

# Fibre Channel Driving Storage Innovation, Speed and Convergence

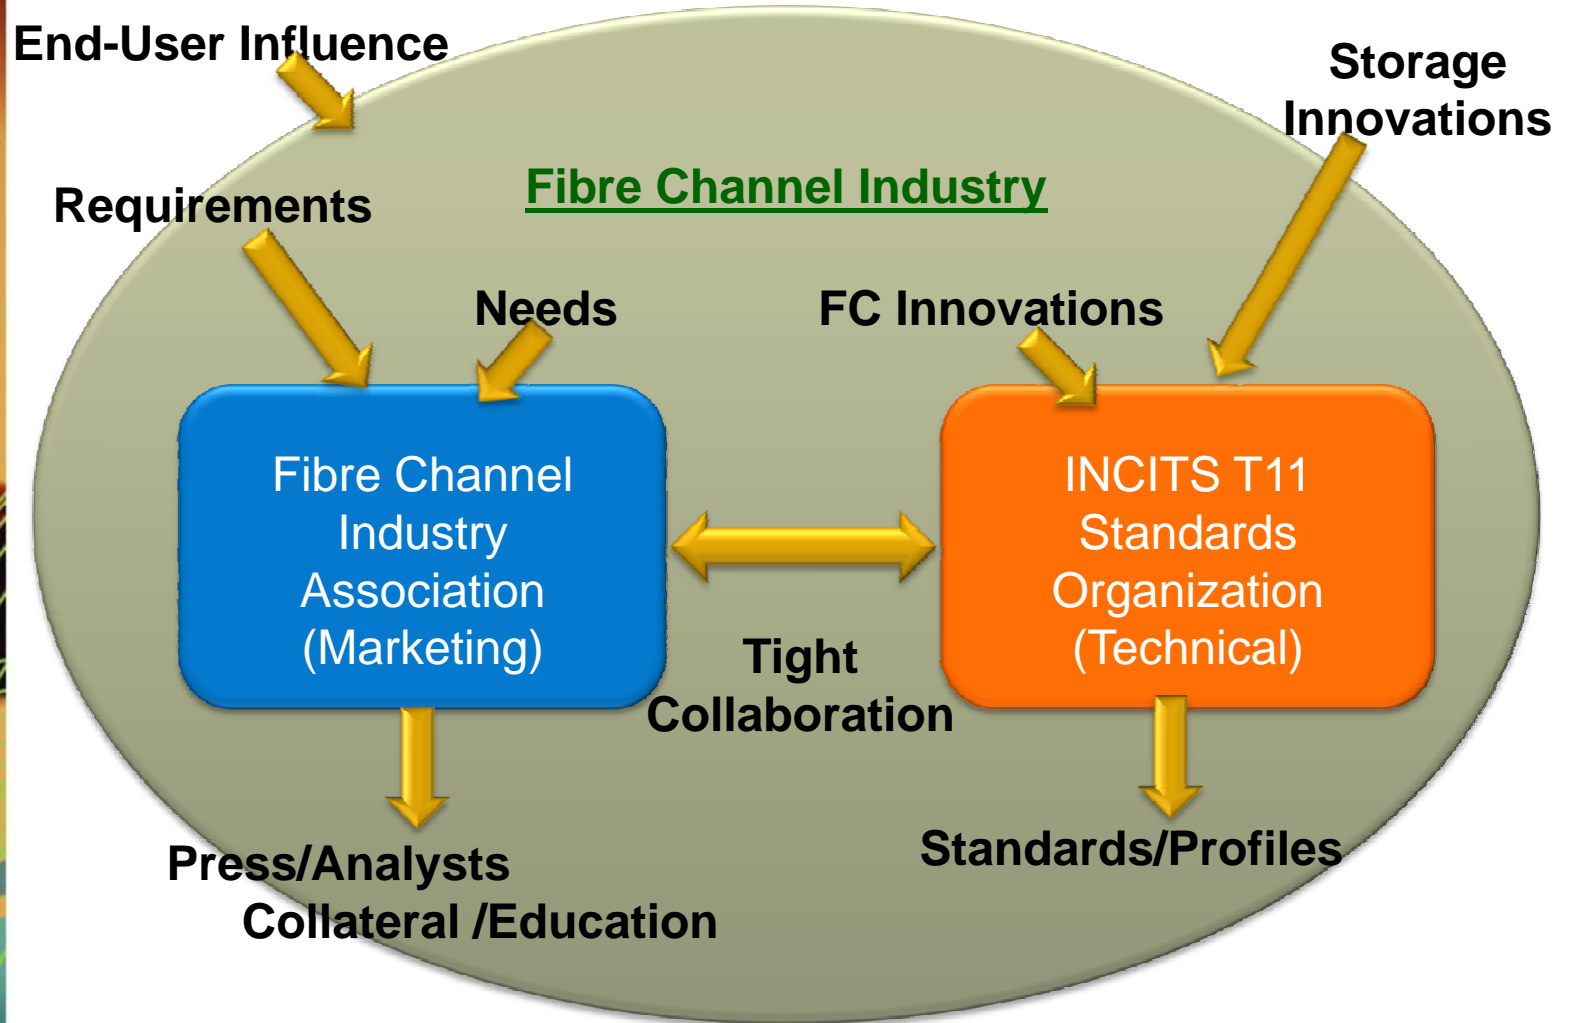
Tom Hammond-Doel  
Vice-Chairman  
FCIA

**FIBRE CHANNEL DRIVING STORAGE INNOVATION,  
SPEED AND CONVERGENCE**



## Agenda


- Update on the Fibre Channel Industry
- The Fibre Channel Roadmaps
- Innovations Abound
- Summary



## **End User Influence on Fibre Channel**

- Fibre Channel is in a unique position due to its large install base and longevity
- This distinguishes Fibre Channel from other storage transport technologies
- Many requirements being addressed by the INCITS T11 FC Standards Committee are based on REAL end user experience and input
- This has lead to the development of many new Fibre Channel features and enhancements

## Normalized Switch & HBA Port Count & Market





## Fibre Channel Speed Roadmap - v13 - 2/4

FC

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	2006	2008
16GFC	3200	14.025	2009	2011
32GFC	6400	28.05	2012	2014
64GFC	12800	TBD	2015	Market Demand
128GFC	25600	TBD	2018	Market Demand
256GFC	12800	TBD	2021	Market Demand
512GFC	25600	TBD	2024	Market Demand

"FC" used throughout all applications for Fibre Channel infrastructure and devices, including edge and ISL interconnects. Each speed maintains backward compatibility at least two previous generations (i.e., 8GFC backward compatible to 4GFC and 2GFC)

†Line Rate: All "FC" speeds are single-lane serial stream

‡Dates: Future dates estimated

**Backward and Forward Compatibility!**

## Fibre Channel Speed Roadmap - v13 - 3/4



**ISL**  
(Inter-Switch Link)


Product Naming	Throughput (MBps)	Equivalent	Spec Technically	Market
10GFC	2400			
20GFC	4800			
40GFC/FCoE	9600			
100GFC/FCoE	24000			
400GFC/FCoE	96000			
1TFC/FCoE	240000			

Now THAT's a fast Yorkie!

**Yes, a Terabit**

† Some solutions are Pre-Standard Solutions: There are several methods and/or data stream lines to achieve the core bandwidth necessary for compatibility guidelines. Refer to the FCoE page 4 for 40GFCoE and 100GFCoE.

## Fibre Channel Speed Roadmap - v13 - 4/4



FCoE

Product Naming	Throughput (MBps)	Equivalent Line Rate (GBaud) <sup>†</sup>	Spec Technically Completed (Year) <sup>‡</sup>	Market Availability (Year) <sup>‡</sup>
10GFCoE	2400	10.3125	2008	2009
40GFCoE	9600	41.225	2010*	Market Demand
100GFCoE	24000	103.125	2010*	Market Demand

Fibre Channel over Ethernet tunnels FC through Ethernet. For compatibility all 10GFCoE FCFs and CNAs are expected to use SFP+ devices, allowing the use of all standard and non standard optical technologies and additionally allowing the use of direct connect cables using the SFP+ electrical interface. FCoE ports otherwise follow Ethernet standards and compatibility guidelines.

<sup>‡</sup>Dates: Future dates estimated

\* It is expected that 40GFCoE and 100GFCoE based on 2010 standards will be used exclusively for Inter-Switch Link cores, thereby maintaining 10GFCoE as the predominant FCoE edge connection



## Innovations – Compliments of Fibre Channel


- The FCIA is helping to extend the FC protocol through close cooperation with the T11 Standards Organization through innovation in multiple areas:

FC-EE

FCoE

16GFC+

Security



Configuration &  
Management

Virtualization  
Technologies


**FC-EE**

## **Energy Efficient Fibre Channel**



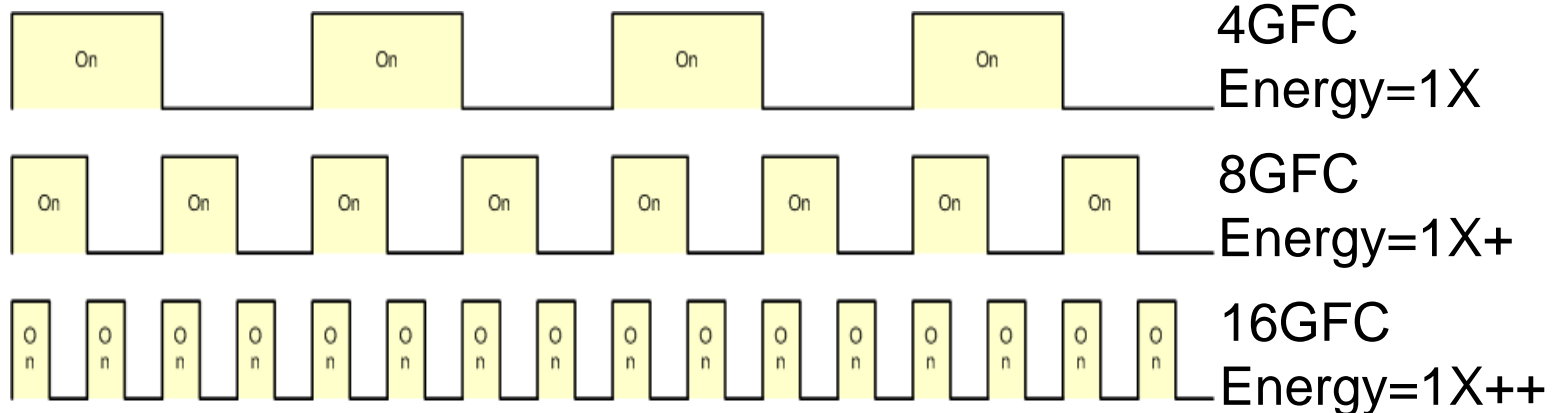
## Fibre Channel Powering **Green Storage**!

- Tiered storage is an element of **Green Storage**
- Fibre Channel offers the high performance **Green Storage** requires to obtain the best Efficiency/Watt ratings
- The FCIA has a **Green Initiative** in place with the T11 Standards organization
- **FC-EE** is exploring different possibilities in committee



## Fibre Channel Powering **Green Storage!**

“NRZ” means speed is relatively power independent



FC-EE Brings many different ways of lowering power



**Just One Example of FC-EE Possibilities**



**FCoE**

# Fibre Channel over Ethernet



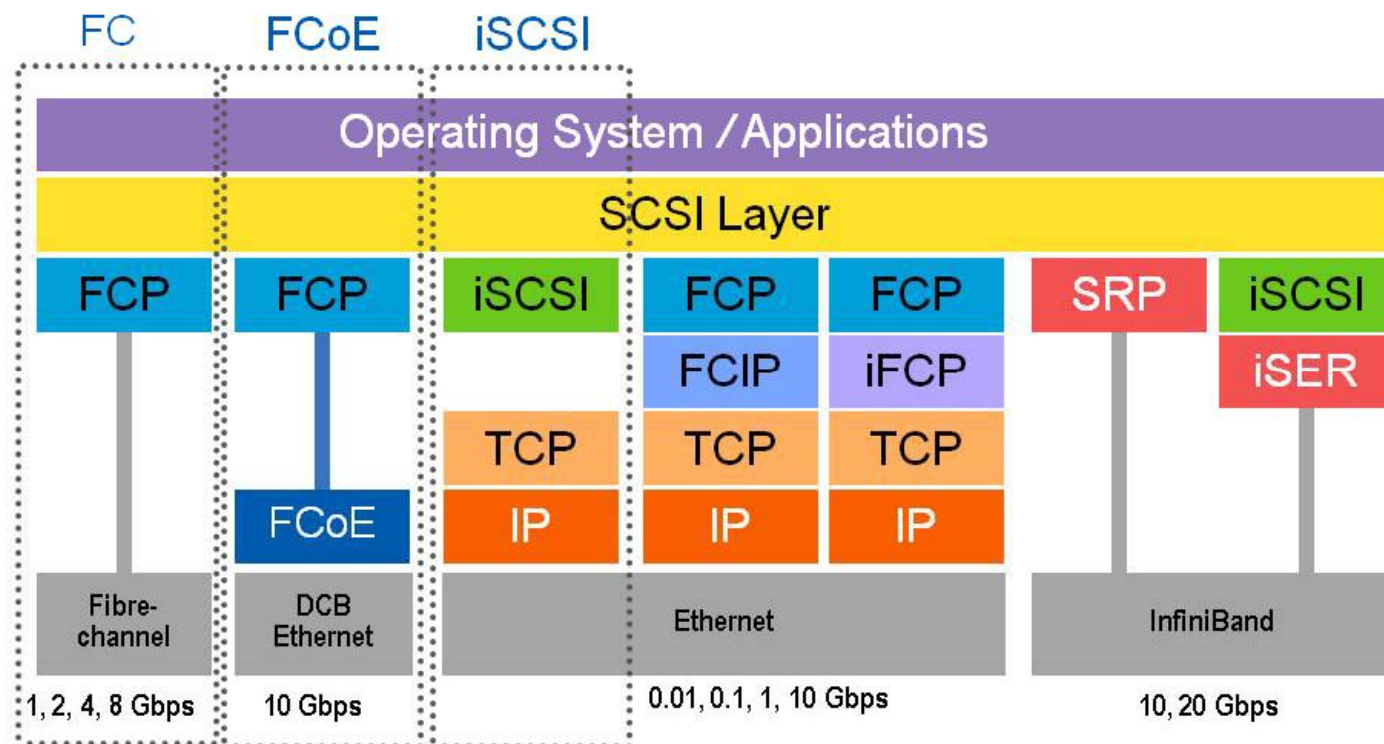


## Innovations – FCoE

- **Data Center Bridging (DCB)**
  - Set of IEEE Standards
    - Makes Ethernet more robust and enables lossless environments
  - Enables 10GbE I/O consolidation
  - Consists of Priority Flow Control, Congestion Notification, and Enhanced Transmission Selection
- **FC over Ethernet (FCoE)**
  - Ethernet encapsulates FC; another upper-layer protocol
  - Managed like FC at initiators, switches, storage systems
    - Central design goal for FCoE
  - Same cabling (SFP+) for 8GFC, 16GFC and 10G FCoE

# Innovations – FCoE

# Fibre Channel over Ethernet



## **6th FCoE Plugfest (May '11)**

- Testing the interaction of FCoE end devices with DCBx switches. This will validate that FCoE devices operate properly with PFC (Pause) and other Ethernet traffic in a Data Center Ethernet environment
- Using only the standard BB5 frame format and standardized that all devices must use FIP
- Enhanced Transmission Selection is now in the test plan. Manages the bandwidth of the links
- The 7<sup>th</sup> FCoE Plugfest is already scheduled for October 2011 to include 16GFC

**Evolution of a WINNING I/O Technology!**

**16GFC+**

# **16GFC+ Fibre Channel**





## 16GFC Overview

- What is driving 16GFC
  - Prolific applications, server virtualization, multi-core processors, more memory, solid state drives, PCIe Gen 3, traffic aggregation, VDI
- Benefits of 16GFC
  - Higher performance leads to fewer links, easier cable management, less power consumption per bit
- Defined by the following standards
  - FC-P1-5, Complete in T11, in final processing at INCITS
  - FC-MSQS, Complete in T11, in final processing at INCITS
  - FC-FS-3, Complete in T11, in final processing at INCITS

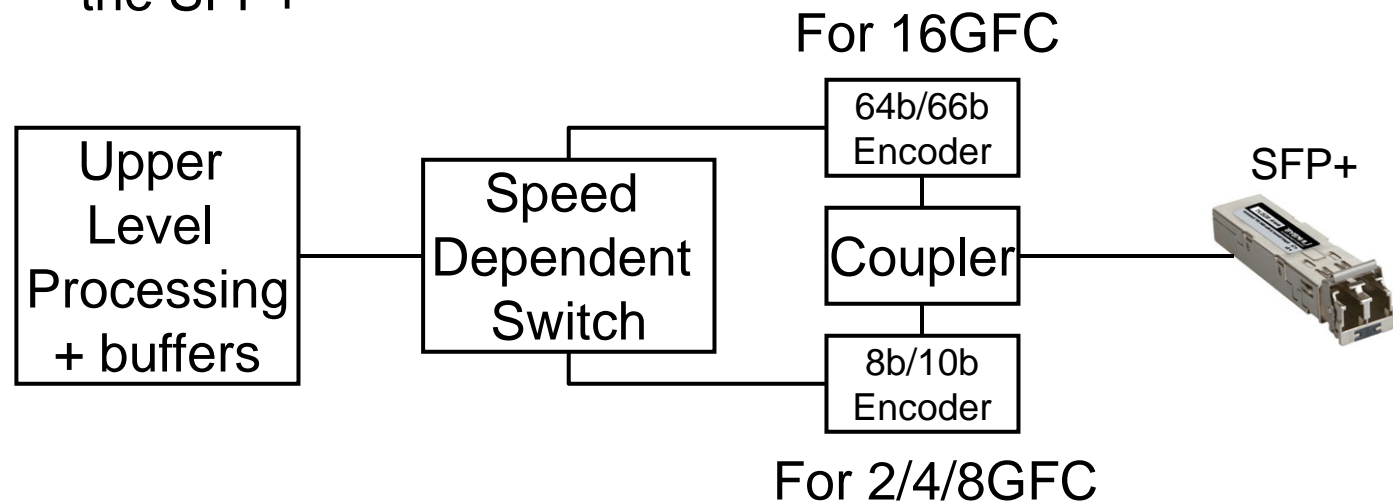


## **16GFC Backward Compatibility**

- 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 can 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

## 16GFC 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+



## **Innovation – 32GFC**

- **Highlights of FCIA 32GFC MRD for T11 standard**
  - Stay serial and single-lane; 28.05Gbaud
  - T11.2 FC-PI-6: Stable by August '11, release Feb '12
  - 70 to 100 meters on OM3 optics, 7 meters on copper
  - $\leq 50\%$  Watts/Port of 40GE and  $\leq 50\%$  \$/port of 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
      - IB will have 25G per lane option in 2014/2015
- **8GFC you buy today will work with 32GFC**

**Now THAT's a Safe Investment!**

**Security**

## **Fibre Channel Security Enhancements**



## **Innovations – FC-SP-2**

- FC-SP-2 builds on the existing FC-SP security model
  - Threat Model drove FC-SP-2 Architecture work
- Non-NULL DH-CHAP w/RSA-2048 mandatory to implement
  - Original SP-SF required only NULL DH-CHAP
- Key Management Services under active consideration
  - KMIP Profile: FC-SP-2 Annex under consideration
- Last new technical input Apr 2011
  - Current schedule: Ballot in Oct 2011



## Configuration & Management

## Fibre Channel Configuration Enhancements



## **FC-SCM Behaviors**

- **Management Tools**
  - The set of behaviors provided by a host, switch, or storage device that a management tool running locally in that host, switch, or storage device may use to manage an SCM environment
- **Hosts**
  - The set of behaviors required to allow a Host, including a Host Bus Adapter, to be used to manage a SCM environment and be managed in an SCM environment
- **Fabrics**
  - The set of behaviors required to allow a Fabric to be used to manage an SCM environment and be managed in an SCM environment
- **Storage**
  - The set of behaviors required to allow a Storage device to be used to manage an SCM environment and be managed in an SCM environment

## Some FC-SCM Examples

- An SCM Switch shall support 256 entries in the name server database
- An SCM Switch shall support 256 entries in the Fabric Configuration Server database
- An HBA shall provide the following Name Server attributes for each of its ports:
  - Symbolic Port and Node Names,
  - Port Types,
  - FC-4 Types and Features
- A switch shall default its configuration to operate in Enhanced Zoning mode.

## **Fabric Device Management Interface (FDMI)**

- Obviates the need to have management clients on every host
- A single client on a single host may access information about all HBA's in the Fabric
- Simplifies management
- FC-GS-6

## Enhanced Commit Service

- Makes the Fabric more reliable and available
- Two phase commit protocol used to update Fabric databases
  - Examples include the Zone and Policy (Security) Databases
- Provides Additional Function
  - Application Level Locking
    - Application rather than fabric basis
  - Advanced Error Recovery
  - Transaction Semantics
    - Atomic operations
- FC-GS-6



## Inter-Fabric Routing

- Allows fabrics to communicate without merging
- Limits notifications (like RSCNs)
- Functionality allowing limitations on inter-device communications
- This is achieved using the Inter-Fabric Zoning protocol (see FC-IFR)

## Virtualization Technologies

## Fibre Channel Virtualization Enhancements



## **N\_Port ID Virtualization**

- Allows Multiple Virtual N\_Ports to exist behind a physical N\_Port
- Normal FLOGI Occurs for the first Login
- Subsequent Logins Request FC-IDs using the FDISC ELS
- Enables Virtual Machine Environments
  - Each Virtual Machine is able to have its own WWN and FC\_ID
- Enables N\_Port Virtualizers

## **Virtual Fabrics (Frame Tagging)**

- Applies the Virtual Machine Paradigm to the Switch
- Provides for Multiple Virtual Switches in a Single Physical Switch
- Each Virtual Fabric (SAN) Runs an Instance of the FC Services
- Provides for Isolation by Carving Up the Physical Fabric into Multiple Virtual Fabrics
- Routing Tables are Indexed by the 12 bit Fabric ID Field in the Frame

# Virtual Channels

- Multiple Virtual Channels on ISLs
- Allows Traffic Differentiation
- Three Assignment Schemes
  - Simple, Two Virtual Channels, Class F and All Others
  - Fixed, Based on D\_ID with class F and Broadcast assigned
  - Variable, Based on D\_ID with class F and Broadcast not assigned
- Each VC Has Own Set of BB Credits
  - Negotiated During ELP



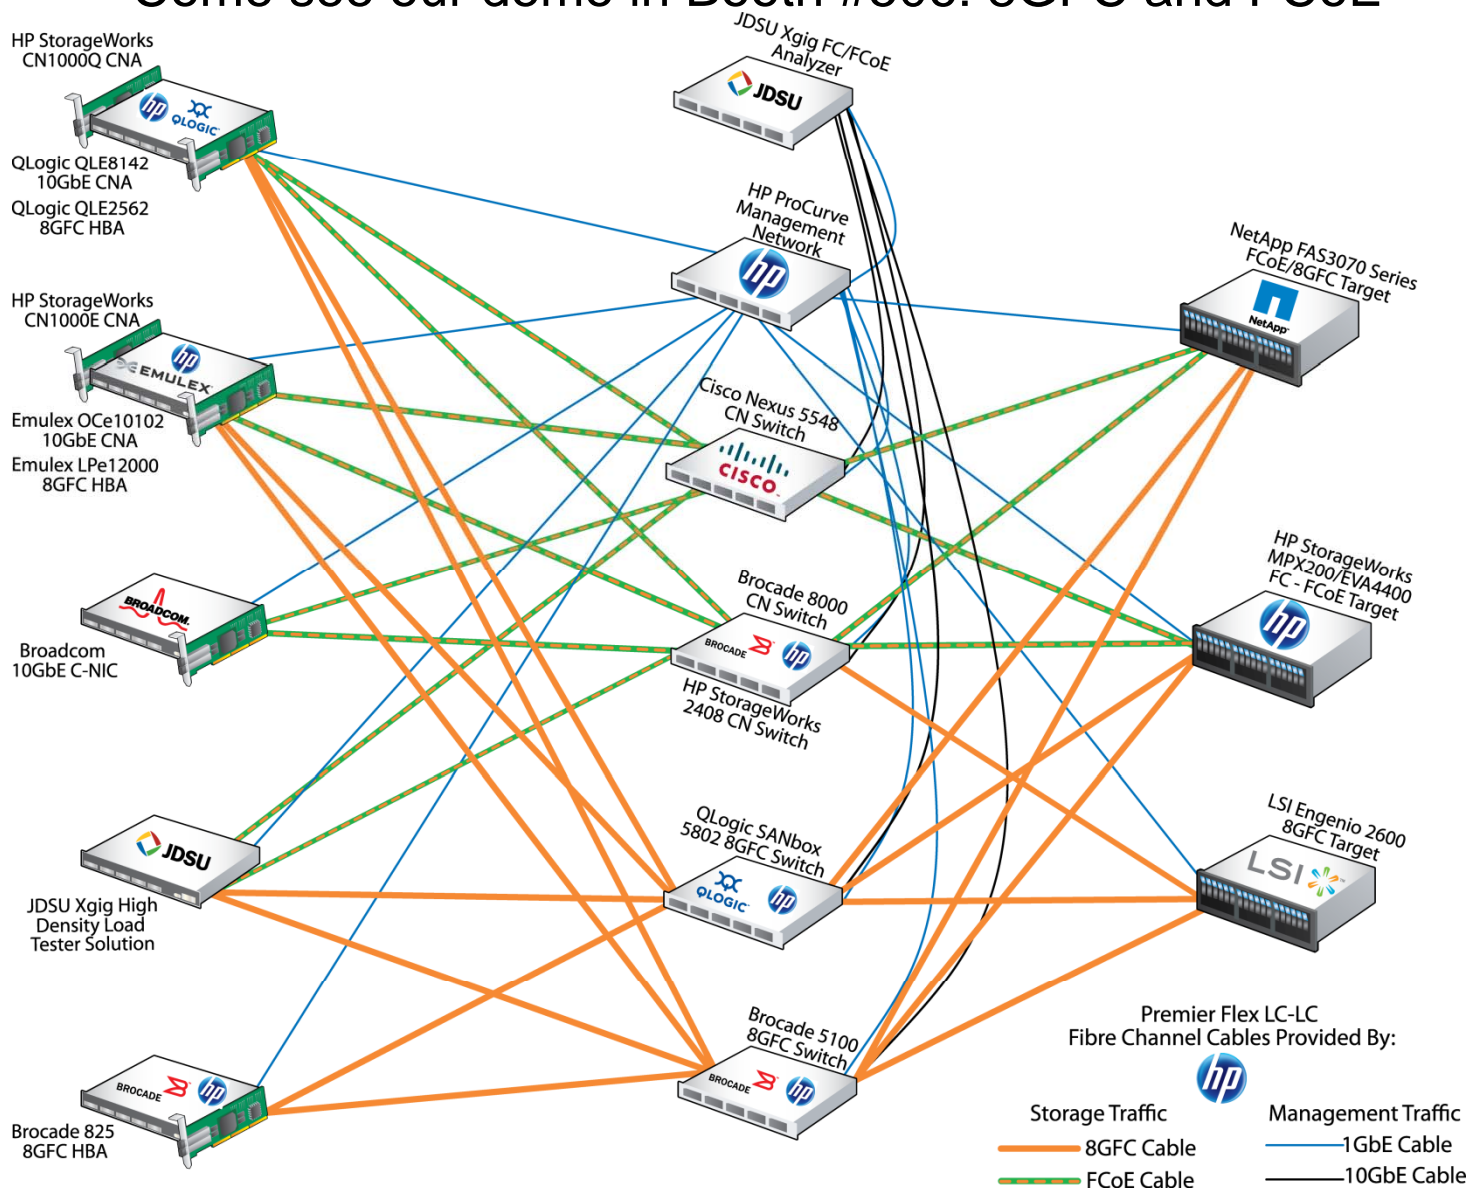
## Updated FC Port Model

- Drivers such as virtualization have driven the definition of a new FC port model
  - The new model is defined in FC-SW-5 and FC-FS-3
- The new model modifies the FC-2 level
- FC-2 defines the basic link level functionality
- FC-2 is now divided into 3 new sub-levels

## Summary: Fibre Channel...

- *Dominates* the SAN market
- Well understood
- Easy to learn, use and implement
- Protects and future-proofs storage investments
- Provides comprehensive solutions
- Vision for extending into additional markets
- Continuous speed & Bandwidth/\$ improvements
- Listening to customer needs
- Aggressively pursuing **Energy Efficiency**

Come see our demo in Booth #506: 8GFC and FCoE



## Thanks...and REMEMBER

**Please fill out our evaluations!**



This One

5 = "Aw, shucks. Thanks!"

4 = "Mighty kind of you!"

3 = "Hope you enjoyed this!"

2 = "Better luck next time."

1 = "We won't go here"

**Questions?**

Tom (T2) Hammond-Doel  
Thomas.Hammond-Doel@LSI.com





# **Driving Innovation Through the Information Infrastructure**

**SPRING 2011**