SNIA

Education

Fibre Channel Technologies "Current & Future"

Dr. M. K. Jibbe, LSI Corporation (ESG)

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 use cases on how Fibre Channel Structure, Speed and Feed are used with other technologies
  - The goals of the 8GFC use cases are to highlight its speed and performance and to demonstrate its role in increased processing density, faster server I/O buses, device virtualization, and storage product expansions
- Project the market outlook and roadmap of Fibre Channel
- Share what is New in Fibre Channel Standards for Protocols APIs, and Management.
  - FCOE → FC from a SAN concept
  - Fabric Services,
  - Virtualization,
  - Operational Flexibility

What can FC provide today?

1. Flexible, Scalable
   - Flexible Topologies,
   - High Speed,
   - High Performance,
   - Long Distance,
   - Concurrent Node connectivity,
   - Low cost.
2. Communication
   - Framing Structure,
   - Data Phases,
   - Ordered Communication,
   - Reliable Access Control
3. What 's about Data Overhead?
   - Low Latency,
   - High Efficiency,
   - Flexible Routing Control,
4. Redundancy, Availability, and Failover,
5. Applicability in SAN with large IT User Base

FC Topologies

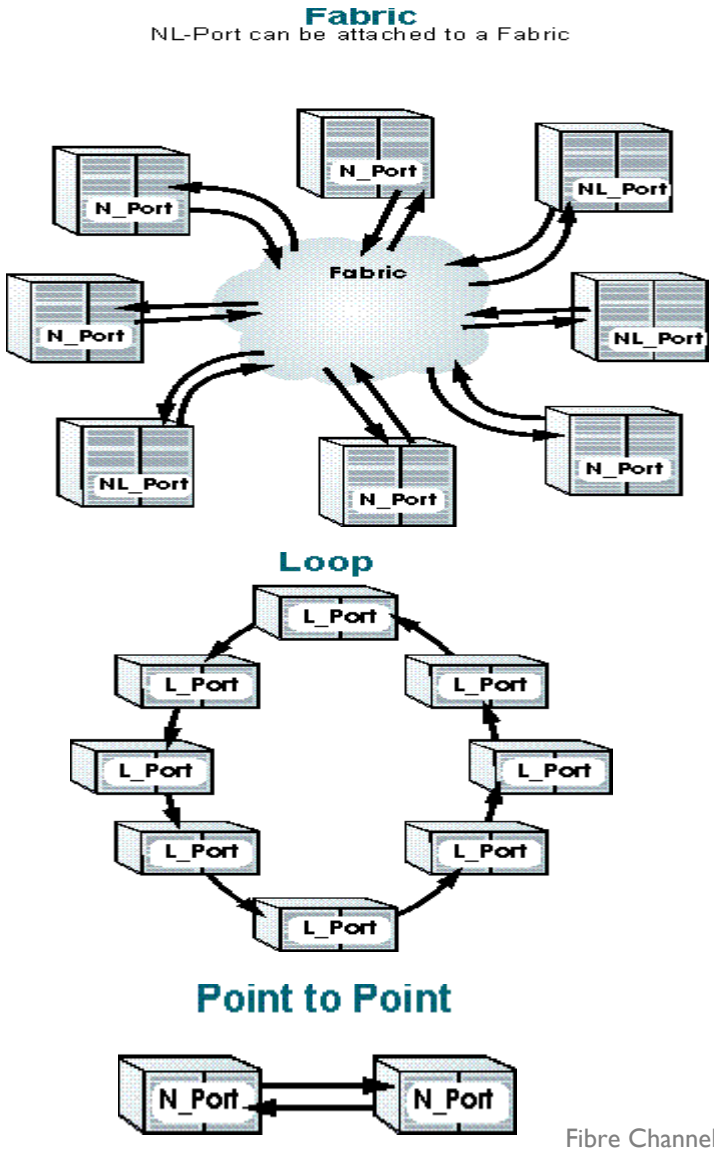

Switched Fabric

- Fibre channel supports a 24-bit address space (2^24 ports in a switched interconnect)
- Multiple concurrent communications for high aggregate throughput
- FC routing is done based on the Domain ID FC
- Device ports are uniquely identified by a WWPN

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^{24}
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

Speed, Flexibility and Scalability

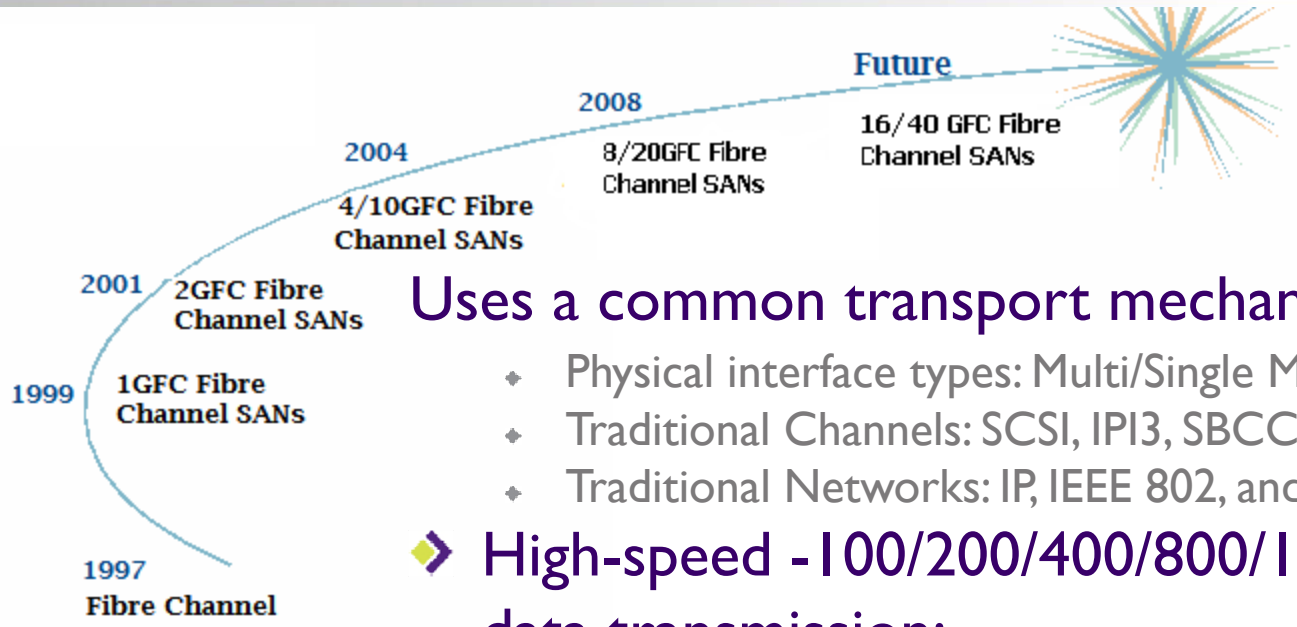

Uses a common transport mechanism to support:

- Physical interface types: Multi/Single Mode Fiber, and Copper
- 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
- 16G will be technically stable in Dec 2009.

Flow Control: Access Control, Latency, and Efficiency

Login Buffer to Buffer
- Node to Fabric
- Fabric to Node

Login Node to Node

Login Fabric to Fabric

Flow control is credit based
- Each R_RDY received increments the available BB Credit value
- Control pace of frame transmission

Latency
- Across a single switch, average latencies can be as low as 400 nsec

Exchanges and Frames
- Maximize bandwidth

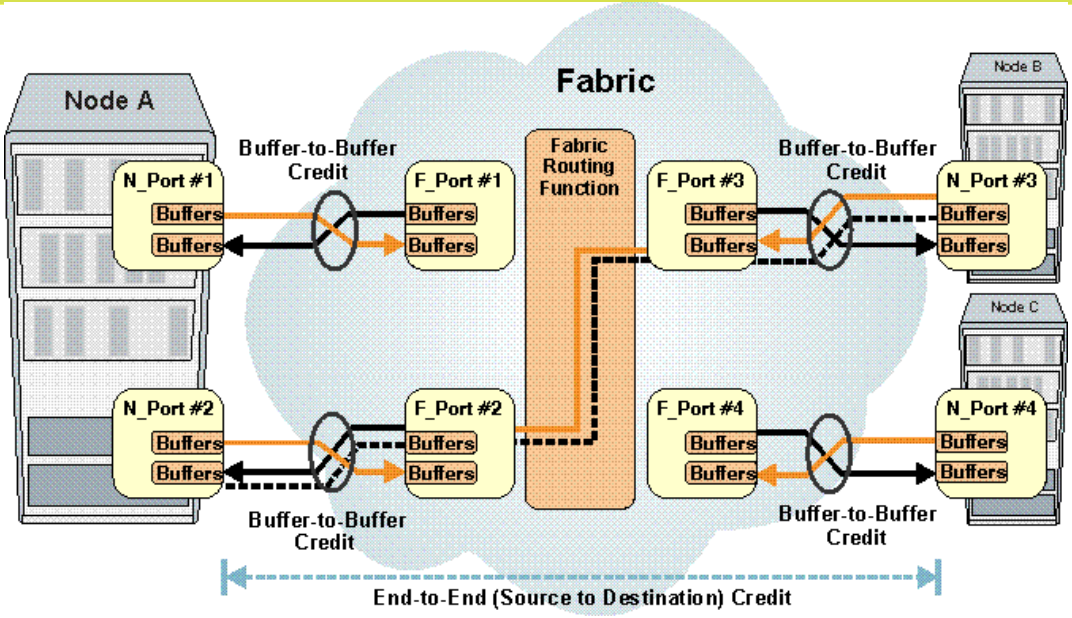

FC Routing

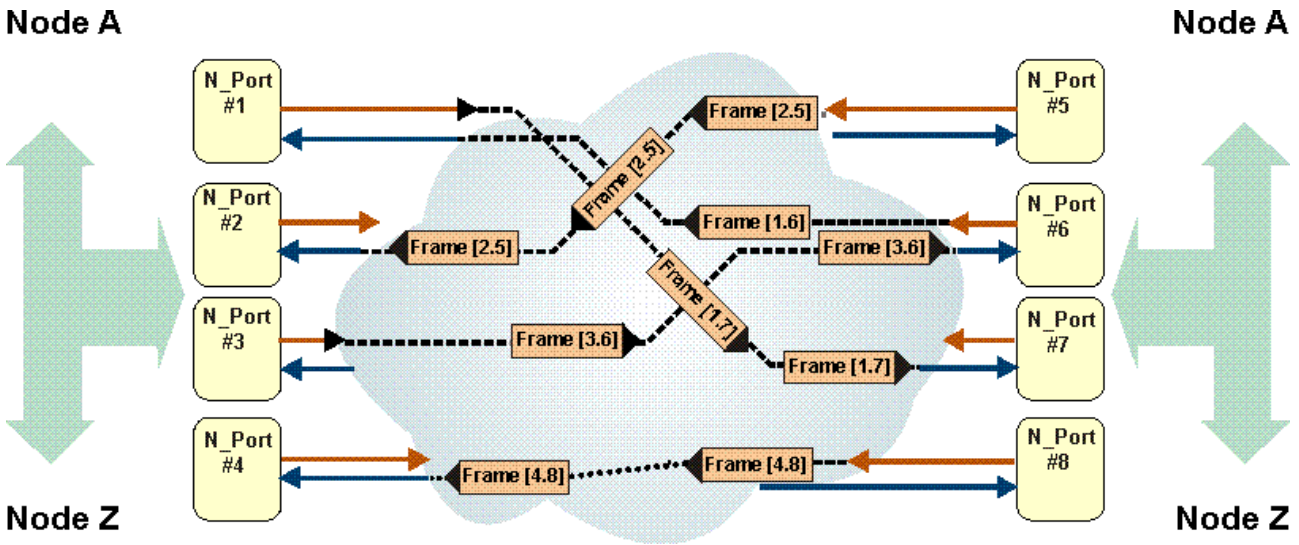

1. **Connect Any to Any**
2. **Maximize Connectivity**
3. **Simplex and Duplex**

Use case 1: FC for data access Control

1. Local switching – Lower switching latency and higher full-speed port density.
2. Trunking – Better reliability, throughput and availability of data travelling between switches.
3. Flexible zoning capabilities - FC-SP and FC-GS-4 fabric security features
4. Data Migration Feature – Management features plus increase data migration efficiency while reducing risk and operational costs.
5. **Inter-Fabric Routing**
   - Allows devices on two separate Fabrics to communicate without a merge unlike ISL
   - SAN level zoning determines which devices may communicate since IFR connect the Fabric clouds.
   - IFR functions are in the Routers between the Fabrics

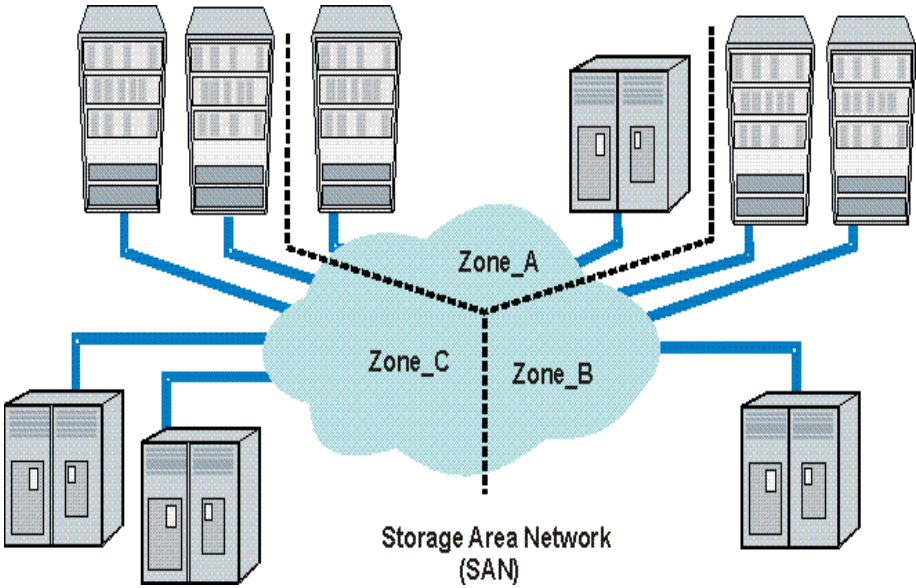

Provide Accessibility
- **Soft Zoning**
- **Hard Zoning**
- **Virtual SAN**

Inter-Fabric Routing

Fabrics and Inter-Fabrics

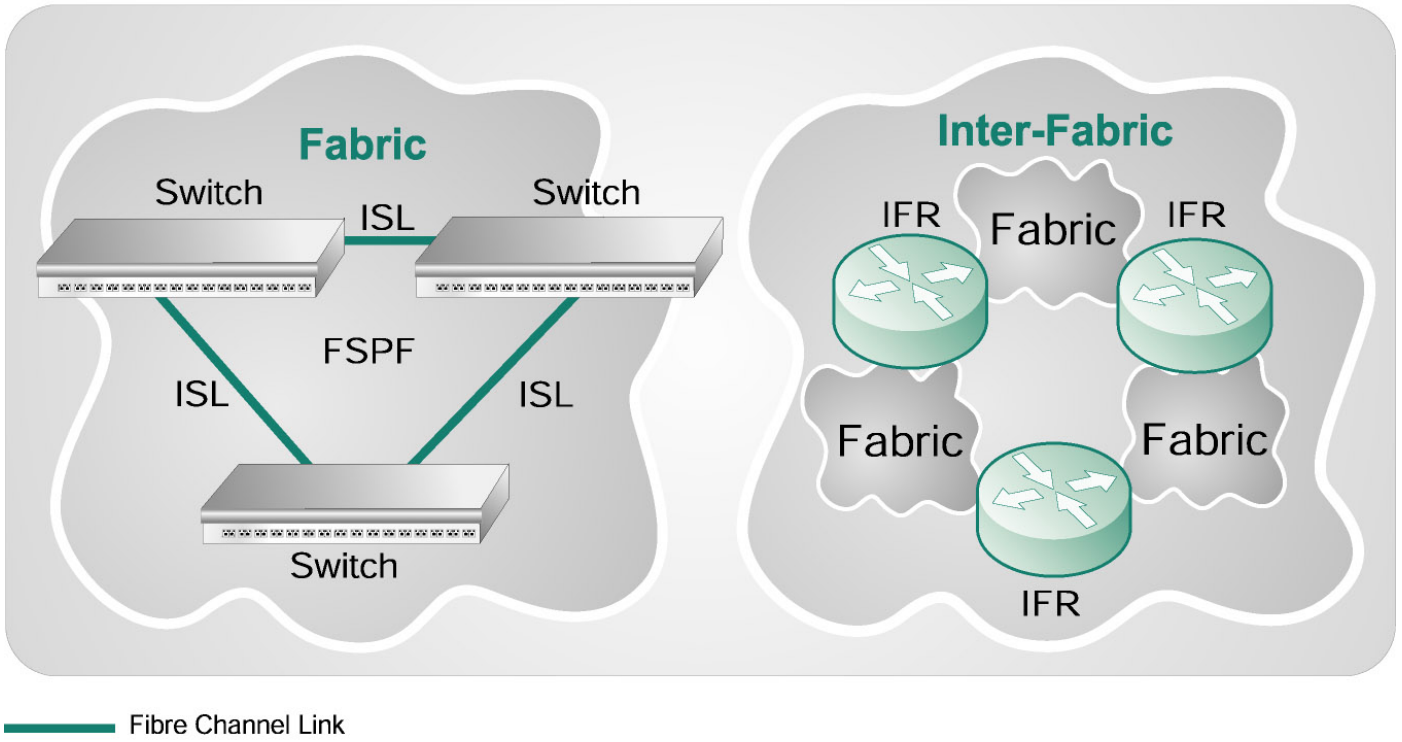

Layer 2 – Switching
FSPF = Fabric Shortest Path First
ISL = Inter-Switch Link

Layer 3 – Routing
IFR = Inter-Fabric Router
Simple Routing

Use Case 2: Customer requires Redundancy, Availability, and Failover in SAN

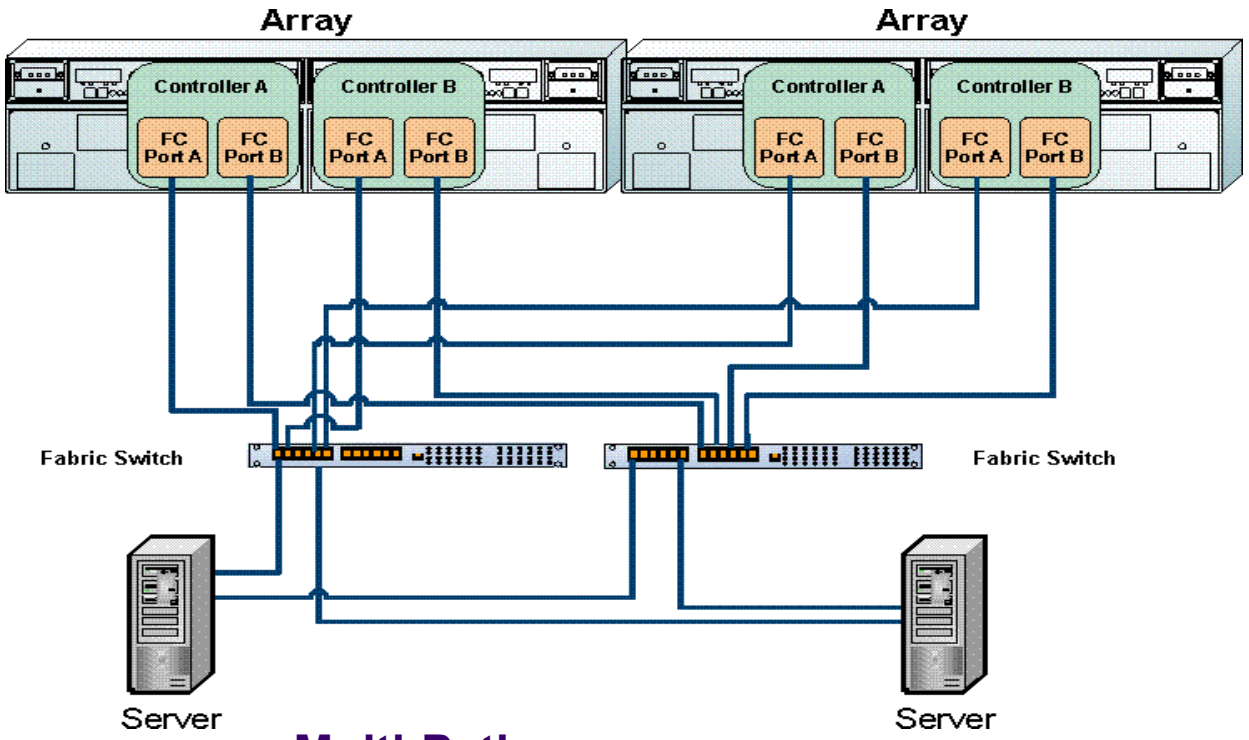

- Multi-Path
- Full Redundancy
-Path Failover
-Active-Active / Active-Passive
-Reliable Failover drivers in an FC SAN
-NPIV

Use Case 3: FC Coexist with different Technologies in a Enterprise Datacenter

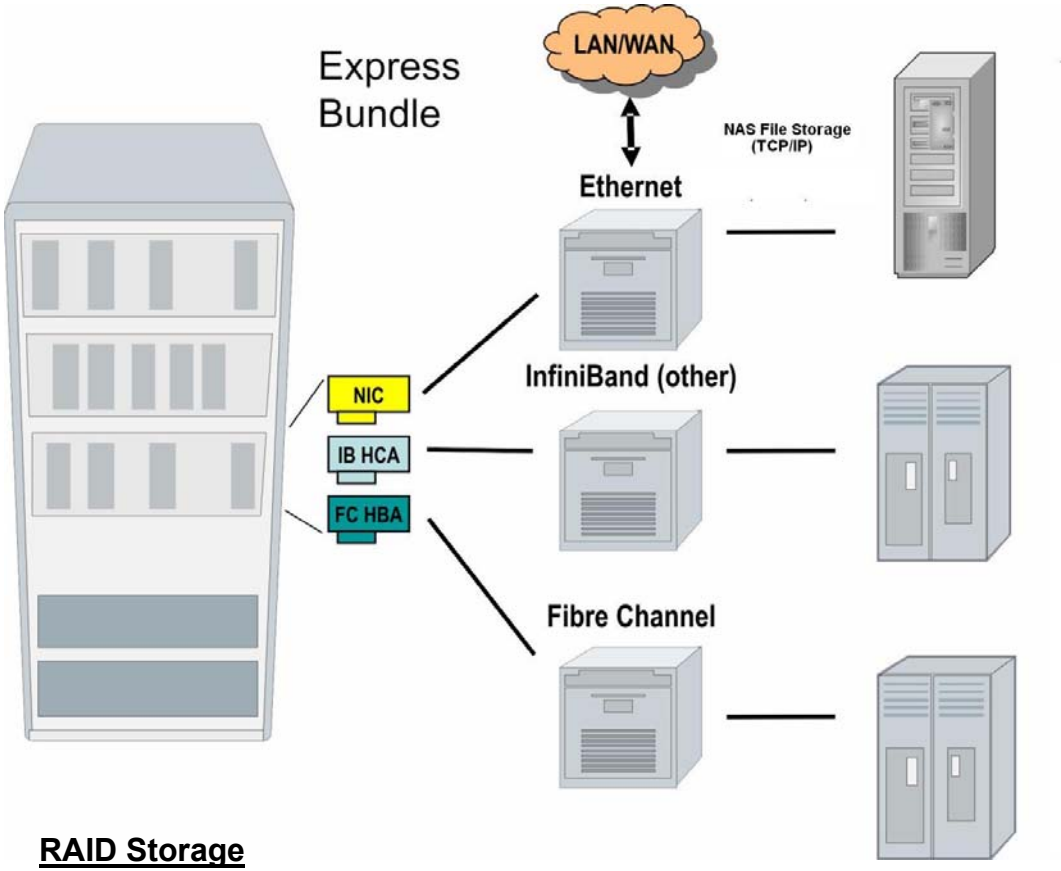

Product	Host Interface	Drive Interface
NAS	Ethernet	FC
FC	FC/iSCSI HBA	SAS/SATA
Infiniband	HCA	FC

- NAS , Infiniband and FC coexist in SAN Today
- Every interface has its own express bundle
- Uses different interface/protocols for Host & Media Connectivity
- Merits of co-existence of NAS , FC & IB
  - Block Storage and File system storage
  - Various RAS functions
  - Tunnels
  - Bandwidth

RAID Storage

- FC Dominated Front End
- IB Opportunity in HPC / DB
- FC Dominated Back End
  - ✓ SATA HAD with FC IO
  - ✓ FC-SATA for Tier 2 Apps
  - ✓ FC Front end – SAS Backend
  - ✓ Mixed

Use Case 4: Customer's Budget dictates Array Module with Different Drive Technologies

Project budget and Application Data sensitivity determine the type of drive technologies to use in a SAN.

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

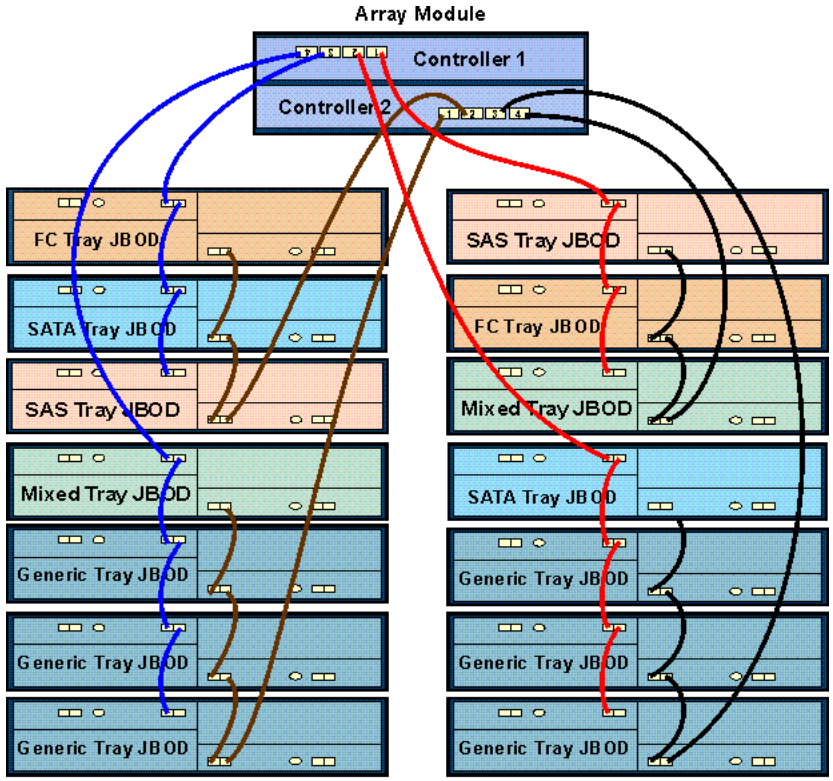

Array Module with Different Drive Tray Types

Entry Cost Comparison Table

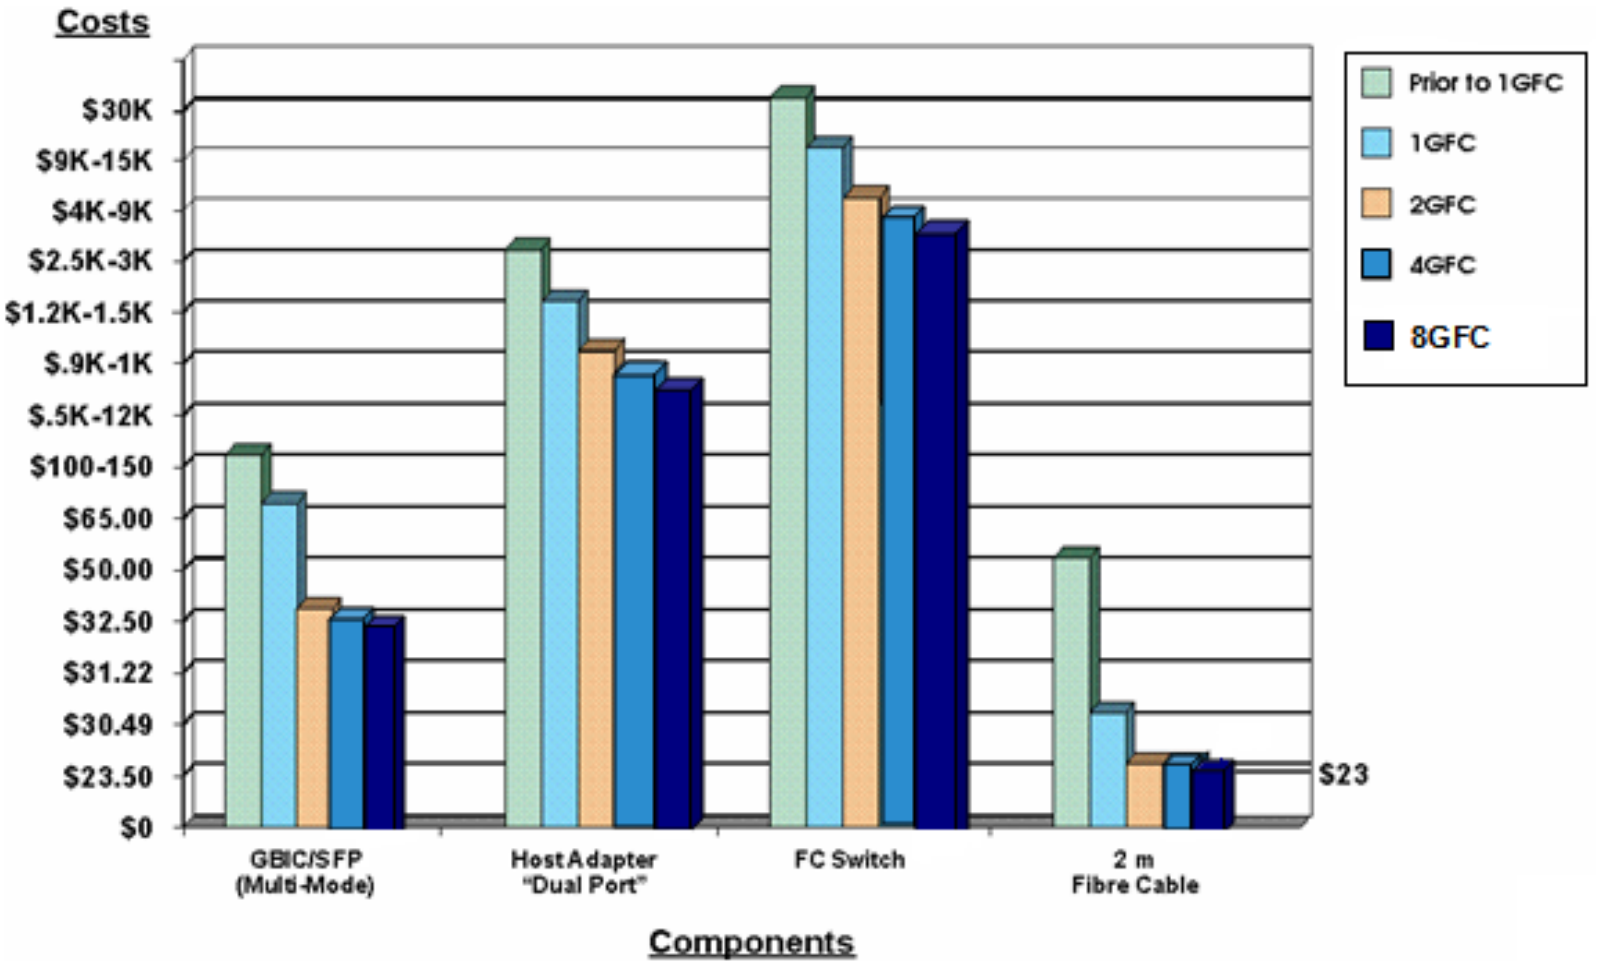

FC Product Performance – IOPS: Host Interface – Drive Interface

		Array Host Interface		
	Drive Type	Dual 4 GFC	Quad 4 GFC	Current Dual 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	32K	32K	70K
	SATA	7K	7K	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 for benchmark calculation and code thread	FC, SAS, / SATA	96D / 8T	96D / 8T	96D / 8T

- **FC continues to evolve with different technologies**
- **4K- R5 -> 4K I/O size and RAID 5**
- **Benchmark numbers are calculated using RAID models**

Use case 5: FC SAN-based storage virtualization

1. **SAN-based storage virtualization holds the promise of addressing SAN complexity**
    1. Single point of administration
2. **Nondisruptive data migration:**
3. **Information lifecycle management (ILM)**
4. **Improved allocation efficiencies** Storage managers know that improving asset utilization is a quick way to lower the total cost of ownership (TCO) for their department.
5. **Heterogeneous replication** One of the huge challenges associated with maintaining agnosticism among disk array vendors is disaster recovery replication.

Efficiency with Virtualization

1. Capacity and costs fail to meet expectations
2. **Server failure.** Large-scale consolidation may put many key processes, applications and services in the same proverbial basket.
3. **Over-provisioning**. Starting consolidation without a clear picture of an application's function, workload or profile
4. **Service levels.** Virtualization technology requires new skills;
5. Hardware bandwidth to accommodate the increased virtual load
6. Platform to support allocating resources for each virtual load
7. Hardware downtime should not cause downtime for entire virtual platforms on top of it.
8. Backup and disaster recovery implementation gets more complex and rapid recovery becomes more important in a virtualized environment

FC Storage Virtualization Illustration

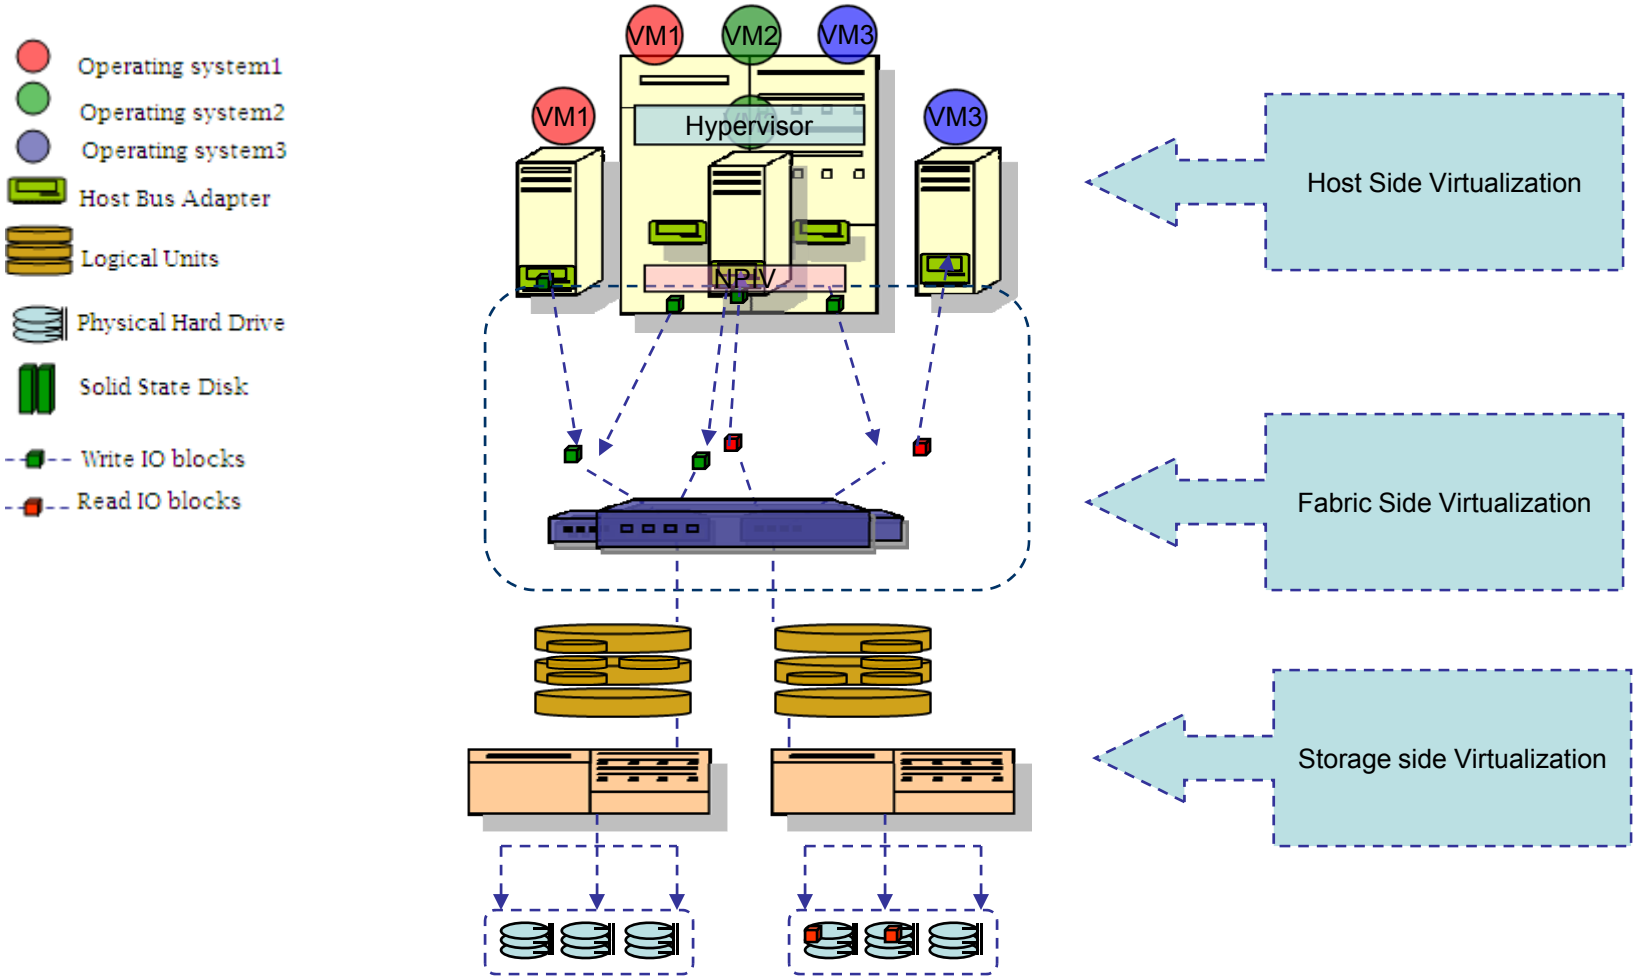

18

SAN-based storage virtualization Components

1. The modern data center is faced with ever-growing demands

2. 8Gb FC HBAs and Fabric offer a multitude of hardware-assist virtualization enhancements that help provide a complete HBA and Fabric virtualization solution
   - Simplify management and provide port-level QoS
   - Reduce system overhead and improve reliability
   - Achieve higher, deterministic I/O performance with QoS functionality
     - Multiple device I/O queues
     - Unique requester IDs
     - Improved device interrupt routing.

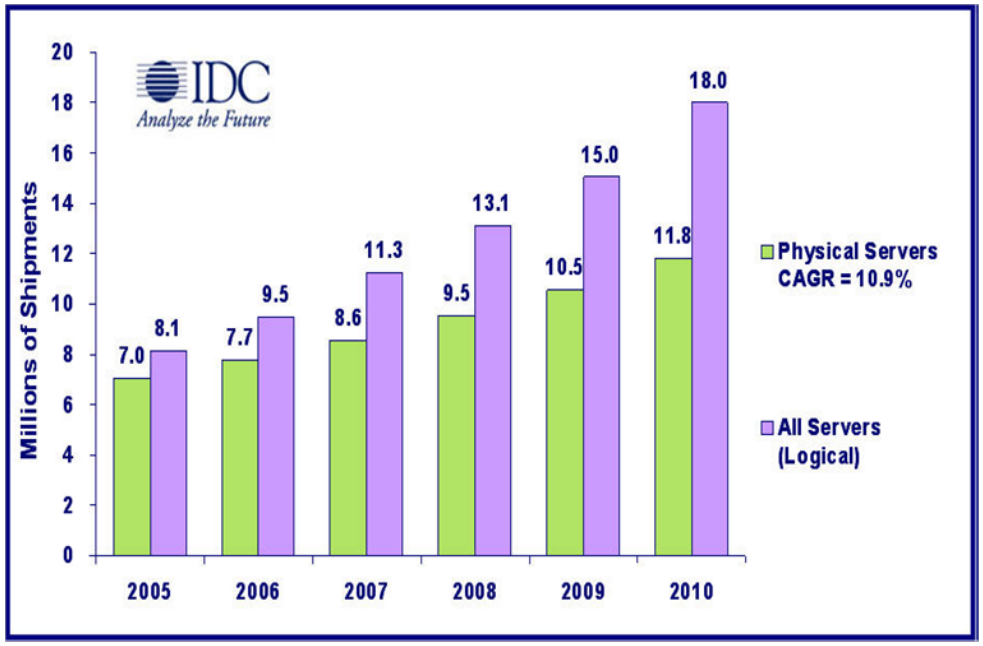

1. SSDs or file cache accelerators can be used as a Storage resource

2. SSDs have historically been used to successfully respond to
   - I/O bottlenecks affecting application performance on a single server
   - Single application basis

3. SSD as a dedicated device for the broader spectrum of applications has been historically difficult

4. Leading SSD devices
   - Support up to 8 or more Fibre Channel ports for switched SAN environments
   - Partitionable into numerous logical unit numbers (LUNs)
   - Can be managed using standard off-the-shelf SAN and SRM tools.

SSD in SAN Virtualization - Illustration

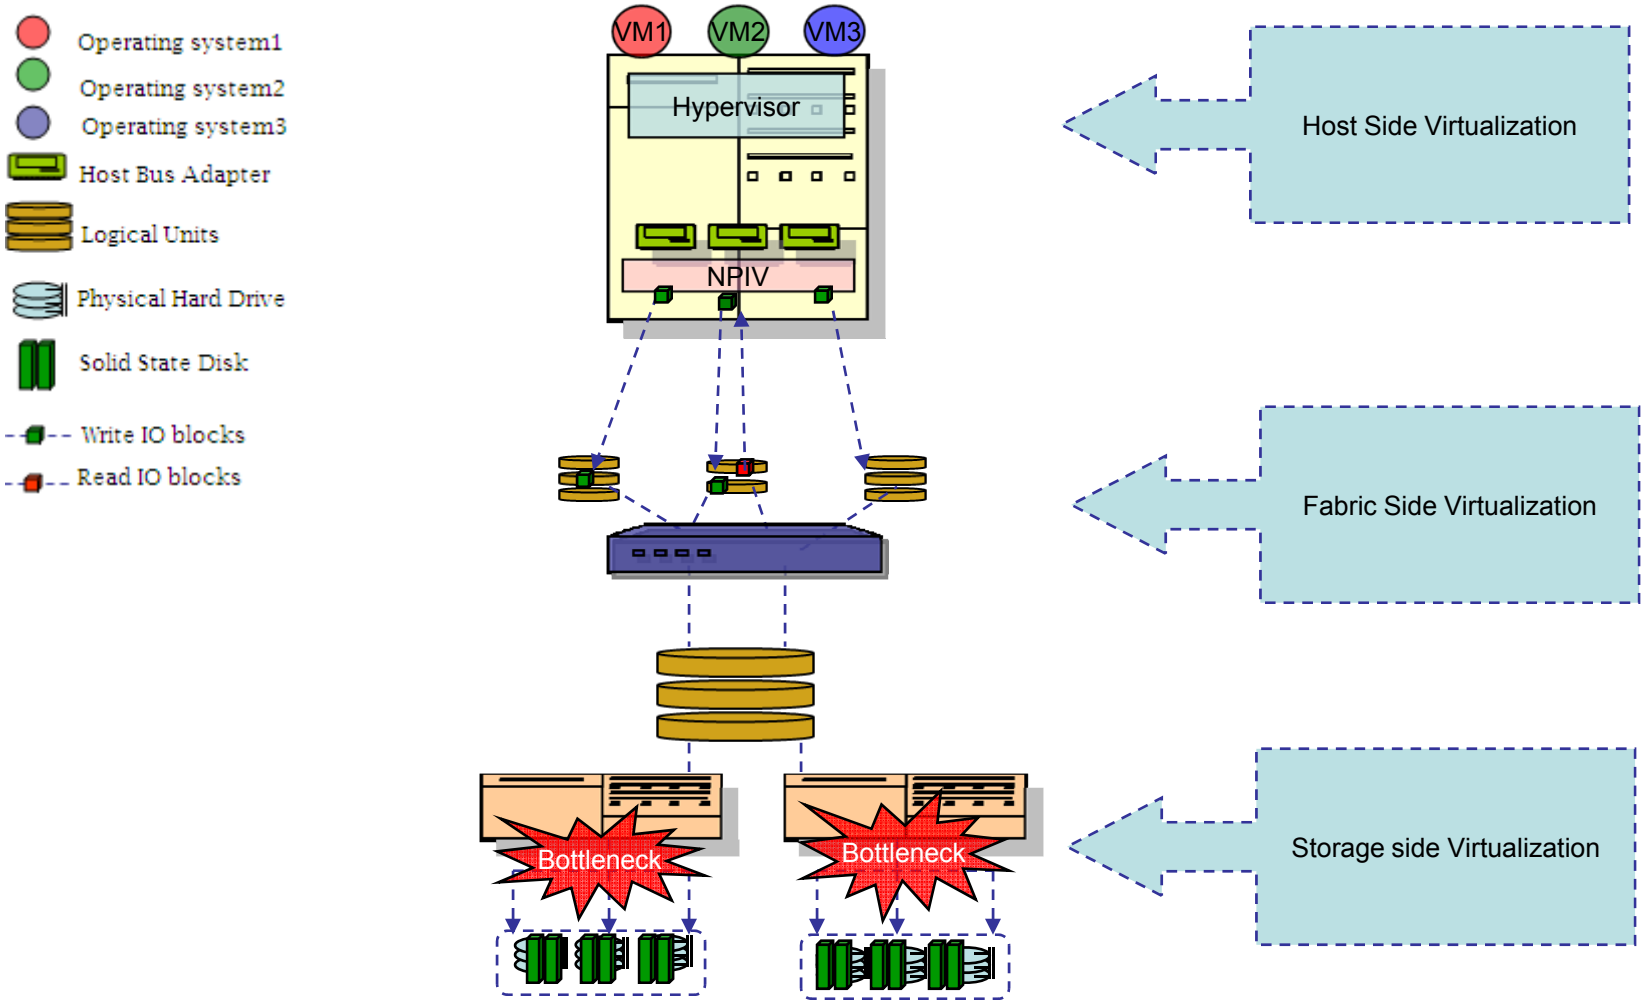

Security

- FC-SP is An Approved INCITS Standard
- Provides a Security Framework for FC Environments
  - Infrastructure (Passwords, PKI, Secrets)
  - Authentication (FCAP, DH-CHAP, FCPAP)
  - Authorization (Security Policies)
  - Data Integrity (Hash, Keyed-Hash, Signatures, ESP)
  - Confidentiality (ESP)
  - Policy Distribution

Check out SNIA Tutorial:
An Introduction to Storage Security

FC features for users

- 8 GFC is plug-compatible with 1, 2 and 4GFC (devices auto negotiate w/o user intervention)
  - Trunking
  - Flexible zoning capabilities with available secure fabric connection utilizing standards-based FC-SP and FC-GS-4 fabric security features
  - Data Migration Feature –RASUI
  - Coexistence with other technologies in DC.
  - SAN base storage virtualization
- With innovative storage virtualization tools that continue to evolve to better help manage storage environments, companies can focus on tuning the SAN for maximum benefit to all applications.
- Designing solid-state disk into the virtualized SAN provides a powerful and cost-effective resource that can be used to significantly improve performance, productivity, and profitability across the enterprise.
- Storage virtualization in all of its forms has been providing benefits to businesses for years; yet there are still some challenges that remain.
- SSD Storage and virtualization enable a data center to take advantage of the high speed FC

FC features for users Cont'd

- It's easy to see that SAN-based virtualization over FC holds a lot of promise for the industry.

- Discuss these value propositions with your customers to open their minds to the possibilities of SAN virtualization
- Optimization areas in Virtualization to fit a storage architecture with SSD and 8Gb FC
  - Improve Storage capacity and resource utilization
  - Simplify management of Storage resources
  - Dynamic Storage allocation
  - Uninterrupted online migrations for Storage administrations
- Storage vendor are investing and advertising FC SSD much more than SATA or SAS SSD.
- FC SSD are leveling the performance marks for storage products at the back end (drive side "user data and Meta data updates").
- FC storage technology shows that the combination of 8g FC, Virtualization, and SSD in a SAN improves performance, productivity, and profitability across the enterprise.

Fibre Channel is Here to Stay.

What's New?

- A Roadmap to the Future!

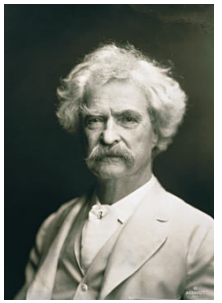

The report of my death was an exaggeration.

FCIA Fibre Channel Speed Roadmap

Education SNIA

	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	14.025	2009	2011
	32GFC	6400	28.5	2012	Market Demand
	64GFC	12800	57	2016	Market Demand
	128GFC	25600	114	2020	Market Demand
Base10**	10GFC	2400	10.52	2003	2004
	20GFC	4800	21.04	TBD	2008 (Vendor)
	40GFC	9600	TBD	TBD	Market Demand
	100GFC	24000	TBD	TBD	Market Demand
FCoE***	10GFCoE	2400	10.3125	2008	2009
	40GFCoE	9600	41.225	TBD	Market Demand
	100GFCoE	24000	100.3125	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., 8GFC backward compatible to 4GFC and 2GFC)

****Base10** is for ISLs, core connections, and other high speed applications demanding maximum bandwidth. Except for 100GFC (which follow Ethernet standards and compatibility guidelines), each Base10 speed is expected to be compatible at least one previous generation.

*****FCoE**: Fibre Channel over Ethernet tunnels FC through Ethernet and is compatible with all existing Fibre Channel fabric environments. FCoE ports follow Ethernet standards and compatibility guidelines.

†Line Rate: All Base2 speeds are single-lane serial stream. Base10 and FCoE rates listed are equivalent data rates for serial stream methodologies. However, final output speed is generated with aggregated methodologies.

‡Dates: Future dates estimated

Maximum Distance

Education SNIA

Application Standard Speed	Transceiver Type	Baud Rate	Wavelength	Link Length (Maximum Distance)[6]					
				OM1 (62.5/125)	OM2 (50/125)	OM3 (50/125)	OM3 Bendable[5] (50/125)	OM4[1] (50/125)	SMF (9/125)
				200 MHz-km	500 MHz-km	2000 MHz-km	>3000 MHz-km	3500-4700 MHz-km	>25000 MHz-km
Fibre Channel									
1 Gbps Short Wave (SW)[3]	SFP[2]	1.0625 Gbps	850 nm VCSEL	300 m	500 m	860 m	> 860 m	1100 m - 1500 m	n/a
1 Gbps Long Wave (LW)	SFP	1.0625 Gbps	1310 nm FP	n/a	n/a	n/a	n/a	n/a	10 km
2 Gbps Short Wave (SW)	SFP	2.1250 Gbps	850 nm VCSEL	150 m	300 m	500 m	> 500 m	650 m - 700 m	n/a
2 Gbps Long Wave (LW)	SFP	2.1250 Gbps	1310 nm FP	n/a	n/a	n/a	n/a	n/a	10 km
4 Gbps Short Wave (SW)	SFP	4.2500 Gbps	850 nm VCSEL	70 m	150 m	270 m	>270 m	350 m	n/a
4 Gbps Long Wave (LW)	SFP	4.2500 Gbps	1310 FP	n/a	n/a	n/a	n/a	n/a	4 km
4 Gbps Long Wave (LW)	SFP	4.2500 Gbps	1310 DFB	n/a	n/a	n/a	n/a	n/a	10 km
8 Gbps Short Wave (SW)	SFP+	8.5000 Gbps	850 nm VCSEL	21 m	50 m	150 m	> 150 m	180 m - 200 m	n/a
8 Gbps Long Wave (LW)	SFP+	8.5000 Gbps	1310 nm FP	n/a	n/a	n/a	n/a	n/a	1.4 km
8 Gbps Long Wave (LW)	SFP+	8.5000 Gbps	1310 nm DFB	n/a	n/a	n/a	n/a	n/a	10 km
Ethernet									
1 Gbps Short Reach (SX)[4]	SFP	1.25 Gbps	850 nm VCSEL	275 m	550 m	860 m	> 860 m	1100 m - 1500 m	n/a
1 Gbps Long Reach (LX)	SFP	1.25 Gbps	1310 nm FP	n/a	n/a	n/a	n/a	n/a	10 km Dark Fiber
10 Gbps Short Reach (SR)	SFP+, XFP	12.5 Gbps	850 nm VCSEL	33 m	82 m	300 m	> 300 m	380 m – 420 m	n/a
10 Gbps Long Reach (LX)	SFP+, XFP	12.5 Gbps	1310 nm DFB	n/a	n/a	n/a	n/a	n/a	10 km Dark Fiber

Last updated: 28 February 2008

FCIA "Condensed" Roadmap
(Speed GFC or GFCoE)

- **FC-Base2 (Edge, Backend, and ISL)**
  - *1GFC, 2GFC, 4GFC, 8GFC* shipping today
  - 16GFC will be technically stable in Dec 2009.
  - 32GFC, 64GFC, 128GFC
- **FC-Base10 (ISL)**
  - 10GFC shipping today
  - 20GFC Under consideration, some vendor implementations shipping
  - 40GFC, 80GFC, 160GFC
    - 100GFC under study (leverage IEEE 802.3 work)
- **FCoE (Datacenter)**
  - Uses any standard 10 Gb/s Ethernet link technology with Lossless Ethernet Bridges
  - Most applications will use the SFP+ being developed in the SFF Committee.
  - 10GFCoE shipping today
  - 1GFCoE, 40GFCoE and 100GFCoE based on market demand

Continuous SAN Market Growth

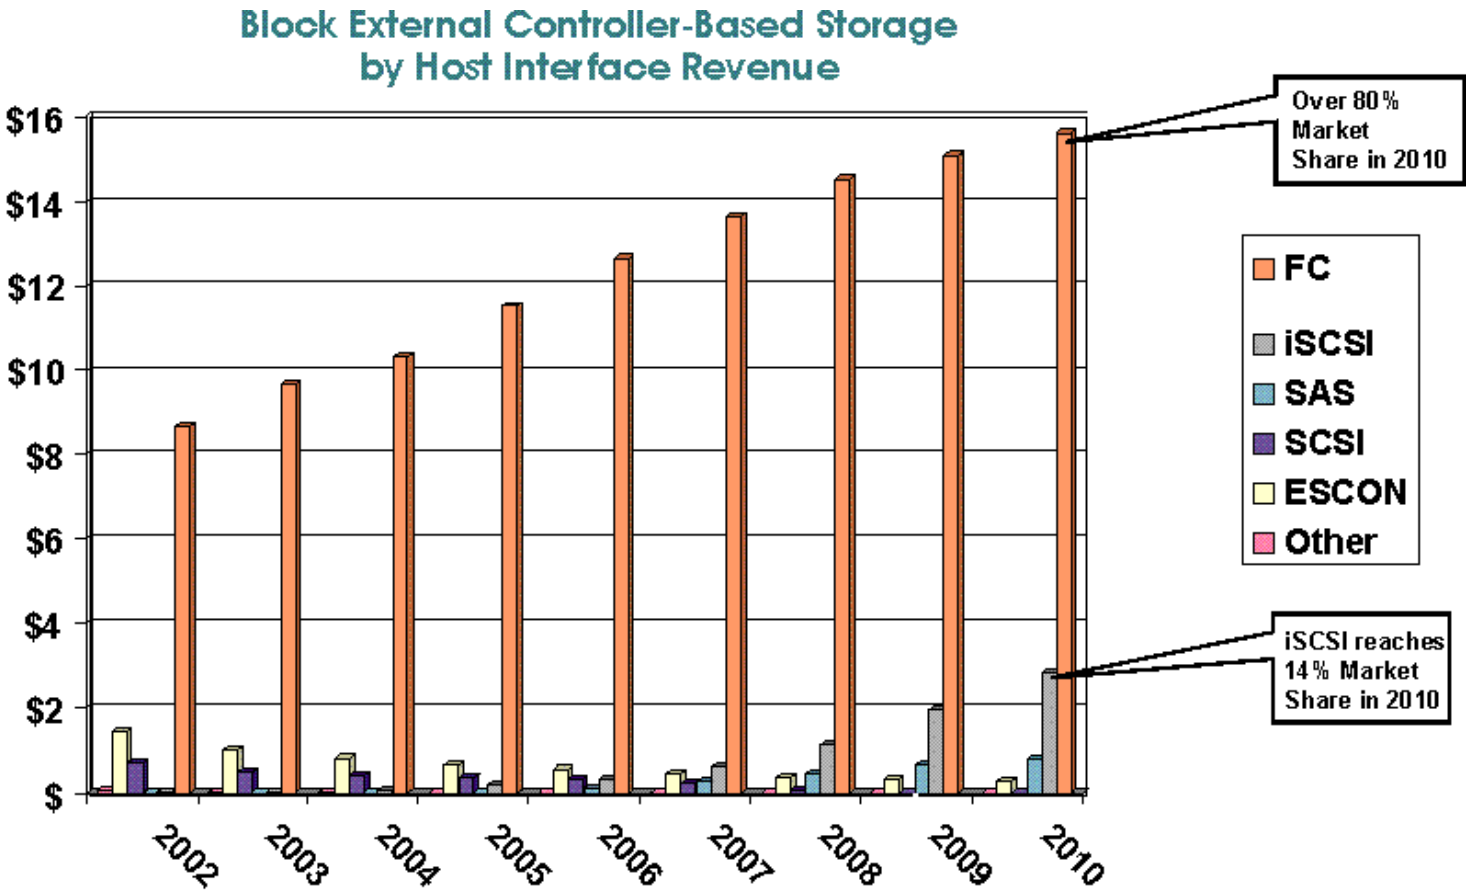

Source: Gartner External Controller-Based Disk Storage WW 2006-2010, 6 October 2006 (FICON included in Fibre Channel)

Continuous Back-End Market Growth

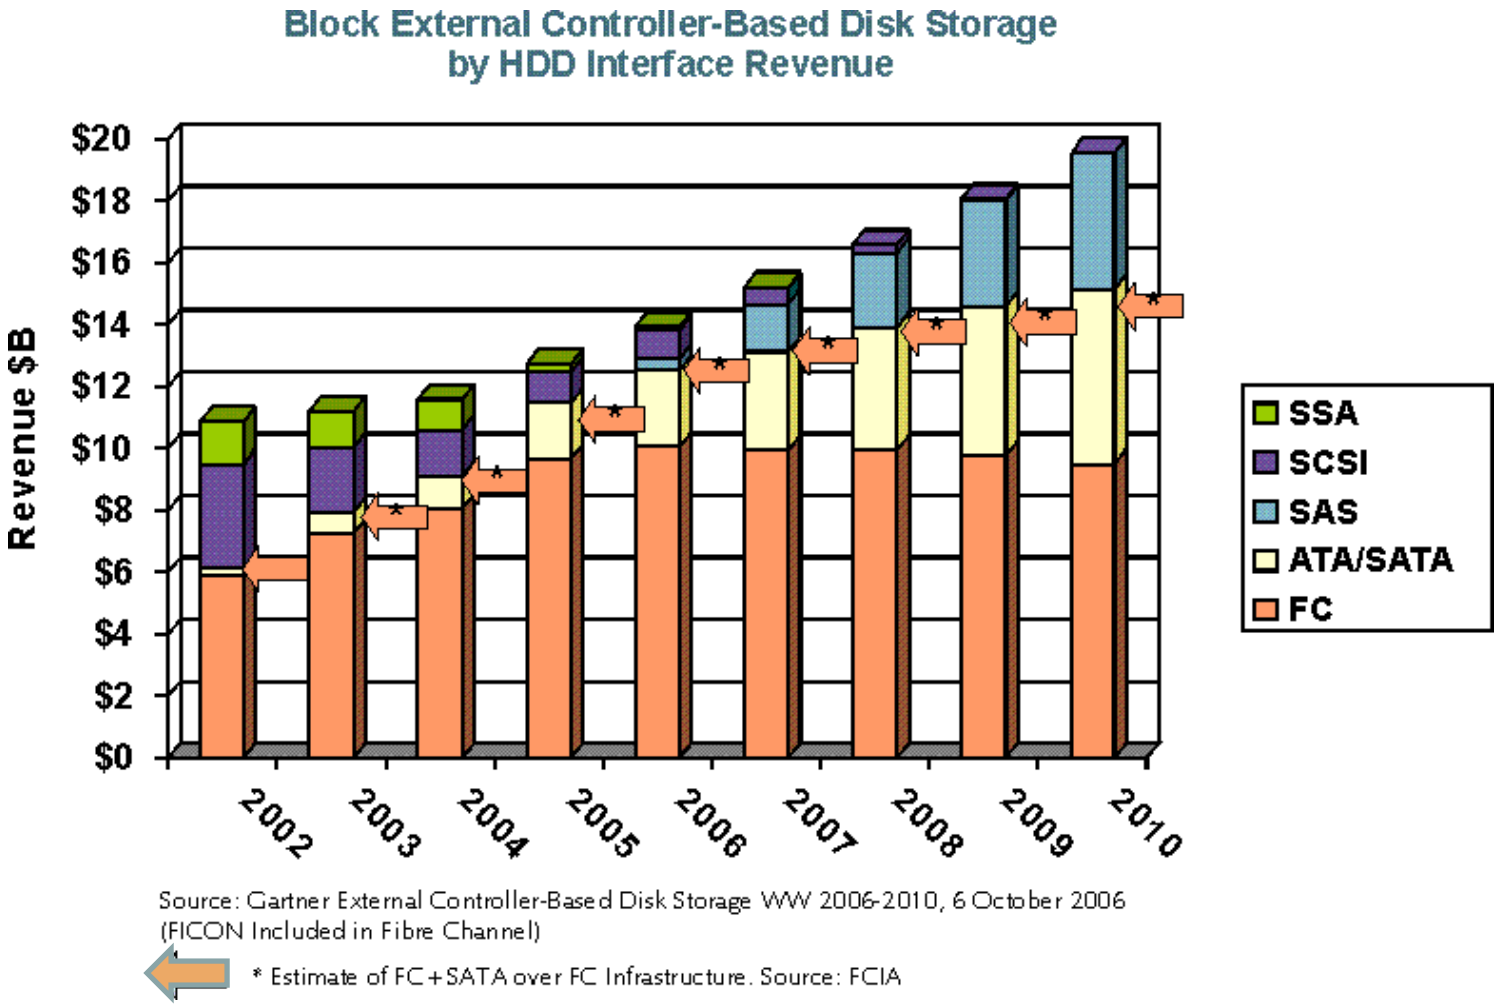

Fibre Channel's Continuous Evolution

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: FCOE
- Lower cost solutions
- Simplified solutions (Plug-n-play)
  - FC-SCM

Fibre Channel Is Being Improved According To Real Customer Requirements

What is New in Fibre Channel Standards for Protocols APIs, and Management?

- FCOE → FC from a SAN concept
- Fabric Services,
- Virtualization,
- Operational Flexibility

FCoE

- One should think about FCoE as placing the FC protocol on a new physical link
  - New Lossless Ethernet links instead of physical FC links
  - But it is still Fibre Channel
- The protocol is defined by the INCITS Fibre Channel (T11) technical committee
- The new Lossless Ethernet links was defined in the IEEE 802.1
- FC-BB-5 (FCoE) completed at T11: Document has been forwarded to INCITS. Public review completed with no comments, next step is publication.

FCoE Fabrics

- FCoE fabrics must inter-operate seamlessly with FC Fabrics
- Existing Fabric Services continue to operate in an FCoE environment
- FCoE supports all Fibre Channel advanced features (e.g. virtual fabrics, IFR, security, etc.)
- **FCoE will not require changes to existing SAN management software**

FCoE Relation to ISO Layers

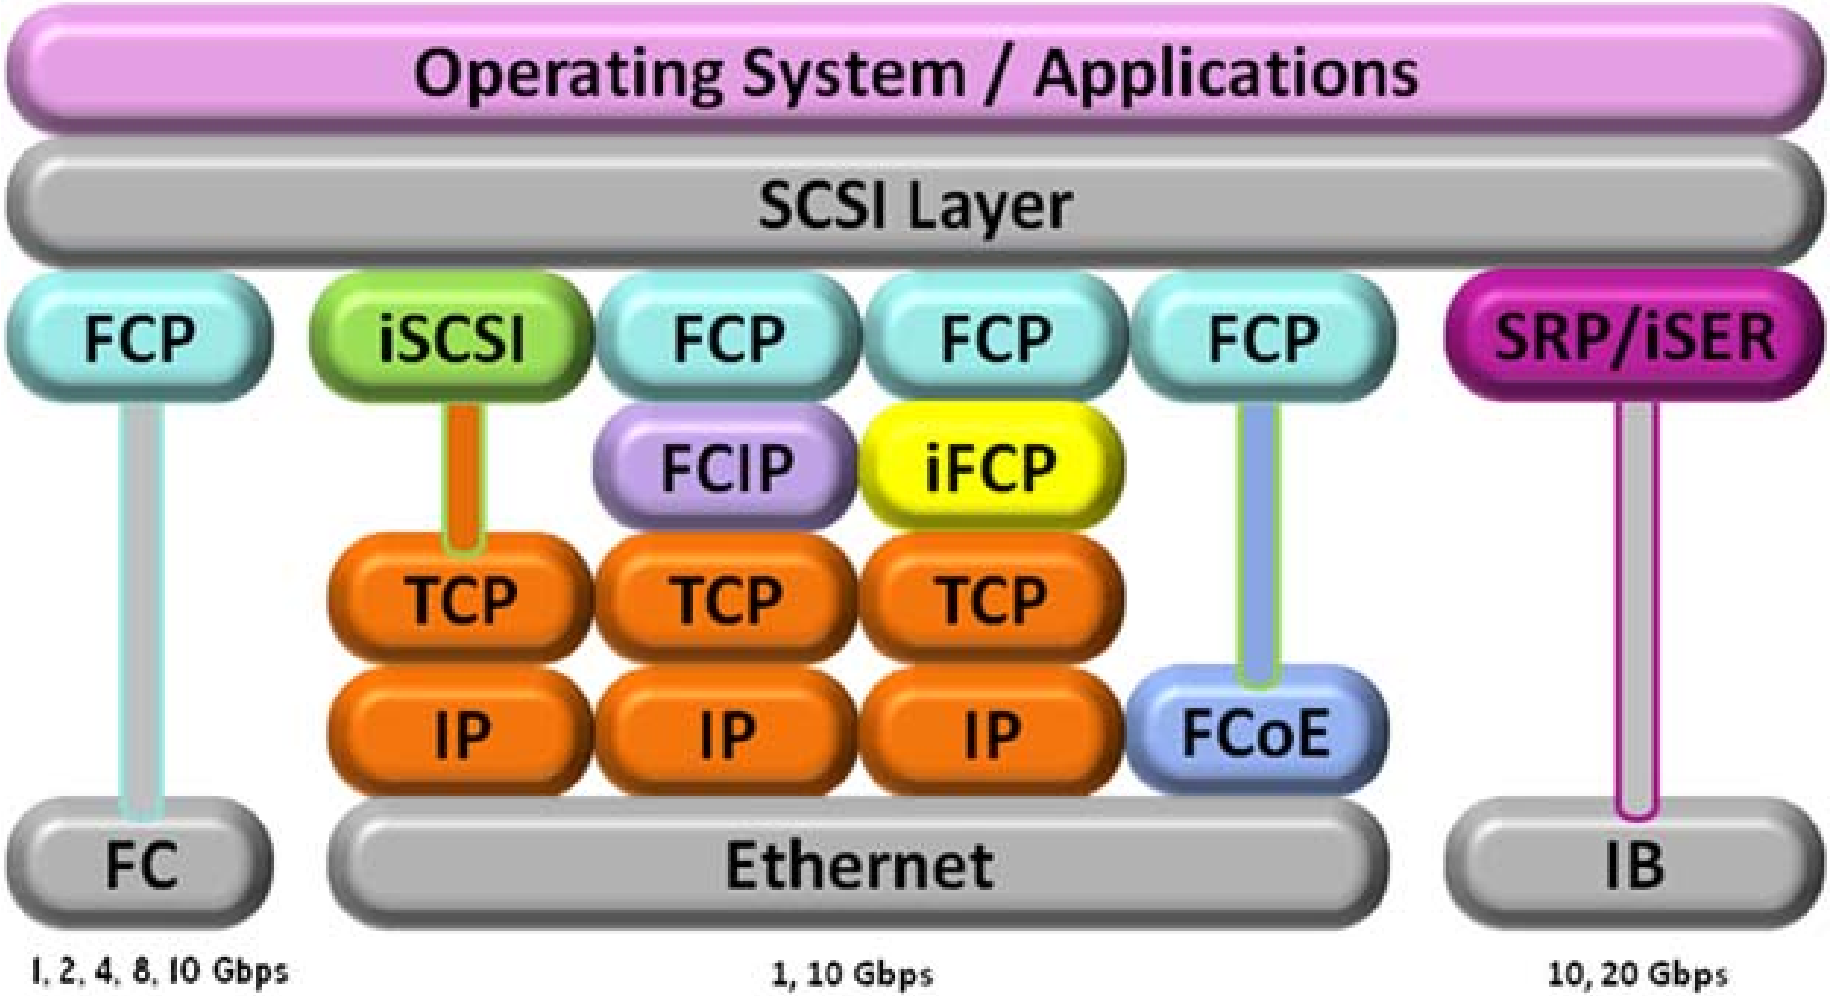

FCoE – Network Convergence

- Server or host system has a single adapter replacing multiple different types of adapters
- The datacenter has a single network instead of two or three different networks
- The number of cables and connections is dramatically reduced
- According to a recent press release by an FCIA member, a converged network can provide up to:
  - 28% savings on switches, adapters and rack space,
  - 42% savings on power and cooling, and
  - 80% savings on cabling

Fabric Services

- **More than just the Name Server and Basic Zoning**
  - Distributed Services
  - Discovery and Management Services
  - Diagnostic Services

Distributed Fabric Services

- All services appear as a single logical service to the client
- Client access using Well-Known address
- Entry switch may forward requests to other switches or resolve the request locally
- Enhanced Commit Service provided for Fabric database consistency (e.g., Zoning, Policy Database)
- For each request, switches may assume multiple roles:
  - Entry, Managed, Managing
- Fabric may utilize caching mechanisms to boost performance

Discovery and Management Services

- Fabric Configuration Server
  - Provides Fabric Topology Information
  - Includes the Platform Database
- Fabric Device Management
  - Includes both static and dynamic information
  - Registration protocol for static information
  - Redirection for dynamic information
- Unzoned Name Server
  - Provides Name Server information irrespective of zoning
- Event Server
  - Notifies Nx_Ports and Switches of events
  - Improvement over RSCN by providing more granular event notifications

Diagnostic Services

- Trace Route
  - Finds the route in the Fabric between two Fx_Ports
- FC Ping
  - Verifies path to an Nx_Port
  - May also be used to verify a path to switch via the Domain Controller

Virtualization

- **Virtual Fabrics**
  - Frame Tagging
- **Virtual Channels**
  - Enables QoS on the link
- **Fabric Based Virtualization**
  - FAIS Technology

Frame Tagging

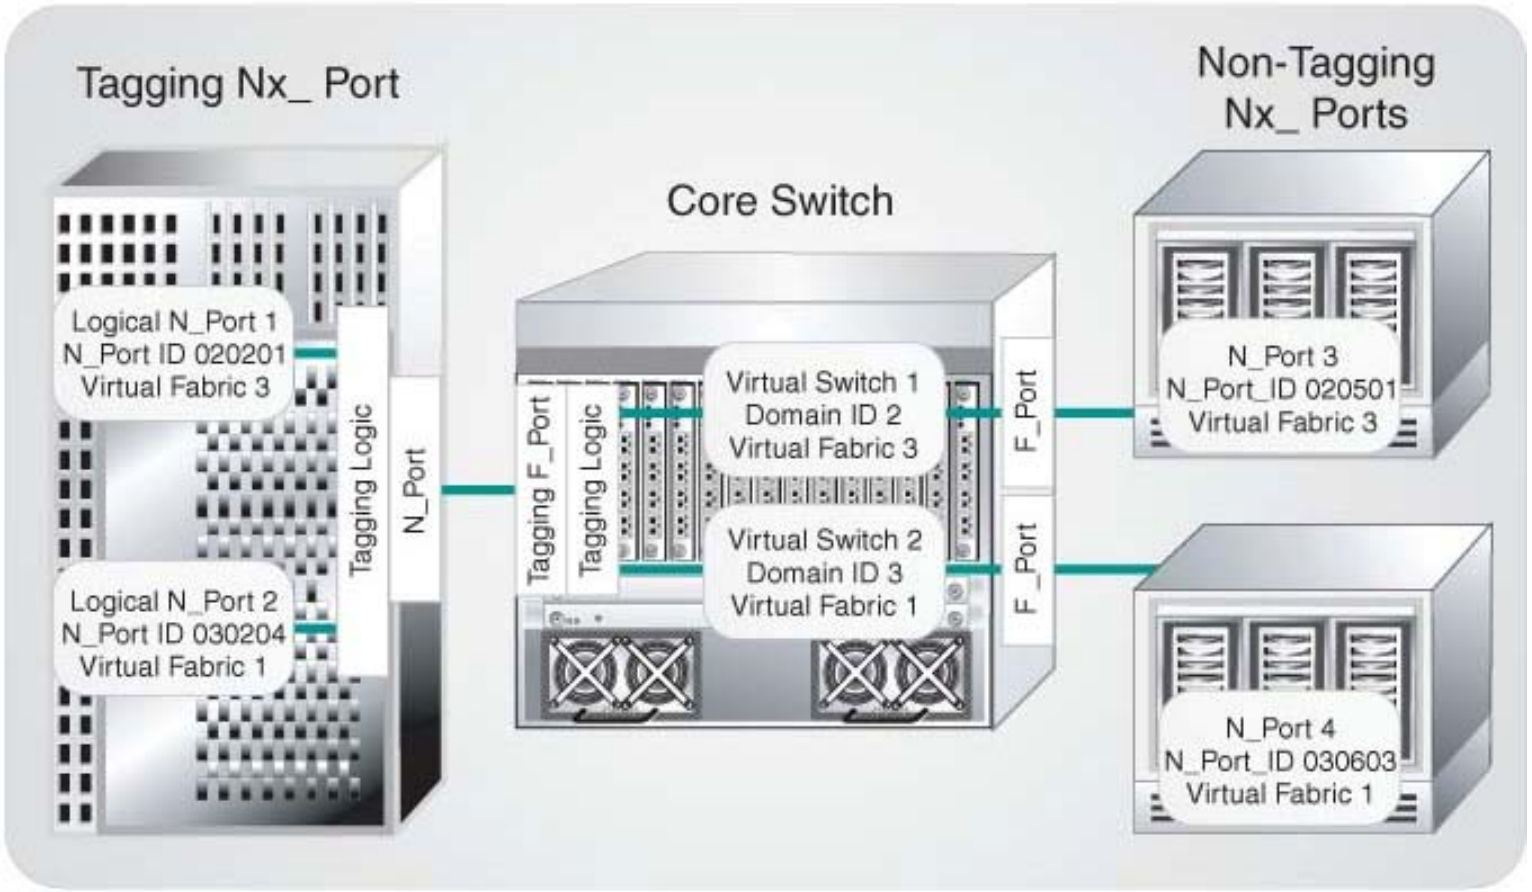

Virtual Channels

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

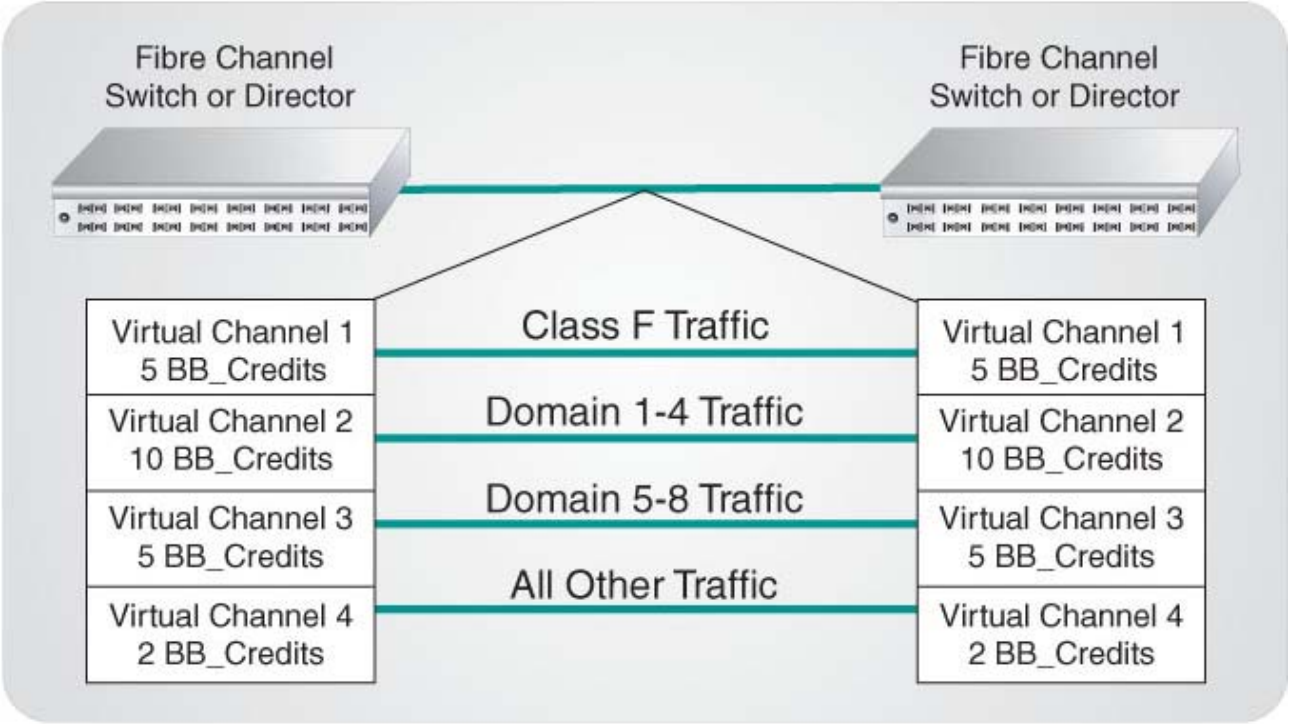

Fabric Based Virtualization

- FC-Fabric Application Interface Specification (FAIS)

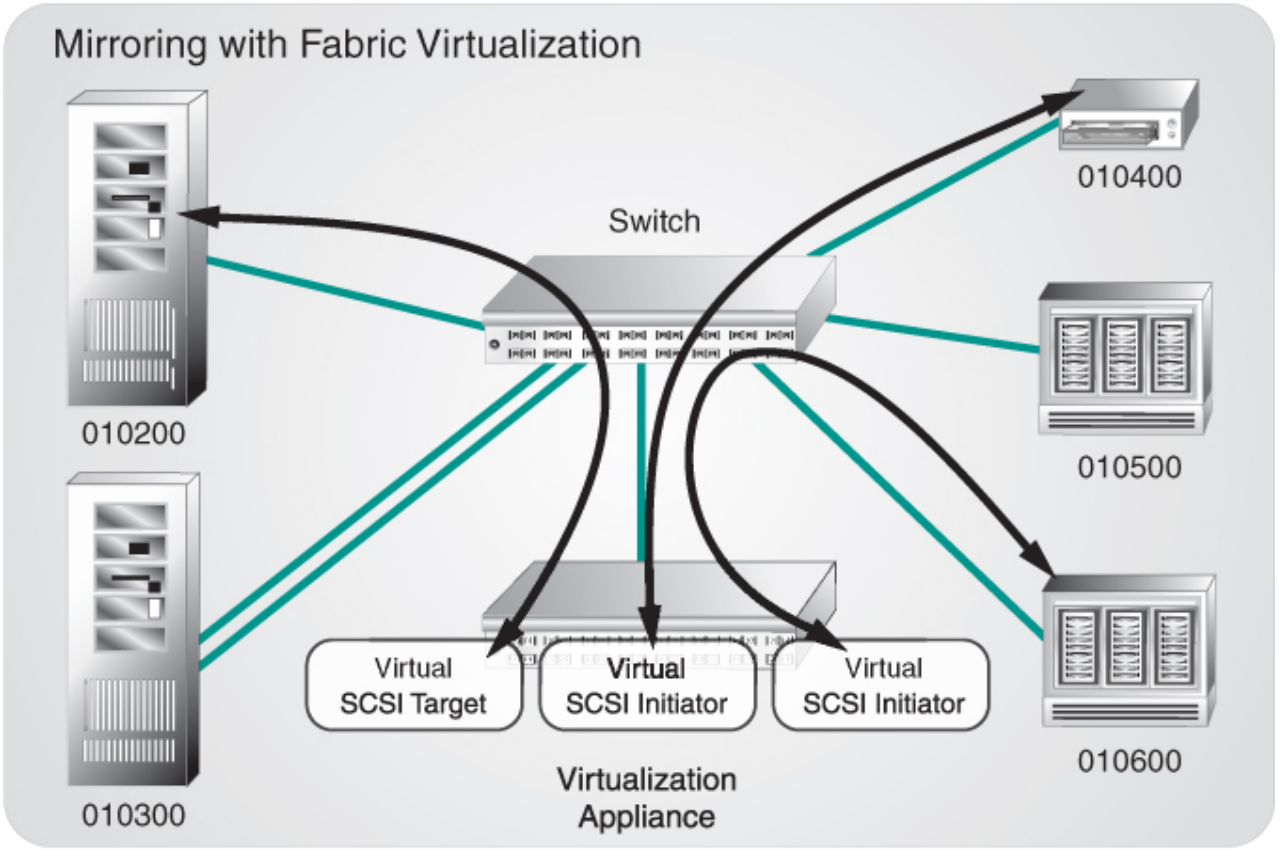

Operational Flexibility

Simplified Configuration Management

- Defines behaviors to simplify interoperability is small environments
  - Management Tools
  - Hosts
  - Fabrics
  - Storage
- A Technical Report, not a standard

Fibre Channel: The Storage of Business

Education SNIA

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

The combination of 8 GFC, virtualization, SSD in a SAN improves performance, productivity, and profitability across the enterprise.

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
Steve Wilson
Joseph L. White
Skip Jones
Tom Hammond-Doel