

16GFC Sets The Pace For Storage Networks

Scott Kipp – Brocade

Mark Jones – Emulex

August 30th, 2011

To be presented to the Block Storage Track
at 1:30 on Monday September 19th

❑ What is 16GFC?

❑ What is Driving 16GFC

- ❑ Prolific applications, server virtualization, multi-core processors, more memory, solid state drives, PCIe 2.0, traffic aggregation, VDI

❑ Benefits of 16GFC

- ❑ Higher performance leads to fewer links, easier cable management, less power consumption per bit

❑ Application of 16GFC

- ❑ 16GFC is needed for ISLs, high bandwidth applications, data migration and VDI

Generations of Fibre Channel

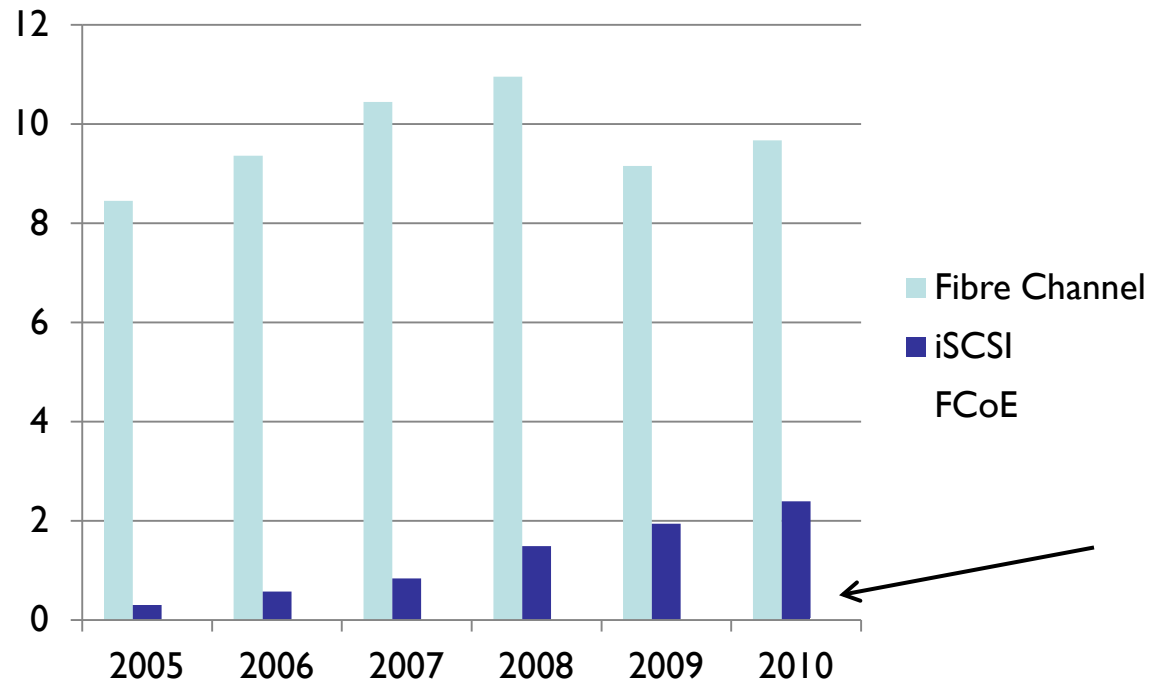
The newest speed in Fibre Channel - Keep it Serial Stupid (KISS)

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+
Data Rate(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



FC External Storage Dominates

External
Storage
Sales (\$B)

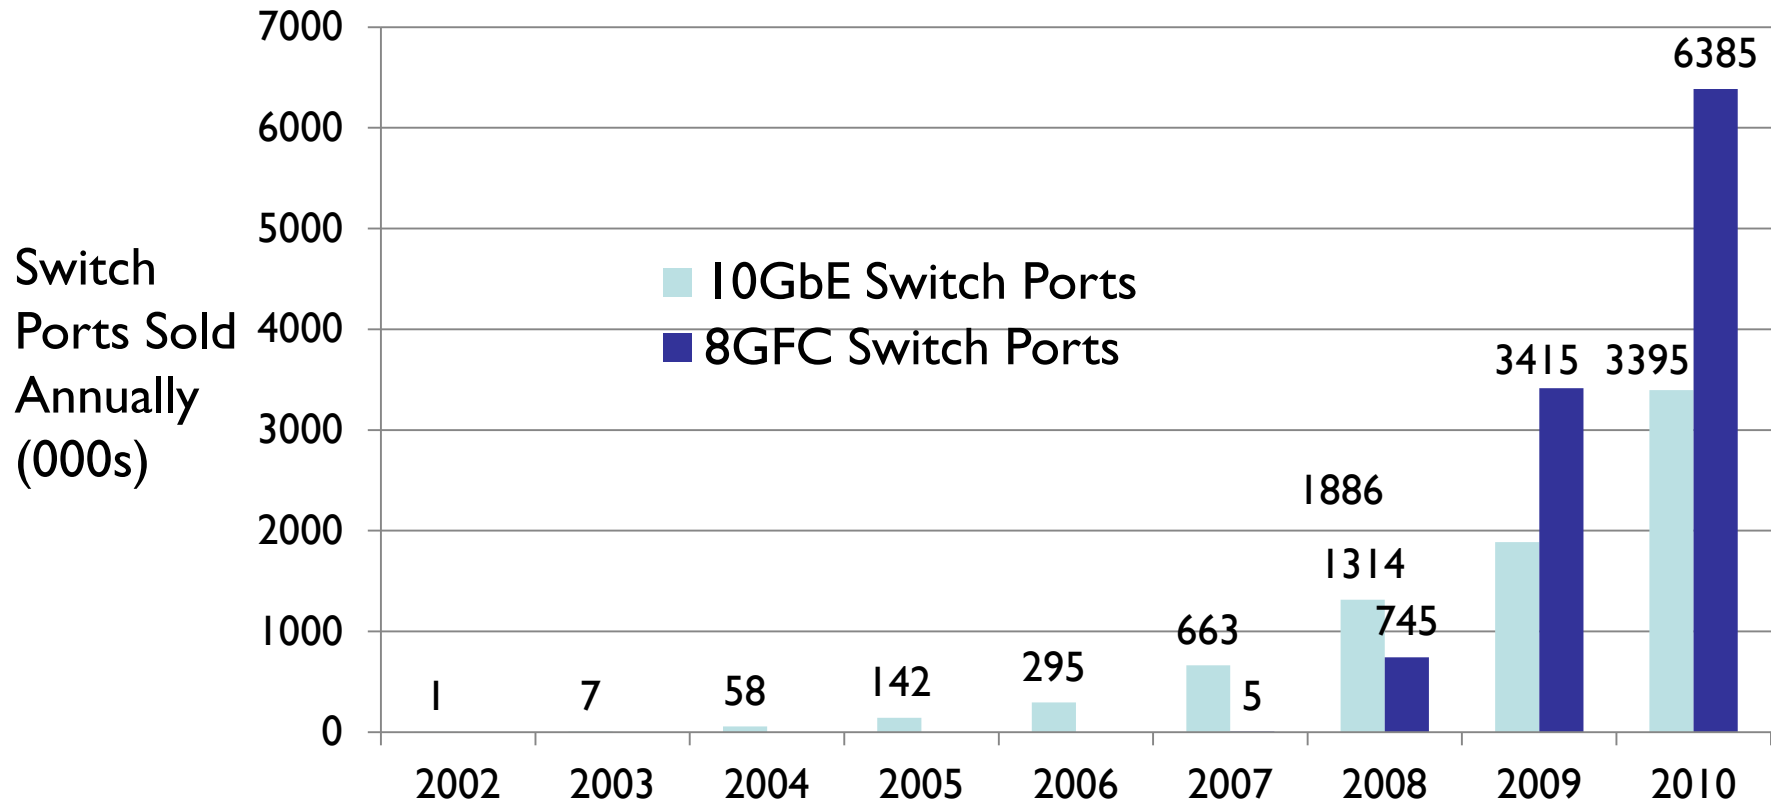


Source: IDC

FCoE storage
barely visible in
2010

8GFC Switch Ports Outsell 10GbE

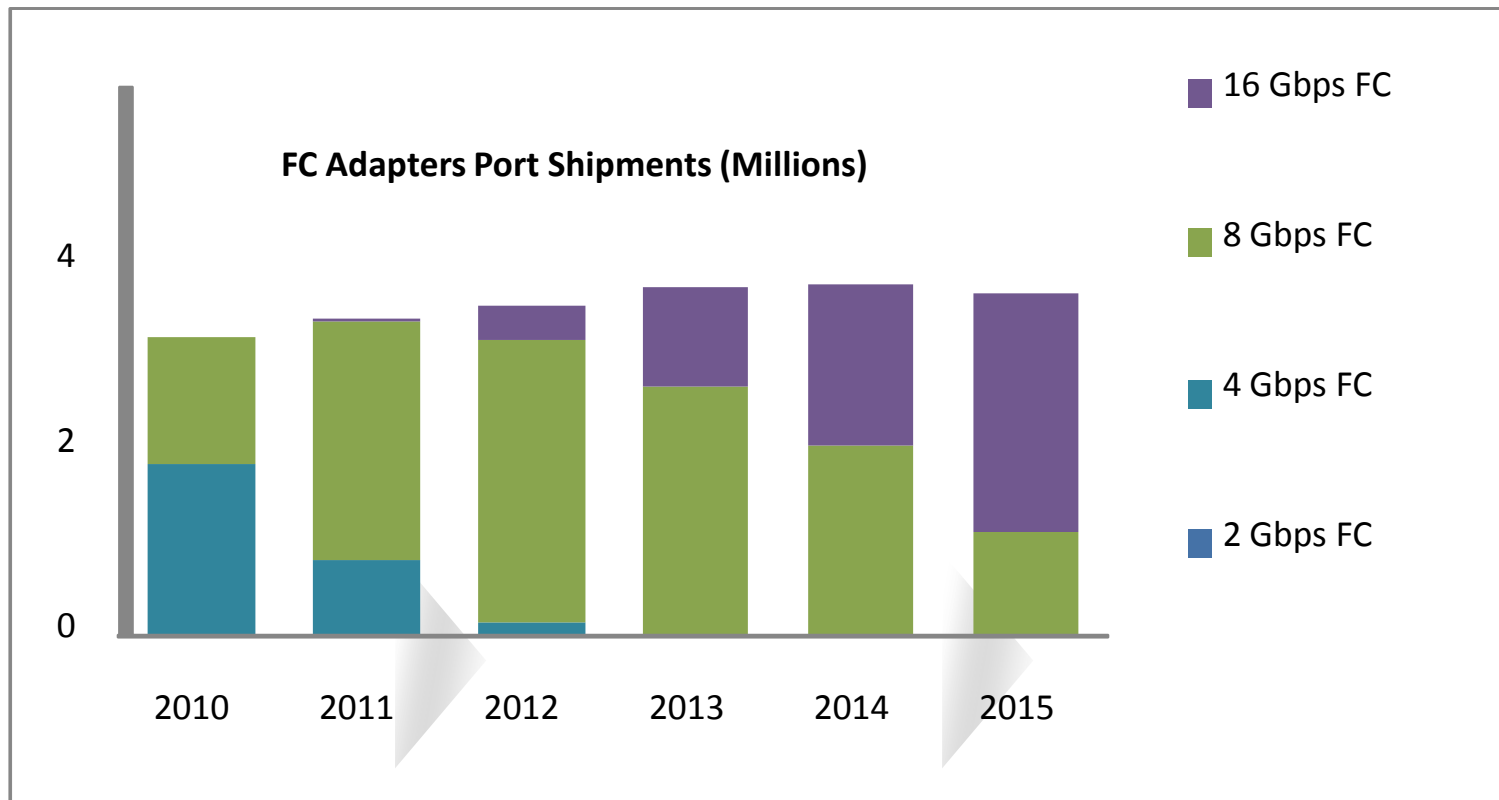
10GbE standard released in 2002, 5 years before 8GFC in 2007



Source: Dell'Oro

2011 Storage Developer Conference. © Fibre Channel Industry Association.
All Rights Reserved.

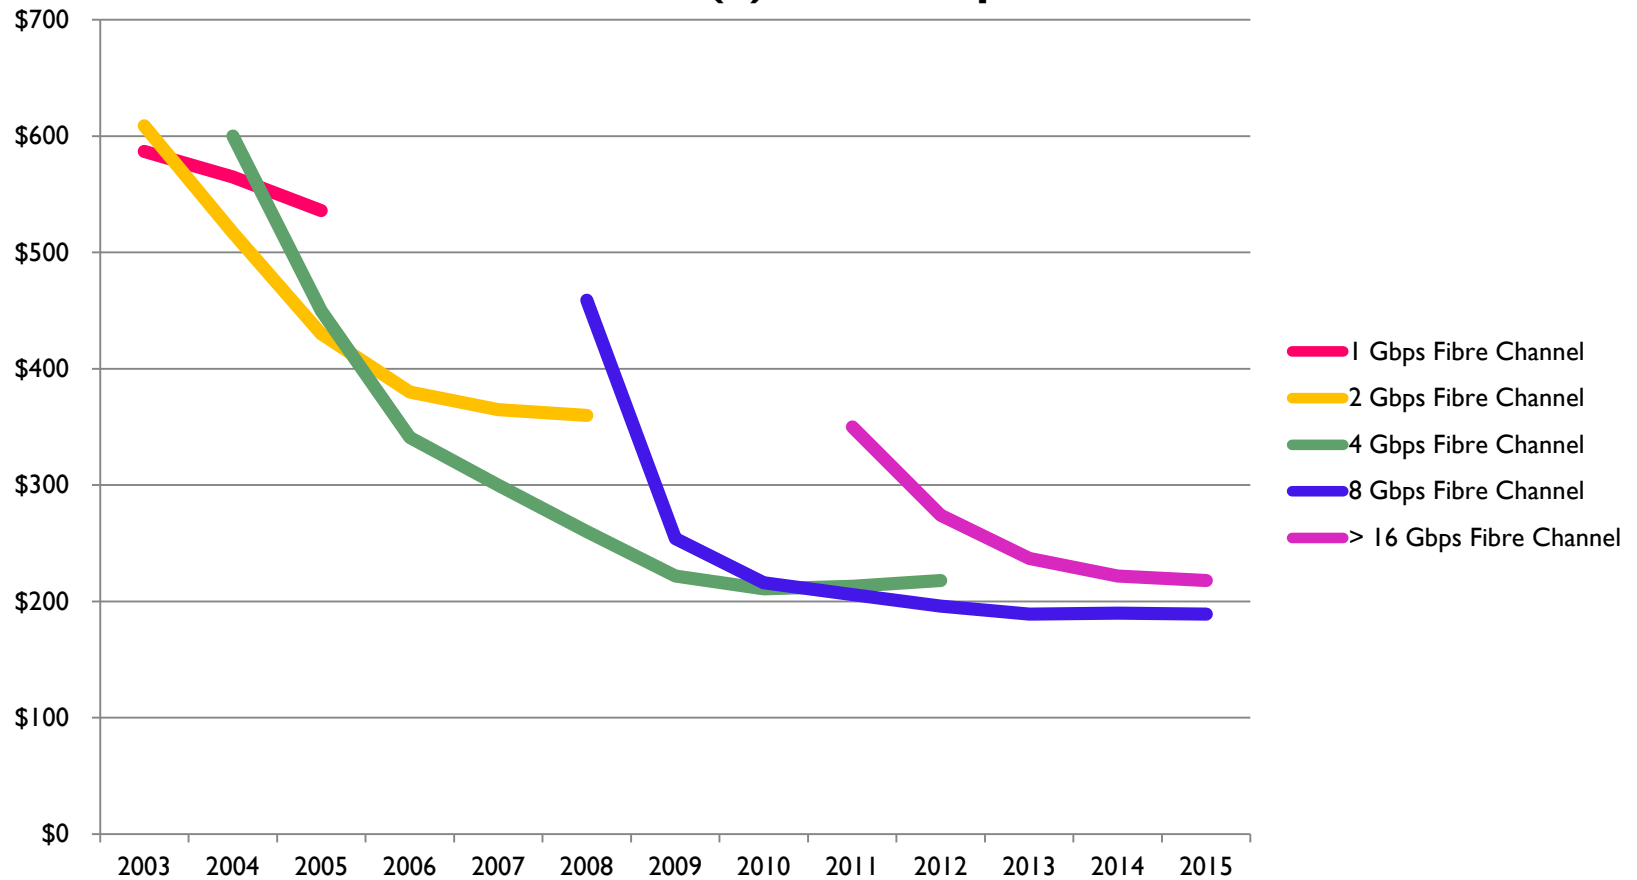
16Gb FC Adapters to Dominate



Source: Dell'Oro Group SAN Forecast Report, July 2011

16Gb FC Adapters Price Decline

World Wide FC Adapter Market MFG Port ASP (\$) – All Adapters



Source: Dell'Oro Group SAN Forecast Report, July 2011

Characteristics of 16GFC

Double the throughput over backplanes, 100 meters and 10 kilometers

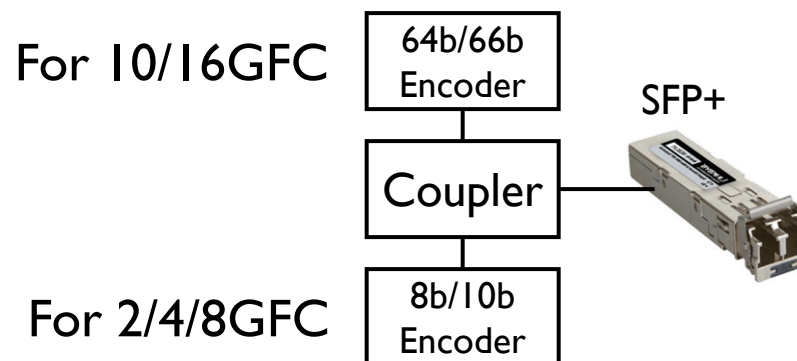
- Fibre Channel Physical Interfaces 5 (FC-PI-5) standardized 16GFC

Speed Name	Throughput (MB/sec)	Retimers in the module	Transmitter Training	OM1/2/3/4 Link Distance (meters)
1GFC	100	No	No	300/500/860/*
2GFC	200	No	No	150/300/500/*
4GFC	400	No	No	50/150/380/400
8GFC	800	No	No	21/50/150/190
10GFC	1200	Yes	No	33/82/300/*
16GFC	1600	Yes	Yes	15/35/100/125

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

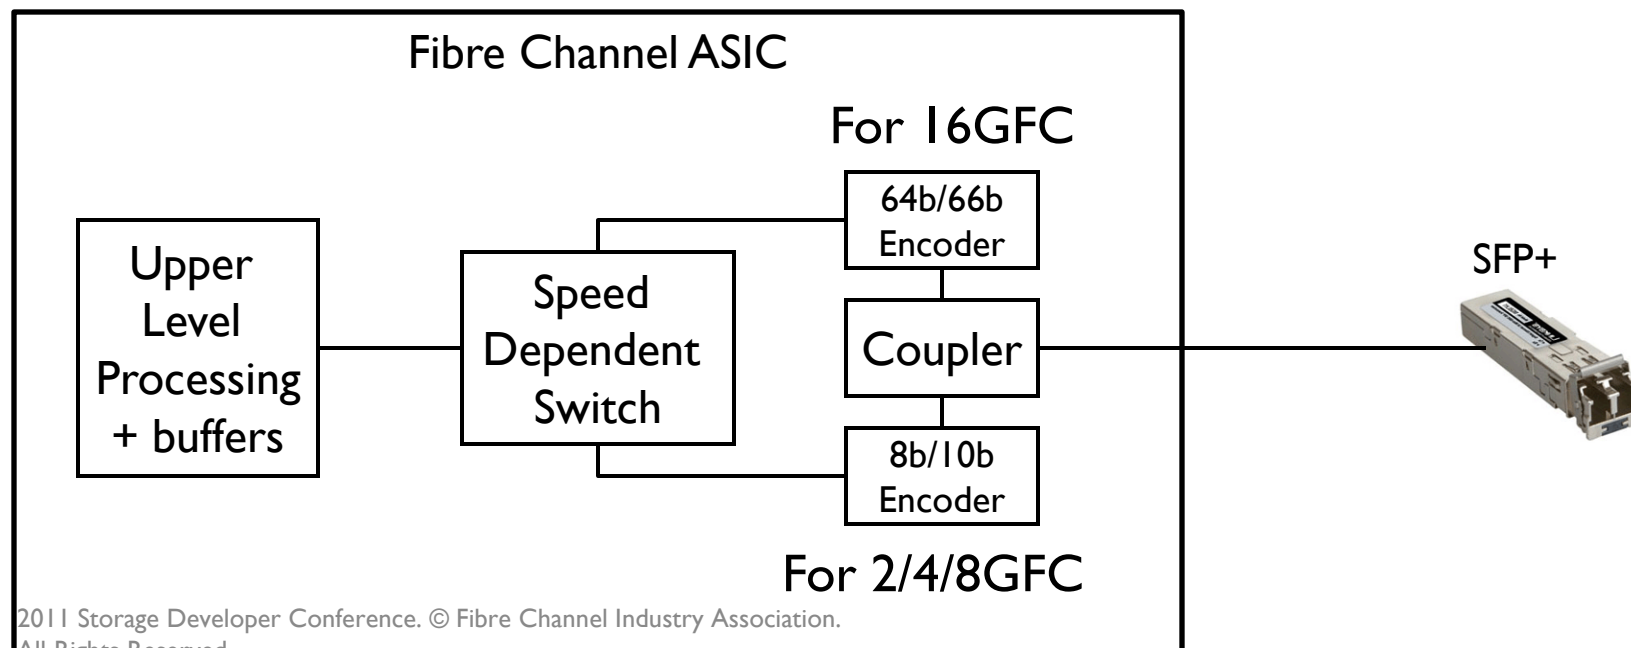
Dual Codecs for Backward Compatibility

- ❑ To improve the efficiency of the protocols, 16GFC only uses 64b/66b coding that is 98% efficient
 - ❑ 8b/10b codes used for 2/4/8GFC are 80% efficient
 - ❑ 16GFC signals do not use the 8b/10b encoders
- ❑ To be backward compatible with 2/4/8GFC, 16GFC ASICs need to support both 8b/10b and 64b/66b coder/decoders (codec) on each link
- ❑ During speed negotiation, the transmitter and receiver switch back and forth between the speeds (and the corresponding codecs) until the fastest speed is reached for a given link



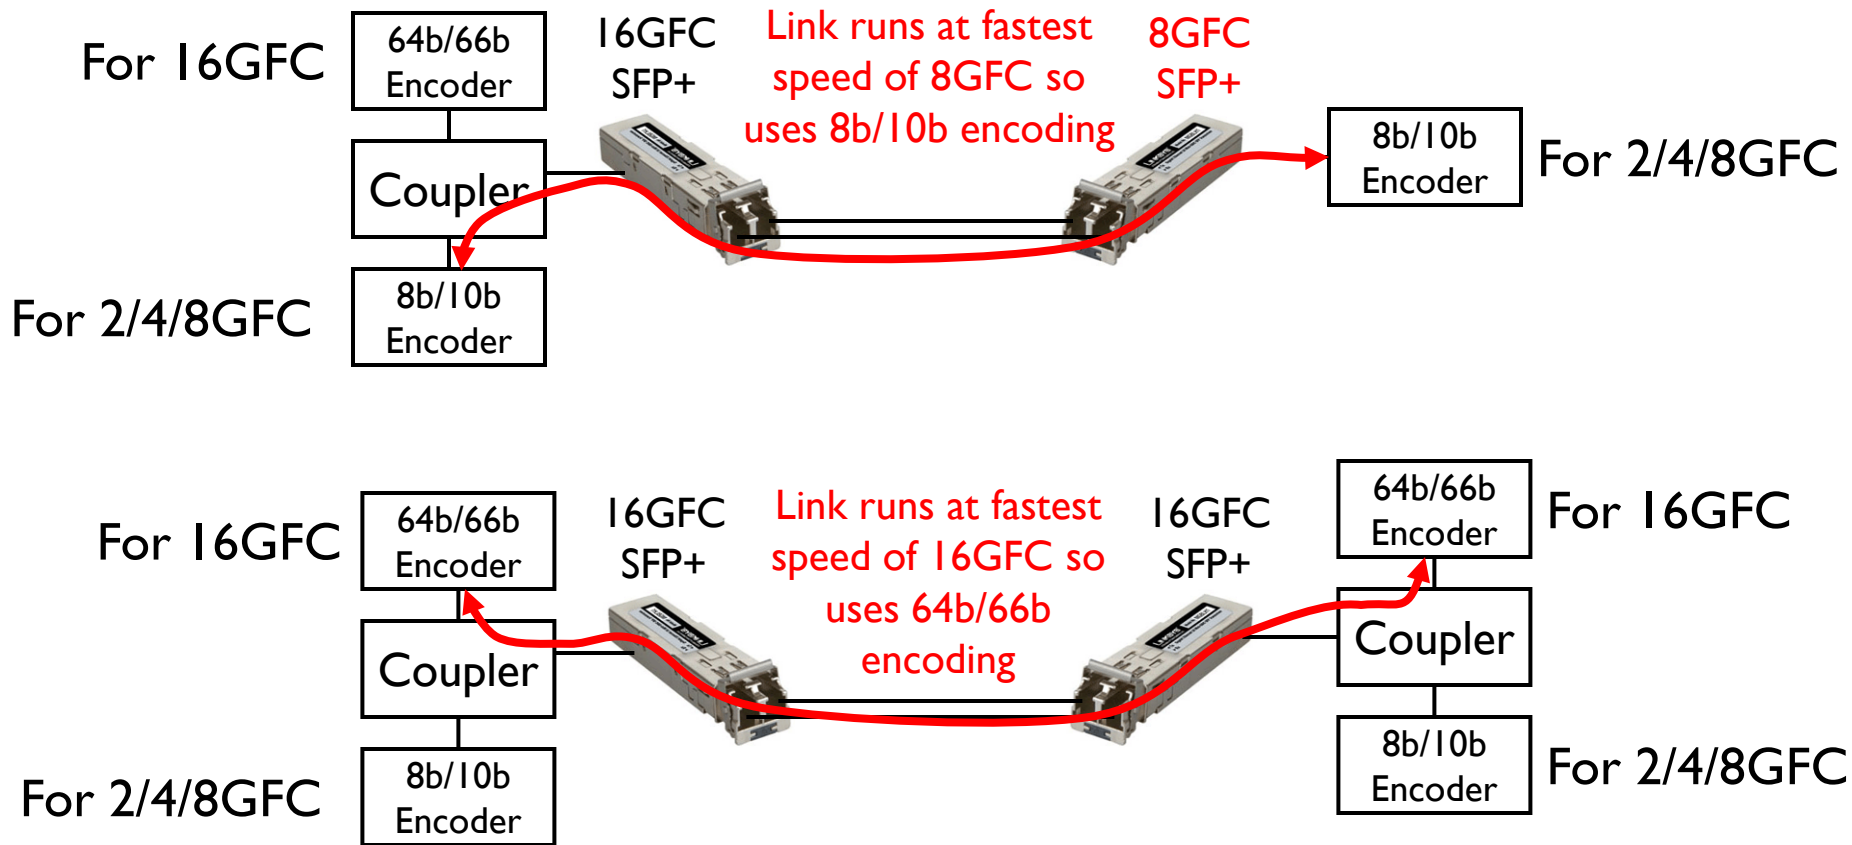
Speed Negotiation

- During speed negotiation, the speed dependent switch routes the initialization sequence to the appropriate encoder
 - 64b/66b for 16GFC
 - 8b/10b for 2/4/8GFC
 - The coupler sends the signals from one the encoders to the SFP+



Speed Negotiation Examples

After Speed Negotiation, only one encoder works at the fastest supported speed

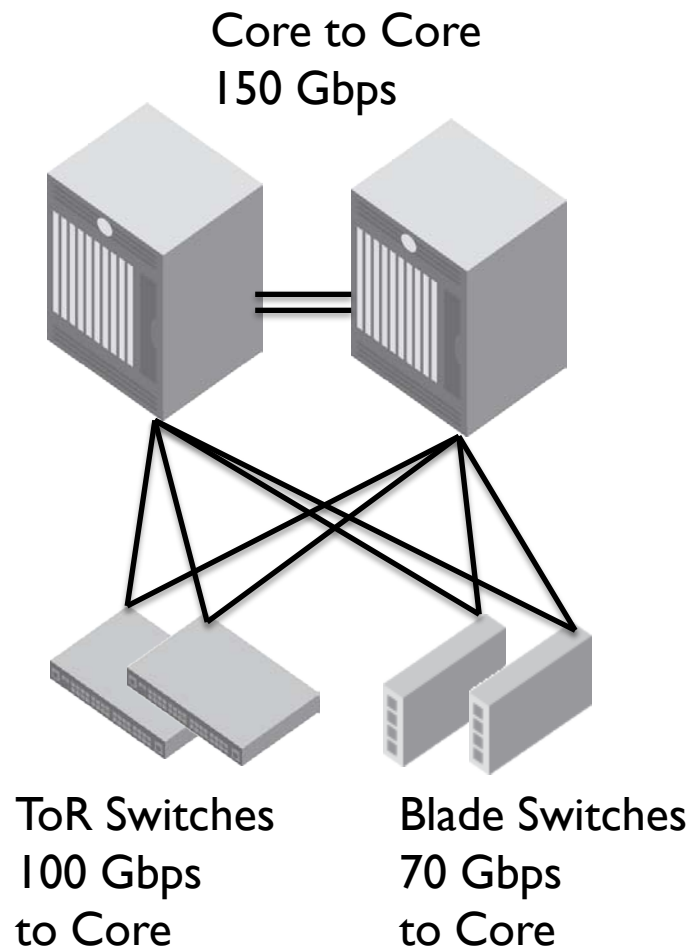


The Benefits of 16GFC

- ❑ 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

- ❑ While processors are considered multi-tasking, each core is still serial and switch between tasks rapidly while they are waiting
- ❑ If individual tasks like reading and writing data are done faster, then the processor doesn't have to wait as long
- ❑ Similar to how SSDs speed up performance, faster links transfer data quicker and decrease latency


The Benefits of 16GFC – Fewer Links



	8GFC Links	16GFC Links
ISLs from ToR Switch to Core	16	8
ISL from Blade Switch to Core	10	5
Core to Core	24	12
Total ISLs	50	25
Total Ports	100	50

Easier Cable Management

- ❑ Fewer links lead to less cables to manage
- ❑ With structured cabling costs up to \$300/port, each cable reduction saves money and reduces opex and the amount of infrastructure to manage
- ❑ Multiple vendors provide dense cable Management solutions



Less Power Consumption

- ❑ 16GFC ports consumes 40% more power than an 8GFC port while delivering 100% more bandwidth
- ❑ Instead of two ports at 8GFC where the power would increase by 100%, a 16GFC port can deliver twice the performance with less than 40% increase in power
- ❑ 16GFC reduces the power per bit that must be expended for each data transfer and reduces operational expenses (opex)

Summary of Benefits

2 x 8GFC Links
= 4 x 0.5 W/sfp+
= 2.0 Watts



1 x 16GFC Links
= 2 x 0.75 W/sfp+
= 1.5 Watts

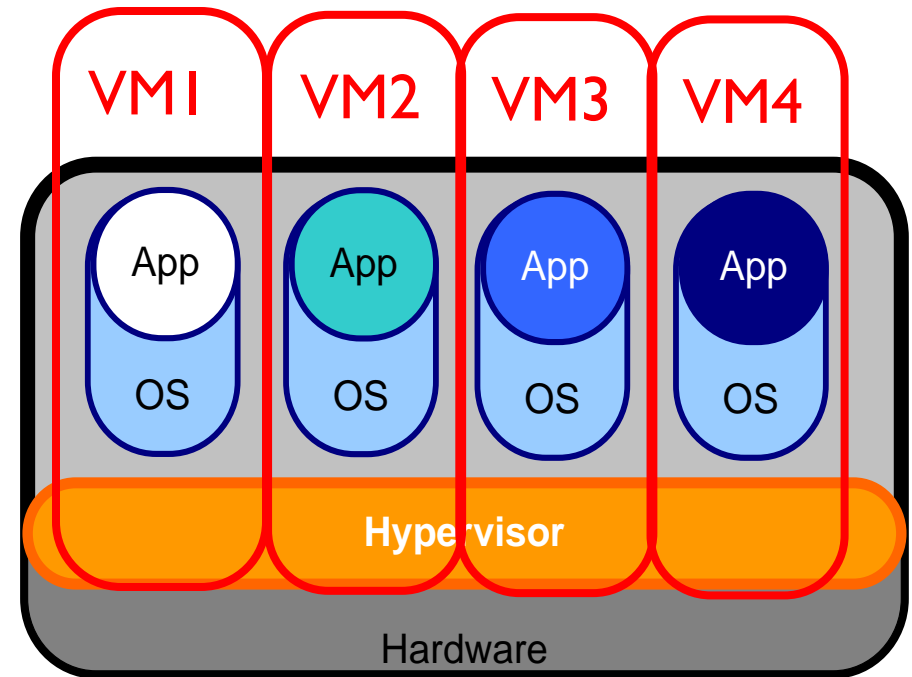


What is driving 16GFC?

- ❑ Multiple factors require Fibre Channel to advance
 - ❑ Prolific applications that increase in scale and number
 - ❑ Server virtualization
 - ❑ Multi-core processors
 - ❑ Large Memory increases
 - ❑ Solid State Disks (SSD)
 - ❑ Faster PCIe rates
 - ❑ Aggregation of traffic on ISLs
 - ❑ Intra-chassis bandwidth
 - ❑ Competition with iSCSI and FCoE

Prolific Applications and Server Virtualization

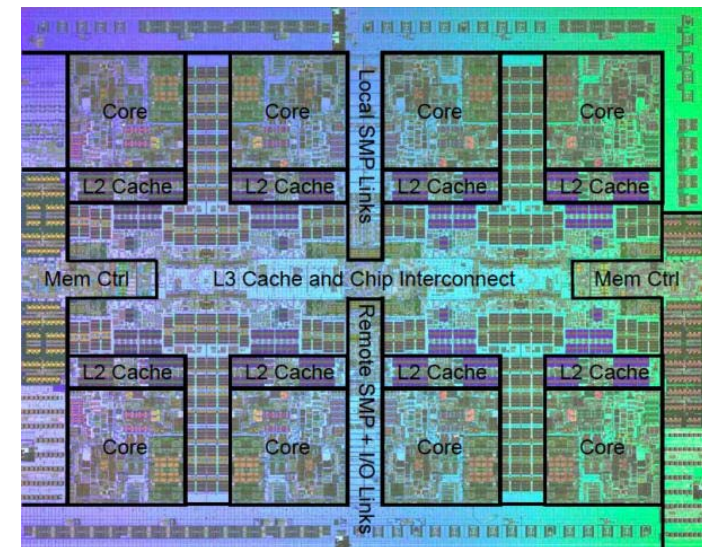
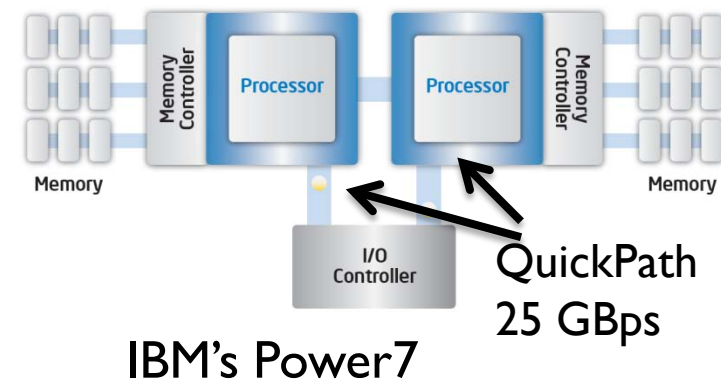
- ❑ Applications keep growing in number and breadth
- ❑ Multiple servers need to access data from shared storage
- ❑ Database applications drive bandwidth
- ❑ Server virtualization creates multiple virtual machines (VMs) for each application so each physical server is producing more Input/Output (I/O)



Faster Processors

- ❑ NehalemEX has 8 cores and 16 threads and uses Intel QuickPath Interconnect at 6.4 Gigatransfers/s delivering up to 25 GigaBytes/second (GBps)
- ❑ IBM has the Power7 that has 8 cores and supports 50 GBps or peak IO and directly interconnects 32 of these processors on a server
- ❑ 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

Nehalem-EX in two-chip configuration



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
 - ❑ 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
- ❑ 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



Increased Memory in Servers


32GB RDIMM



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

PCIe Continues Ramping

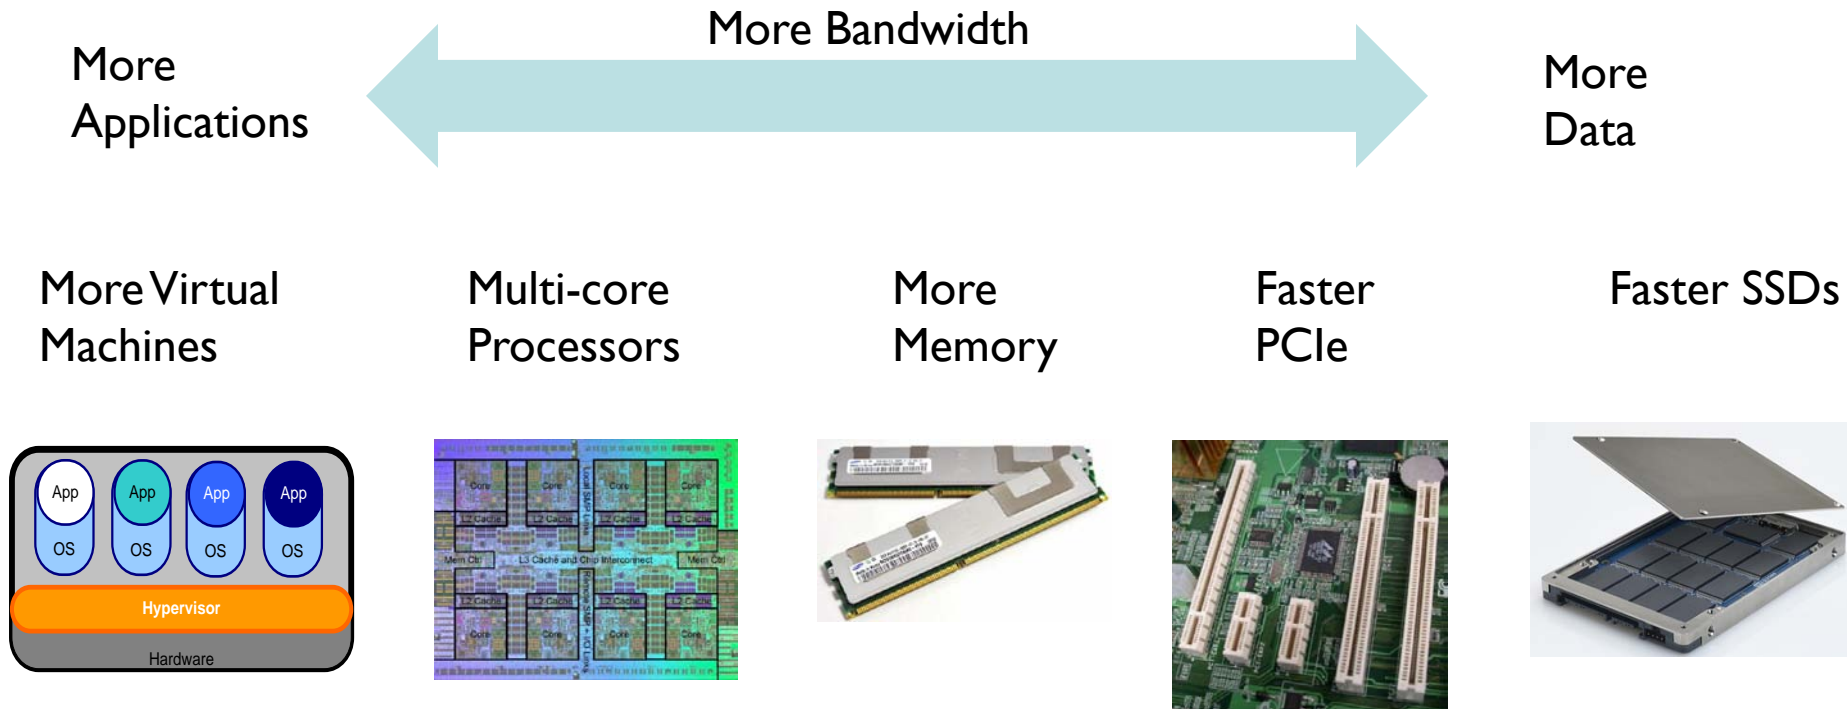
- ❑ PCIe 2.0 increases in speed to support dual ported 16GFC HBAs
- ❑ PCIe 3.0 will support quad ported 16GFC HBAs
- ❑ PCIe4.0 expected in



	Number of Lanes	Speed per Lane (MBps)	Directional Bandwidth (Gbps)	Ports Supported
PCIe -1.0	8	250	16	2 – 8GFC
PCIe -2.0	8	500	32	2 – 16GFC
PCIe -3.0	8	1000	64	4 – 16GFC
PCIe-4.0	8	2000	128	4 - 32GFC

More Applications Drive more Bandwidth

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



How is 16GFC better than 40GE?

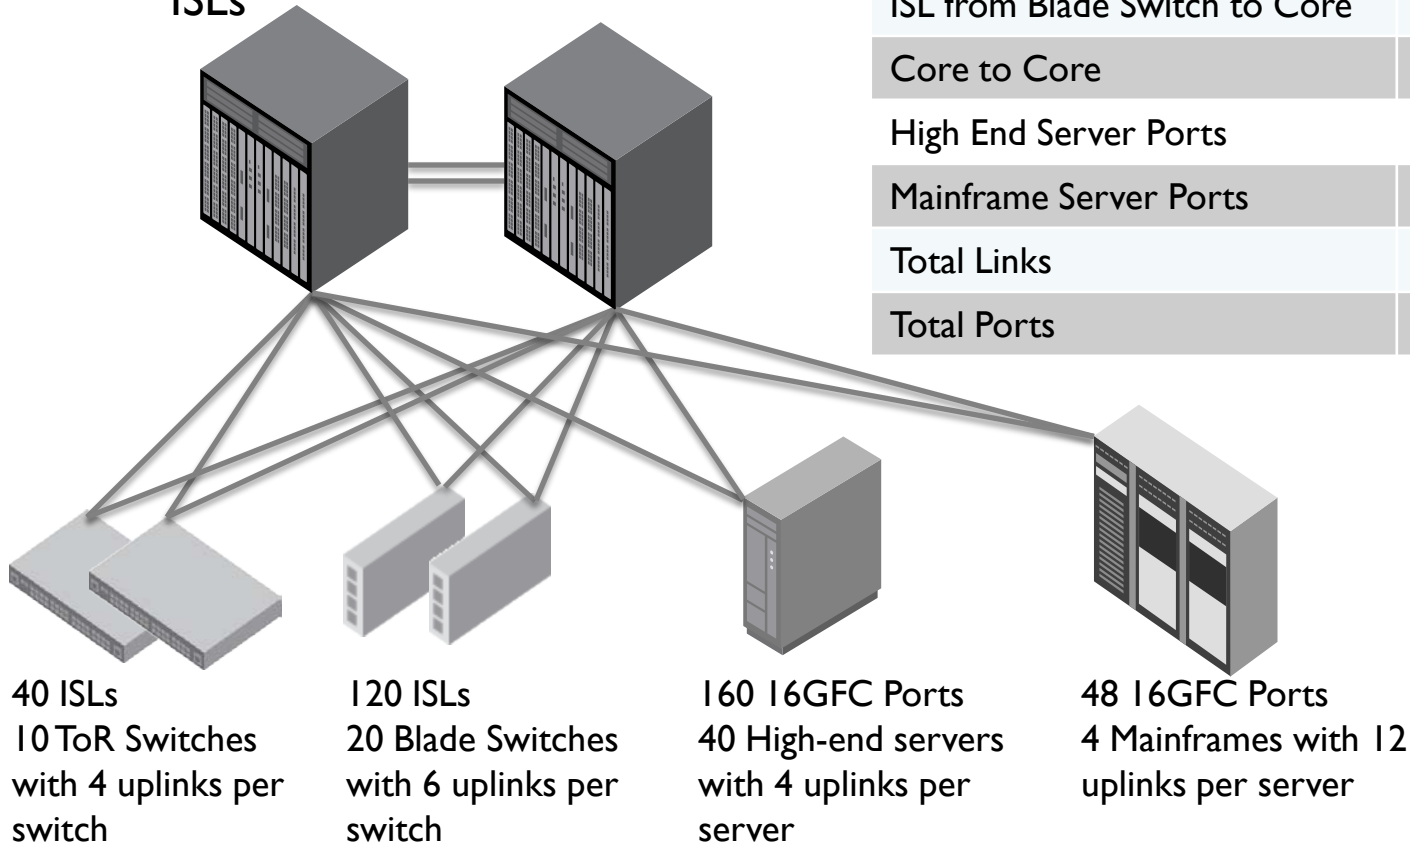
- ❑ 40 Gigabit Ethernet (40GE) is basically 4 lanes of 10GE
- ❑ With Ethernet, there are only steps of 10GE, 40GE and 100GE – with each lane running at 10.3125 Gbps and 4X25.7Gbps in 2014
- ❑ Each lane of 16GFC outperforms each lane of 10GE by 40% and is more efficient in terms of power consumption/bit
- ❑ 3 lanes of 16GFC yield more bandwidth than 40GE
- ❑ Lanes of 16GFC can be trunked in a more flexible way than 40GE because FC lanes can be trunked in steps of 16GFC to effectively produce 32GFC, 48GFC, 64GFC, 80GFC, 96GFC, 112GFC and 128GFC

The Application of 16GFC

- ❑ ISLs – Only one 16GFC ISL is needed instead of two 8GFC ISLs or four 4GFC ISLs to exchange the same amount of data
- ❑ High Bandwidth Applications – applications that use large amounts of bandwidth can perform better with 16GFC
- ❑ Data Migration – Data migrations can be done in less time when the limiting factor is the link speed
- ❑ Virtual Desktop Infrastructure

Sample Fabric

12 Core to Core
ISLs



	8GFC Links	16GFC Links
ISLs from ToR Switch to Core	80	40
ISL from Blade Switch to Core	240	120
Core to Core	24	12
High End Server Ports	320	160
Mainframe Server Ports	96	48
Total Links	760	380
Total Ports	1,520	760

High Bandwidth Applications

- Different applications have varying bandwidth requirements
- I6GFC is designed to support a wide range of applications and is mainly applicable for aggregation points, high end servers and backup servers

Max I/O per Application	Applications per Server	I/O per Server	Servers per Rack	I/O per Rack
10 Mbps	30	0.3 Gbps	40	12 Gbps
100 Mbps	15	1.5 Gbps	30	45 Gbps
400 Mbps	8	3.2 Gbps	24	76.8 Gbps
1.5 Gbps	12	18 Gbps	4	72 Gbps
8 Gbps	12	96 Gbps	1	96 Gbps
4 Gbps	4	16 Gbps	20	320 Gbps

Bursty Data Traffic

- ❑ Data traffic is bursty in nature and the system performance is based on how fast each burst can be processed or transmitted
- ❑ Faster links lead to better performance

Size of Data Exchange	Latency of GE at 125 MBps (Seconds)	Latency of 4GFC at 400 MBps (Seconds)	Latency of 8GFC at 800 MBps (Seconds)	Latency of 10GE at 1250 MBps (Seconds)	Latency of 16GFC at 1600 MBps (Seconds)
100 MB	0.8	0.25	0.125	0.08	0.0625
1 GB	8	2.5	1.25	0.8	0.625
10 GB	800	250	125	80	62.5
100 GB	800	250	125	80	62.5
1 TB	8000	2500	1250	800	625
10 TB	80000	25000	12500	8000	6250
100 TB	800000	250000	125000	80000	62500
1 PB	8000000	2500000	1250000	800000	625000

16GFC
delivers a TB
in 10 minutes

I/O Aggregation

- Small streams of data can add up to large rivers
- I6GFC will minimize the number of links by maximizing the throughput of each link

	Max I/O Requirements (MBps)	Number of Applications	Aggregate Bandwidth	16 GFC Links
Very Low-End Application	10	500	5000	3.125
Low-End Application	100	80	8000	5
Mid-range Application	1000	17	17000	10.625
High-end Application	10000	2	20000	12.5

Data Migration

- When large amounts of data need to be replicated, little processing is involved, but high bandwidth is needed
- Faster links cut the transfer time down

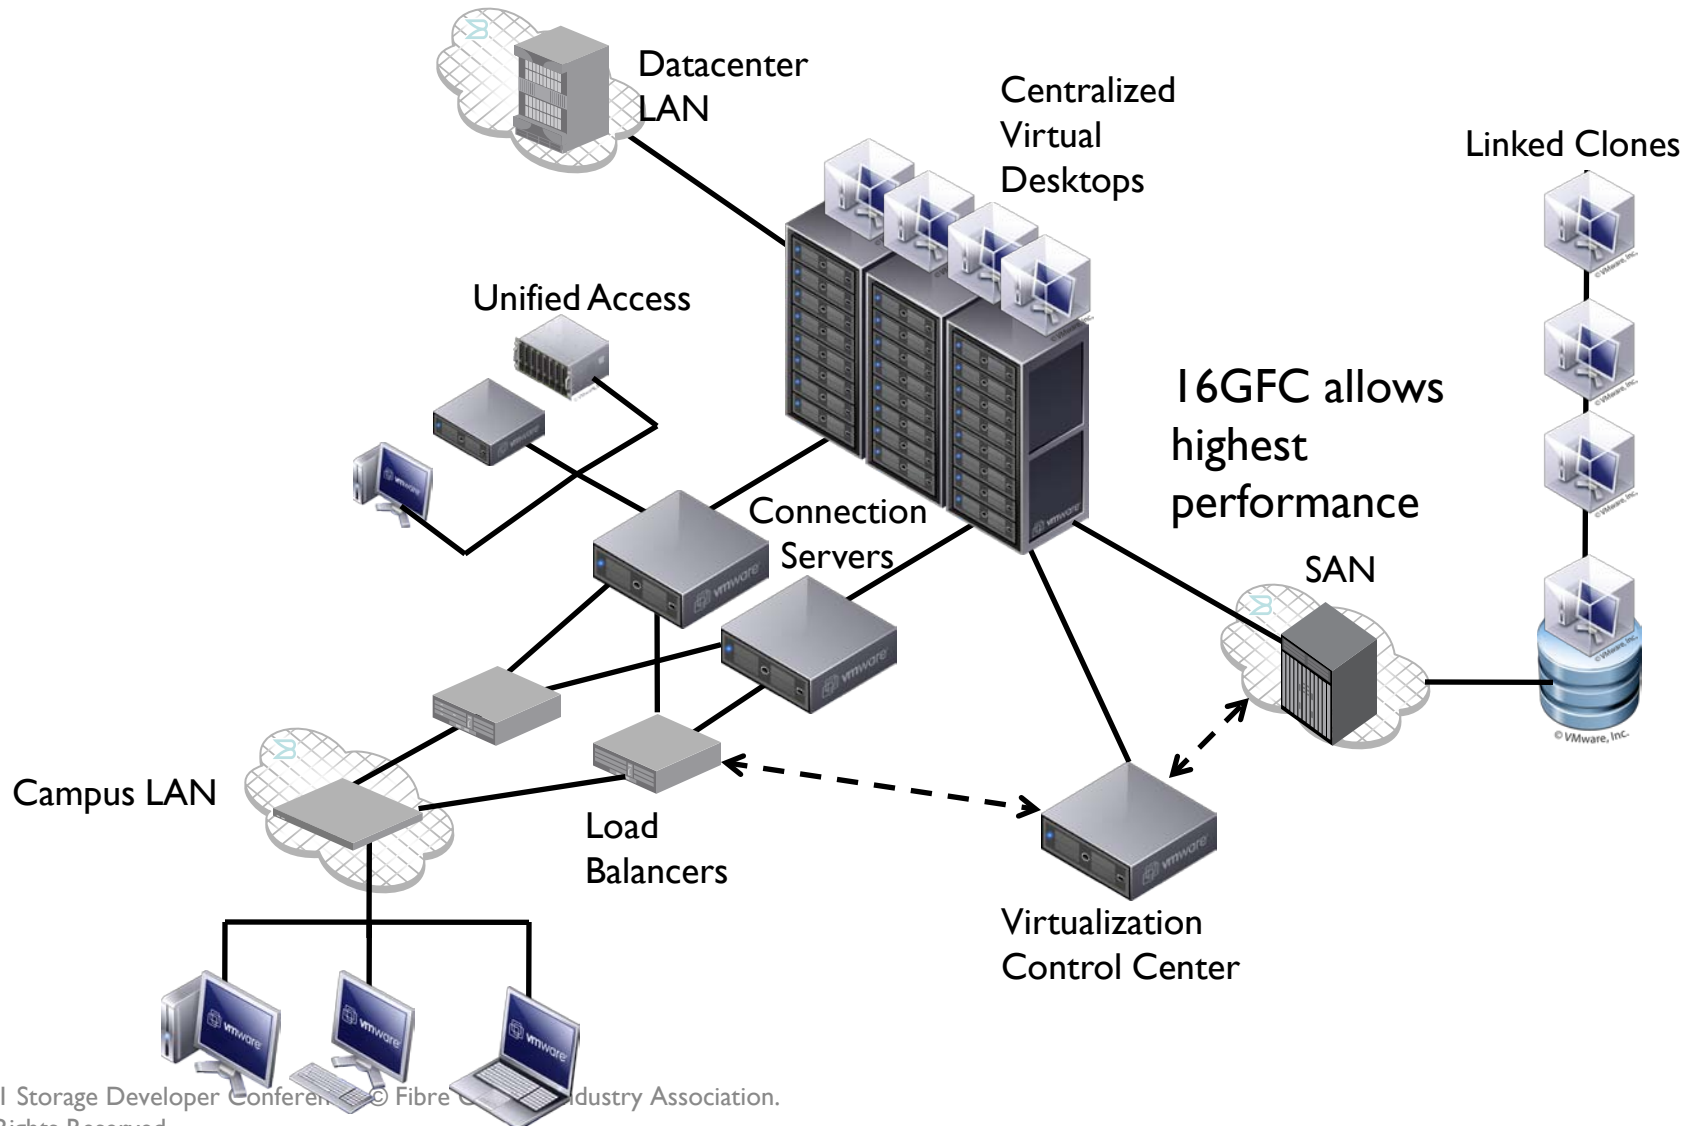
Size of Data Exchange	Latency of 1 GE at 125 MBps (Seconds)	Latency of 4GFC at 400 MBps (Seconds)	Latency of 8GFC at 800 MBps (Seconds)	Latency of 10GE at 1250 MBps (Seconds)	Latency of 16GFC at 1600 MBps (Seconds)
1 GB	8	2.5	1.25	0.8	0.625
10 GB	80	25	12.5	8	6.25
100 GB	800	250	125	80	62.5
1 TB	8000	2500	1250	800	625
10 TB	80000	25000	12500	8000	6250
100 TB	800000	250000	125000	80000	62500
1 PB	8000000	2500000	1250000	800000	625000
10 PB	80000000	25000000	12500000	8000000	6250000

16GFC
delivers a PB
in ~1 week

Virtual Desktop Infrastructure (VDI)

- ❑ VDI lets IT departments run virtual desktops in the datacenter and deliver desktops to employees as a managed service.
- ❑ End users gain a familiar, personalized environment that they can access from any number of devices anywhere throughout the enterprise or from home.
- ❑ Administrators gain centralized control, efficiency, and security by having desktop data in the datacenter.
- ❑ VDI drives bandwidth to the data center when hundreds or thousands of employees work from the data center

VDI Architecture



- ❑ 16GFC is needed to keep pace with other advances
 - ❑ Abundant applications, server virtualization
 - ❑ 12 core processors, 32 GB DIMMs, SSDs, PCIe 2.0
- ❑ Users benefit from higher performance, reduced number of links, easier cable management, and better power efficiency
- ❑ High bandwidth applications, data migrations, VDI and ISLs are perfect applications for 16GFC
- ❑ Fibre Channel will remain the dominant storage interconnect by providing the best performance

Thank You