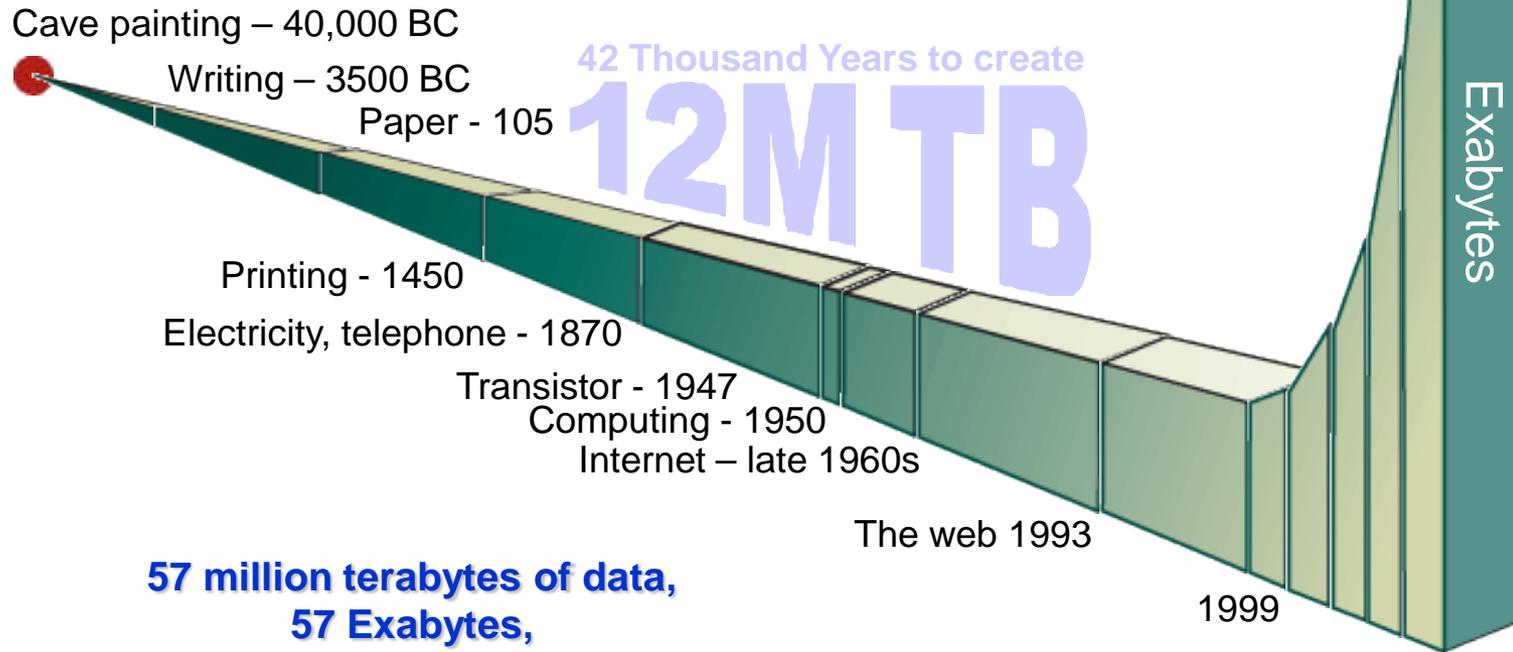
100,000 Years Ago +

Storage and
Optical Networking

TODAY

So The Storage Big Bang Extends Over Time!

The history of storing data from the beginning of human habitation up through 2004.



57
Exabytes
OF INFO
STORED
as of
2004!

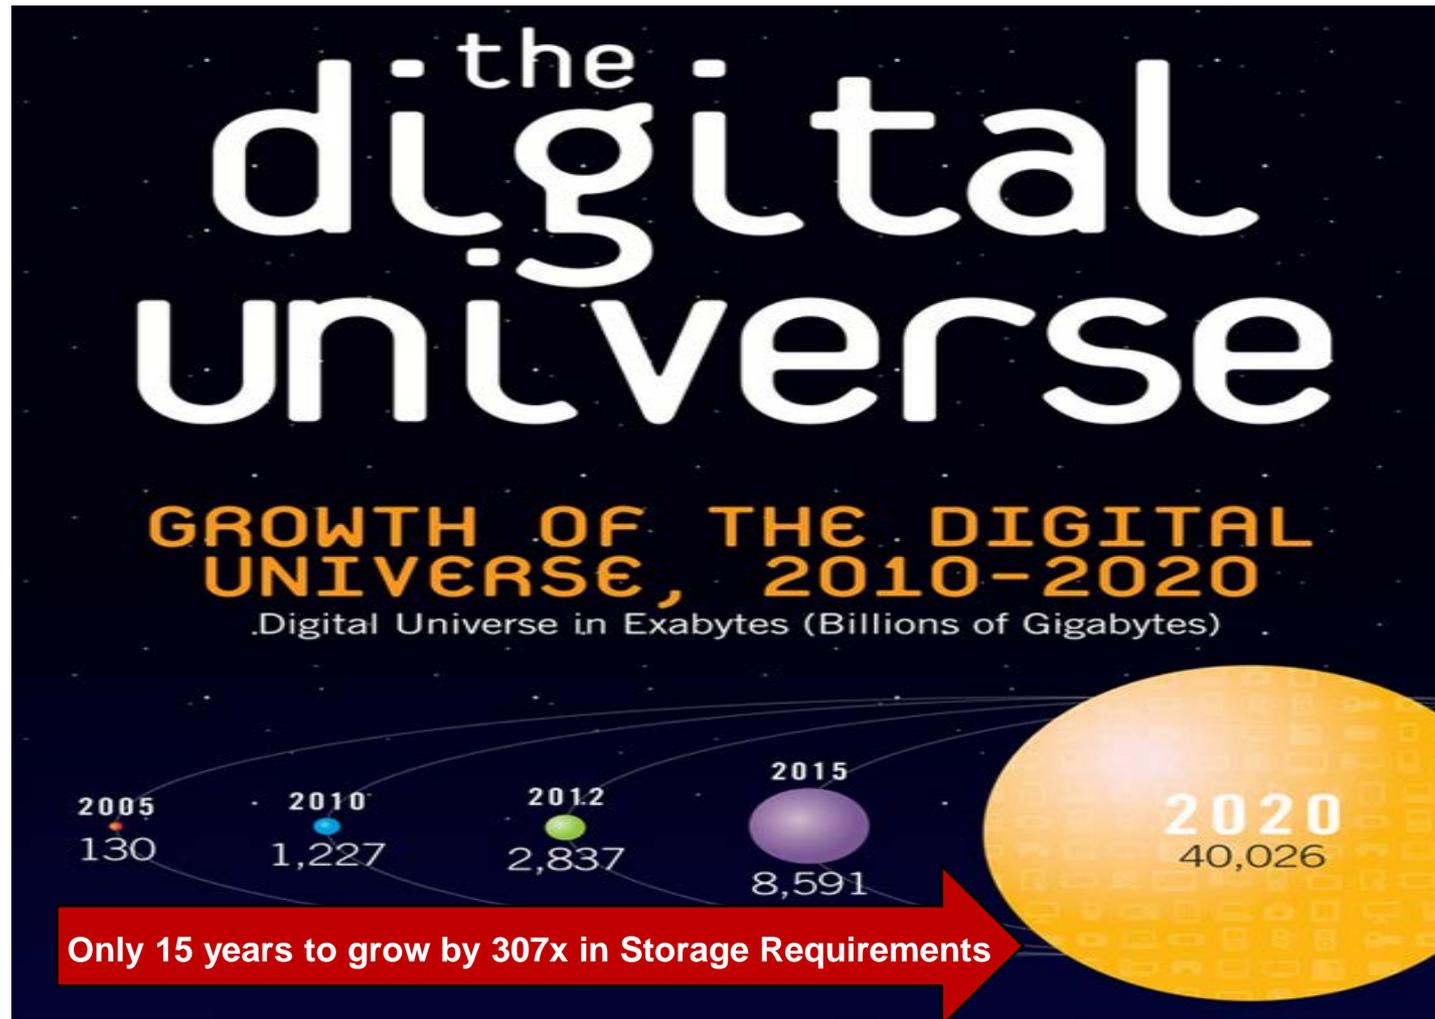
*It has been a
Tsunami
of stored
data and it
does not get
better over
time!*

**57 million terabytes of data,
57 Exabytes,
would consume 2.4 Quadrillion trees
(2,422,500,000,000,000)
if a printer used 12 lpi, single spaced,
to print all of that data**

Source: School of Information Management and Systems, Berkeley.

The Storage Content *Big Bang*

Storage From 2005 Up To The End of 2020



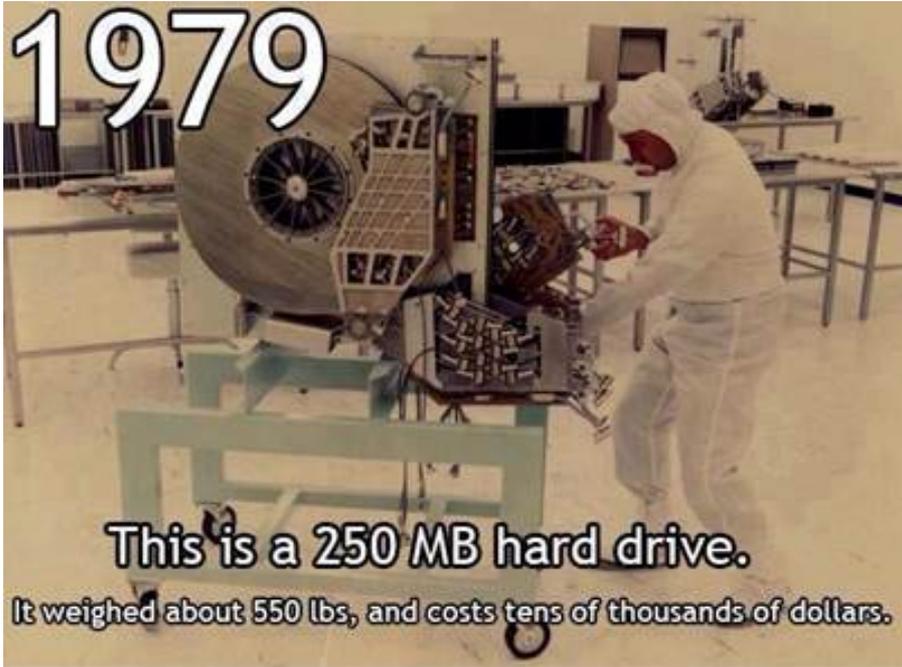
During 2010 we cracked the Zettabyte storage barrier for the first time!

By 2020 users will be creating and replicating 10s of Zettabytes of information every **Year!**

AND GROWING!

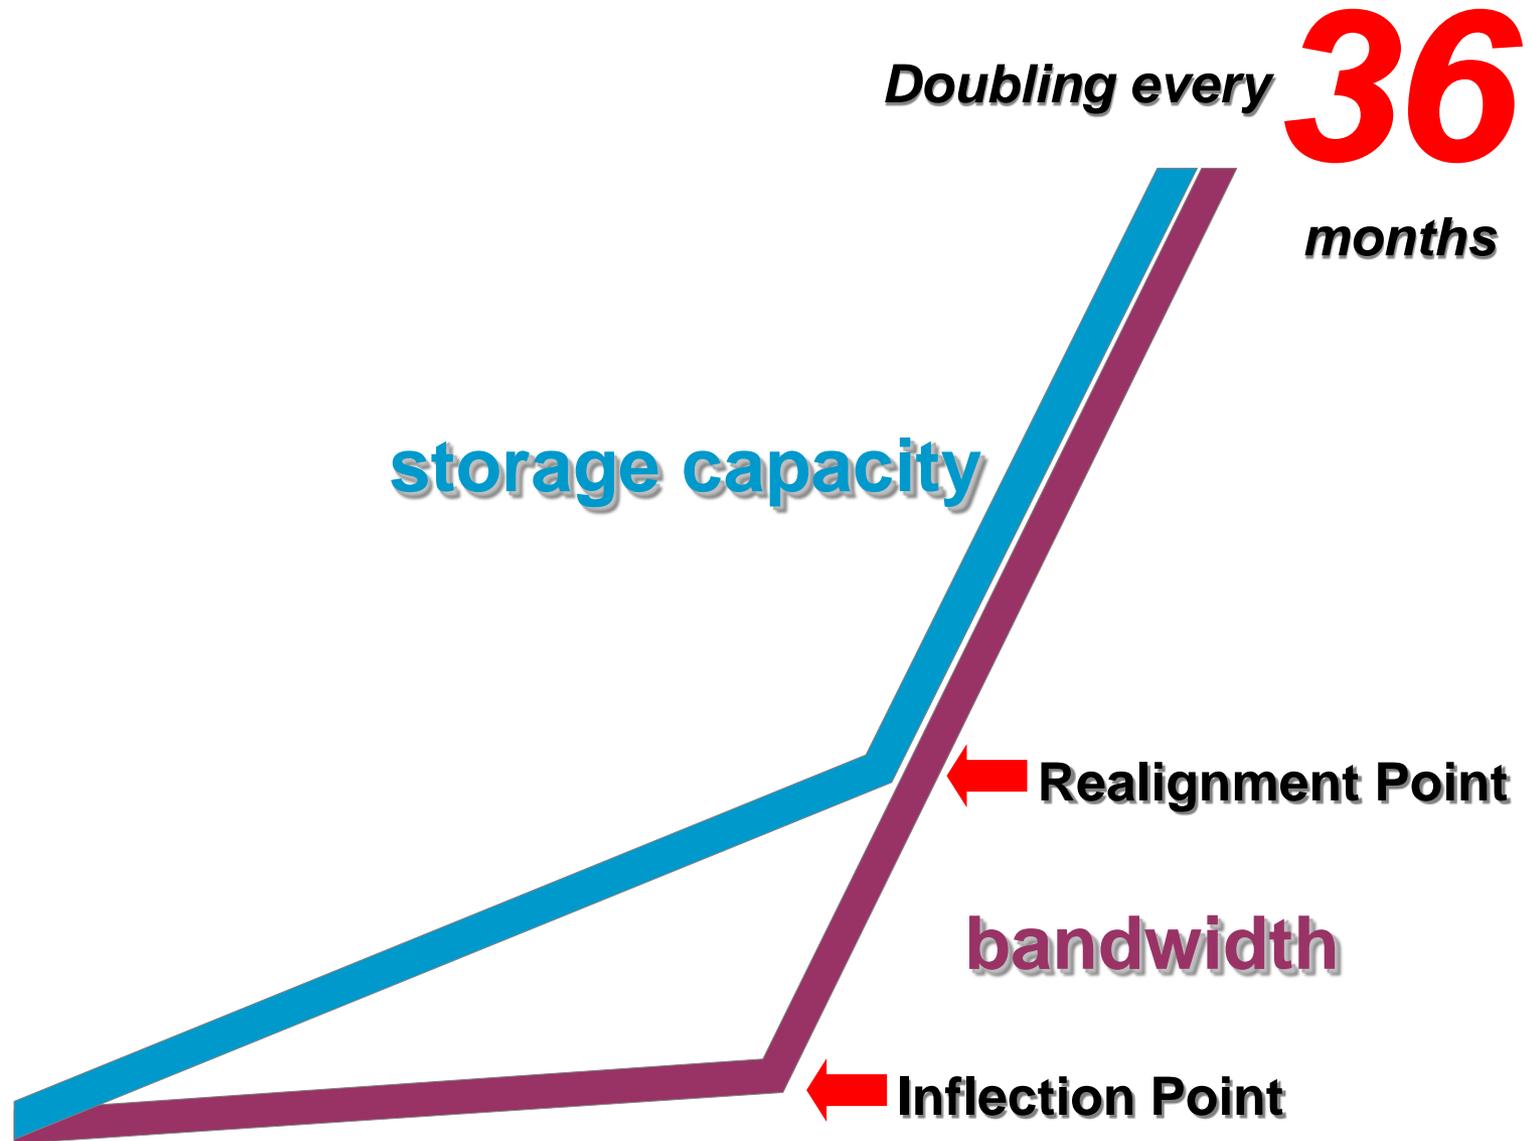
An EMC and IDC Study: <http://www.emc.com/leadership/programs/digital-universe.htm>

This Explosive Growth Of Data...



**...has Provided Us
with Benefits as well
as Challenges!**

But Storage Capacity and Bandwidth Have Now Become Interlocked





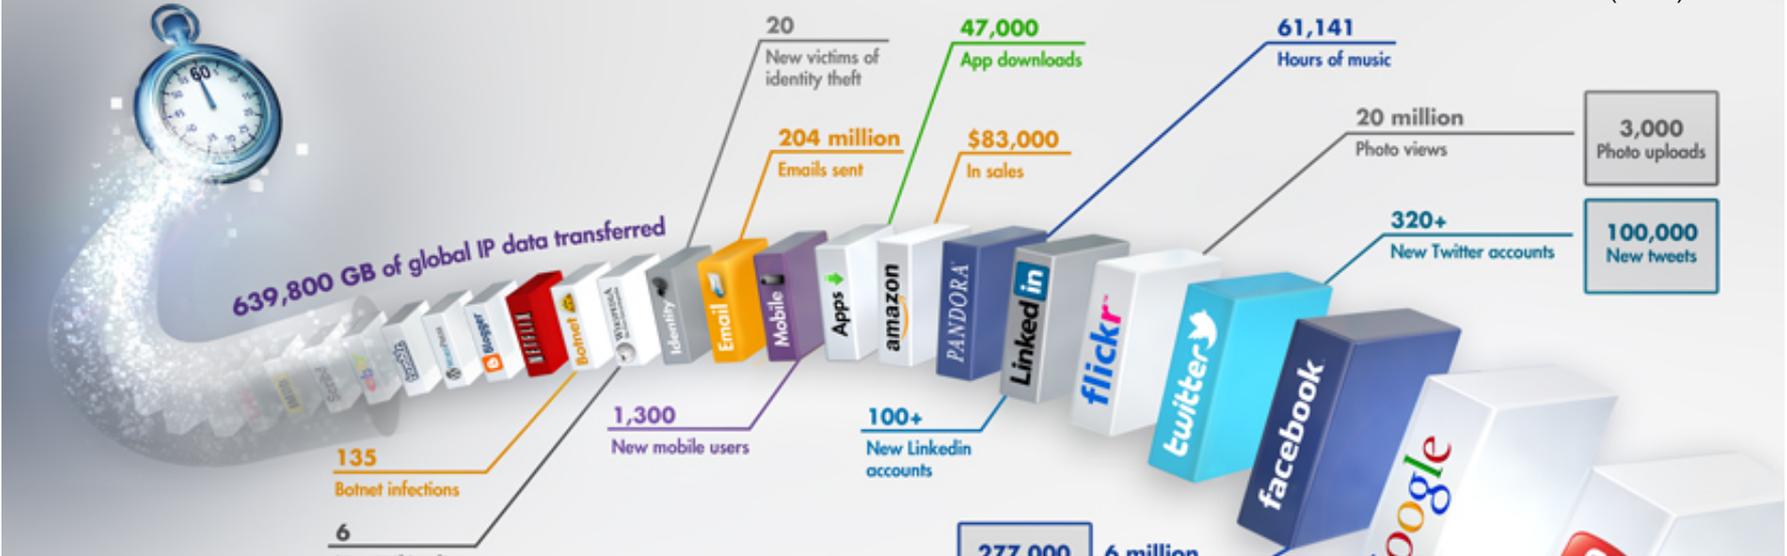
- Bandwidth Drivers
- Fibre Channel Speed Evolution

Trends and Roadmaps

What is driving bandwidth demand?

What Happens in an Internet Minute?

(2012)



And Future Growth is Staggering



Prolific Applications Server Virtualization

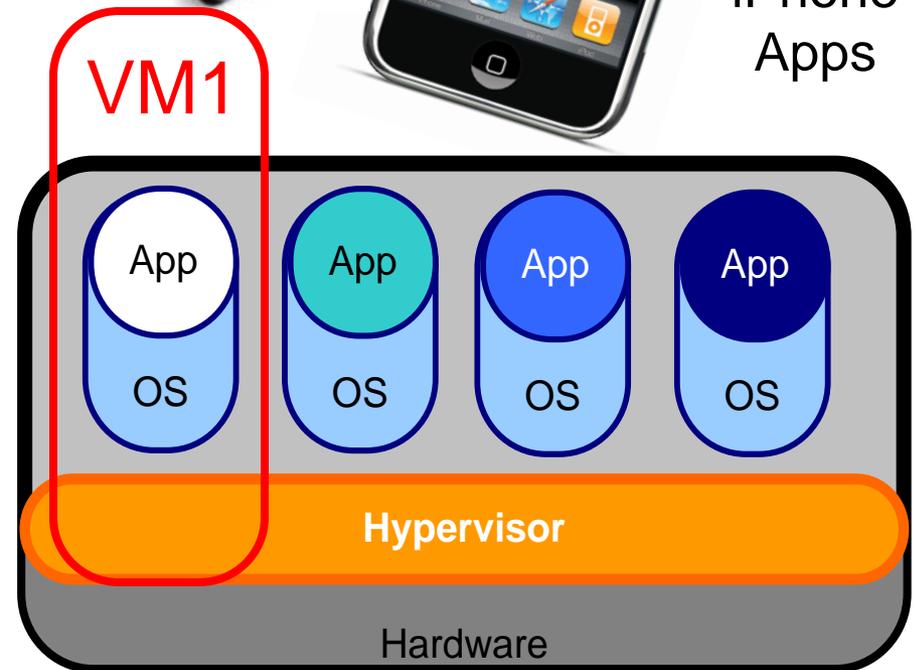
IDC believes that this market is poised for an explosion over the next few years!

http://www.idc.com/getdoc.jsp?containerId=IDC_P10666

- 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)
- Consumerization of IT and the continued explosion of the mobile device market is a huge driver



150,000
iPhone
Apps

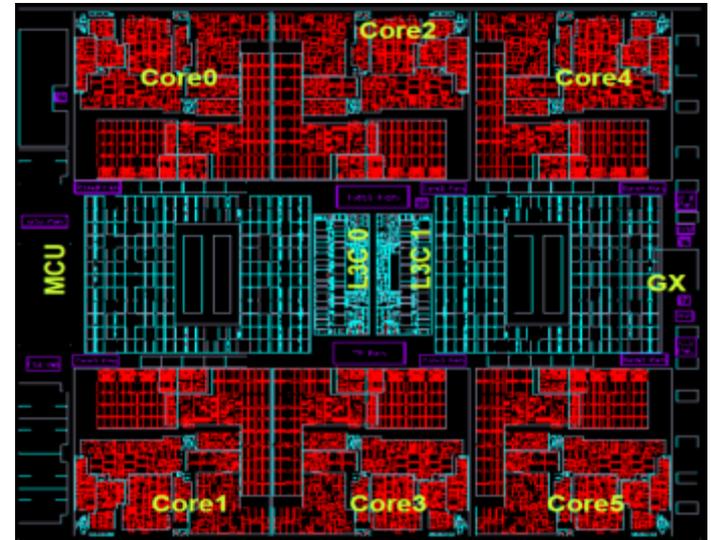


Faster Processors

- 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 1 to 100s of cores

Moore's Law says CPU speed should **double** every two years -- Reality has seen speed doubling **every year!**

IBM's zEC12 CP Chip



System zEC12 deployed a 5.5 GHz, six-core processor, versus the 5.2-GHz, quad-core processor of the z196.

This may be the world's fastest commercial processor.

Increased Memory in Servers

2012 was the year of very cheap computer and consumer electronics memory. **BUT...**
The death of some of the memory makers, and the withdrawal of others, means supply is much more controlled and prices have begun to rebound.

- 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 that demands faster CPUs

32GB RDIMM



Kingston Technology announced a 1TB thumb drive – 1/2013



US\$1,750 for the 512K version

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
 - Traditional spinning disk drive seek times are in the millisecond range
 - SSD seek times are virtually eliminated
 - SSDs are often referred to as Tier-0 storage while disk drives are Tier-1
 - Capacities are 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...

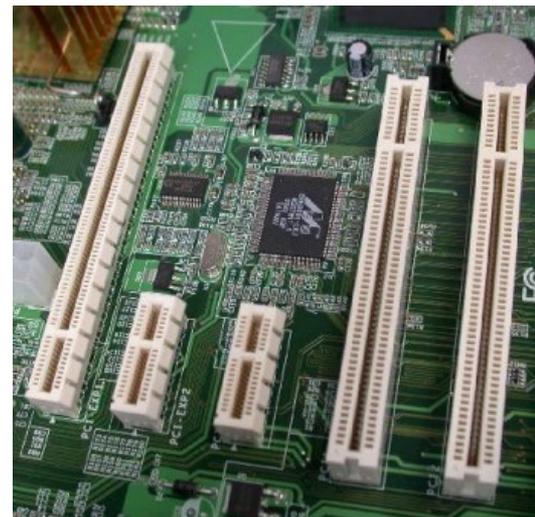
	Latency	Drive IOPS	Raid 0 Array IOPS*
7200 RPM SATA HDD	~12 ms	75 - 100	750-1,000
10K RPM SAS HDD	~9 ms	100 - 140	1,000-1,400
15K RPM SAS HDD	~6 ms	150 - 200	1,500-2,000
SSD	50-250 us**	40k-150k	50k-500k

- BUT...for example, Texas Memory Systems RamSan-630 storage system supports 500,000 IOPS and 8 GBps (64 Gbps) of throughput
 - Assumes 10 drives per array and Raid 0 which does not penalize writes so this is a best possible case
 - ** This is based on Flash memory and multiple parallel processing

PCIe Continues Ramping

System zEC12 has embraced the deployment of PCIe technology with its FICON Express8S cards.

- 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

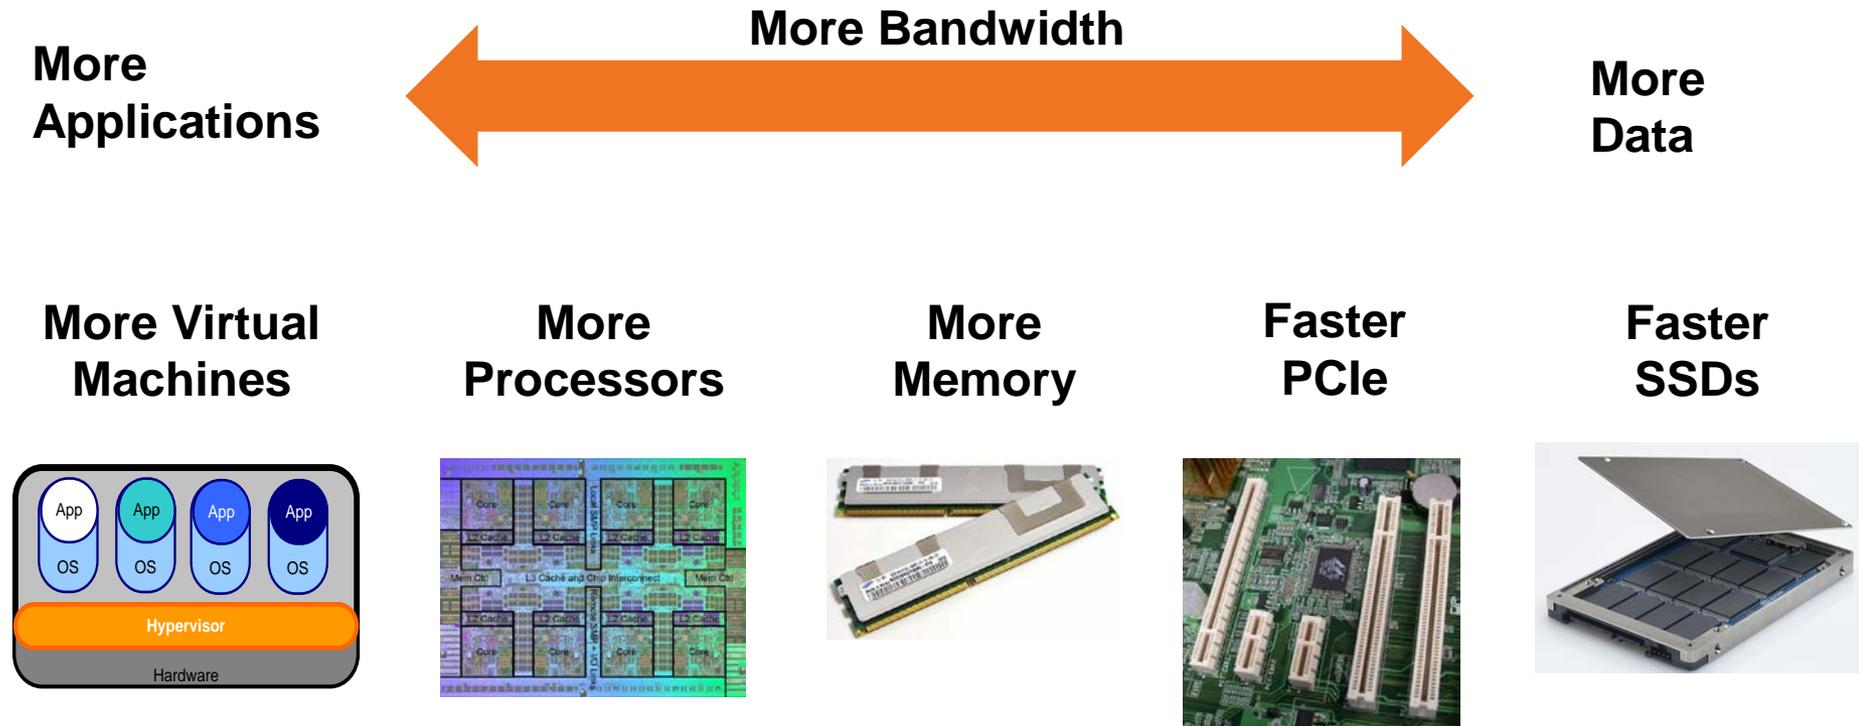


	Number of Parallel Lanes	Speed per Lane (Mbps)	Directional Bandwidth (Gbps)	Ports Supported
PCIe -1.0	4	250	8	1 – 8GFC
PCIe -1.0	8	250	16	2 – 8GFC
PCIe -2.0	4	500	16	1 – 16GFC
PCIe -2.0	8	500	32	2 – 16GFC
PCIe -3.0	4	1000	32	2 – 16GFC
PCIe -3.0	8	1000	64	4 – 16GFC

More Applications Drive more Bandwidth

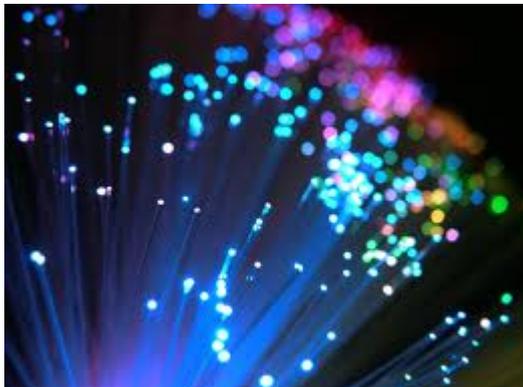
SAN and FICON Switching Devices Already Provide Connectivity At Up To 16Gbps Line Rates

- Gen5, 16G FC was designed for servers over the next few years that will use these technologies



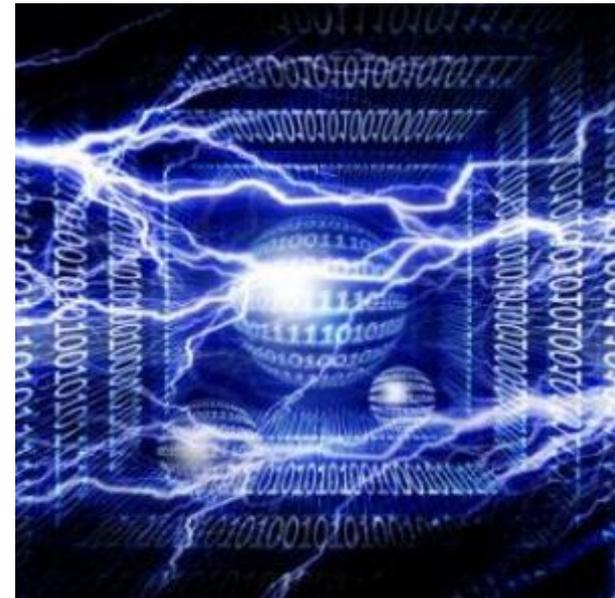
The Evolution of Fibre Channel Speeds

- Five generations of Fibre Channel have been delivered to the market (1Gb, 2Gb, 4Gb, 10Gb and 16Gb)
- Speed doubling about every 3-years
- Fibre Channel dominates the storage market



Even the anti-Fibre Channel bigots are finding out that Fibre Channel is harder to kill than your common-variety Zombie!

Other protocols have only eroded SAN FC dominance by 18%



Generations of Fibre Channel

The newest speed in Fibre Channel - Keep it Serial Stupid

Generation	1 st Gen	2 nd Gen	3 rd Gen	4 th Gen	5 th Gen	6 th 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

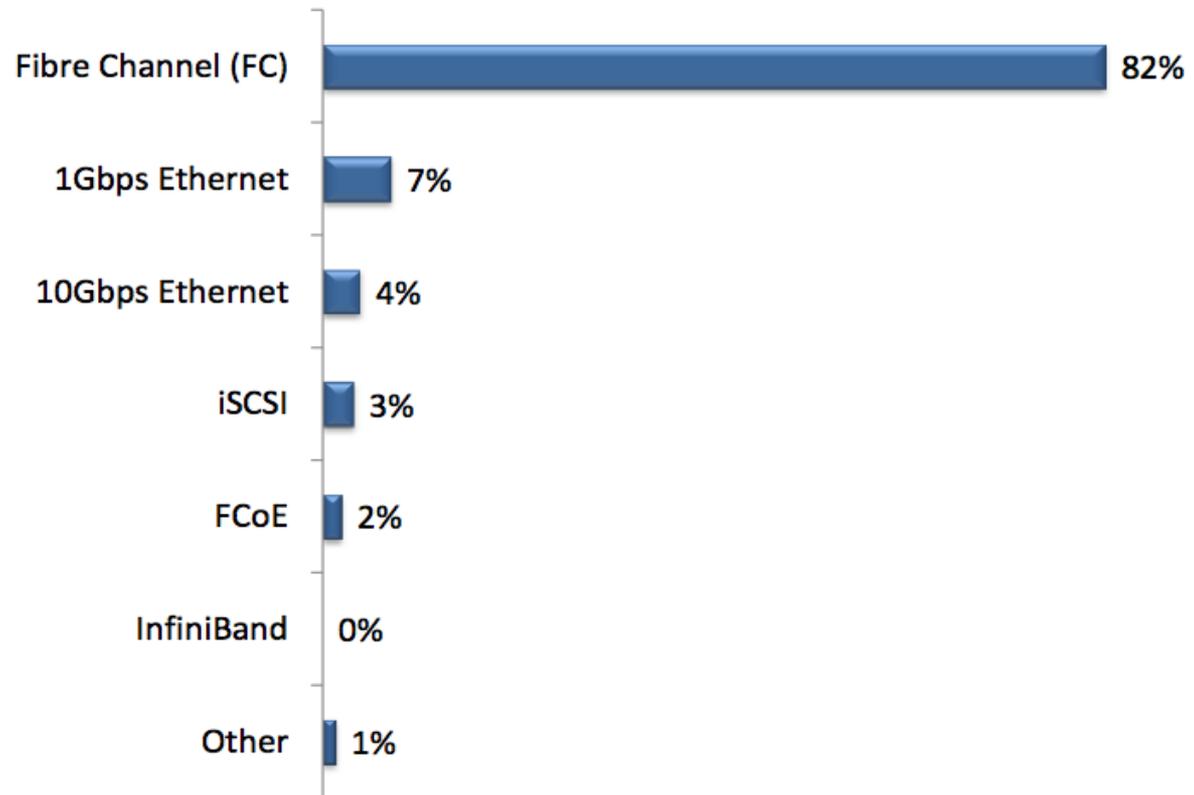


SFP spec provides capabilities for data rates up to 4.25Gbps, while the SFP+ supports up to 16Gbps today.

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%.

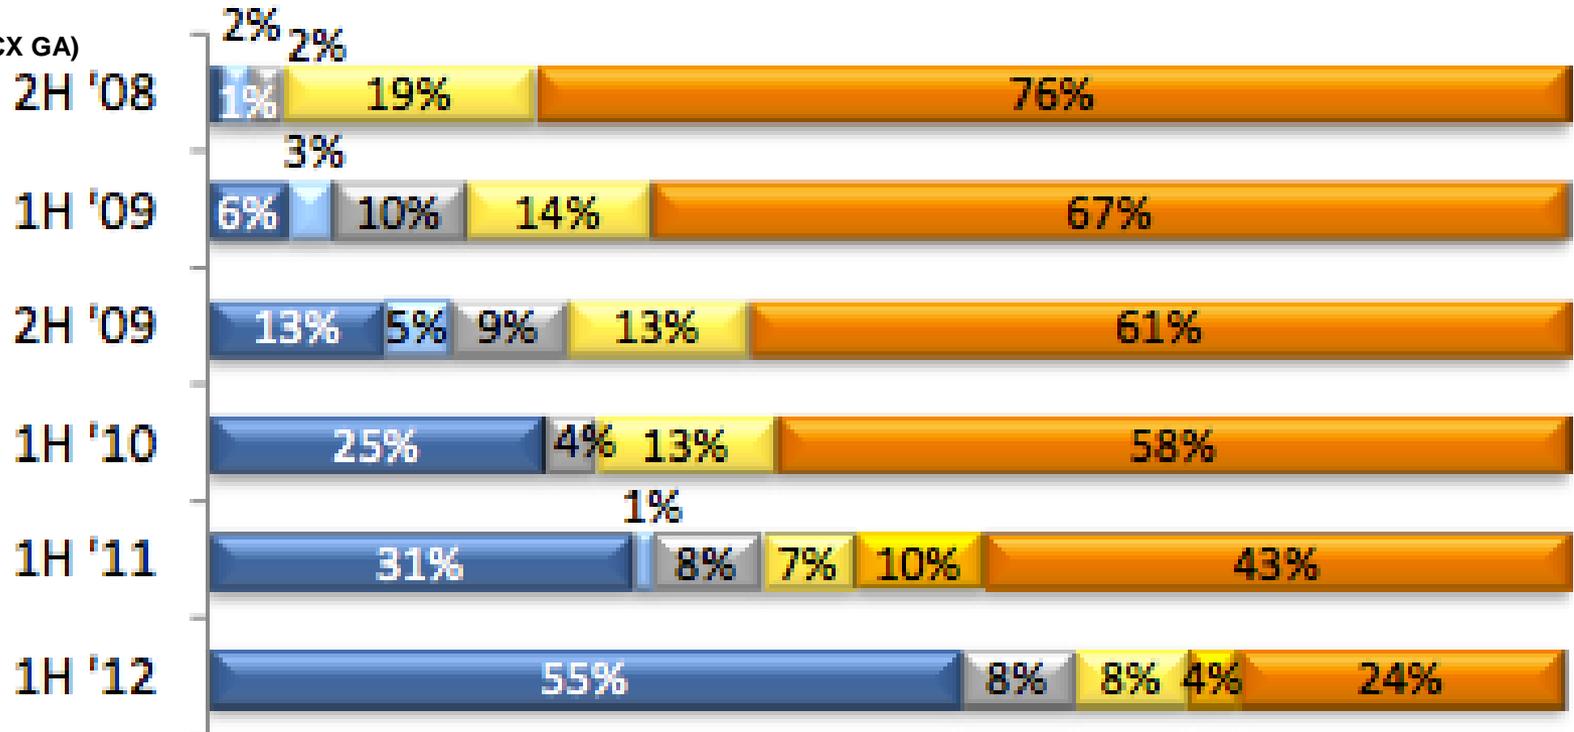


Source:
The InfoPro, 1H CY 2012

8Gbps Fibre Channel Implementation

Implementation Roadmap*

(Brocade DCX GA)

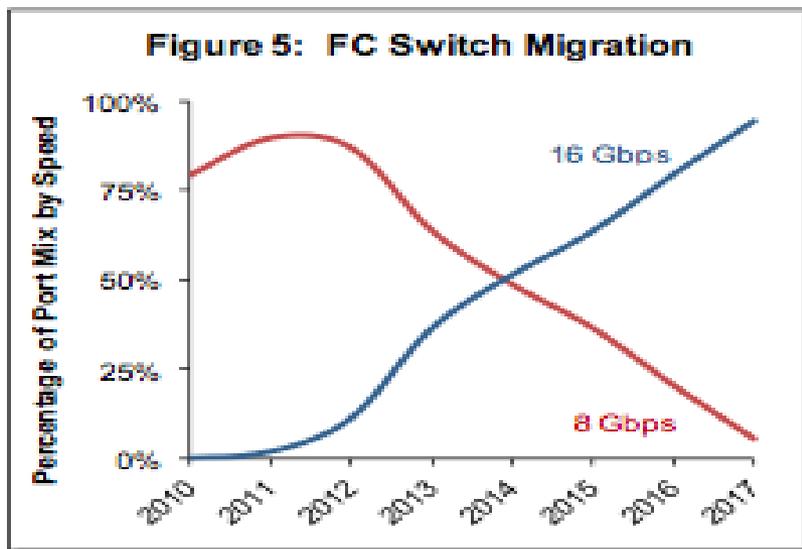


- In Use Now
- Near-term Plan (In Next 6 Months)
- Past Long-term Plan (Later Than 18 Months Out)
- In Pilot/Evaluation (Budget Has Already Been Allocated)
- Long-term Plan (6-18 Months)
- Not in Plan

Source: The InfoPro, 1H CY 2012

Past and Forecast Adoption of 16Gb FC

Source: Dell Oro, Q1 2013



- FC upgrades are in 2x increments, now on its fifth generation in the last dozen years (1, 2, 4, 8 and 16Gb FC)

- From the time that a new speed is ratified at the standard committee, it usually takes about three years for a sizable portion of new sales to be on the new generation, as was seen in the last chart.
- The speed of adoption on switches is a little faster, since the cost and performance of new speeds make sense for inter-switch links, and switch vendors can pull users into the new generation through pricing and end-of-life of previous generations.

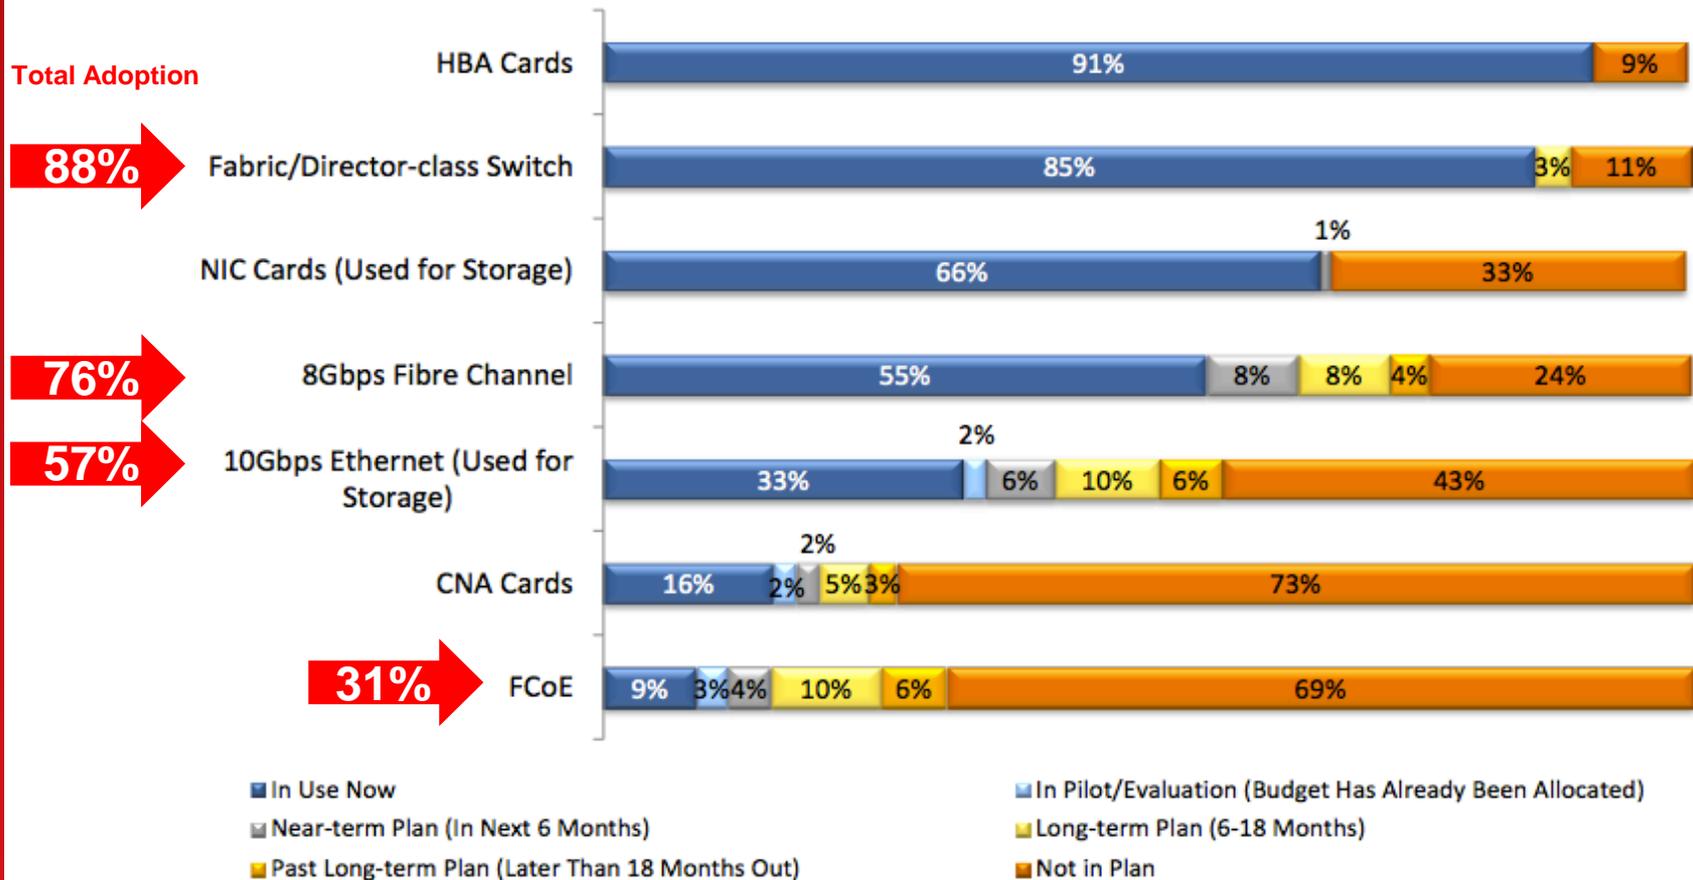


What's Ahead?

- Fibre Channel Roadmap
- FCoE

At 16G and Beyond

Storage Networking: Technology Roadmap



Source:
The InfoPro, 1H CY 2012

The Benefits of Gen5, 16GFC

About 45% of Brocade FC Switch Sales (2012-2013)

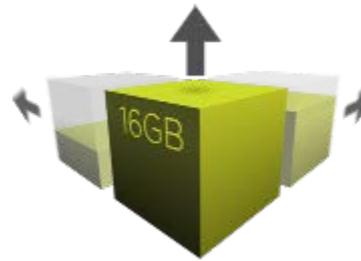
- 16GFC is 100% faster than 8GFC and 40% faster than 10GbE and leads to these benefits:
 - Higher performance lets servers process more data
 - Fewer FC links can do the same job
 - Easier cable and device management
 - Less power consumption per bit



Innovation
(Forward Error Correction)



Performance
(64b/66b)



Scalability
(Inter-Chassis Links)



Environment Friendly
(Less energy per data bit transferred)

Characteristics of Gen5, 16GFC

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

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

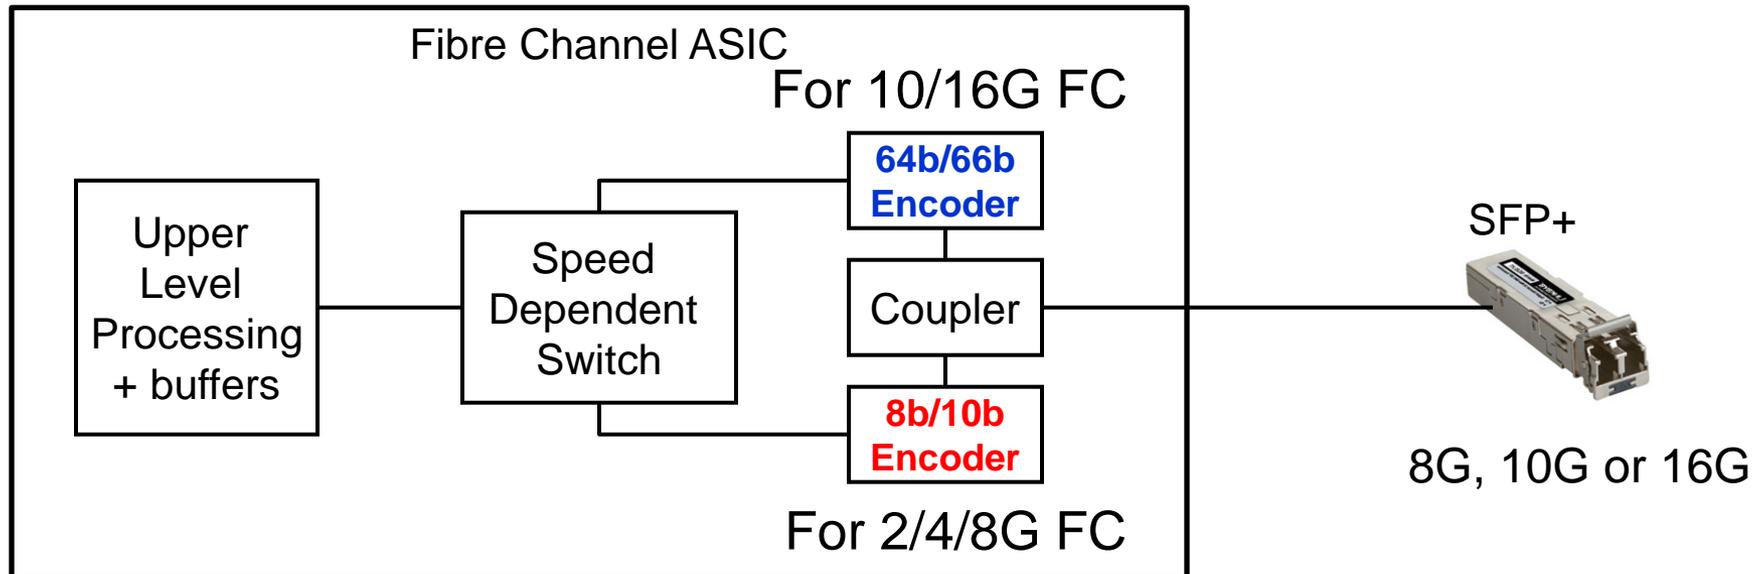
Speed Name	Throughput (MB/sec)	Line Rate (Gbps)	Encoding	Retimers in the module	Transmitter Training	OM1/2/3/4 Link Distance (meters)
1GFC	100	1.0625	8b/10b	No	No	300/500/860/*
2GFC	200	2.125	8b/10b	No	No	150/300/500/*
4GFC	400	4.25	8b/10b	No	No	50/150/380/400
8GFC	800	8.5	8b/10b	No	No	21/50/150/190
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 did not standardize distances for OM4 fiber for 1/2/10Gb FC

Speed Negotiation Is Now Much Smarter

Dual Codics, controlled by ASIC, per port

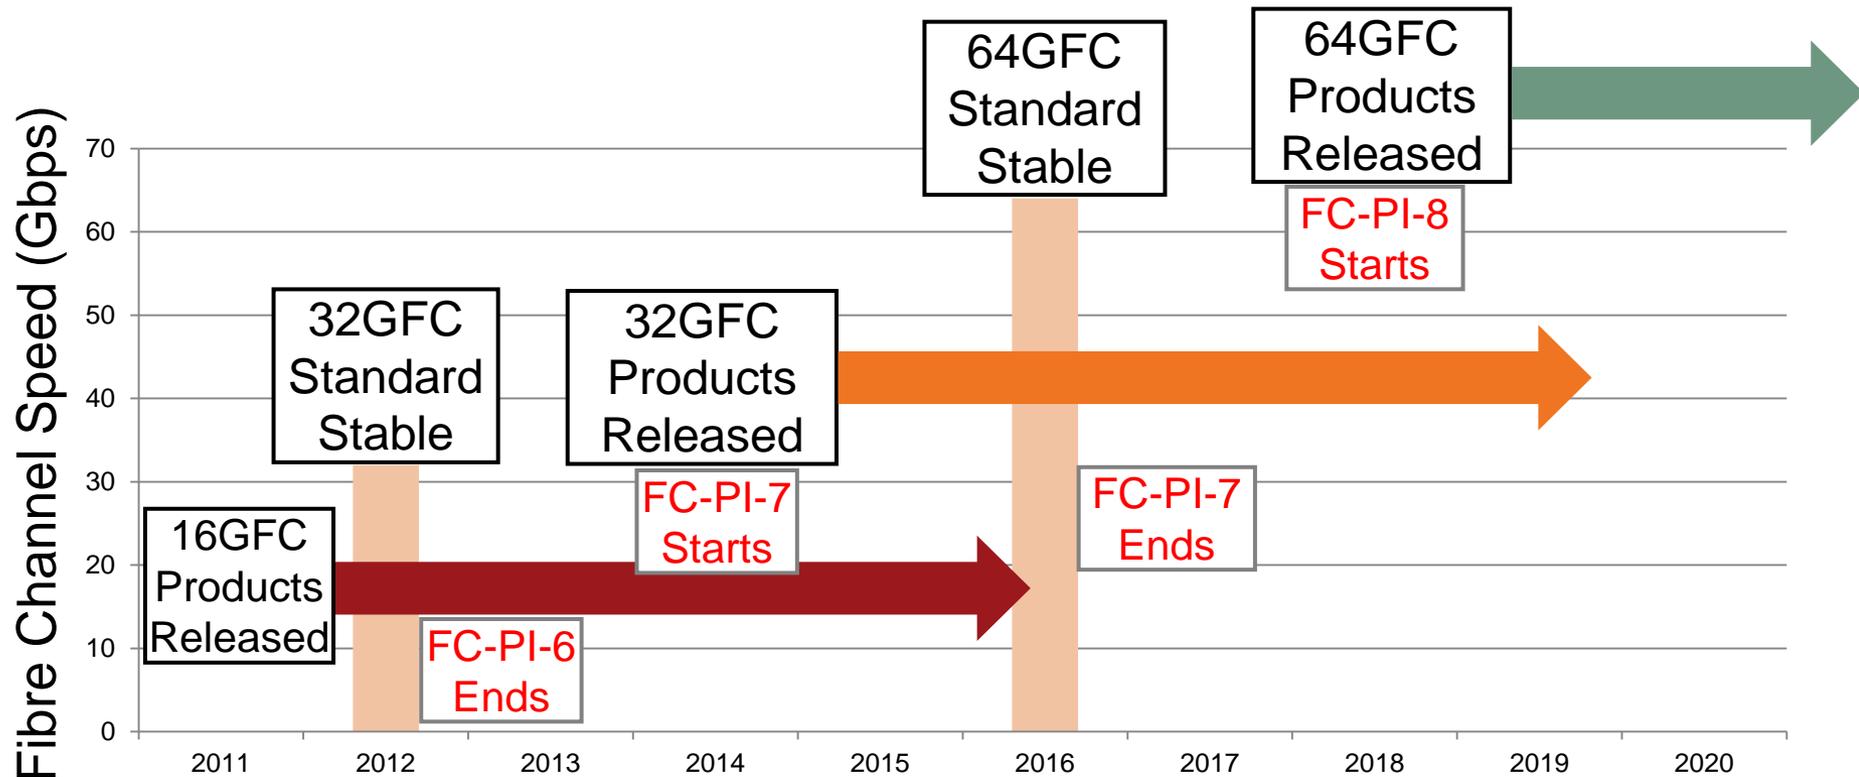
- For the first time, 64b/66b data encoding supports autonegotiation
- During speed negotiation, the speed dependent switch routes the initialization sequence to the appropriate encoder
 - 64b/66b for 10/16G FC
 - 8b/10b for 2/4/8G FC
 - The coupler sends the signals from one of the encoders to the SFP+



Fibre Channel Speeds Double Through 2020

Keeping pace with the Digital Universe

- 32G FC Standard Stabilizing
- 64G FC Standard Starting
- Roadmap up to 128G currently



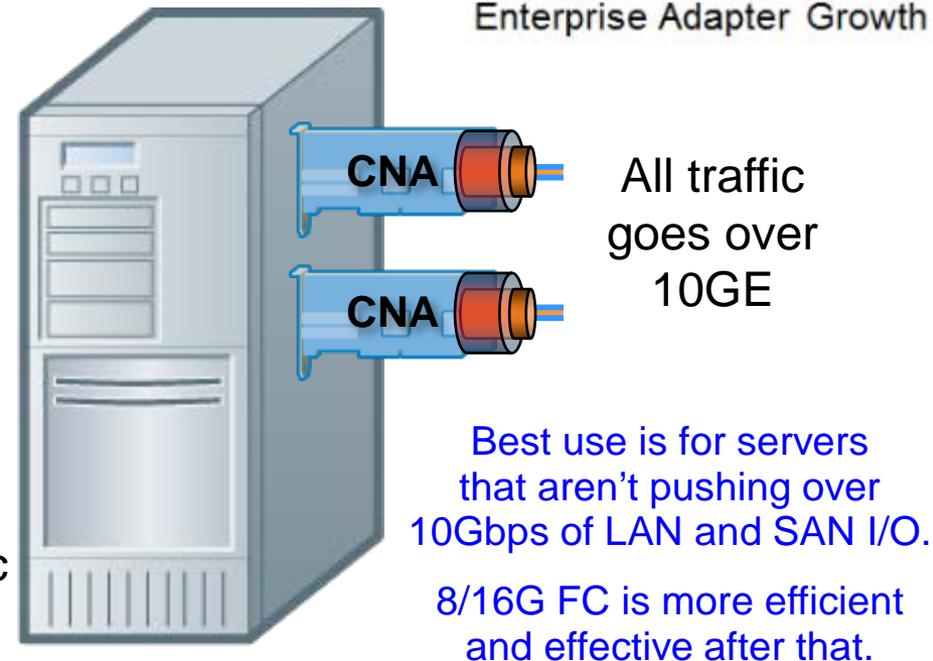
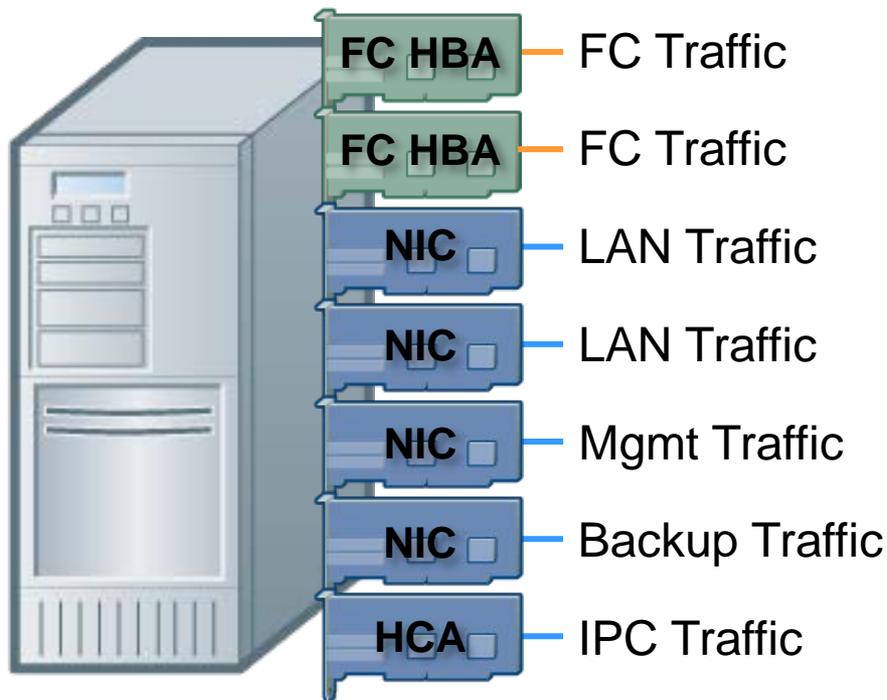
A Mainframe Guy Looks at Converged Fabrics

Why? Reduces Adapters and Cables!

- Fewer CNAs (Converged Network adapters) instead of NICs, HBAs and HCAs
- Limited number of interfaces for Blade Servers / Rack Mounted Servers



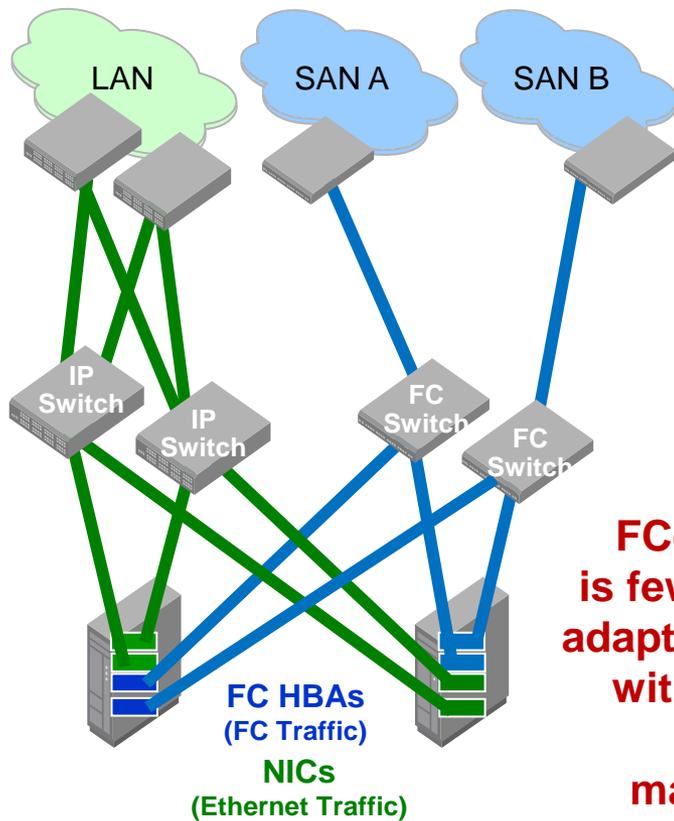
Enterprise Adapter Growth



A Mainframe Guy looks Primary FCoE Use Case

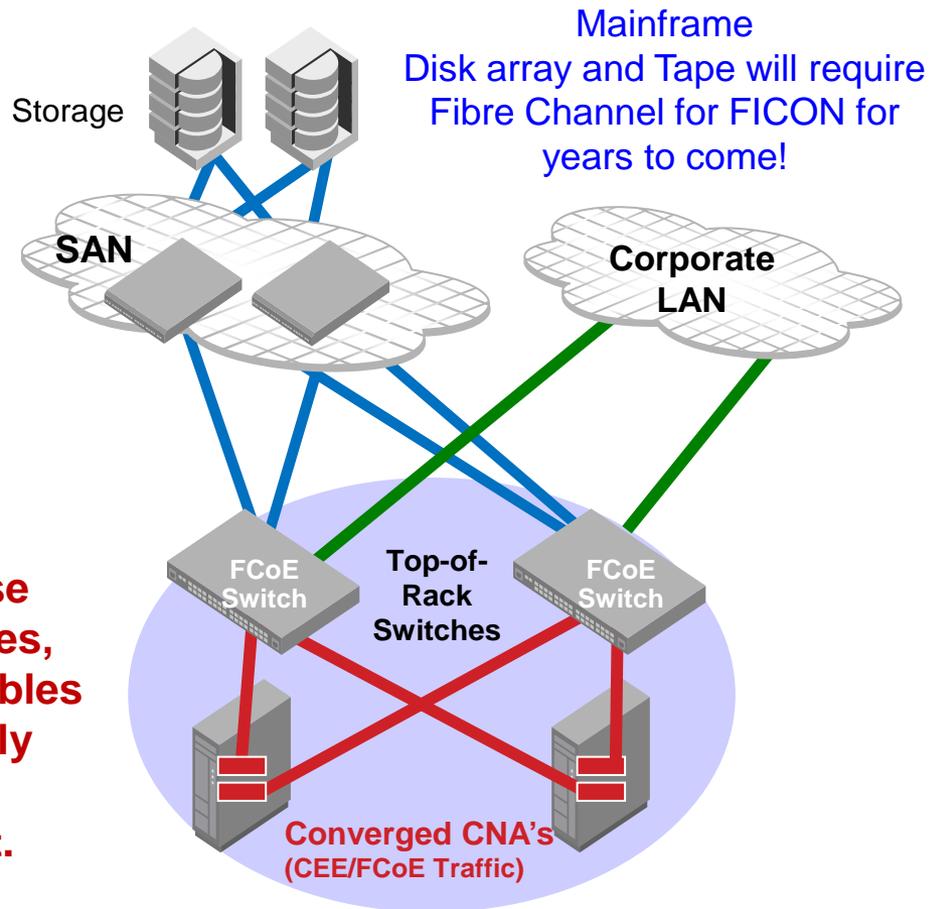
IBM has no plans for FCoE on the mainframe

- Before Unified I/O



FCoE use case is fewer switches, adapters and cables with potentially simpler management.

- After Unified I/O



— FC — Ethernet — FCoE and CEE

Current Platforms Supporting FCoE

Brocade supports FCoE on switching devices

- Blade servers
 - Cisco
 - IBM
 - HP
- Pizza Box PCs
- IBM Power System Announced Feb 2013, (FCoE and 16G FC)¹.
- NetApp Filers
- EMC VMAX (for SRDF only currently)

10G for LAN On Motherboard (LOM) Makes FCoE more attractive

10GbE Deployment Costs

SFP+ Top of Rack

Fiber

\$243 per port



\$778



\$256 per port



\$1277 per port TOR

Direct Attach

\$243 per port



\$135



\$256 per port

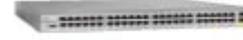


\$634 per port TOR

10GBASE-T Top of Rack

10GBASE-T NIC

\$305 per port



\$5



\$256 per port



\$566 per port TOR

10GBASE-T LOM

\$305 per port



\$5



LOM

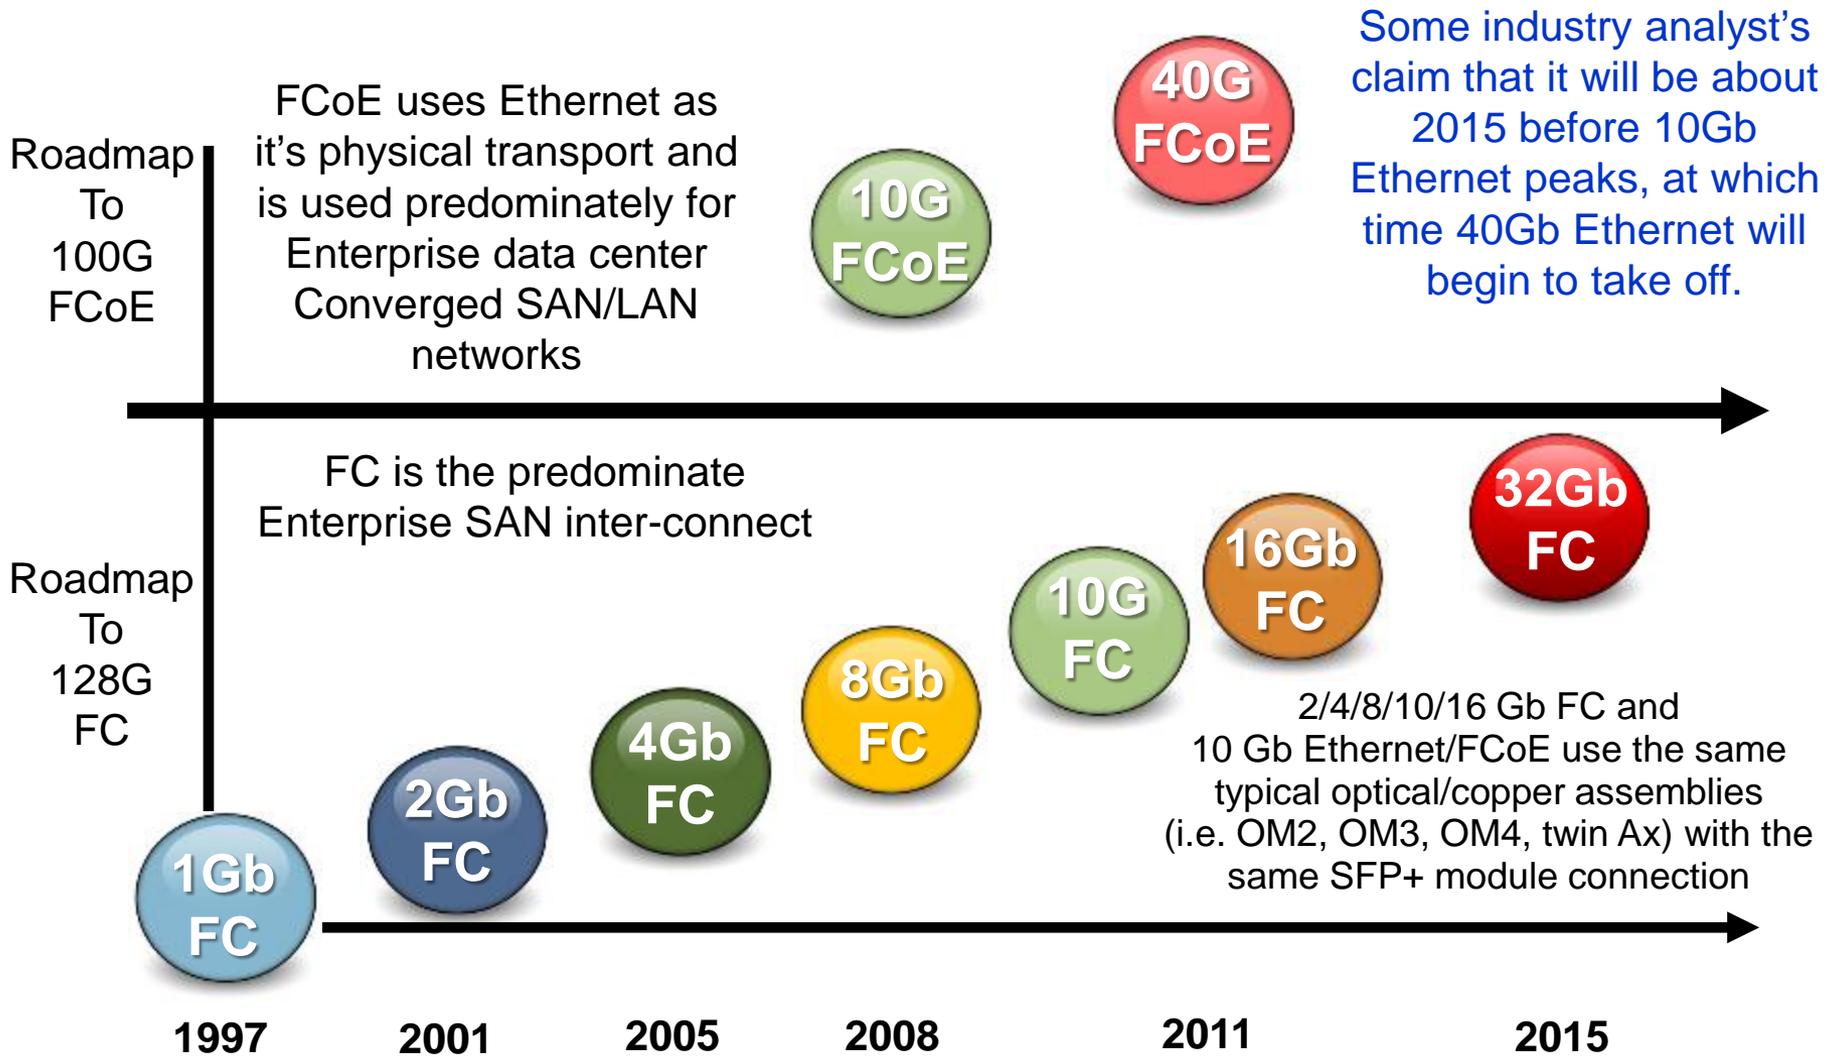


\$310 per port TOR

Source: List prices as of Feb. 2011 from the following sources:
 1. Cisco Technical note SFP+ 10G - hwwp.com
 2. Cisco 2232T1 - CDF.com
 3. Cisco 2232T1 (10GBASE-T patch) - Technonet
 4. Intel 8000 and Intel 8001 - hwwp.com
 5. Multirate cable, Cisco SFP+ transceiver, Fiber cable - CDF.com



FC and FCoE Bandwidth Roadmaps

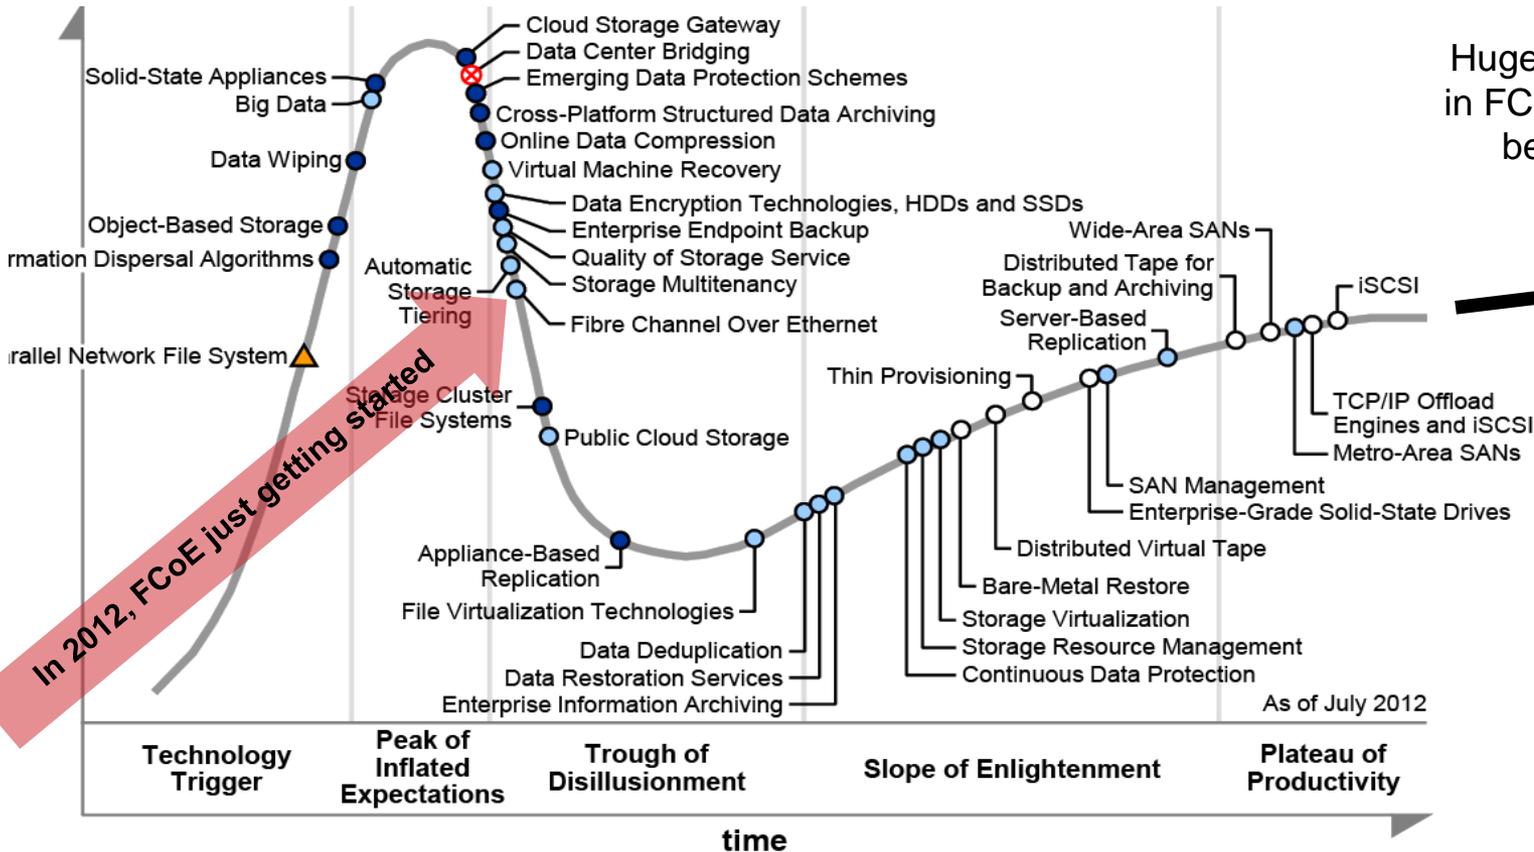


Some industry analyst's claim that it will be about 2015 before 10Gb Ethernet peaks, at which time 40Gb Ethernet will begin to take off.

TOTAL Investment Protection!

Gartner 2012

Storage Technology Hype Cycle Curve



Huge Investments in FC have already been made.

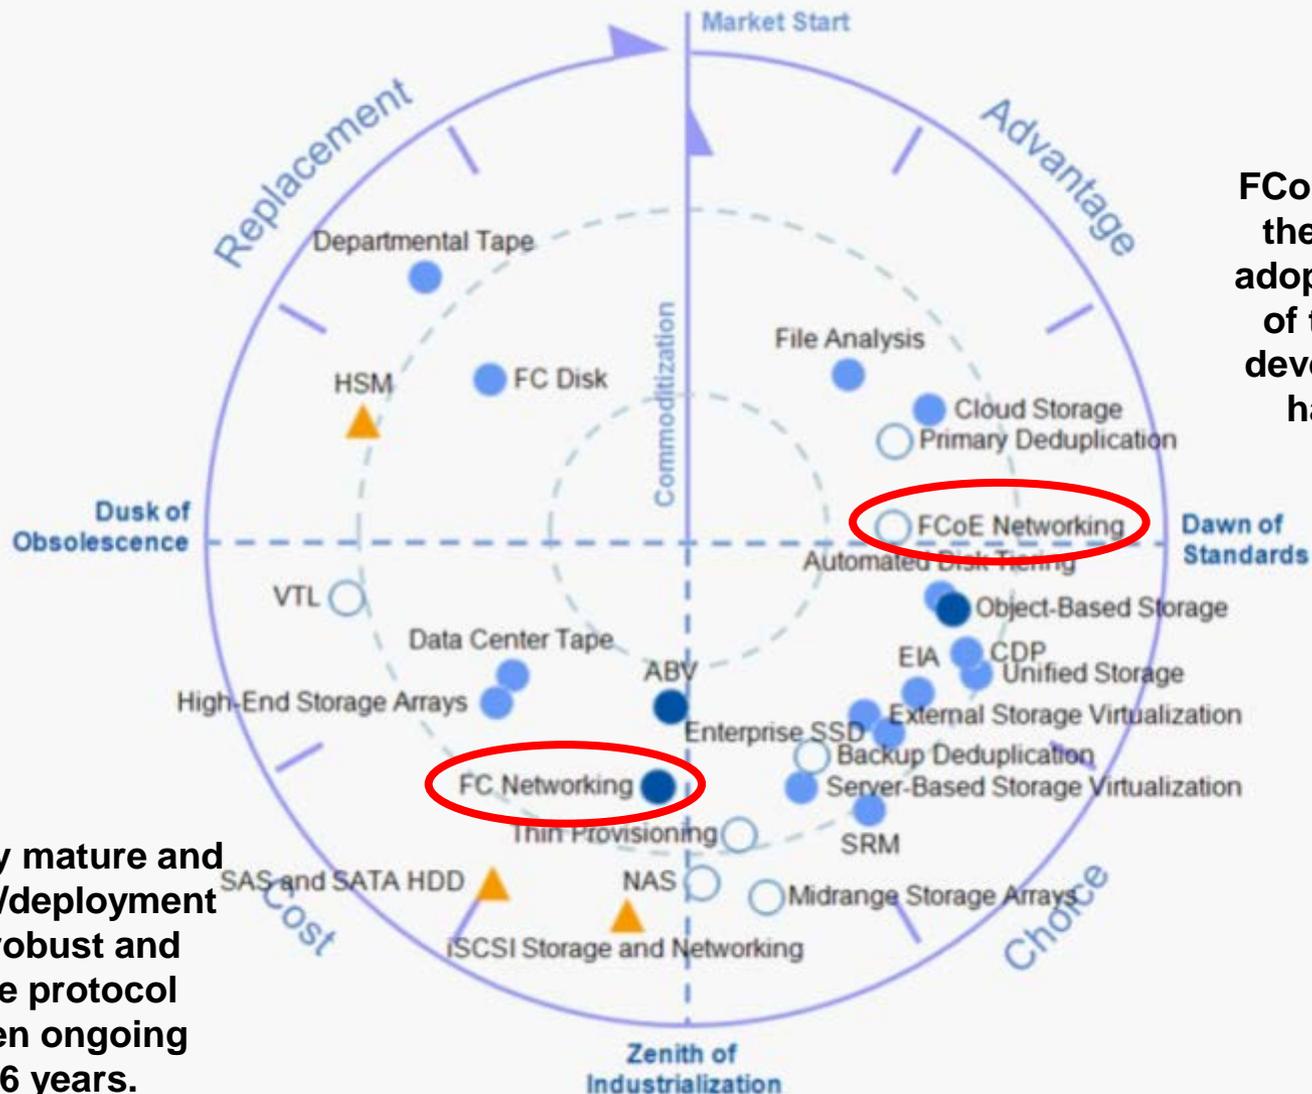


Plateau will be reached in:

○ less than 2 years ● 2 to 5 years ● 5 to 10 years ▲ more than 10 years ⊗ obsolete before plateau

Gartner 2012

IT Storage Market Clock

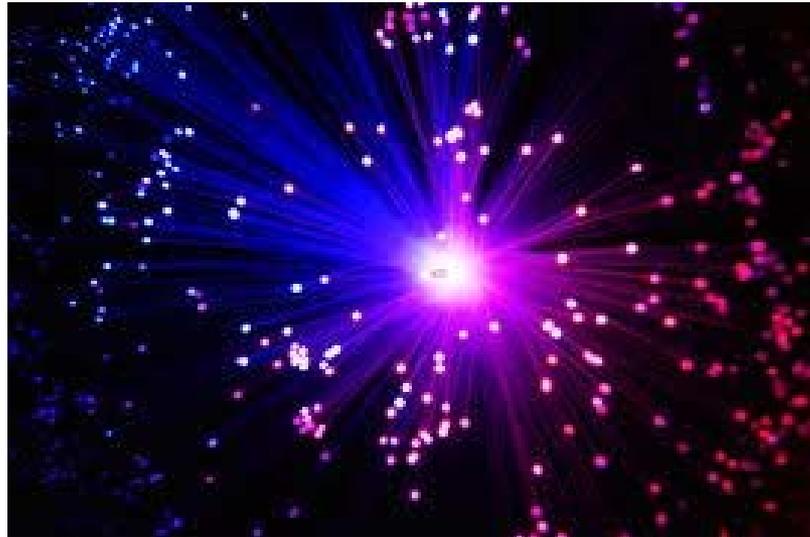


FCoE is just clearing the standards and adoption/deployment of this young and developing protocol has just begun.

FC is very mature and adoption/deployment of this robust and capable protocol has been ongoing for 16 years.

Reflection / Discussion

- These are all interesting data-points but – for your data center consider:
 - When do your applications need greater than 8G ?
 - When will your servers have higher speed availability ?
 - When will your Disks / Tapes / VTLs have higher speeds ?
 - Which Technology(s) will have highly available solutions sooner ?



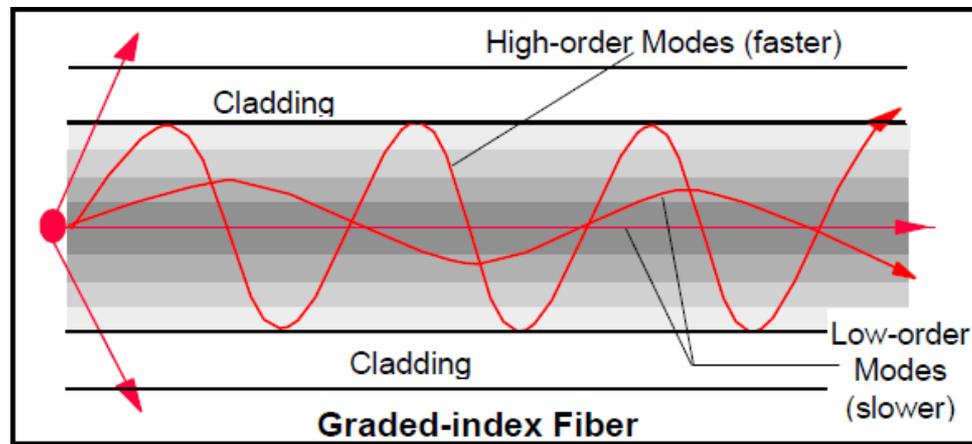
- Modal Dispersion
- Light in Flight
- Measuring Light Signals

Let's talk about light

FC Storage Networking Terminology

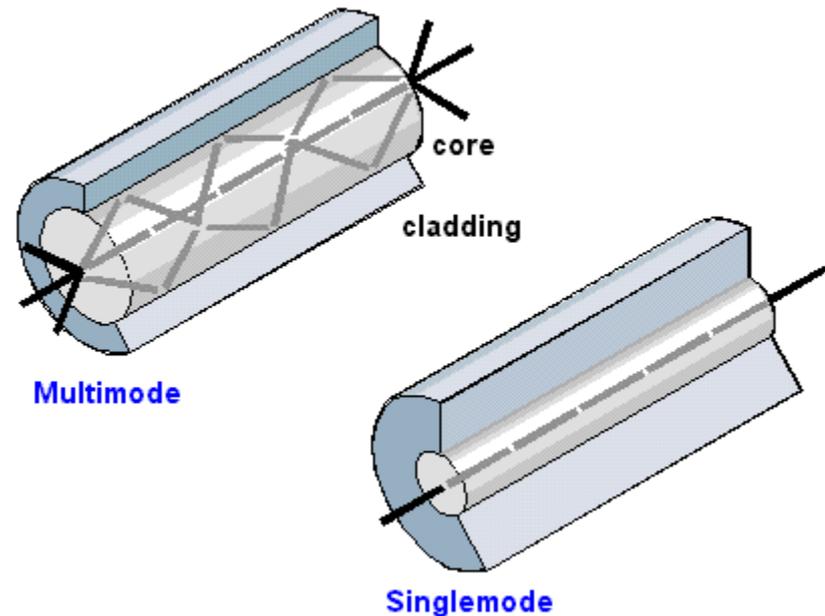
Fiber Channel Links

- **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.



Graded-index Fiber

OM3 and OM4

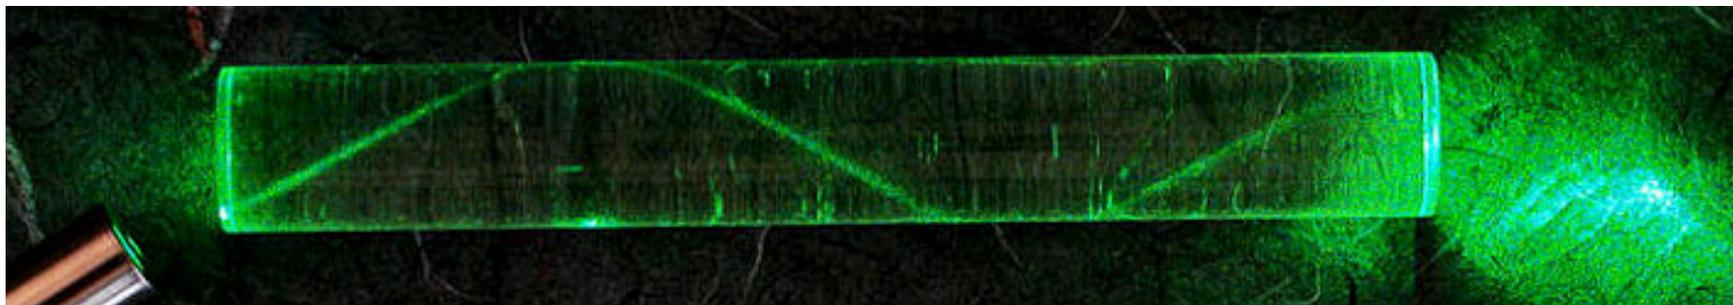


FC Storage Networking Terminology

Fiber Channel Links

- **Photo 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

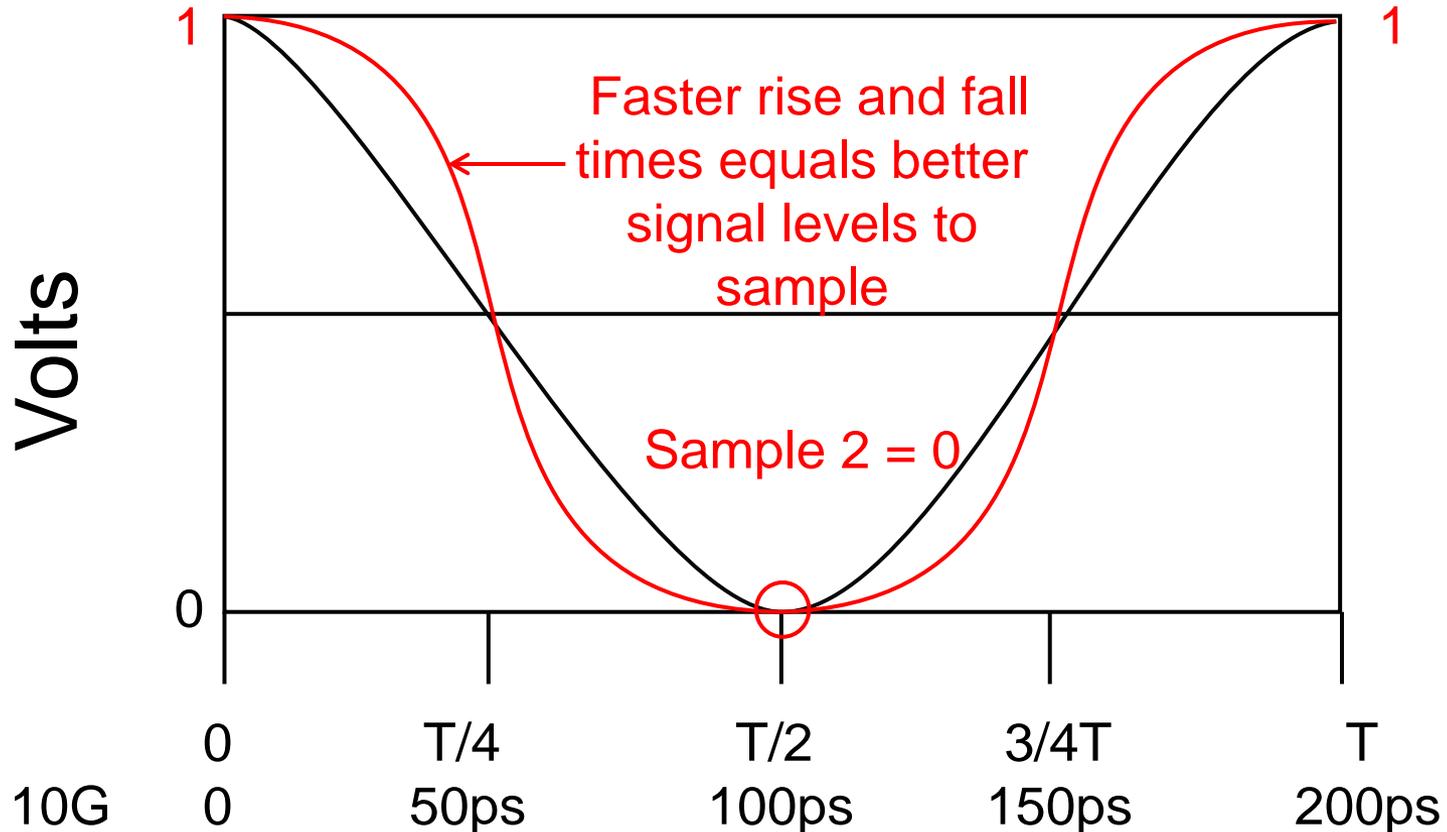
- There is now a camera that can take a trillion frames per second. Below is a photo of light in flight from a laser pointer. The distance of the light shown below is the total distance that light travels in atmosphere in a Femtosecond.
- A femtosecond (10^{-15} seconds) is one quadrillionth, or one 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



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
- Gartner, Inc.:
 - Hype Cycle for Storage Technologies, 2012 (5 July 2012)
 - IT Market Clock for Storage, 2012 (6 September 2012)
 - <http://www.gartner.com/technology/home.jsp>
- TheInfoPro
 - <https://451research.com/about-theinfopro>

Three Great FICON Resources For You To Use

Visit Brocade's Mainframe Blog Page at:

<http://community.brocade.com/community/brocadeblogs/mainframe>

Almost 250,000 hits

Also Visit Brocade's New Mainframe Communities Page at:

http://community.brocade.com/community/forums/products_and_solutions/mainframe_solutions

You can also find us on Facebook at:

<https://www.facebook.com/groups/330901833600458/>

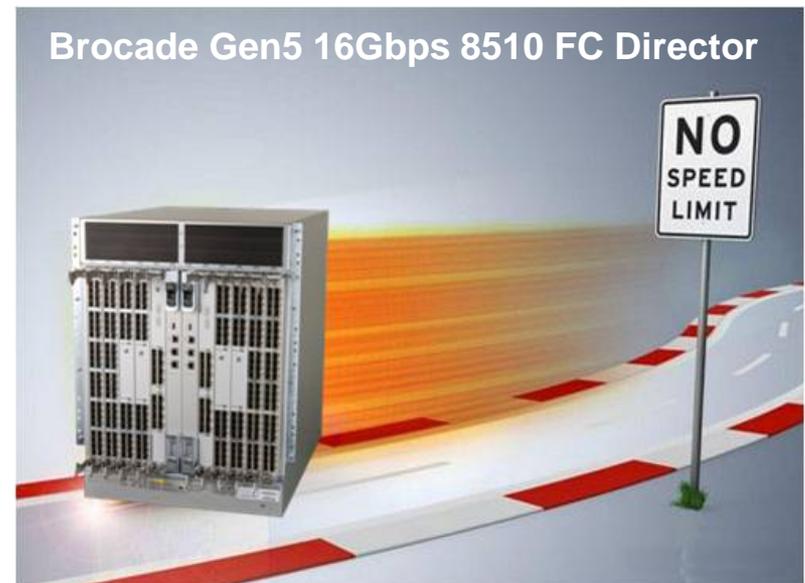
Brocade can provide “in-depth” I/O Training

We Provide Many “Deeper Dive” Seminars

- **Customer Mainframe 101 Fundamentals – 1-day** (great for people new to mainframe)
- **Customer Mainframe 201 Advanced Topics – 1-day** (good for new people and as a refresher)
- **SAN basics for Mainframe Professionals – 1 or 2 days**
- **Advanced FICON Seminar for Customers and Partners – 1-day**
- **FICON Advice and Best Practices Seminar – 1 or 2 days**
- **8G or 16G FICON Cascading Seminar – 1 day**
- **Protocol Intermixed Fabrics – 2 hours to 1-day**
- **FICON Fabric Design Considerations – 2 hours to All Day**
- **Joint IBM-Brocade Mainframe I/O and FICON Performance Training Seminar – 1 / 2 days**
- **CAF200 class: Preparation Class for FICON Certification (BCAF) – 2 days**

Our FICON Certification Class Expects That You Already Have Good Mainframe Experience And IS NOT Designed To Extensively Train You In FICON Concepts!

THANK YOU from all of us at Brocade!



- **Canada CMG Group, Toronto, April 17, 2013**
- Buzz Fibrechannel - To 16G and Beyond!

David Lytle, BCAF

dlytle@brocade.com

www.brocade.com