SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

Fibre Channel Technologies "Current & Future"

Presenters:
Dr. M. K. Jibbe, LSI Logic (ESG)
Tom Hammod-Doel, LSI Logic (ESG)
Steven Wilson, Brocade Communications Systems, Inc.

2/28/2007

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.

Abstract

The objectives of this tutorial are:

- Provide the user with a Primer on Fibre Channel
- Project the market outlook and roadmap of Fibre Channel
- Share what is New in Fibre Channel Standards for Protocols APIs, and Management.

What can FC provide today?

1. Flexible, Scalable relative to Topologies, Speed, Performance, Distance, Node connectivity and Low cost
2. Communication and Data Overhead (Framing, Data Communication, Latency, Efficiency, Routing Control, and Access Control),
3. Redundancy, Availability, and Failover,
4. Applicability in SAN (Easy of use), with large IT User Base

FC Topologies

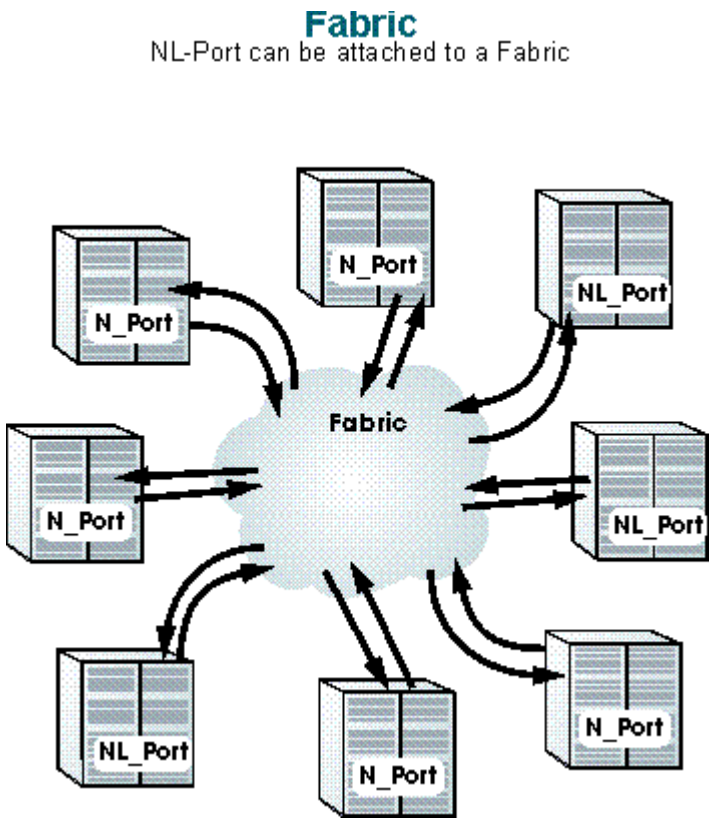

Switched Fabric

Up to 2/24 ports in a switched interconnect

Multiple concurrent communications for high aggregate throughout

- **Fibre channel supports a 24-bit address space**
  - ✓Provides 2 ^24 addresses
  - ✓FC routing is done based on NPort ID assigned on login (24-bit addressing consisting of Domain ID, Area ID, and Device ID)
- **FC Device ports are uniquely identified by a MAC ID (Worldwide Name)**
- **Address lookup is provided by the Fabric Switch using the Name Server portion of Directory Services**

FC Topologies

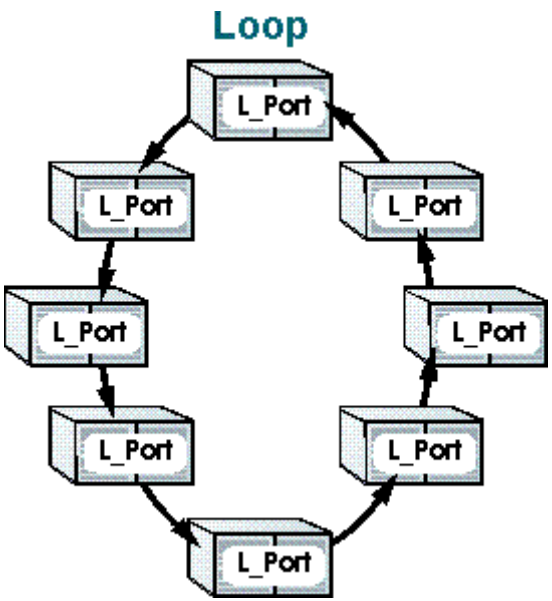

Arbitrated Loop → Up to 127 ports on a shared loop

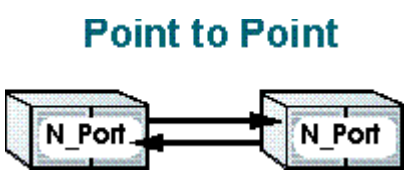

Point-to-Point → Two ports on a dedicated link

Topology Comparison

Attribute	Point to Point	Arbitrated Loop	Switched Fabric
Number of ports	2	2 to 127	Up to 2
Maximum bandwidth	Link rate times 2	Link rate times 2	Link rate times number of ports
Bandwidth allocation	Dedicated	Shared by all loop ports	Managed by fabric
Address assignment	N_Port Login	Loop initialization and Fabric Login	Fabric Login
Number of concurrent circuits	1	1	Number of port pairs (number of ports/2)
Effect of port failure	Point-to-point link fails	Loop fails (port bypass function required)	Link between switch and port fails
Concurrent maintenance	Link is down	May be disruptive to entire loop	Link between switch and port is down
Expansion	Add additional point-to-point links	Attach loop to fabric	Expand fabric
Redundancy/High Availability	Add redundant port and point-to-point links	Use dual loops and dual-ported devices	Use redundant switches
Link rates supported	All	All (all devices on loop must be same rate)	All (fabric may support mixed rates)
Media types supported	All	All	All
Classes of service supported	All	Class-1, -2 -3	All
Frame delivery order	In order	In order	Note 1
Access to interconnect medium	Dedicated	Arbitration	Dedicated
Cost per port	Port cost	Port cost + loop function (+hub if used)	Port cost + fabric port

Note 1: Frame Delivery Ordering is switch implementation dependent

Flexibility and Scalability

- Uses a common transport mechanism to support:
  - Physical interface types: Multi/Single Mode Fiber, and Cu
  - Traditional Channels: SCSI, IPI3, SBCCS, and HIPPI
  - Traditional Networks: IP, IEEE 802, and ATM
- High-speed -100/200/400/800/1200 MB/s, Reliable data transmission:
  - 100/200/400/800/1200 MB/s
  - BER < 10^{-12}
- Provide scalability of performance and cost
- Encourage industry support through open standards
- Designed to fulfill the needs of SANs

Speed

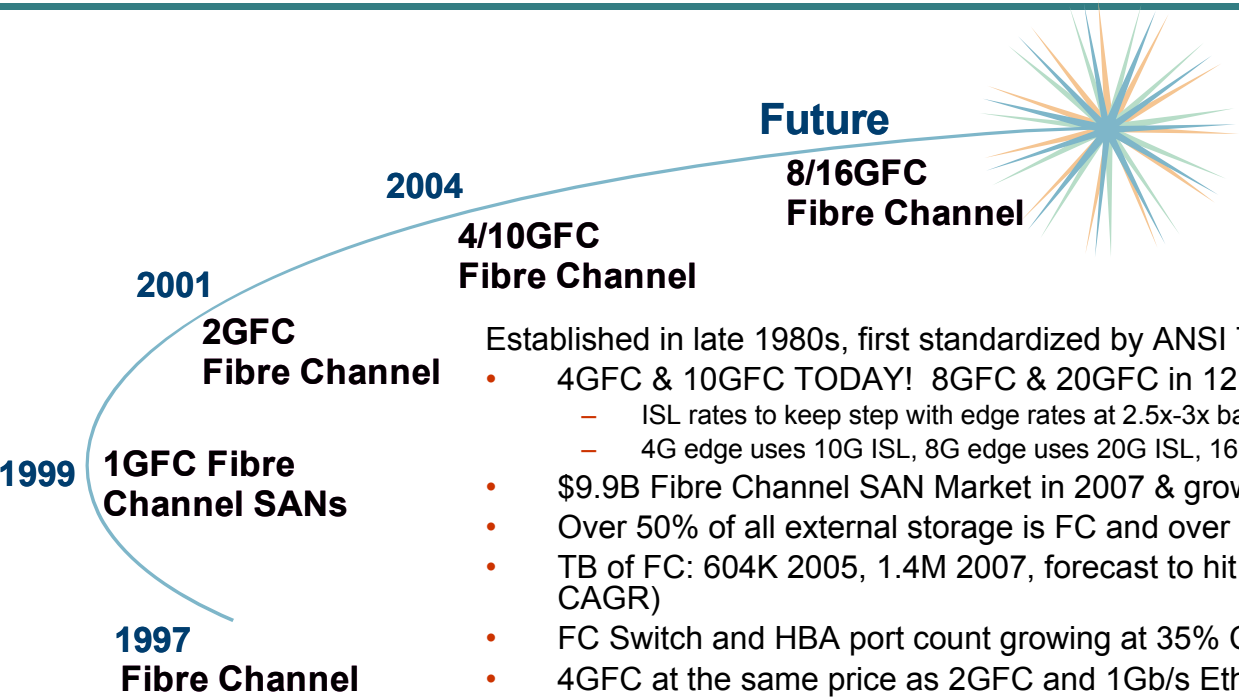

Established in late 1980s, first standardized by ANSI T11 in 1994

- 4GFC & 10GFC TODAY! 8GFC & 20GFC in 12 – 18 months
  - ISL rates to keep step with edge rates at 2.5x-3x bandwidth premium over edge rates
  - 4G edge uses 10G ISL, 8G edge uses 20G ISL, 16G edge uses 40G ISL, etc..
- $9.9B Fibre Channel SAN Market in 2007 & growing faster than total storage market
- Over 50% of all external storage is FC and over 90% of all SAN is FC
- TB of FC: 604K 2005, 1.4M 2007, forecast to hit 4.4M Terabytes of FC storage by 2010 (50% CAGR)
- FC Switch and HBA port count growing at 35% CAGR
- 4GFC at the same price as 2GFC and 1Gb/s Ethernet
- 4GFC is plug-compatible with 1 and 2GFC (devices auto negotiate w/o user intervention)
- Applications are driving higher data rates (i.e. Video, back-up times; more data to backup, less time to do it)
- R/W operations on 4GFC HBAs show dramatic improvements with 98.5% real utilization on saturated lines
- Serial SCSI; FCP protocol
- Minimal error rates for network technology
- Credit-based flow control (no dropped frames)

Maximum Distance

Media Type	Transmitter	Speed	Distance	Variant
Electrical (Differential)	ECL/PECL	400 MB/s	0m – 10m (typical)	400-DF-EL-S
		200 MB/s	0m – 10m (typical)	200-DF-EL-S
		100 MB/s	0m – 30m (typical)	100-DF-EL-S
9 um. Single-Mode Fiber	1550 nm. Long wave Laser	400 MB/s	2m - >50km	400-SM-LL-V
		200 MB/s	2m - >50km	200-SM-LL-V
		100 MB/s	2m - >50km	100-SM-LL-V
	1300 nm. Long wave Laser	400 MB/s	2m - 2km	400-SM-LL-I
		200 MB/s	2m - 2km	200-SM-LL-I
		100 MB/s	2m - 10km	100-SM-LL-L
			2m - 2km	100-SM-LL-VI
50 um. Multi-Mode Fiber	850 nm. Short-wave Laser	400 MB/s	0.5m - 175m	400-M5-SN-I
		200 MB/s	0.5m - 300m	200-M5-SN-I
		100 MB/s	0.5m - 500m	100-M5-SN-I
62.5 um. Multi-Mode Fiber		400 MB/s	0.5m - 70m	400-M6-SN-I
		200 MB/s	0.5m - 150m	200-M6-SN-I
		100 MB/s	0.5m - 300m	100-M6-SN-I

- **2 Km distance with Multi- mode Fibre**
- **10 Km distance with Single Mode Fibre**
- **5000 Km distance with FC over IP**

Frames

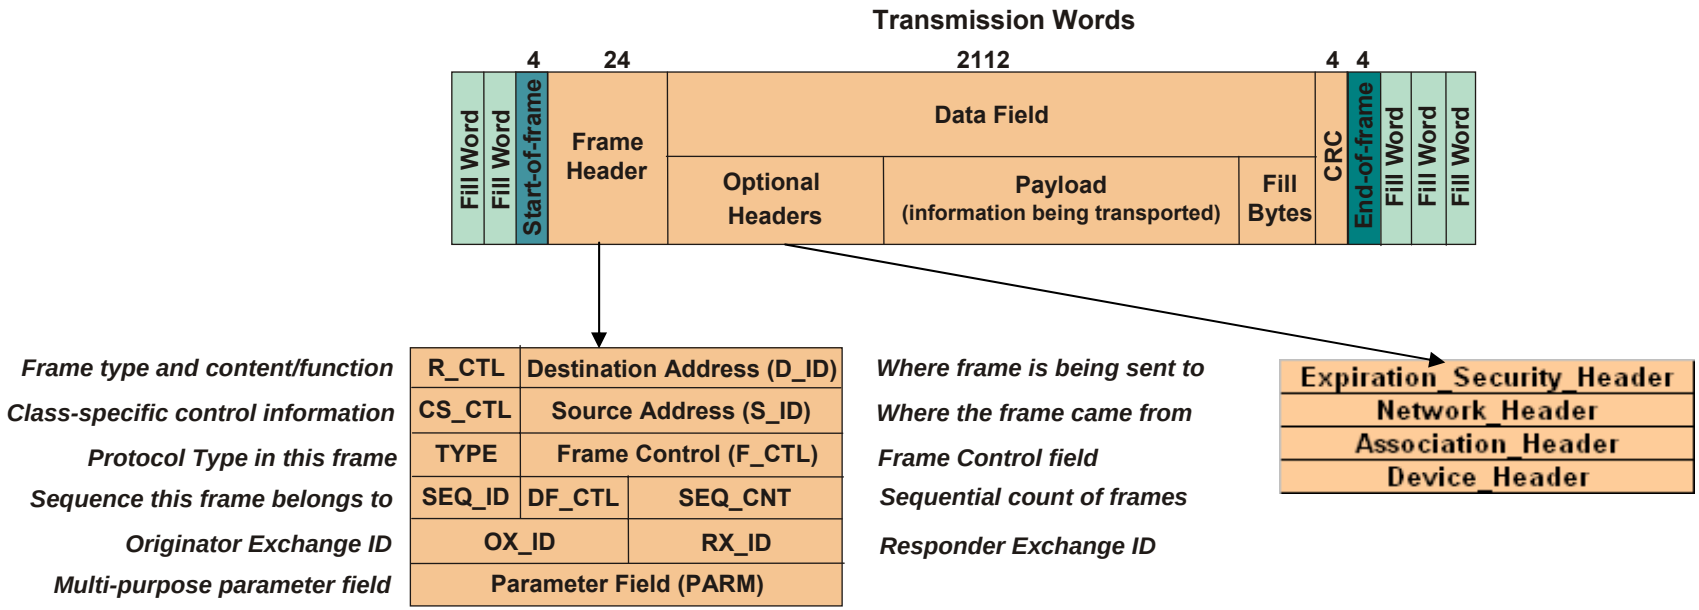

- Flexibility
  - Fly by Frame handling
  - Out of order
- Speed
- Routing

Data Traffic with Exchanges

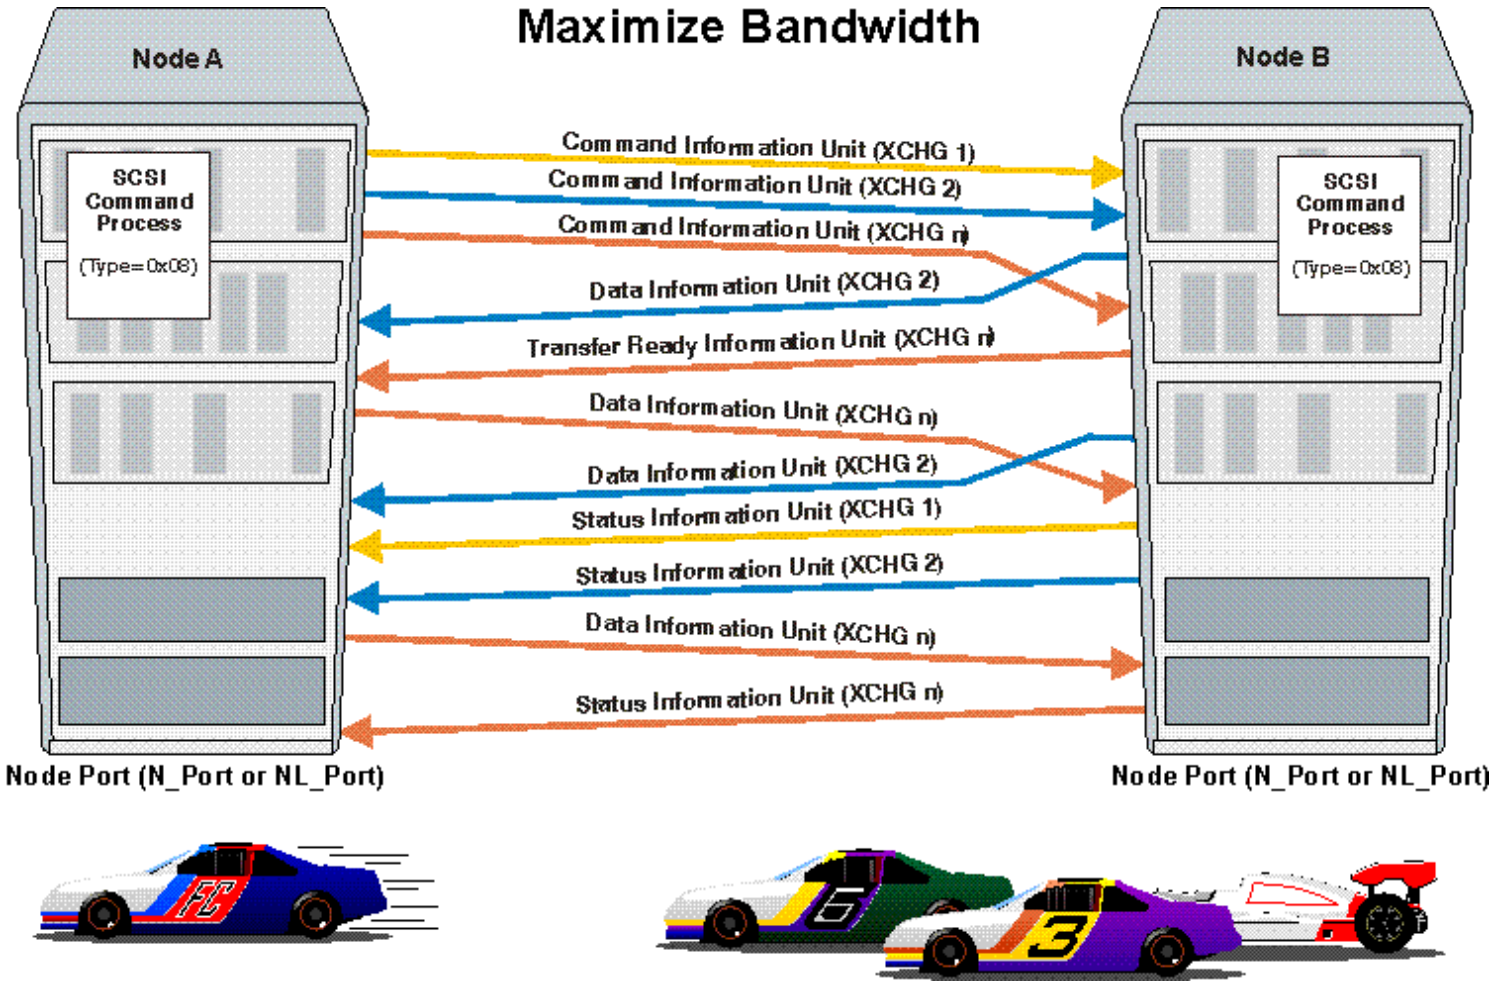

Establishing Operating Environment

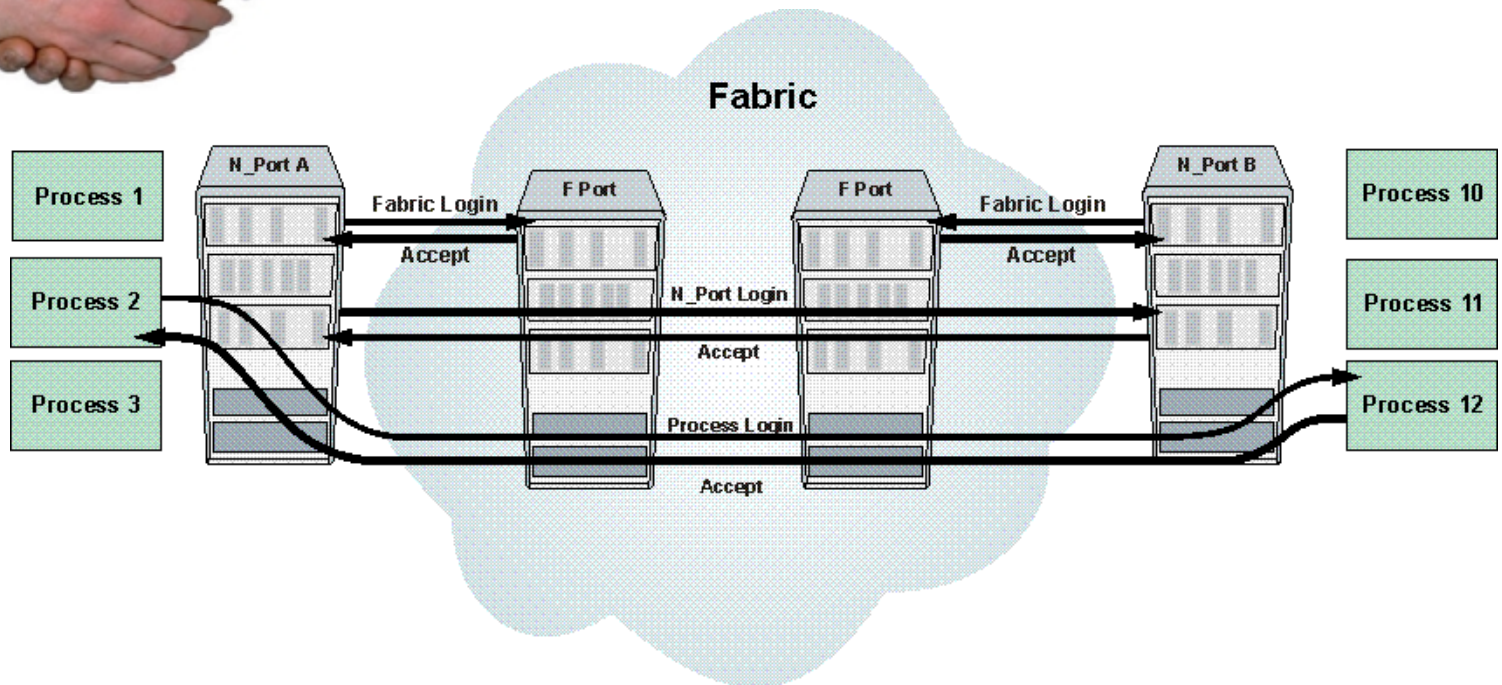

Flow Control: Access Control, Latency, and Efficiency

Login Buffer to Buffer

- Node to Fabric
- Fabric to Node

Login Node to Node

Flow control is credit based

- Buffer-To-Buffer Credit
- Control pace of frame transmission
- Each R_RDY received increments the available BB_Credit value

Latency

- Across a single switch, average latencies are less than 400 nanoseconds.

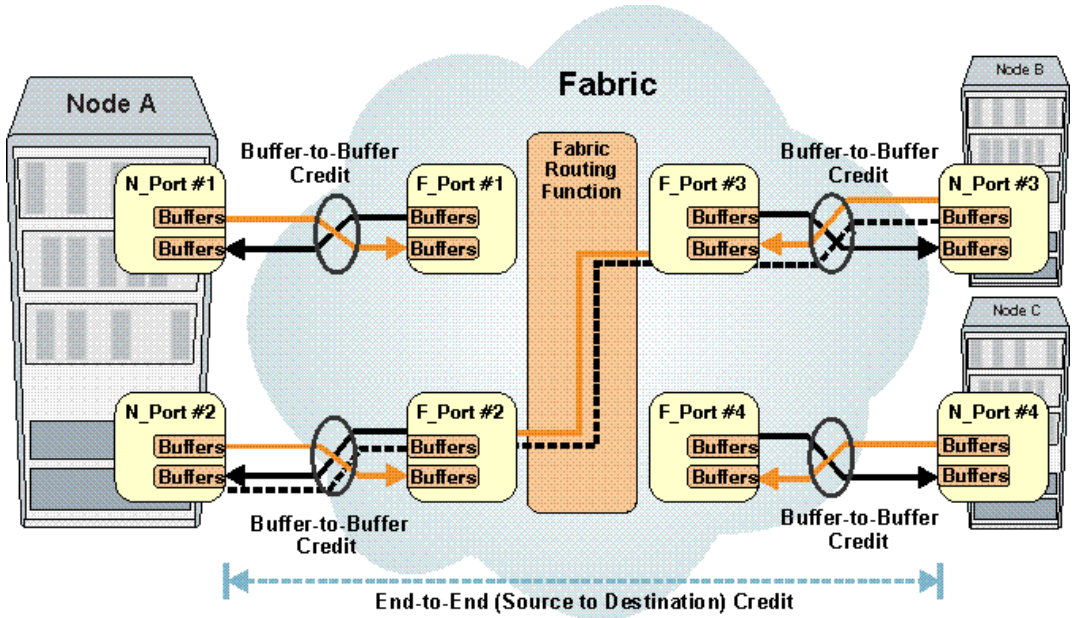

FC Routing

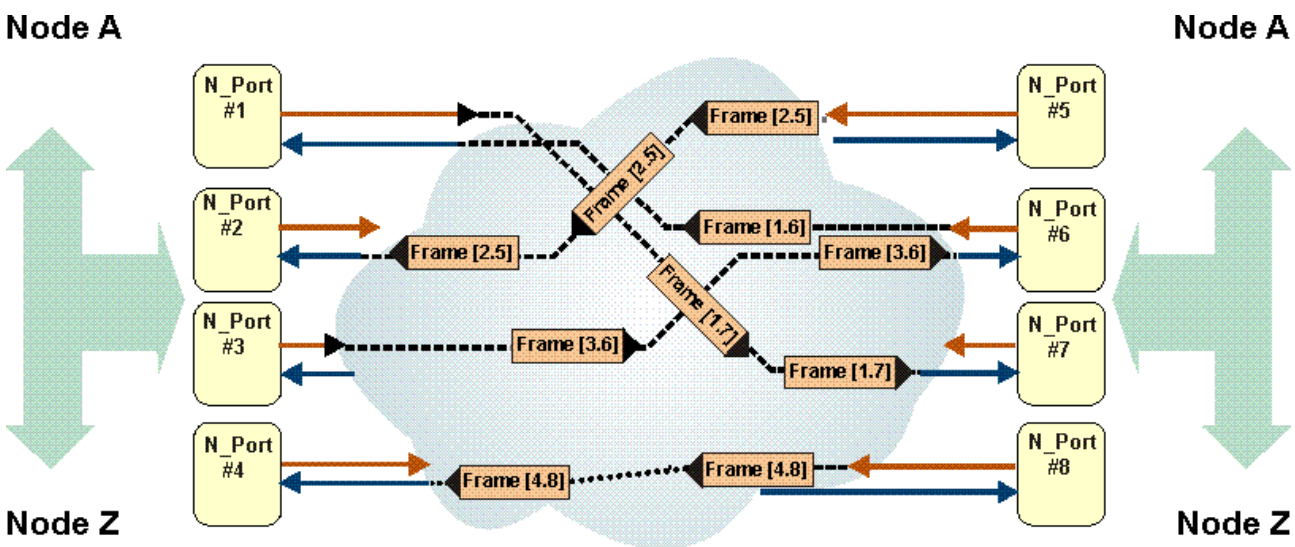

- Connect Any to Any
- Maximize Connectivity
- Simplex and Duplex

FC Access Control

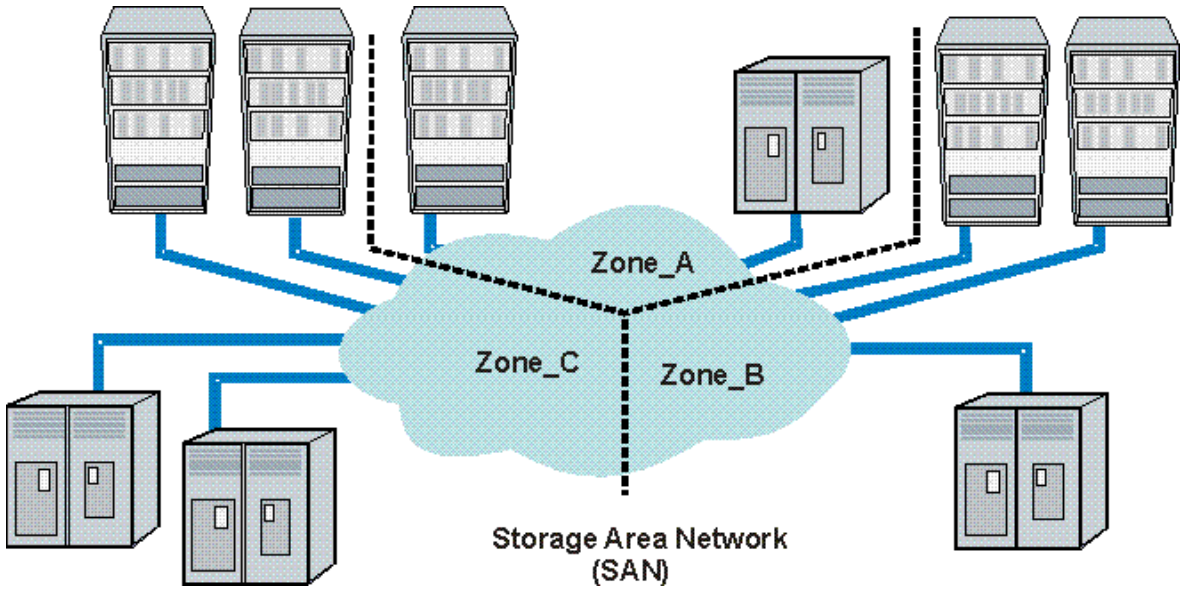

Divide

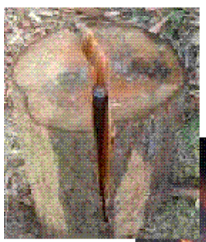

and *conquer*

Provide Accessibility
Soft Zoning: Employs the Name Server to limit the information returned to an initiator in response to a query.
Hard Zoning: Enforced by the Fabric

Redundancy, Availability, and Failover

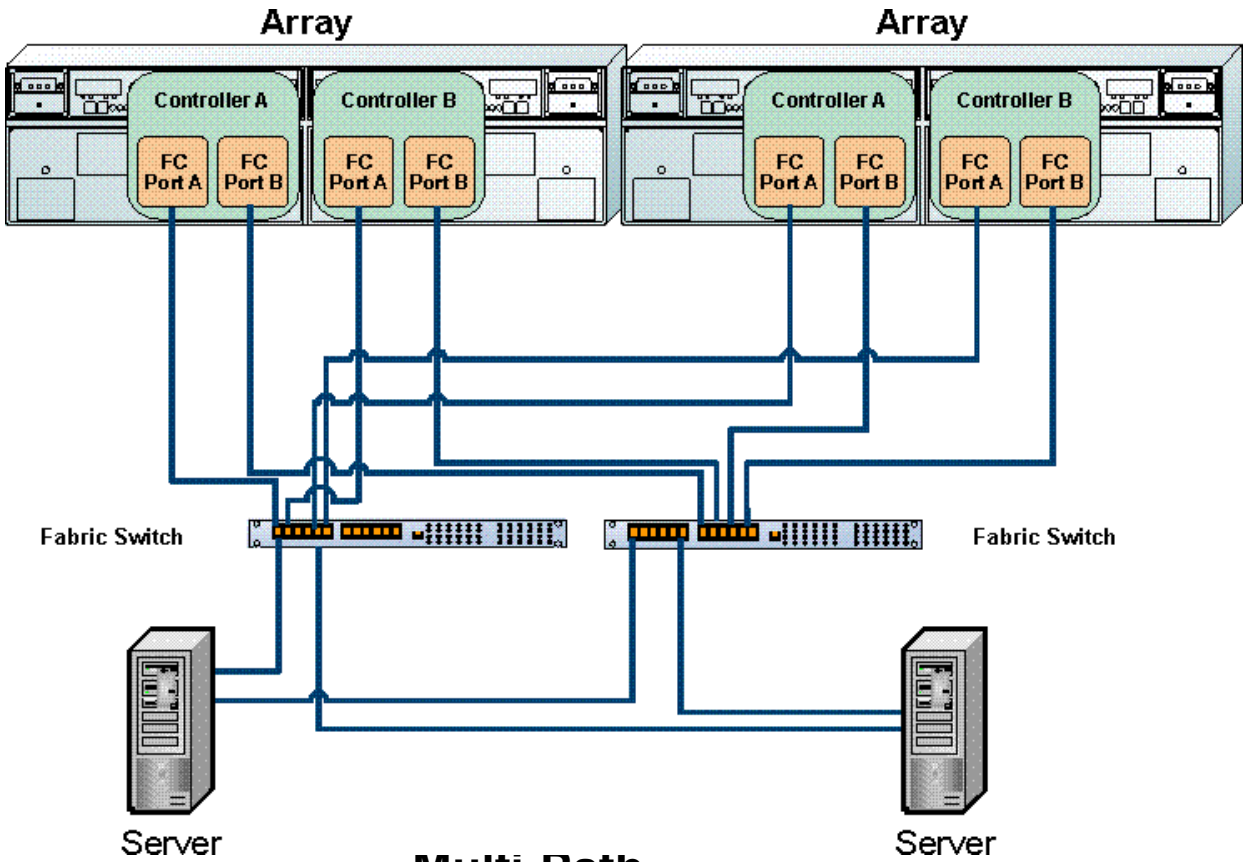

- **Multi-Path**
- **Full Redundancy**
- **Path Failover**

FC Products Dominant in Enterprise Datacenter

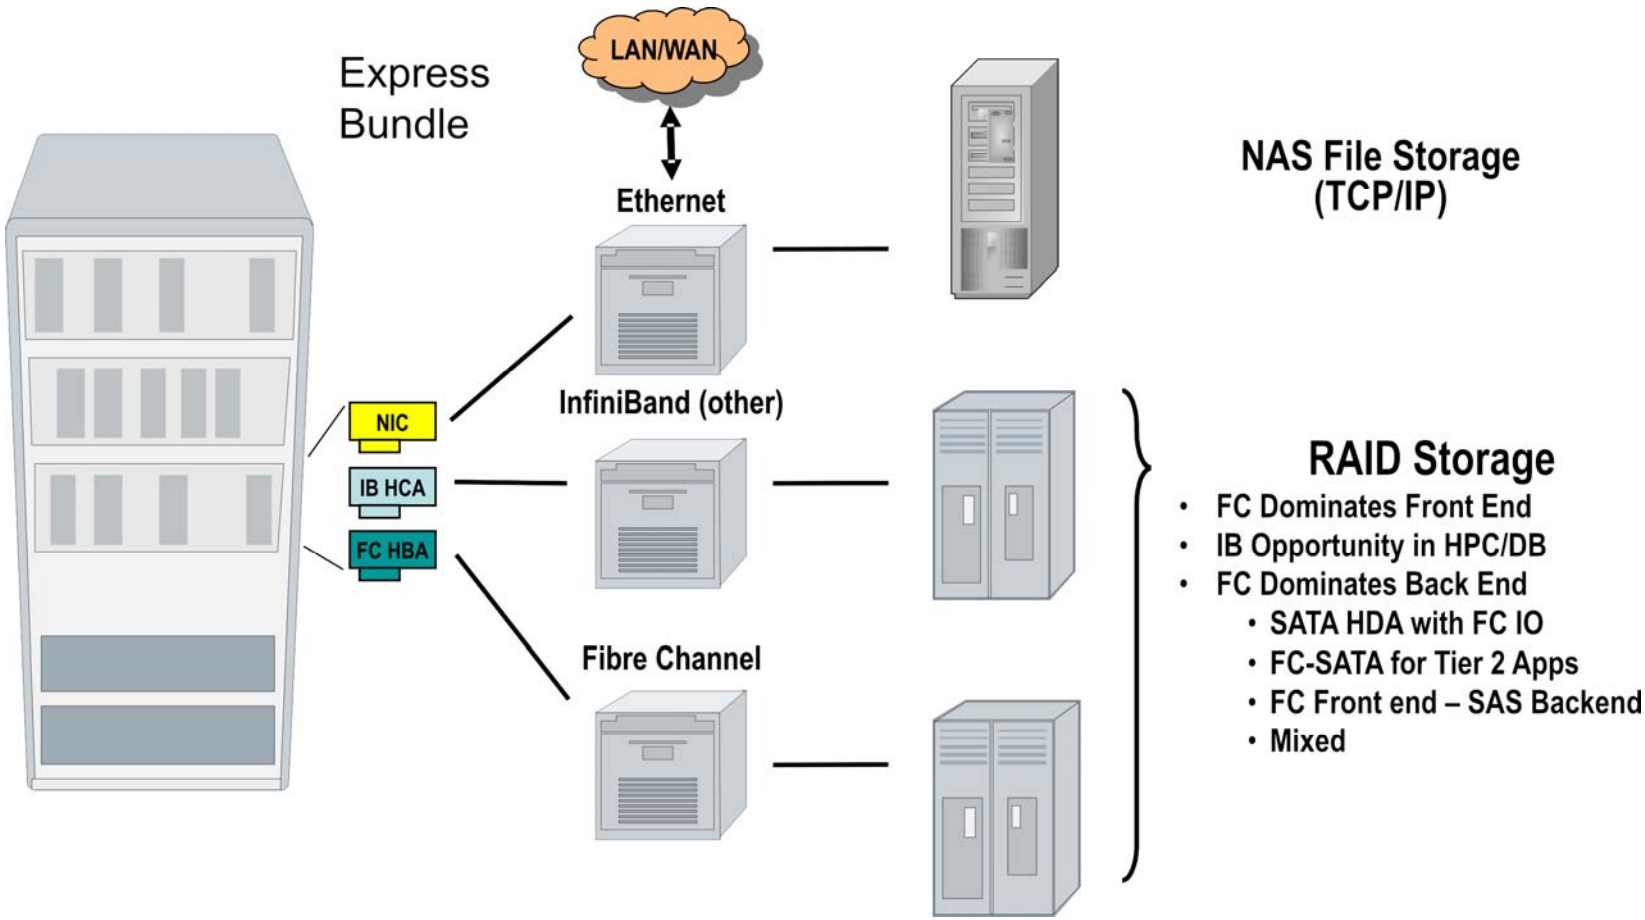

Detailed RAID Storage: Array Module with Different Drive Technologies

SNIA
Storage Networking Industry Association
EDUCATION

There's a lot more under the Fibre Channel hood than appearances reveal:

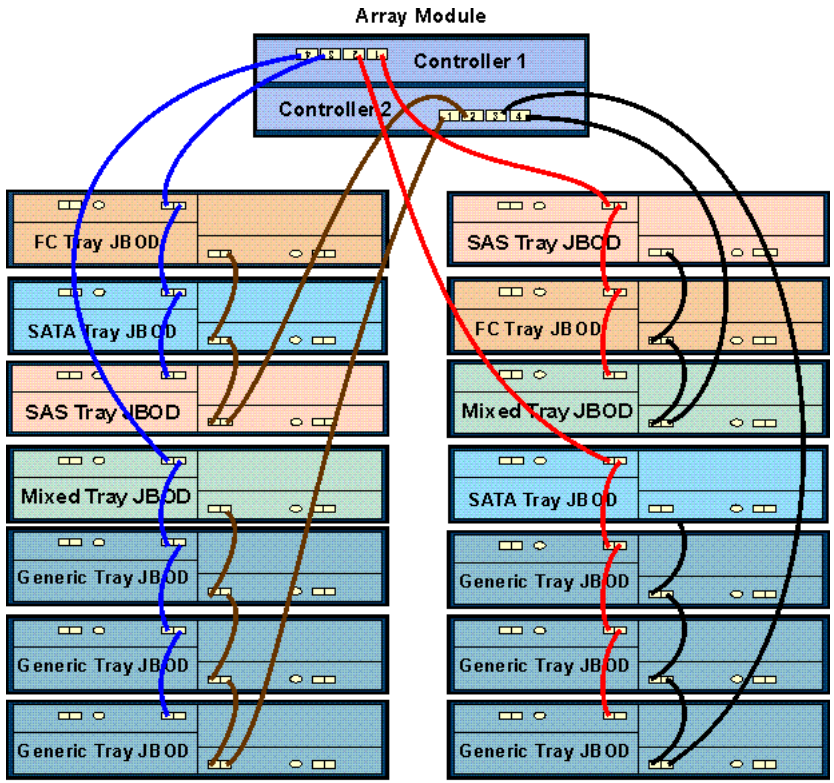

Array Module with Different Drive Tray Types

FC Product Performance – IOPS: Host Interface – Drive Interface

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION
EDUCATION

	Drive Type	Dual 4 GFC	Quad 4 GFC	Future 8 GFC
Burst I/O rate cache reads (512 byte)		125K	125K	200K
Sustained I/O rate disk reads (4k – R5)	FC	40k	40k	80K
	SAS			70K
	SATA			12K
Sustained I/O rate disk writes (4k- R5) - CMD	FC	9k	9k	15K
	SAS	8K	10K	12K
	SATA	2K	2K	4K
Number of drives required for benchmark test and code thread	FC, SAS, / SATA	96D / 8T	96D / 8T	96D / 8T

- FC continues to evolve with different technologies

FC Product Performance – MB/sec

	Drive Type	Dual FC	Quad FC	Future 8 GFC
Sustained throughput cache read (512k)		1600 MB/s	1800 MB/s	3000 MB/s
Sustained throughput disk read (512k)	FC	850 MB/s	850 MB/s	1600MB/s
	SAS	800 MB/s	800 MB/s	1200 MB/s
	SATA	800 MB/s	800 MB/s	900 MB/s
Sustained throughput disk write (512k)	FC	800 MB/s	800 MB/s	1600 MB/s
Cache mirroring disabled	SAS	750 MB/s	750 MB/s	750 MB/s
Cache mirroring disabled	SATA	750 MB/s	750 MB/s	750 MB/s
Number of drives required for benchmark test and code thread	FC	48D / 8T	48D / 8T	48D / 8T

- FC continues to evolve with different technologies

Current Cost Comparison Table

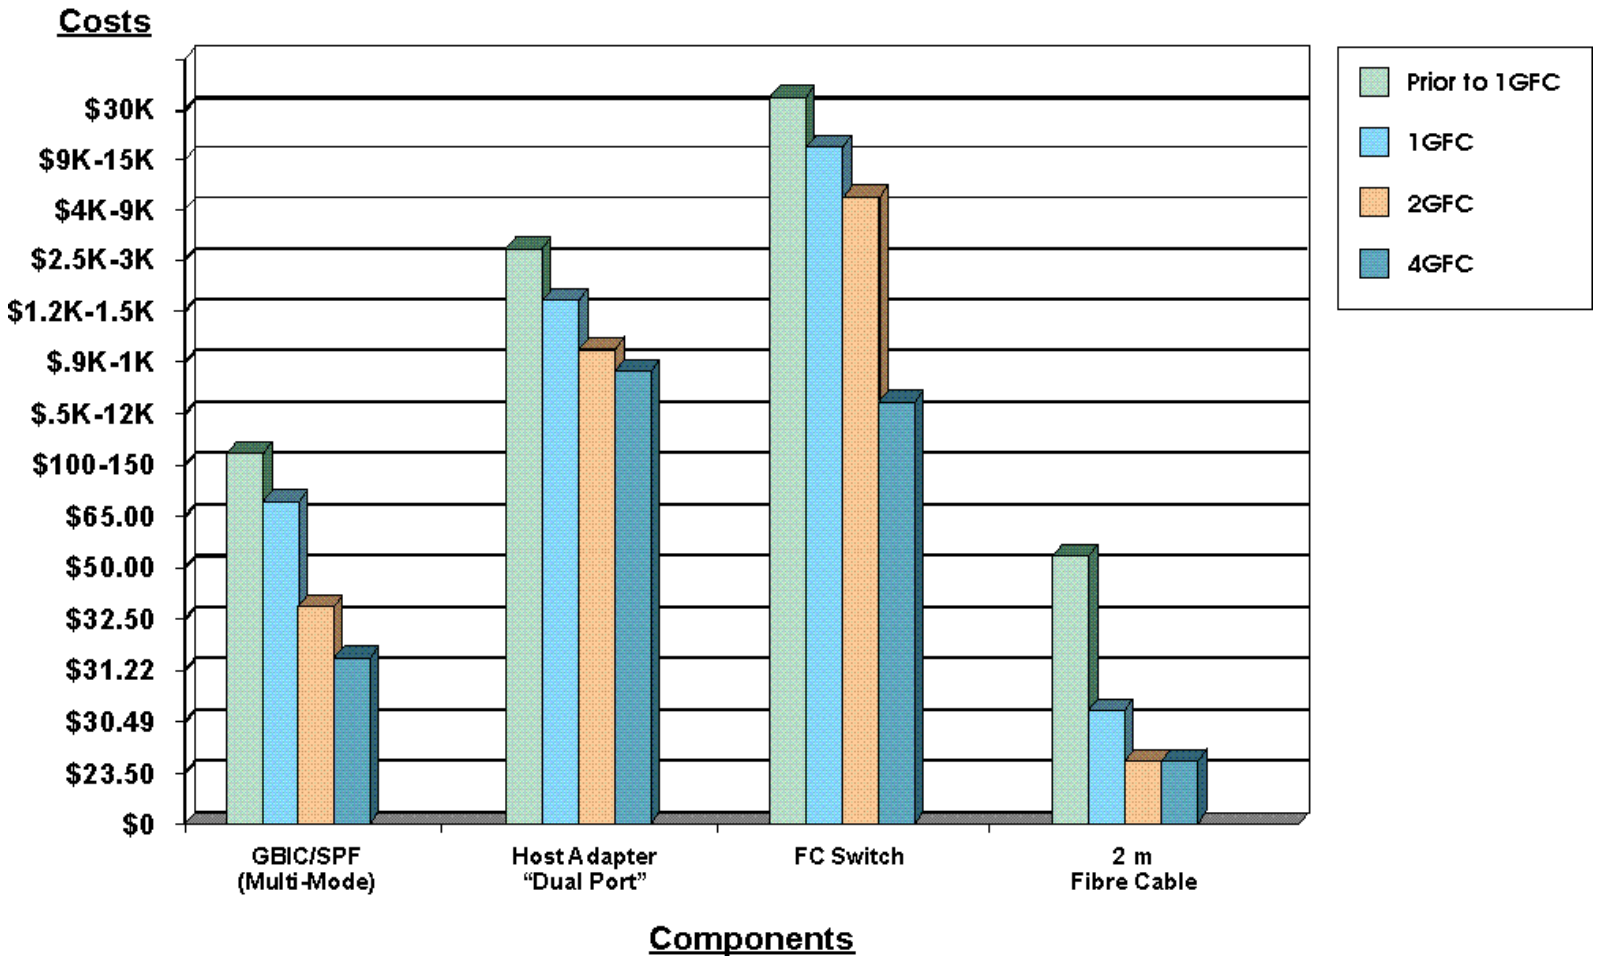

Fibre Channel is Here to Stay.

SNIA
Storage Networking Industry Association
EDUCATION

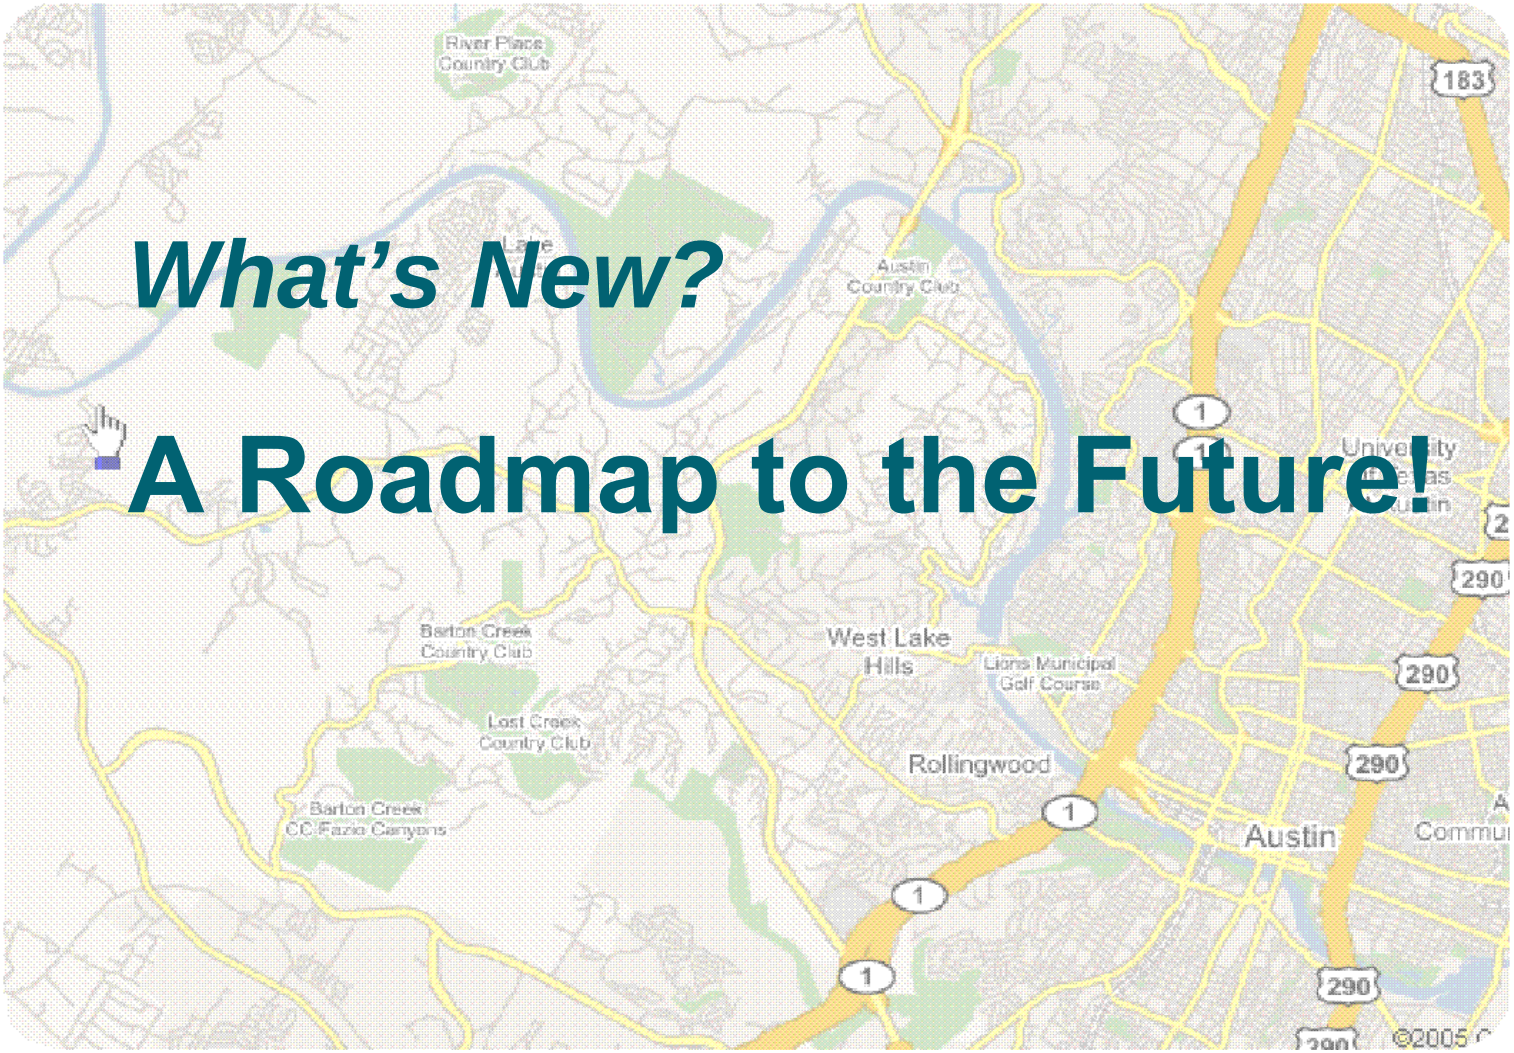

What's New?

A Roadmap to the Future!

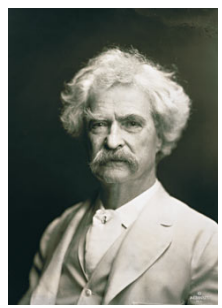

The report of my death was an exaggeration.

Continuous SAN Market Growth

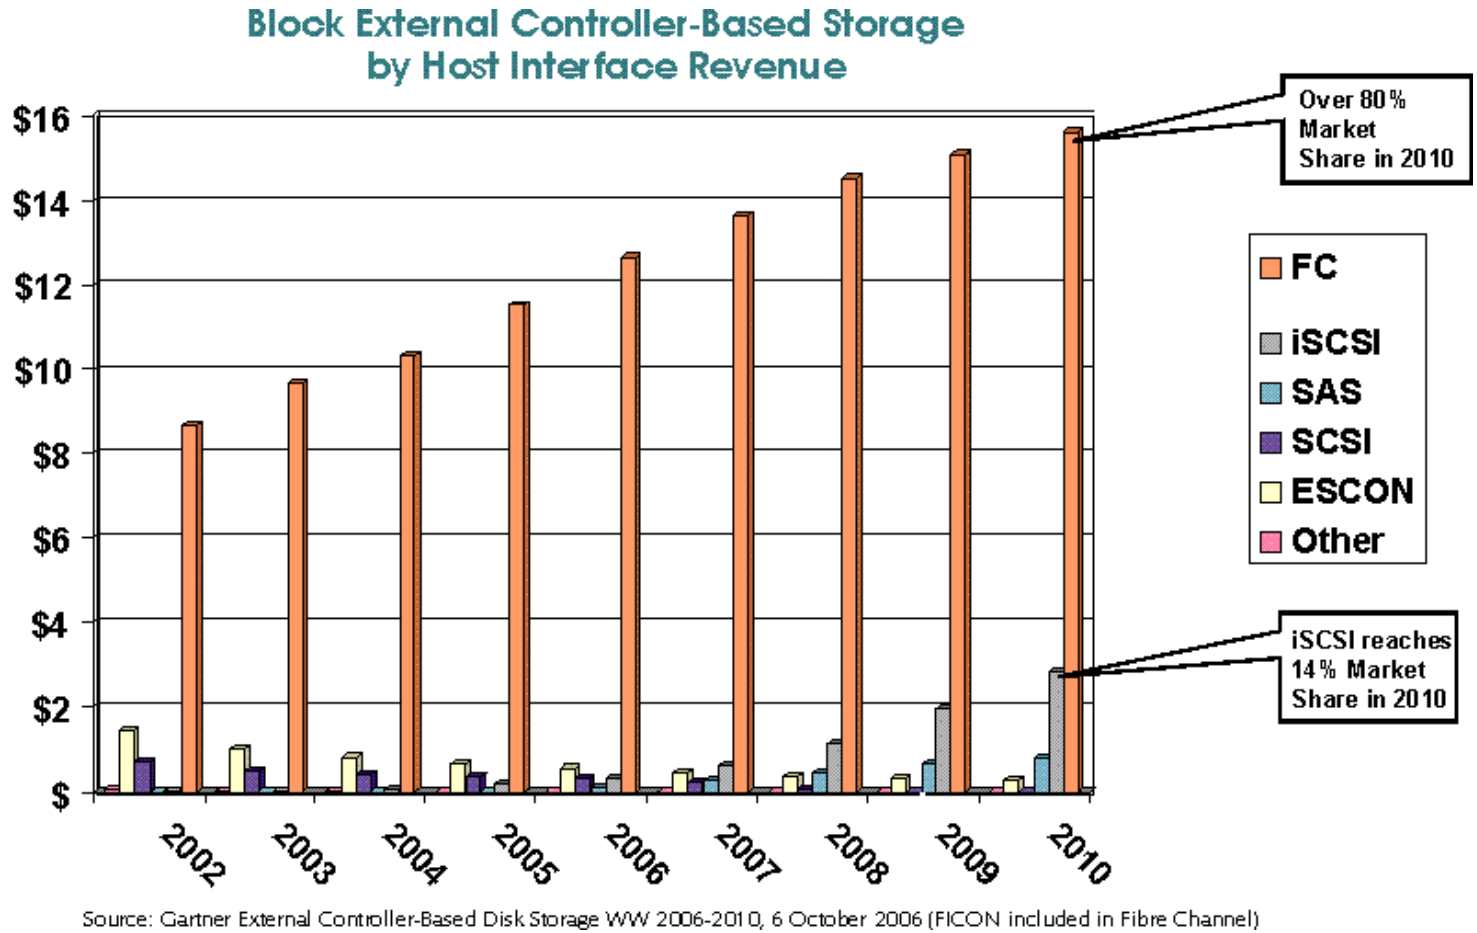

Source: Gartner External Controller-Based Disk Storage WW 2006-2010, 6 October 2006 (FICON included in Fibre Channel)

Fibre Channel's Continuous Evolution

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION
EDUCATION

FC has been the major storage system interconnect since the mid 90s

- FC dominates the SAN and external storage market place

How will FC continue to Meet customers' evolving needs?

- Faster speeds
- Bandwidth/Cost leadership
- Investment protection
- Additional capabilities
- Lower cost solutions
- Simplified solutions (Plug-n-play)

Continuous Back-End Market Growth

SNIA
Storage Networking Industry Association
EDUCATION

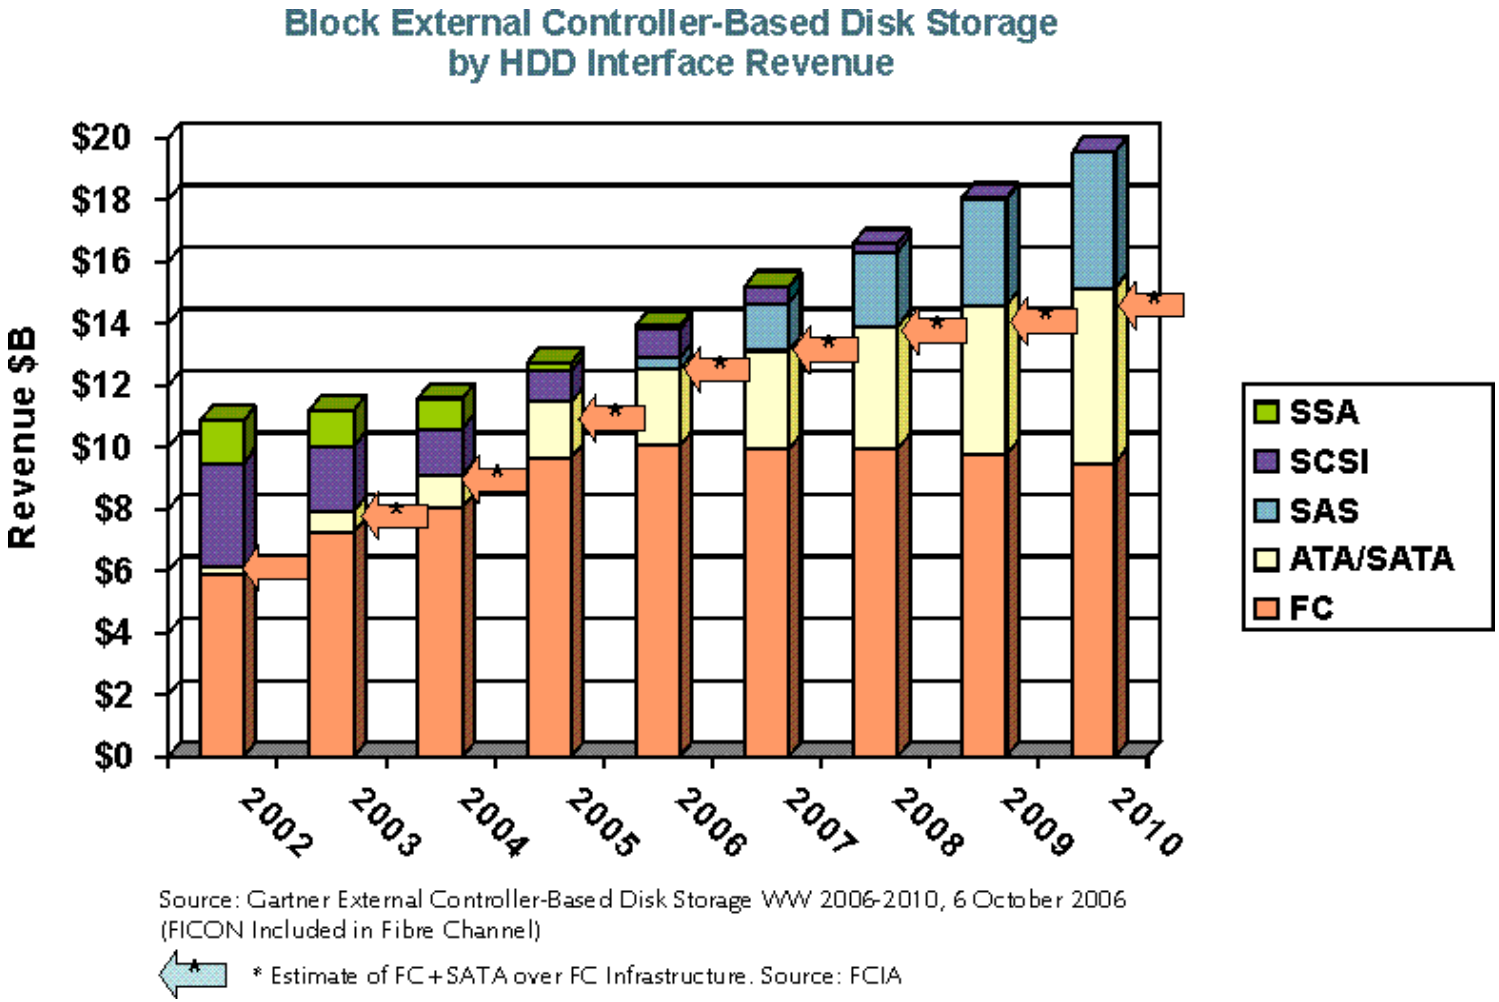

FCIA Fibre Channel Speeds – 3 Connection Types

SNIA
Storage Networking Industry Association
EDUCATION

FC specifies 3 connection types

- FC-Base2
- FC-Base10
- FC-BaseT
- *All speeds of each type Auto-negotiate best speed w/o any user intervention!*
- *Each speed within its connection type is backward compatible 2 generations!*

FC-Base2

- **Predominant FC interconnect**
- **Used for fabric Edge and ISL**
- **Also used for Disk and Tape Drives**
- **All speeds single lane serial** streams
- Optics and copper cabling, SFP/SSF

FC-Base10

- Used for ISL (2.5x-3x bandwidth of edge)
- *4G Edge/10G ISL migrates to 8G Edge/20G ISL migrates to 16G Edge/40G ISL, etc*

FC-BaseT

- FC using Cat5e/6/6a infrastructures
- Copper only (Cat5e/6/6a cables)
- RJ-45 connector
- User can use FC without changing any existing or new Ethernet cabling!

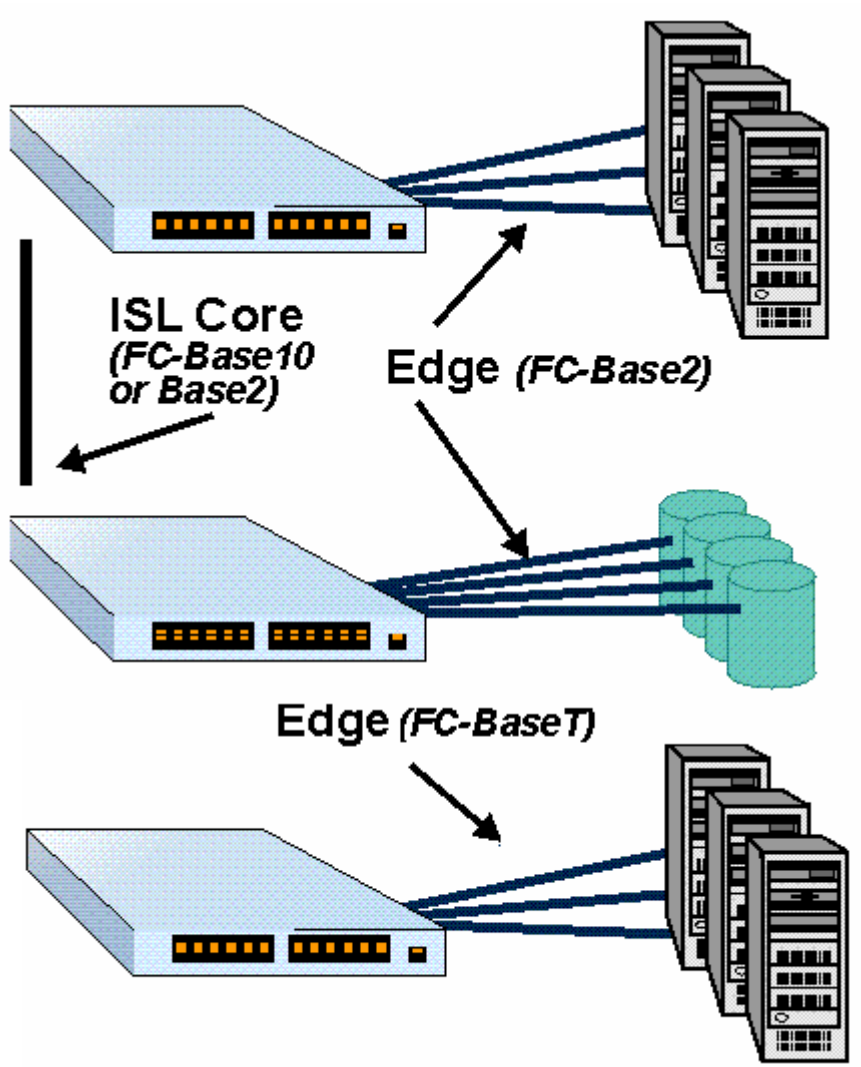

FCIA Fibre Channel Speed Roadmap

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION
EDUCATION

	Product Naming	Throughput (MBps)	Line Rate (GBaud)†	T11 Spec Technically Completed (Year)‡	Market Availability (Year)‡
Base2*	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	17	2009	2011
	32GFC	6400	34	2012	Market Demand
	64GFC	12800	68	2016	Market Demand
	128GFC	25600	136	2020	Market Demand
Base10**	10GFC	2400	10.52	2003	2004
	20GFC	4800	21.04	2007	2008
	40GFC	9600	42.08	TBD	Market Demand
	80GFC	19200	84.16	TBD	Market Demand
	160GFC	38400	168.32	TBD	Market Demand
BaseT***	1GFC	200	1.0625	2006	2007
	2GFC	400	2.125	2006	2007
	4GFC	800	4.25	2006	2007
	8GFC	1600	8.5	TBD	Market Demand
	10GFC	2400	10.52	TBD	Market Demand

***Base2** used throughout all applications for Fibre Channel infrastructure and devices. Each speed maintains backward compatibility at least two previous generations (I.e., 4GFC backward compatible to 2GFC and 1GFC)
****Base10** commonly used for ISLs, core connections, and other high speed applications demanding maximum bandwidth.
*****BaseT** used in common Ethernet copper infrastructures incorporating CAT5e/6/6a cables and RJ-45 connectors
†**Line Rate**: All Base2 speeds are single-lane serial stream ‡**Dates**: Future dates estimated

Courtesy of FCIA

FCIA "Condensed" Roadmap (Speed Gb/s)

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION
EDUCATION

- **FC-Base2 (Edge, Backend, and ISL)**
  - *1GFC, 2GFC, 4GFC* shipping today
  - 8GFC Ships in 12-18 months
  - 16GFC, 32GFC, 64GFC, 128GFC
- **FC-Base10 (ISL)**
  - 10GFC shipping today
  - 20GFC ships in 12-18 months
  - 40GFC, 80GFC, 160GFC
    - 100GFC under study (leverage IEEE 802.3 work)
- **FC-BaseT (Edge)**
  - new 2006 standard for Ethernet RJ45 Cat5/6 copper)
  - 1GFC, 2GFC, 4GFCF, ships in 12-18 months
  - 8GFC, 10GFC
    - 8GFC follows typical FC trend
    - 10G follows typical Ethernet trend

Fibre Channel Is Being Improved According To Real Customer Requirements

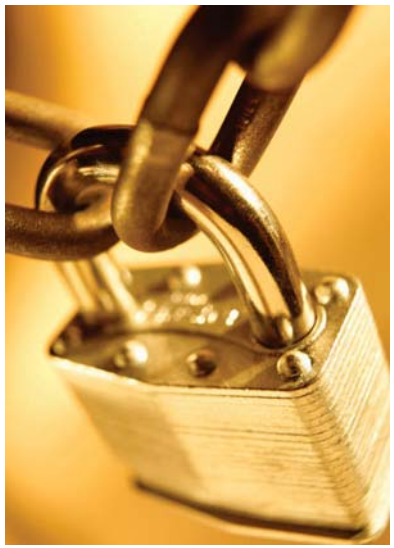

- New Fibre Channel Standards for
  - Management And Ease Of Use
  - Operational Flexibility and Scalability
  - Security

Management Improvements

- **Fabric Device Management Interface**
  - HBA Information Can Be Retrieved From The Fabric
- **Fibre Channel Open Management**
  - SMI-S
  - SNMP MIB Development
- **Improvements to the Fabric Configuration Server**
  - Advanced Topology Discovery and Bulk Data Retrieval
- **Common Transport**
  - Session Semantics Have Been Added
- **Diagnostic Tools**
  - FC Trace Route and Ping

Operational Flexibility

- **FAIS: Fabric Application Interface Specification**
  - Allows fabric to host certain applications
- **Event Server**
  - More Granular Event Registration
- **Virtual Channels**
  - Enables Traffic Differentiation On Links
- **Enhanced Commit Service**
  - Fabric Locking More Granular

Operational Flexibility

- **Frame Tagging**
  - Enables Virtual Fabrics
- **Routing Architectures and Models**
  - Allows Devices On Distinct Fabrics To Communicate Without a Merge
- **FC-SATA: SATA Tunneling over FC**
  - Brings native tiered storage to FC
  - FC SATA: An FC-4 mapping of the Serial ATA storage interface protocol to Fibre Channel

Operational Flexibility

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION
EDUCATION

- What Is NPIV?
  - Acronym for N-port ID virtualization.
  - Additional attribute of an F-port.
  - FLOGI request allocated the base PID 0xddaa00.
  - FDISC(SID=0) requests allocate virtual PIDs: 0xddaa01, 0xddaa02, 0xddaa03 …
  - Used by multiple virtual machines emulated on a physical machine.

NPIV Overview

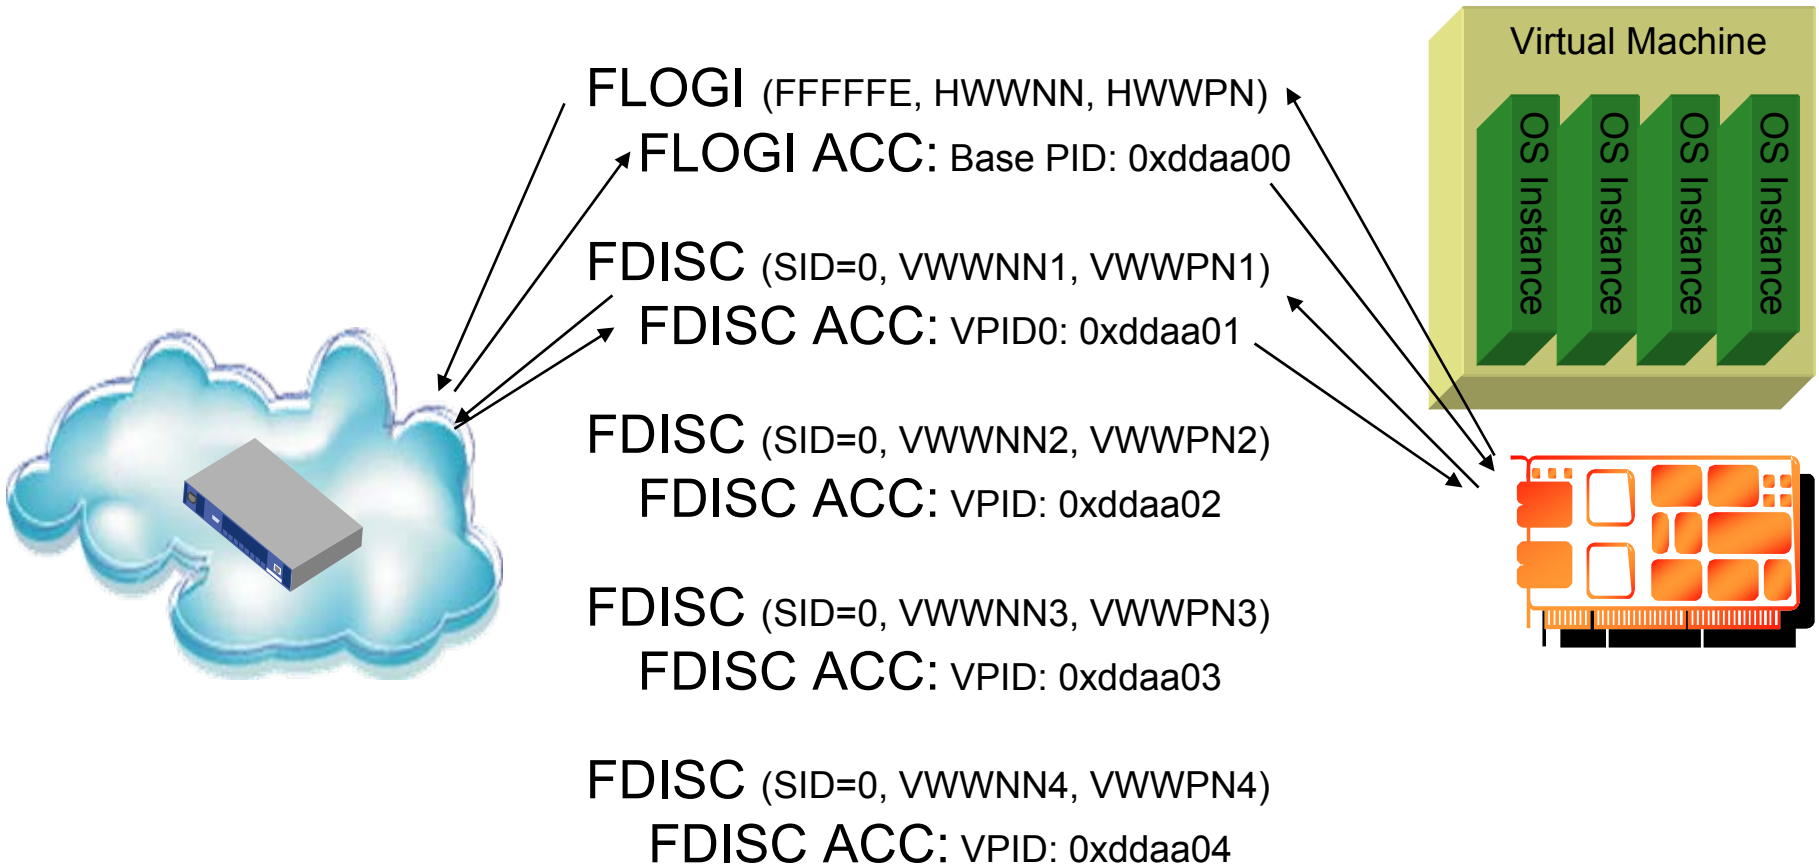

Storage Virtualization

SNIA
Storage Networking Industry Association
EDUCATION

- Three types of storage virtualization

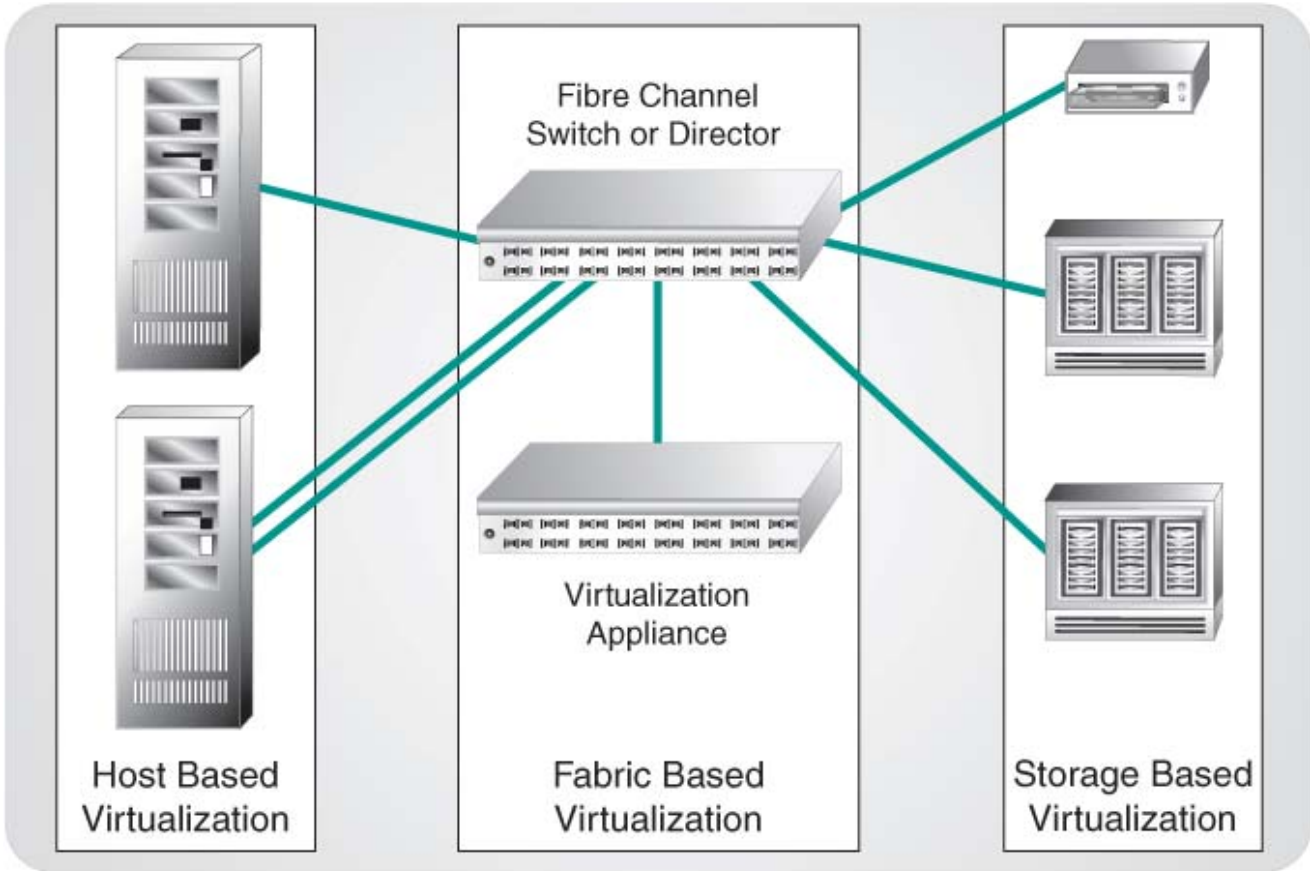

Refer to SNIA Virtualization Tutorials

Fabric Based Virtualization

- FC-Fabric Application Interface Specification (FAIS)

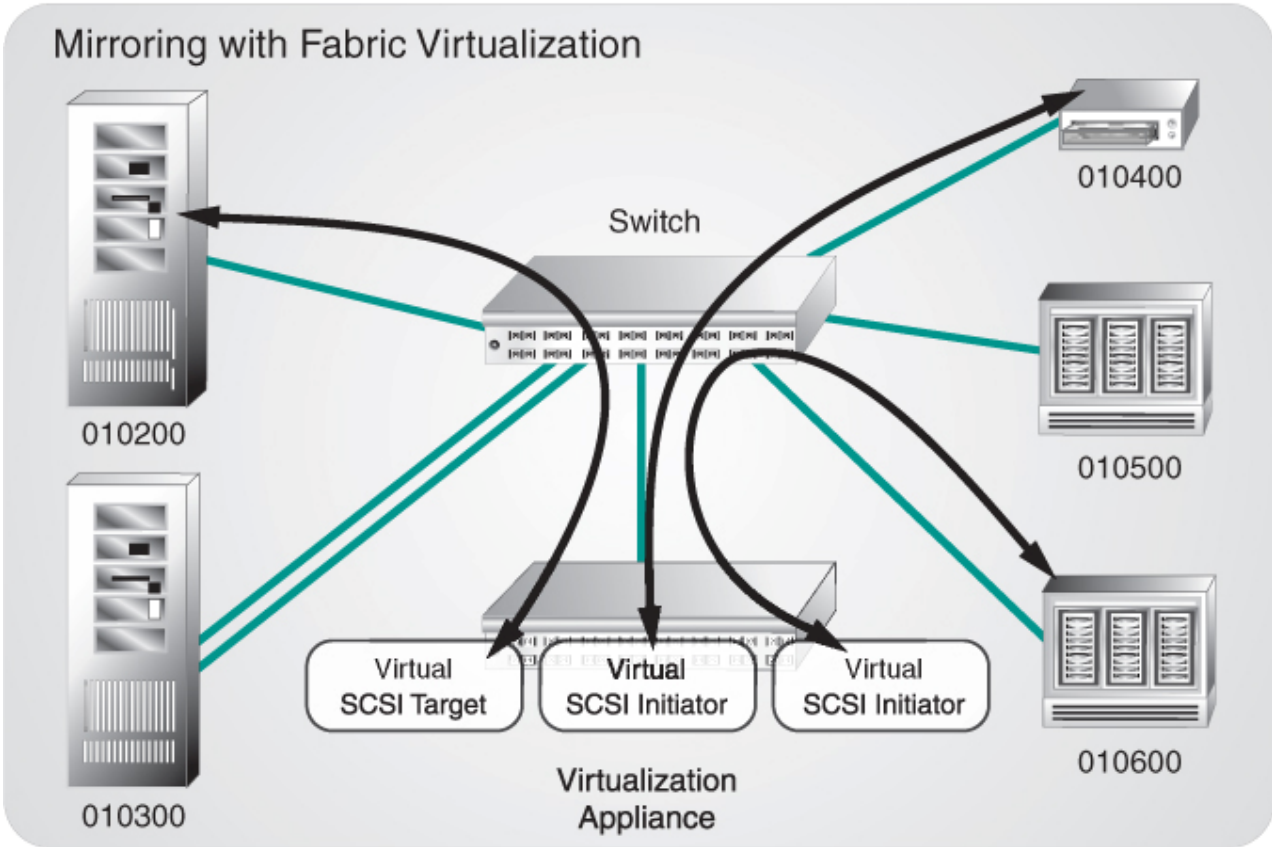

Virtual Channels

- ISL buffer credits are assigned to traffic flows to provide Quality of Service (QoS) between switches

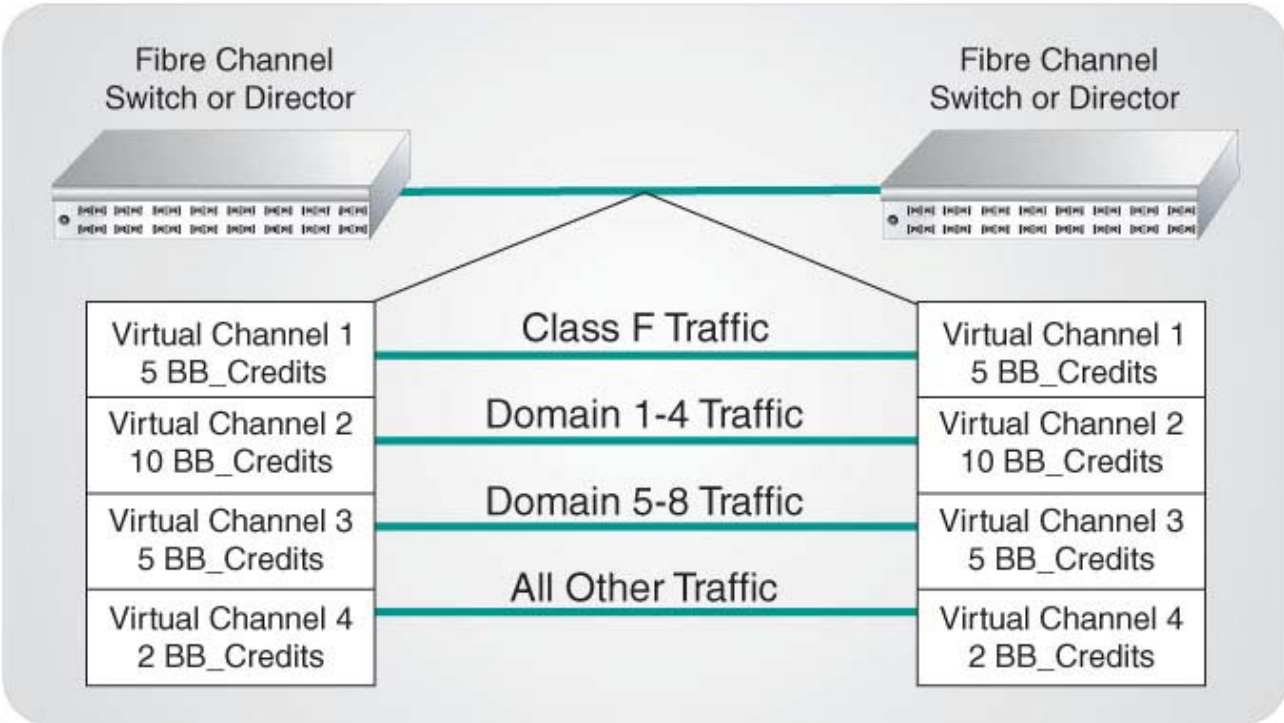

Virtual Fabric Tagging

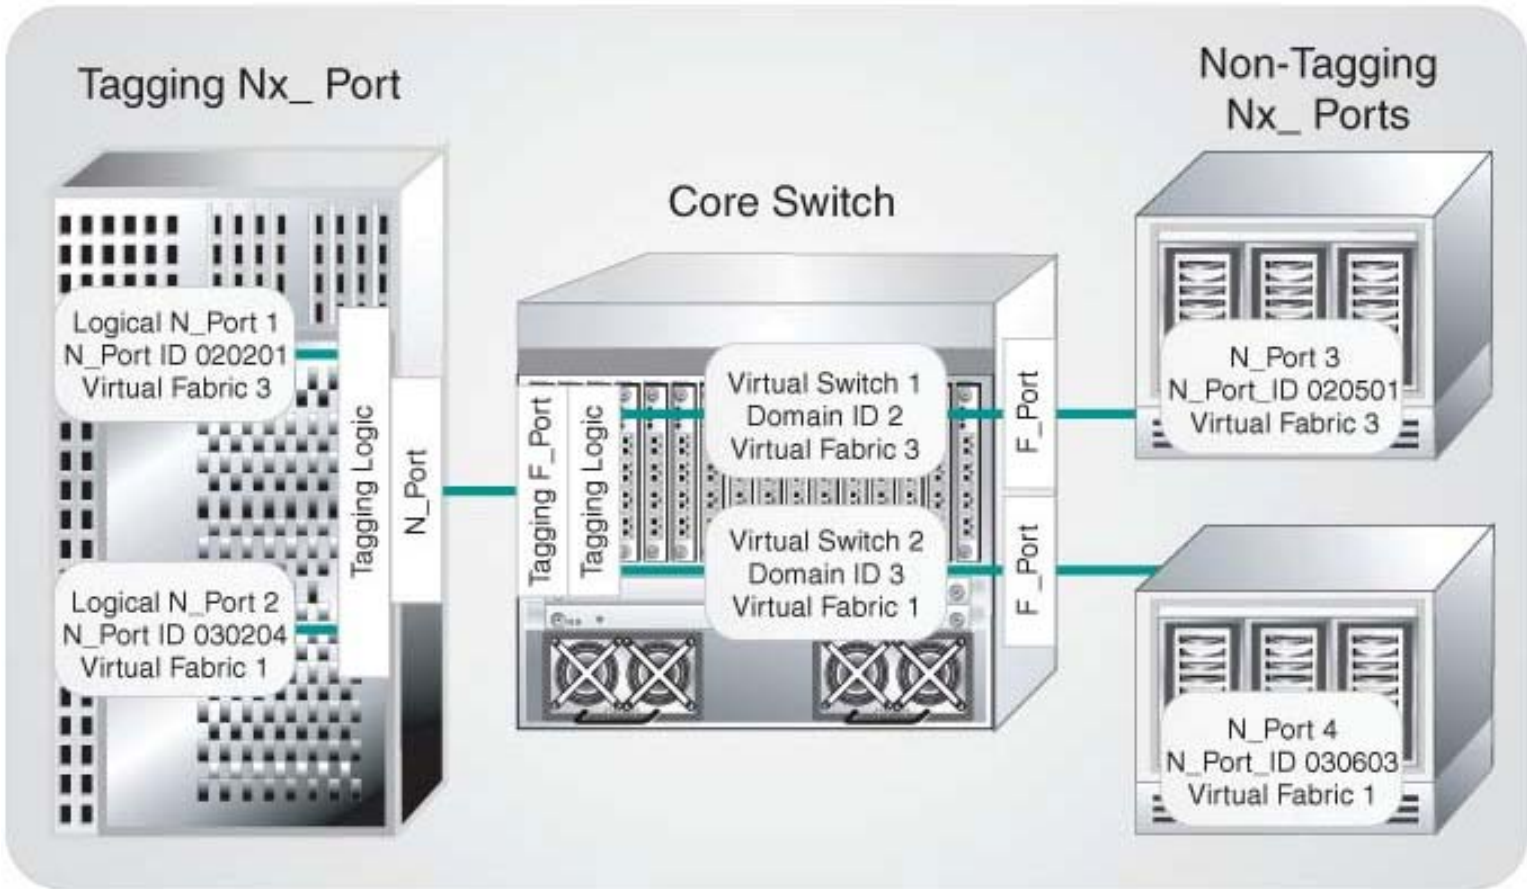

Inter-Fabric Routing

Fabrics and Inter-Fabrics

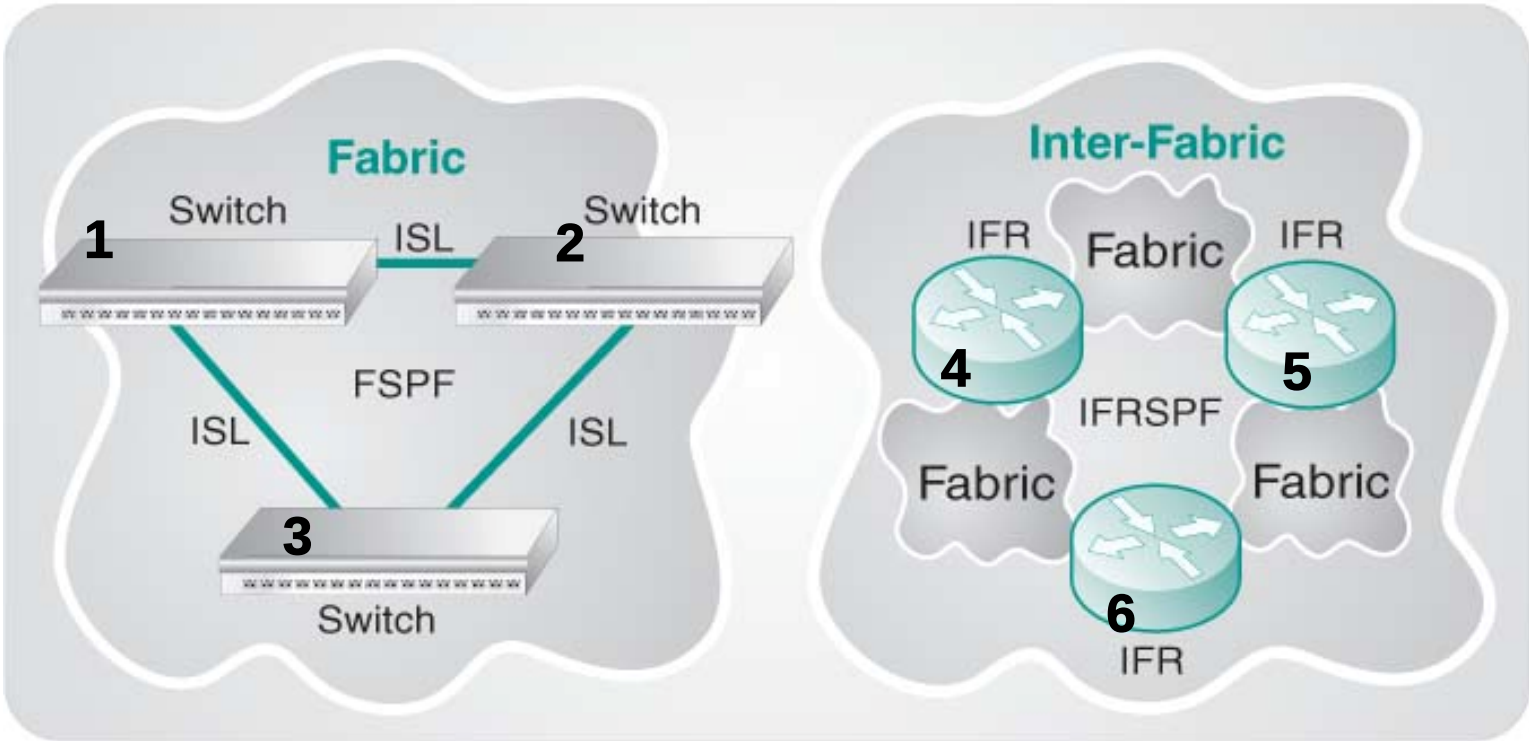

Layer 2 – Switching
FSPF = Fabric Shortest Path First
ISL = Inter-Switch Link

Layer 3 – Routing
IFRSPF – Inter-Fabric Routing Shortest Path First
IFR = Inter-Fabric Router Simple Routing

FC-SATA Configuration

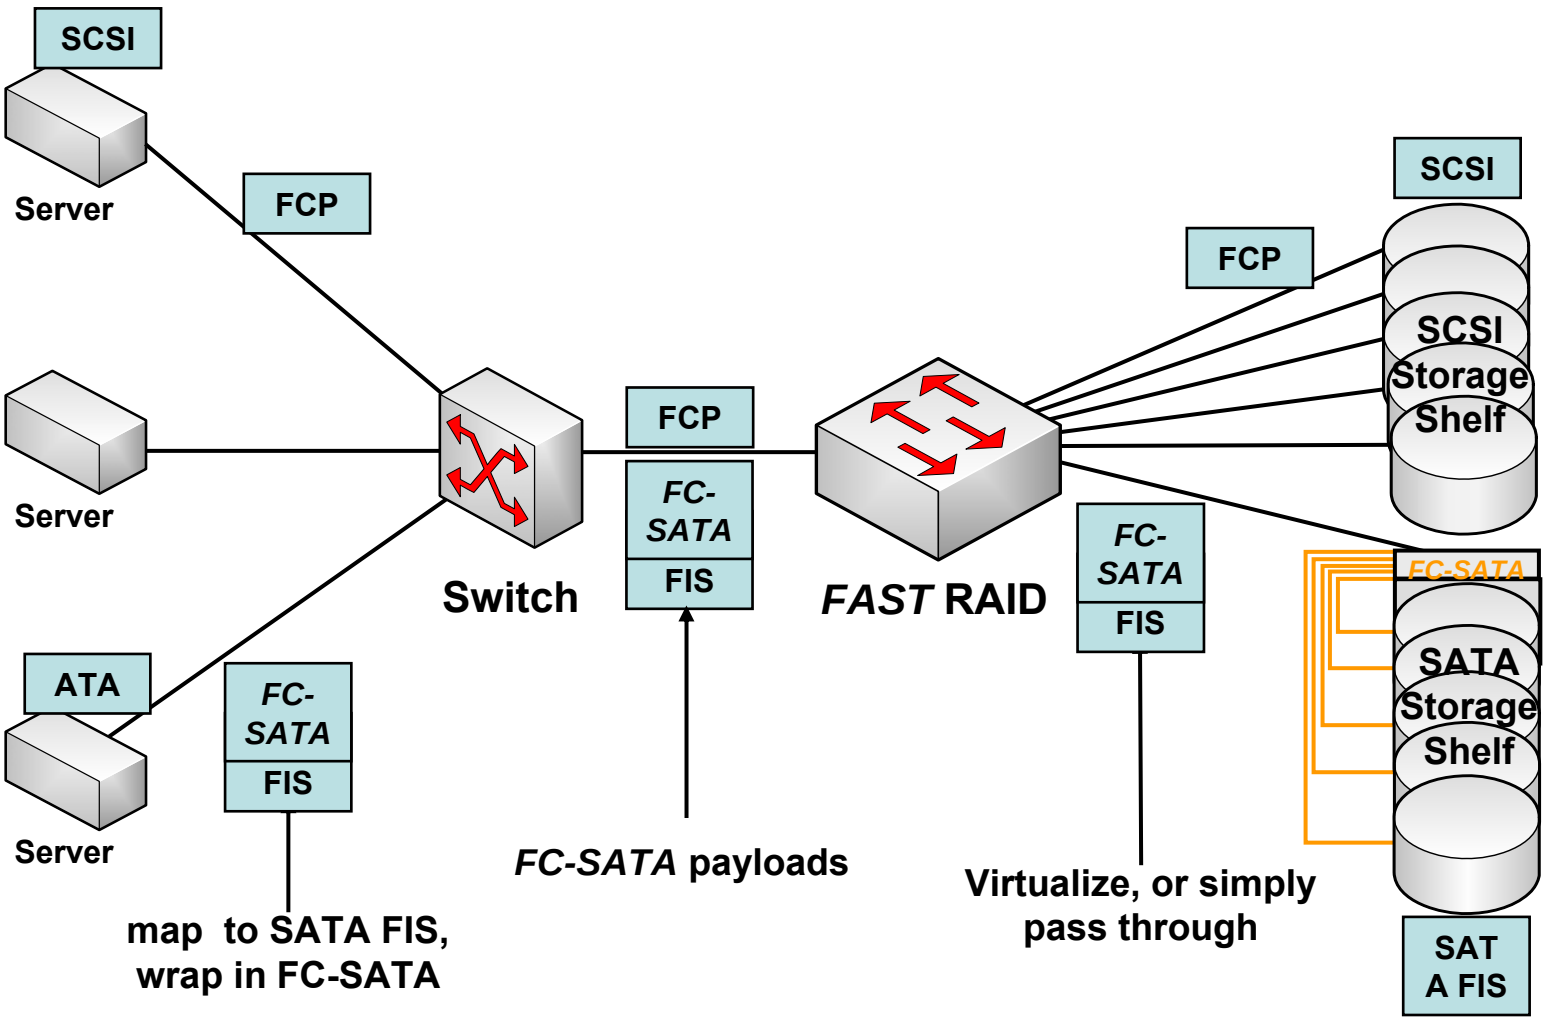

Security

- FC-SP Has Completed Letter Ballot and Will Soon Be A Standard
- Addresses
  - Infrastructure (Passwords, PKI, Secrets)
  - Authentication (FCAP, DH-CHAP, FCPAP)
  - Authorization (Security Policies)
  - Data Integrity (Hash, Keyed-Hash, Signatures, ESP)
  - Confidentiality (ESP)
  - Policy Distribution

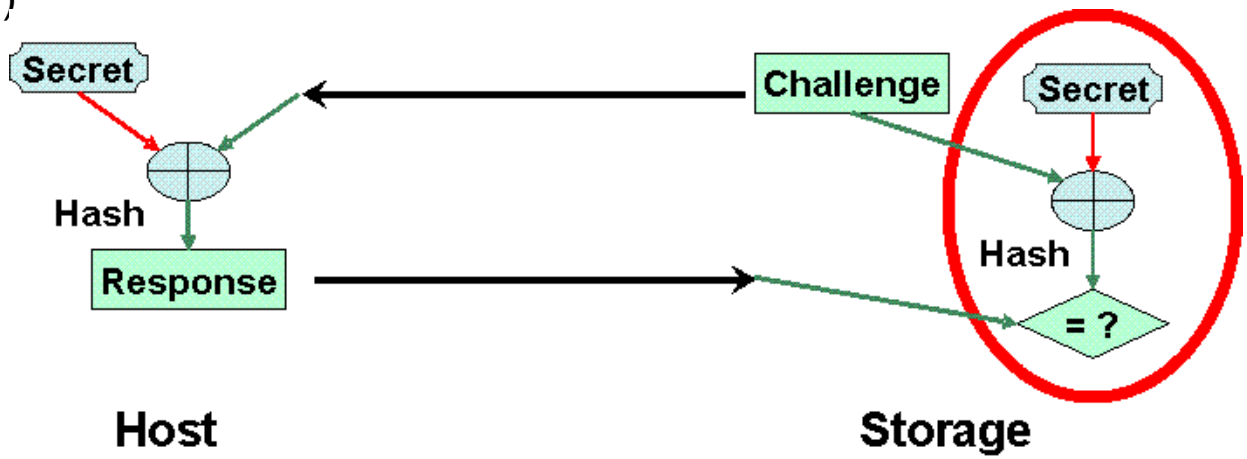

Refer to SNIA FC-SP Tutorial

Fibre Channel: The Storage of Business

Dominates the SAN market today

Fibre Channel has a clear roadmap to provide:

- Higher performance
- Additional capabilities (Security, Tiered Storage, Intelligence…)
- Enablers for new markets

Easy to learn, use and implement

Protects and future proofs storage investments

Comprehensive end to end solution

Fibre Channel Meets the Challenge

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracknetworking@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

SNIA Education Committee

Dr. M. K. Jibbe
Skip Jones
Steve Wilson
Tom Hammod-Doel
Howard Goldstein
Robert Peglar

Continue Your SNIA Education Experience At SNW

SNIA
Storage Networking Industry Association
EDUCATION

- Attend Hands-On Labs in:
  - Data Classification
    Key to Service Level Management
  - Data Security and Protection
    Data Assurance Solutions to Meet Corporate Requirements
  - IP Storage
    iSCSI, Your IP SAN
  - Storage Management
    Manage Storage or Be Managed By It
  - Storage Virtualization
    Increasing Productivity
  - Zero to SAN
    - *Fibre Channel Connectivity in No Time*

Sessions begin Monday afternoon, April 16 and continue through Wednesday, April 18.

All sessions in Emma/Maggie/Annie, 3rd Floor of the Hyatt Manchester.

Registration at the SNW Registration area