



T12

# Fibre Channel Basics and SAN Device Discovery

Sharon P. Wang

IBM @server xSeries  
Technical Conference

Aug. 9 - 13, 2004

Chicago, IL





# Trademarks

- AIX
- AIX 5L
- DFSMS
- Enterprise Storage Server
- ESCON
- eServer
- FICON
- FlashCopy
- iSeries
- Parallel Sysplex
- pSeries
- RS/6000
- Tivoli
- TotalStorage
- xSeries
- z/OS
- zSeries
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of the Microsoft Corporations
- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc,
- TME and Tivoli are trademarks of Tivoli Systems Inc.
- UNIX is a registered trademark licensed through The Open Group
- Other company, product, and service names may be trademarks or service marks of others.

# Channel + Network ==> Fibre Channel


Best characteristics  
of channel and  
bus architecture

Best attributes  
of  
distributed  
networks



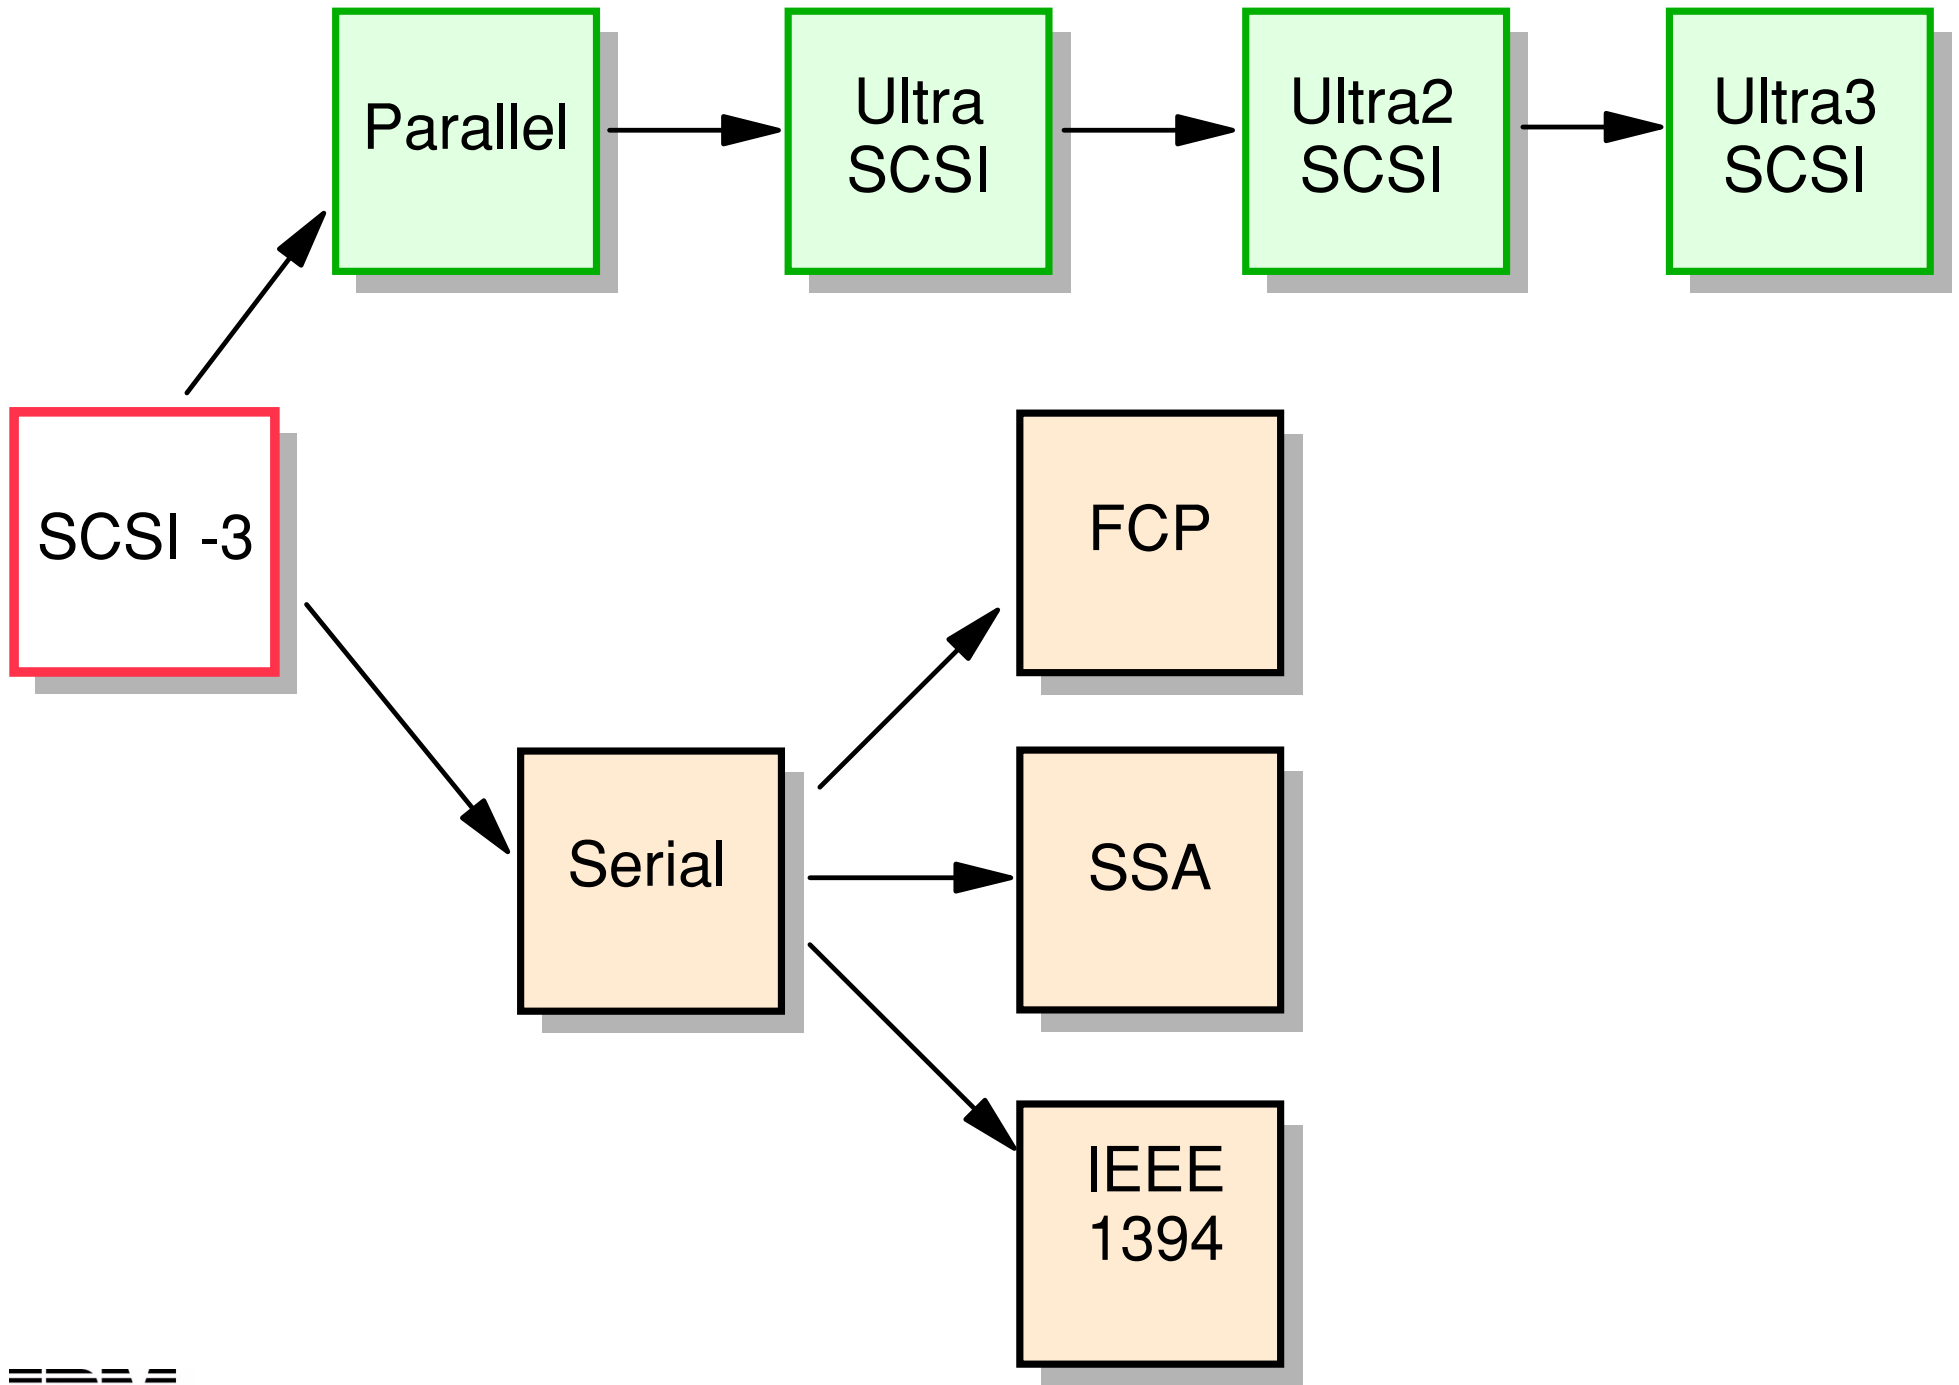
Fibre  
Channel  
Architecture

# Parallel SCSI

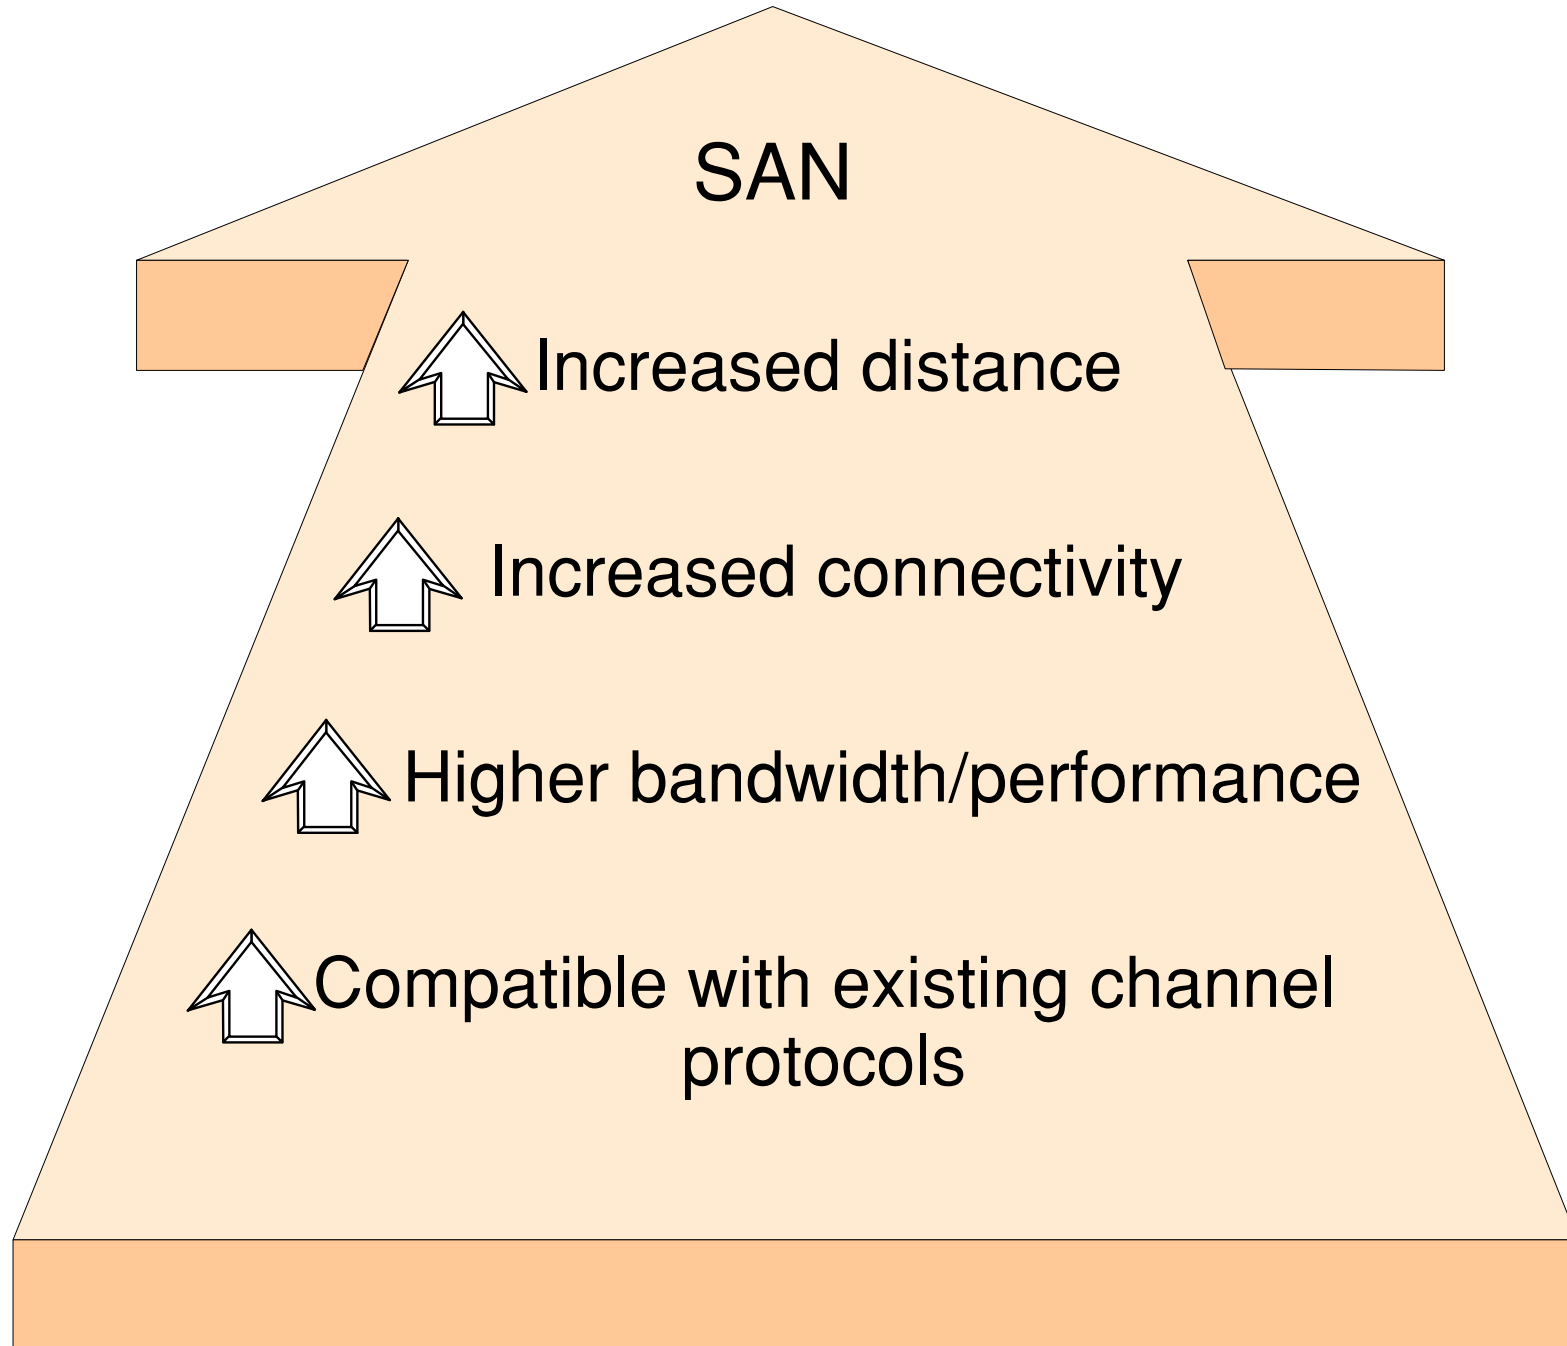


SCSI Bus Type	8-bit Width (Narrow)	16-bit Width (Wide)
SCSI-1	5 MBps	
SCSI-2 Fast	10 MBps	20 MBps
SCSI-3 Ultra	20 MBps	40 MBps
SCSI-3 Ultra2	40 MBps	80 MBps
SCSI-3 Ultra3 (Ultra160)		160 MBps
SCSI-3 Ultra320		320 MBps
SCSI-3 Ultra640		640 MBps

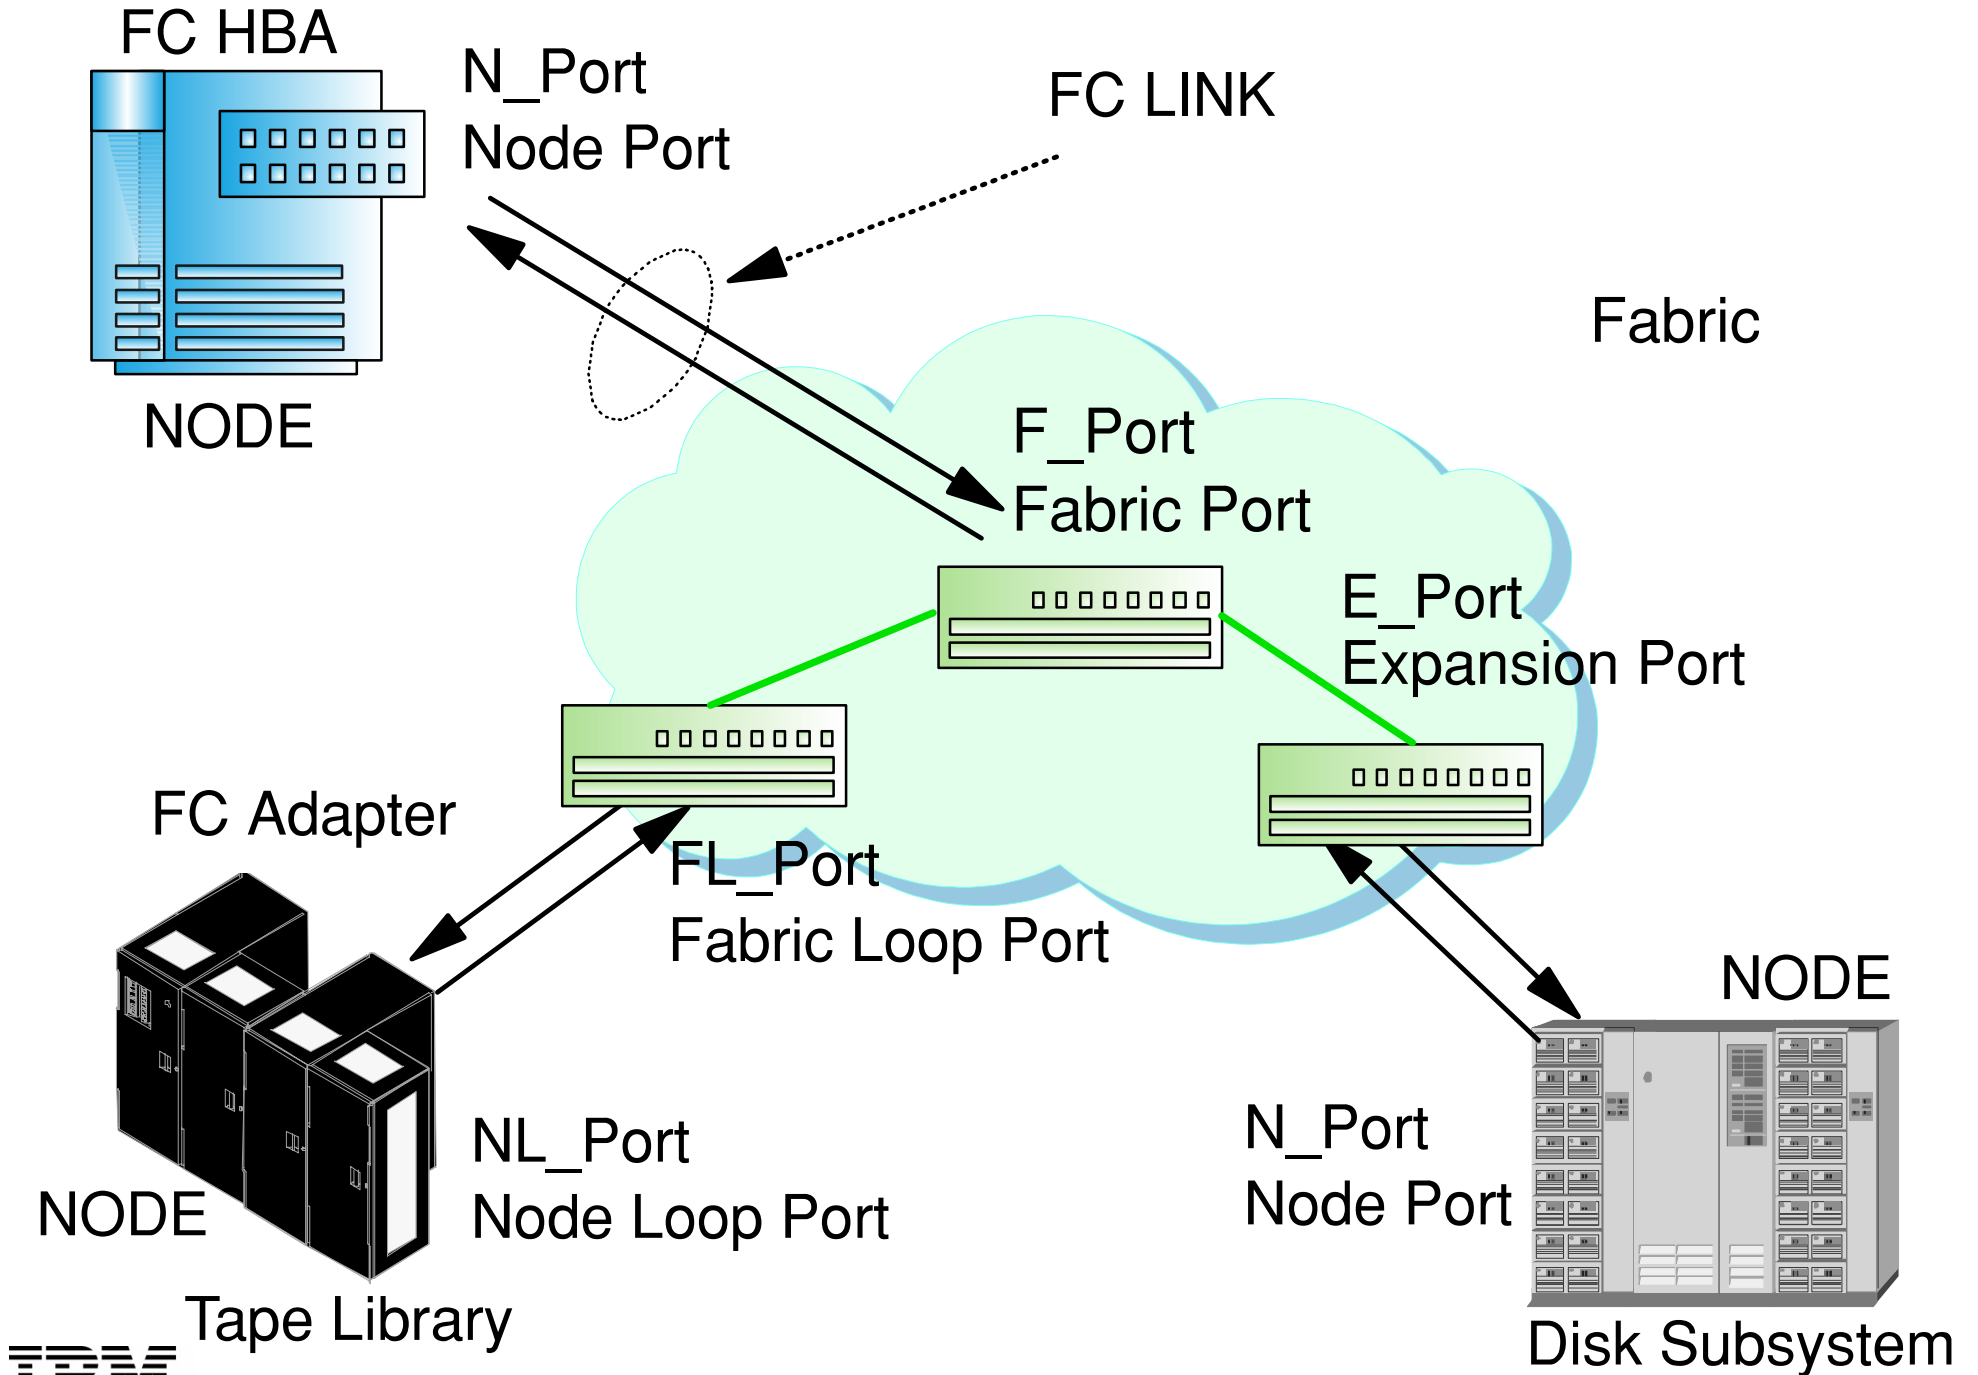
# SCSI-3 Interface Evolution



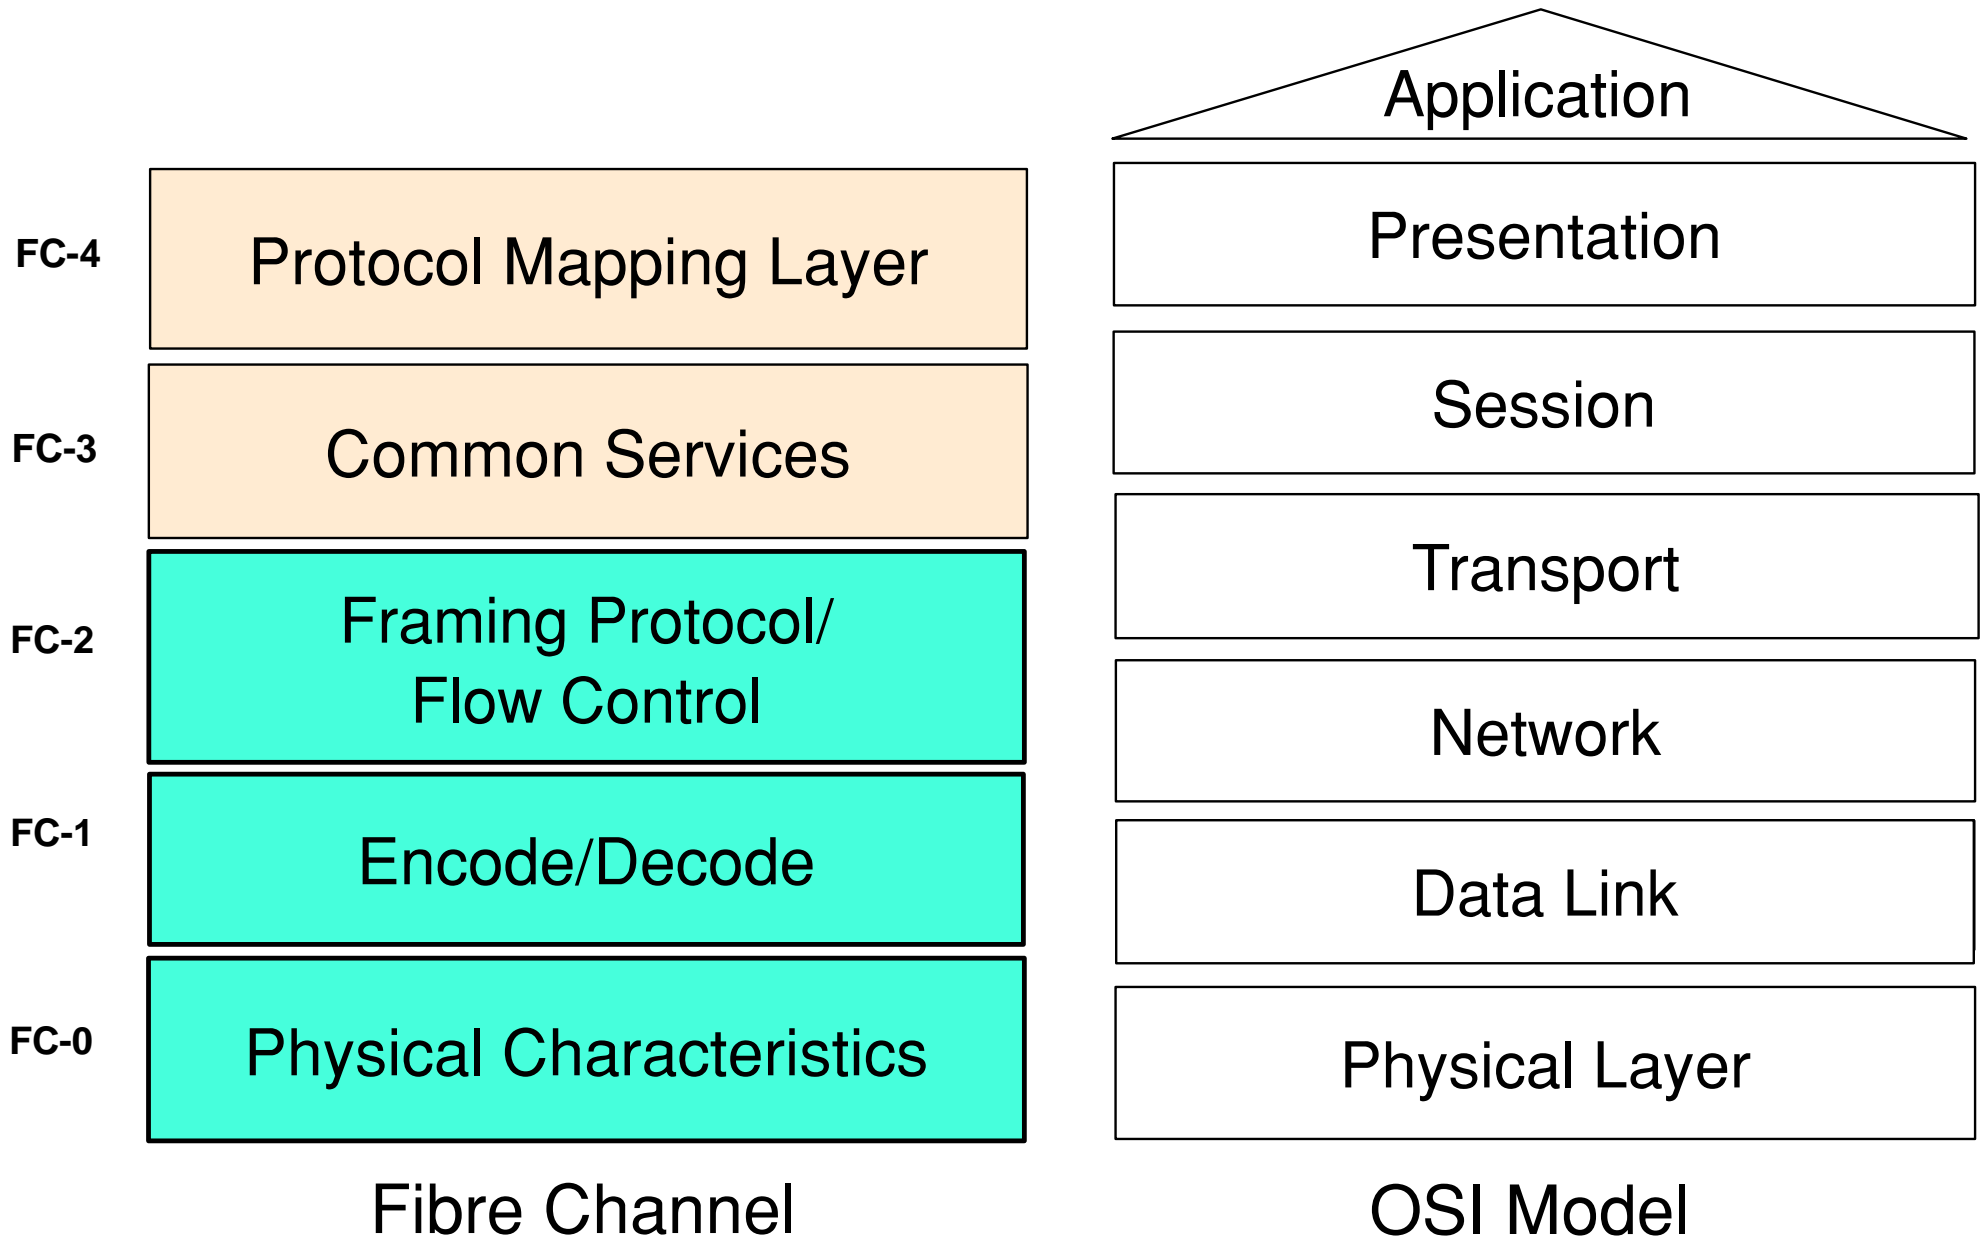
# Why Fibre Channel SAN?



# Fibre Channel Terminology

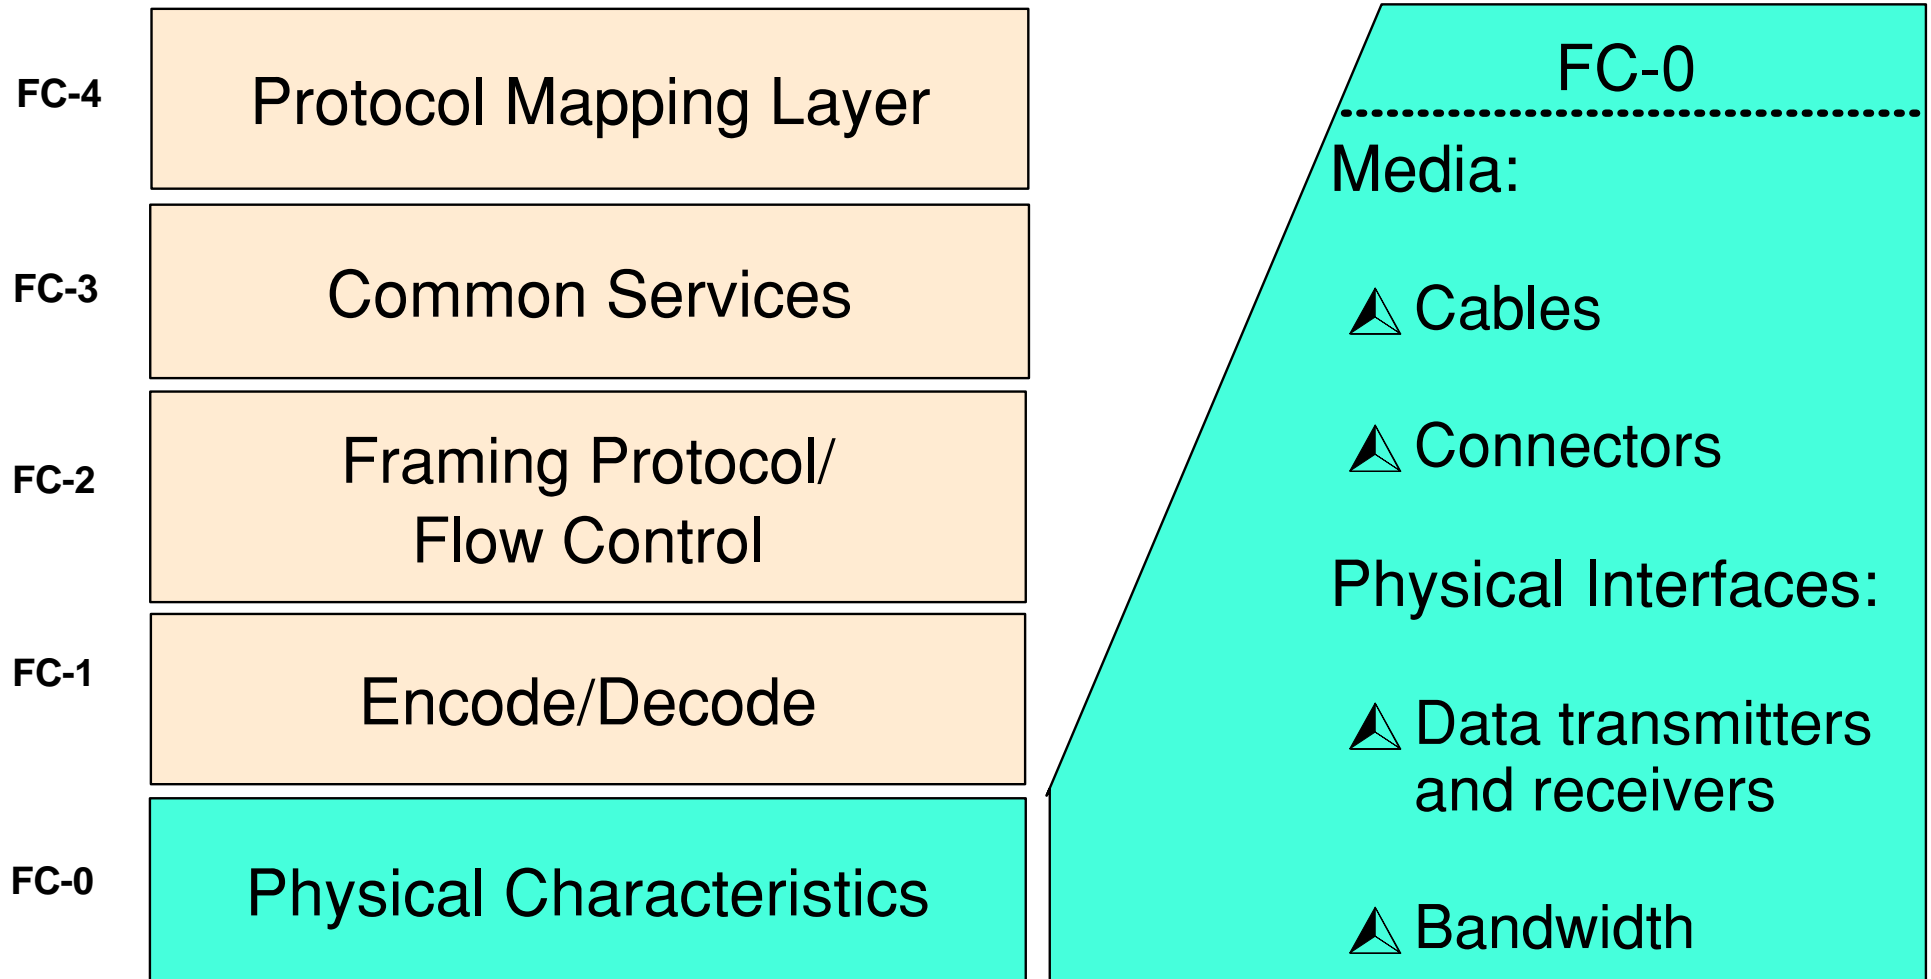


# FC and OSI Structure Comparison








# FC-0: Physical Interface and Media




# Media: Cable Options



# Fiber Optic Technology

## Multimode fiber



Multiple paths for the light to travel



SC Duplex multimode

## Single mode fiber

Single path for the light to travel





SC Duplex single mode



LC Duplex single mode


# Speed Auto Detection and Negotiation

Servers, storage, or switches




Storage

Server



# Physical Interface




Host Bus Adapter

# FC-0: Media Connectors




SC Duplex multimode




GBIC  
(SFP)

Fiber Optic  
Cable



Duplex SC Connector  
(Duplex LC Connector)

(to HBA or  
storage subsystem)

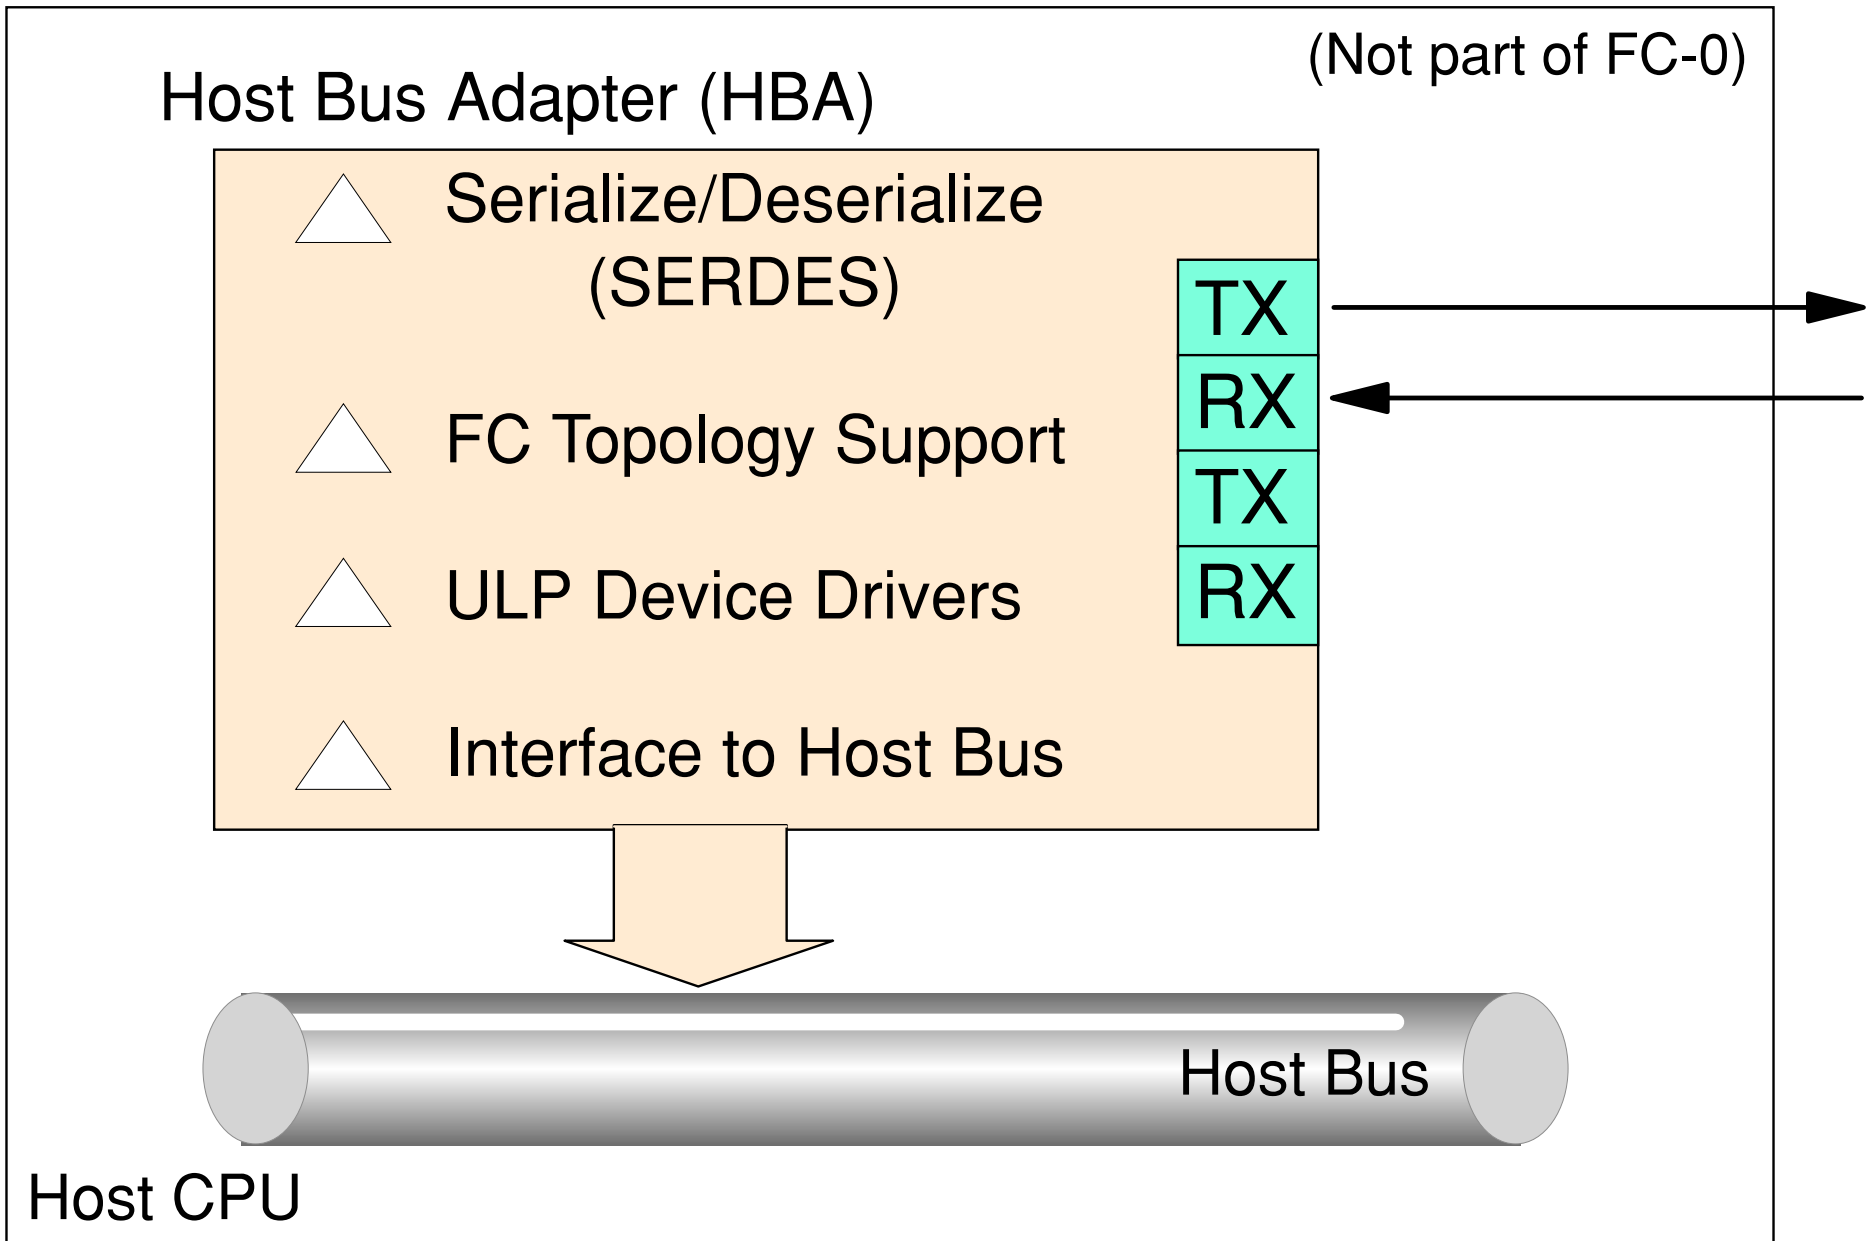


SC Duplex single mode

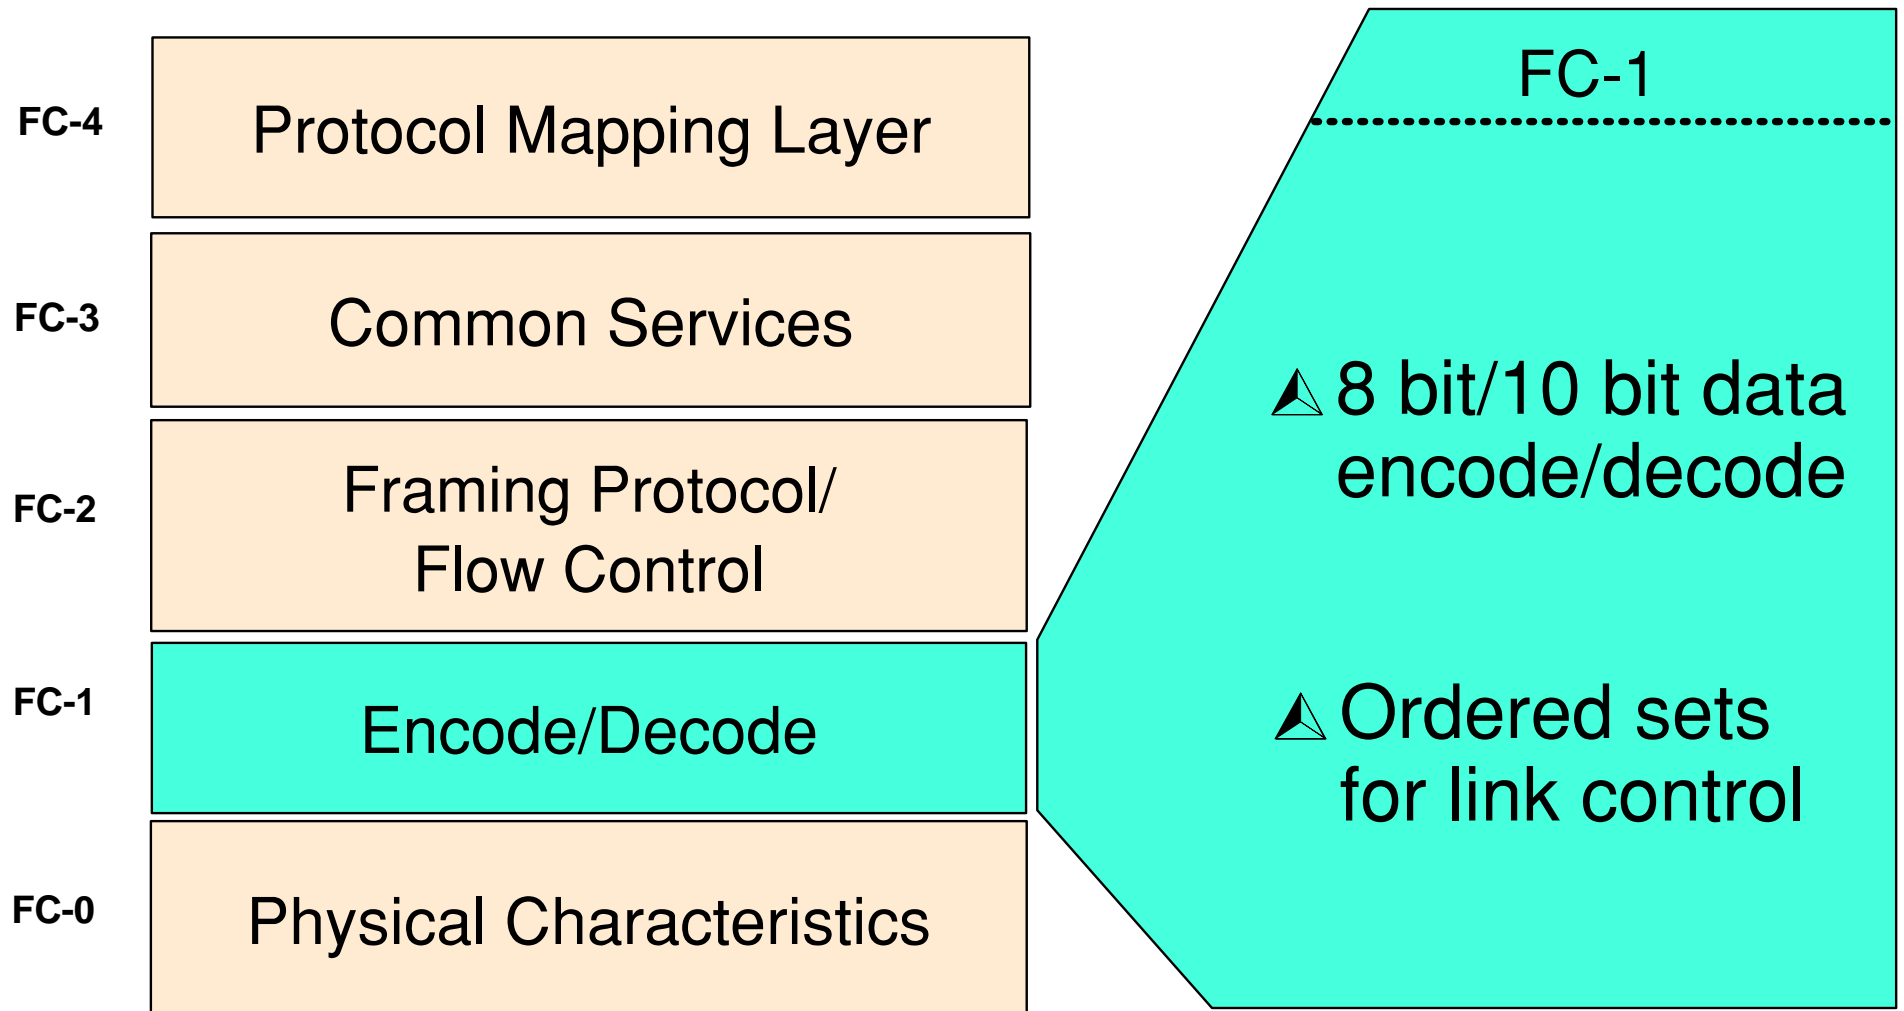


LC Duplex single mode

# HBAs: Interface between Peripherals and Host Bus

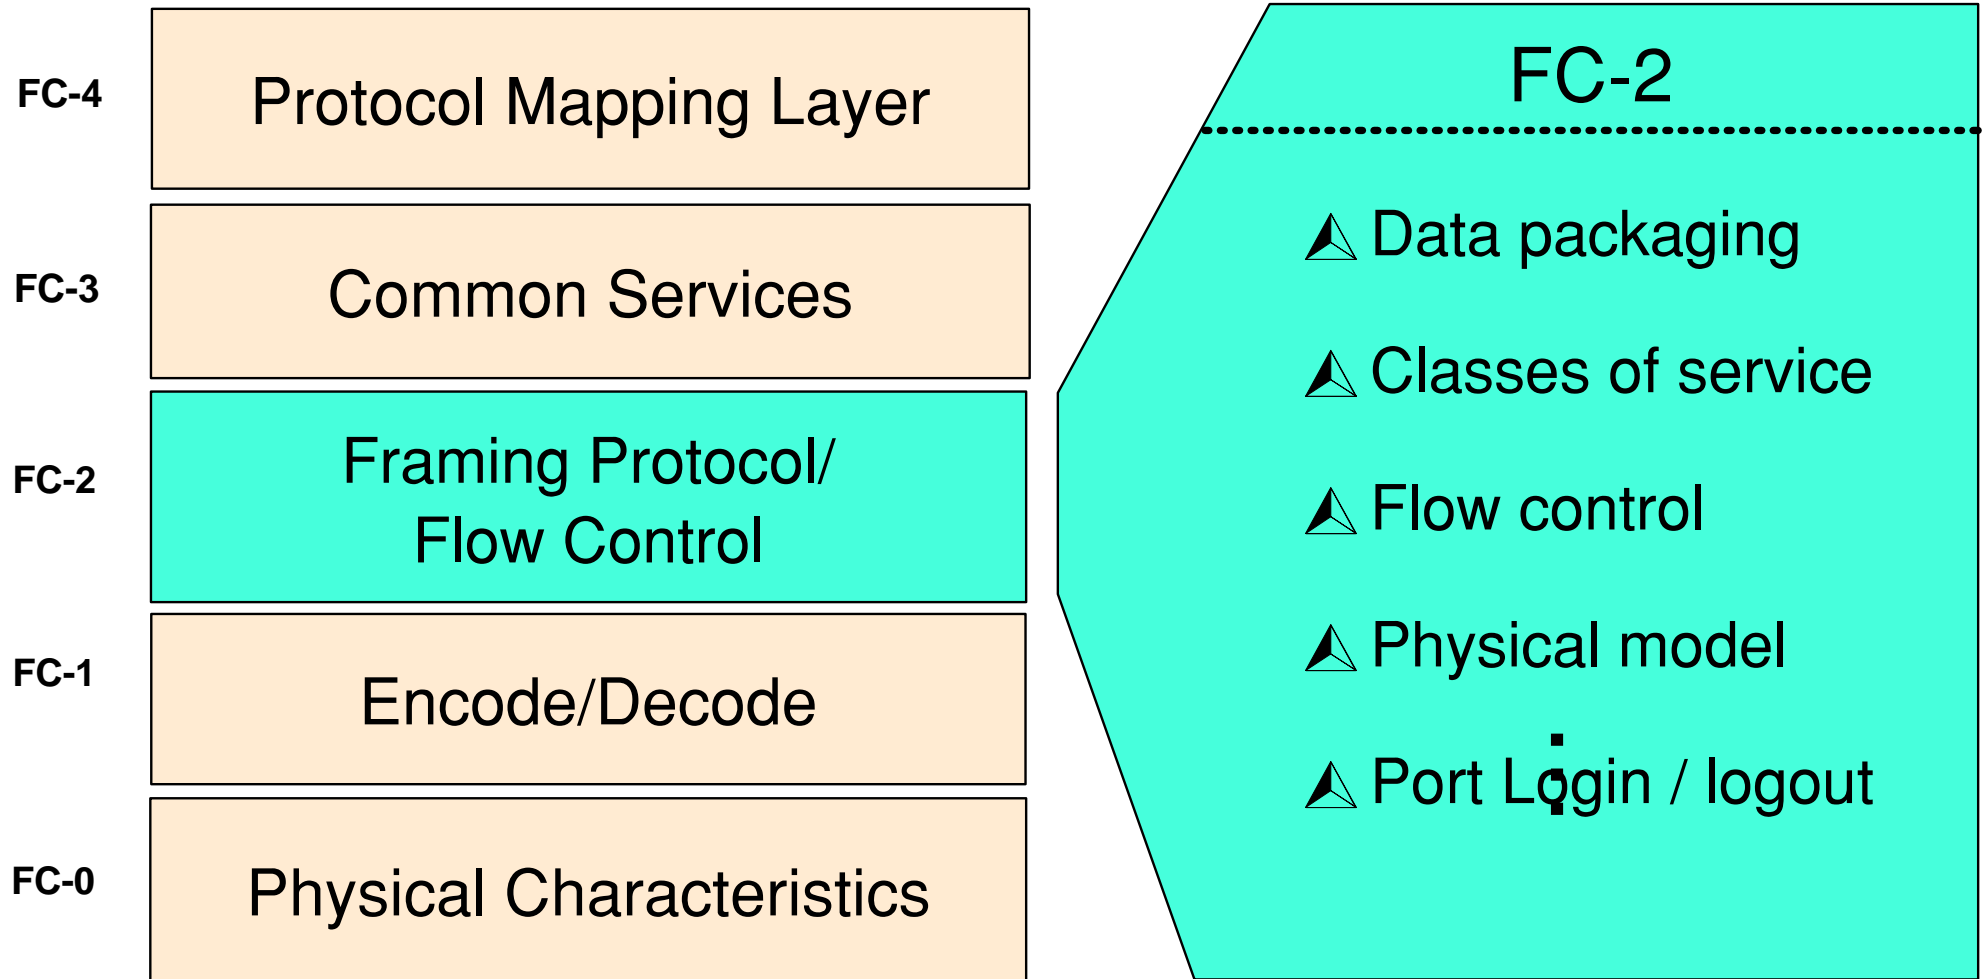


# FC-1: Transmission Protocol






# FC-2: Framing and Signaling Protocol






# Frame Content

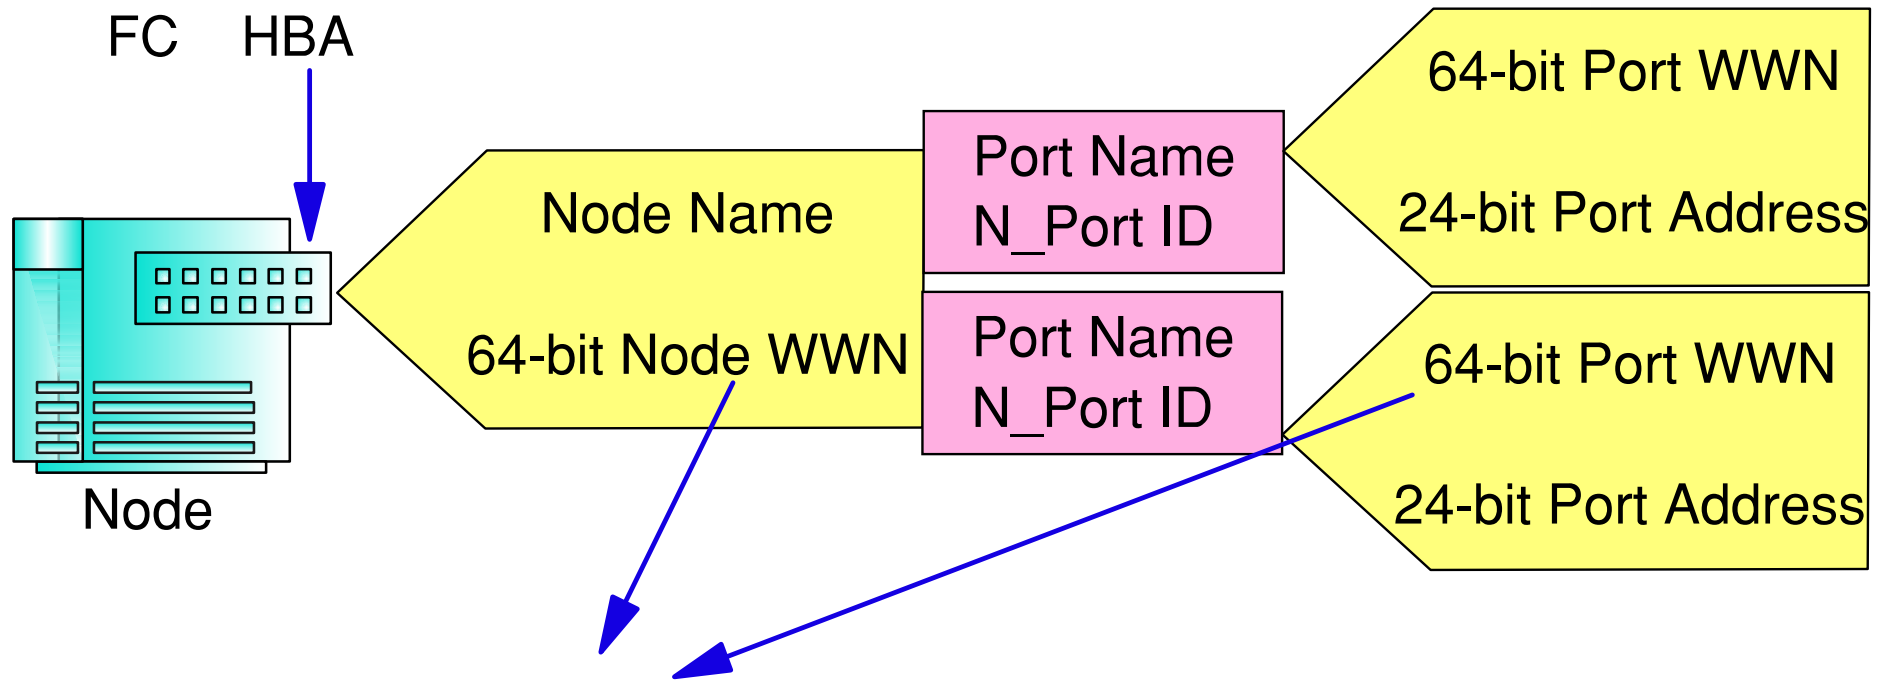


Word	Bits 31-24	Bits 23-16	bits 15-8	bits 7-0
0		Destination ID (Address)		
1		Source ID (Address)		
2				
3				
4				
5				

# Fibre Channel Topologies



# World Wide Name (WWN) & Addressing



64-bit WWN

10 : 00 : 00 : 47 : 11 : 00 : 47 : 11

Format

1000 = standard  
2nnn = extended

Company\_id

Component number

<http://standards.ieee.org/regauth/oui/index.html>

# World Wide Name (WWN)

64-bits



Manufacturer assigned  
component number

IEEE assigned company\_ID

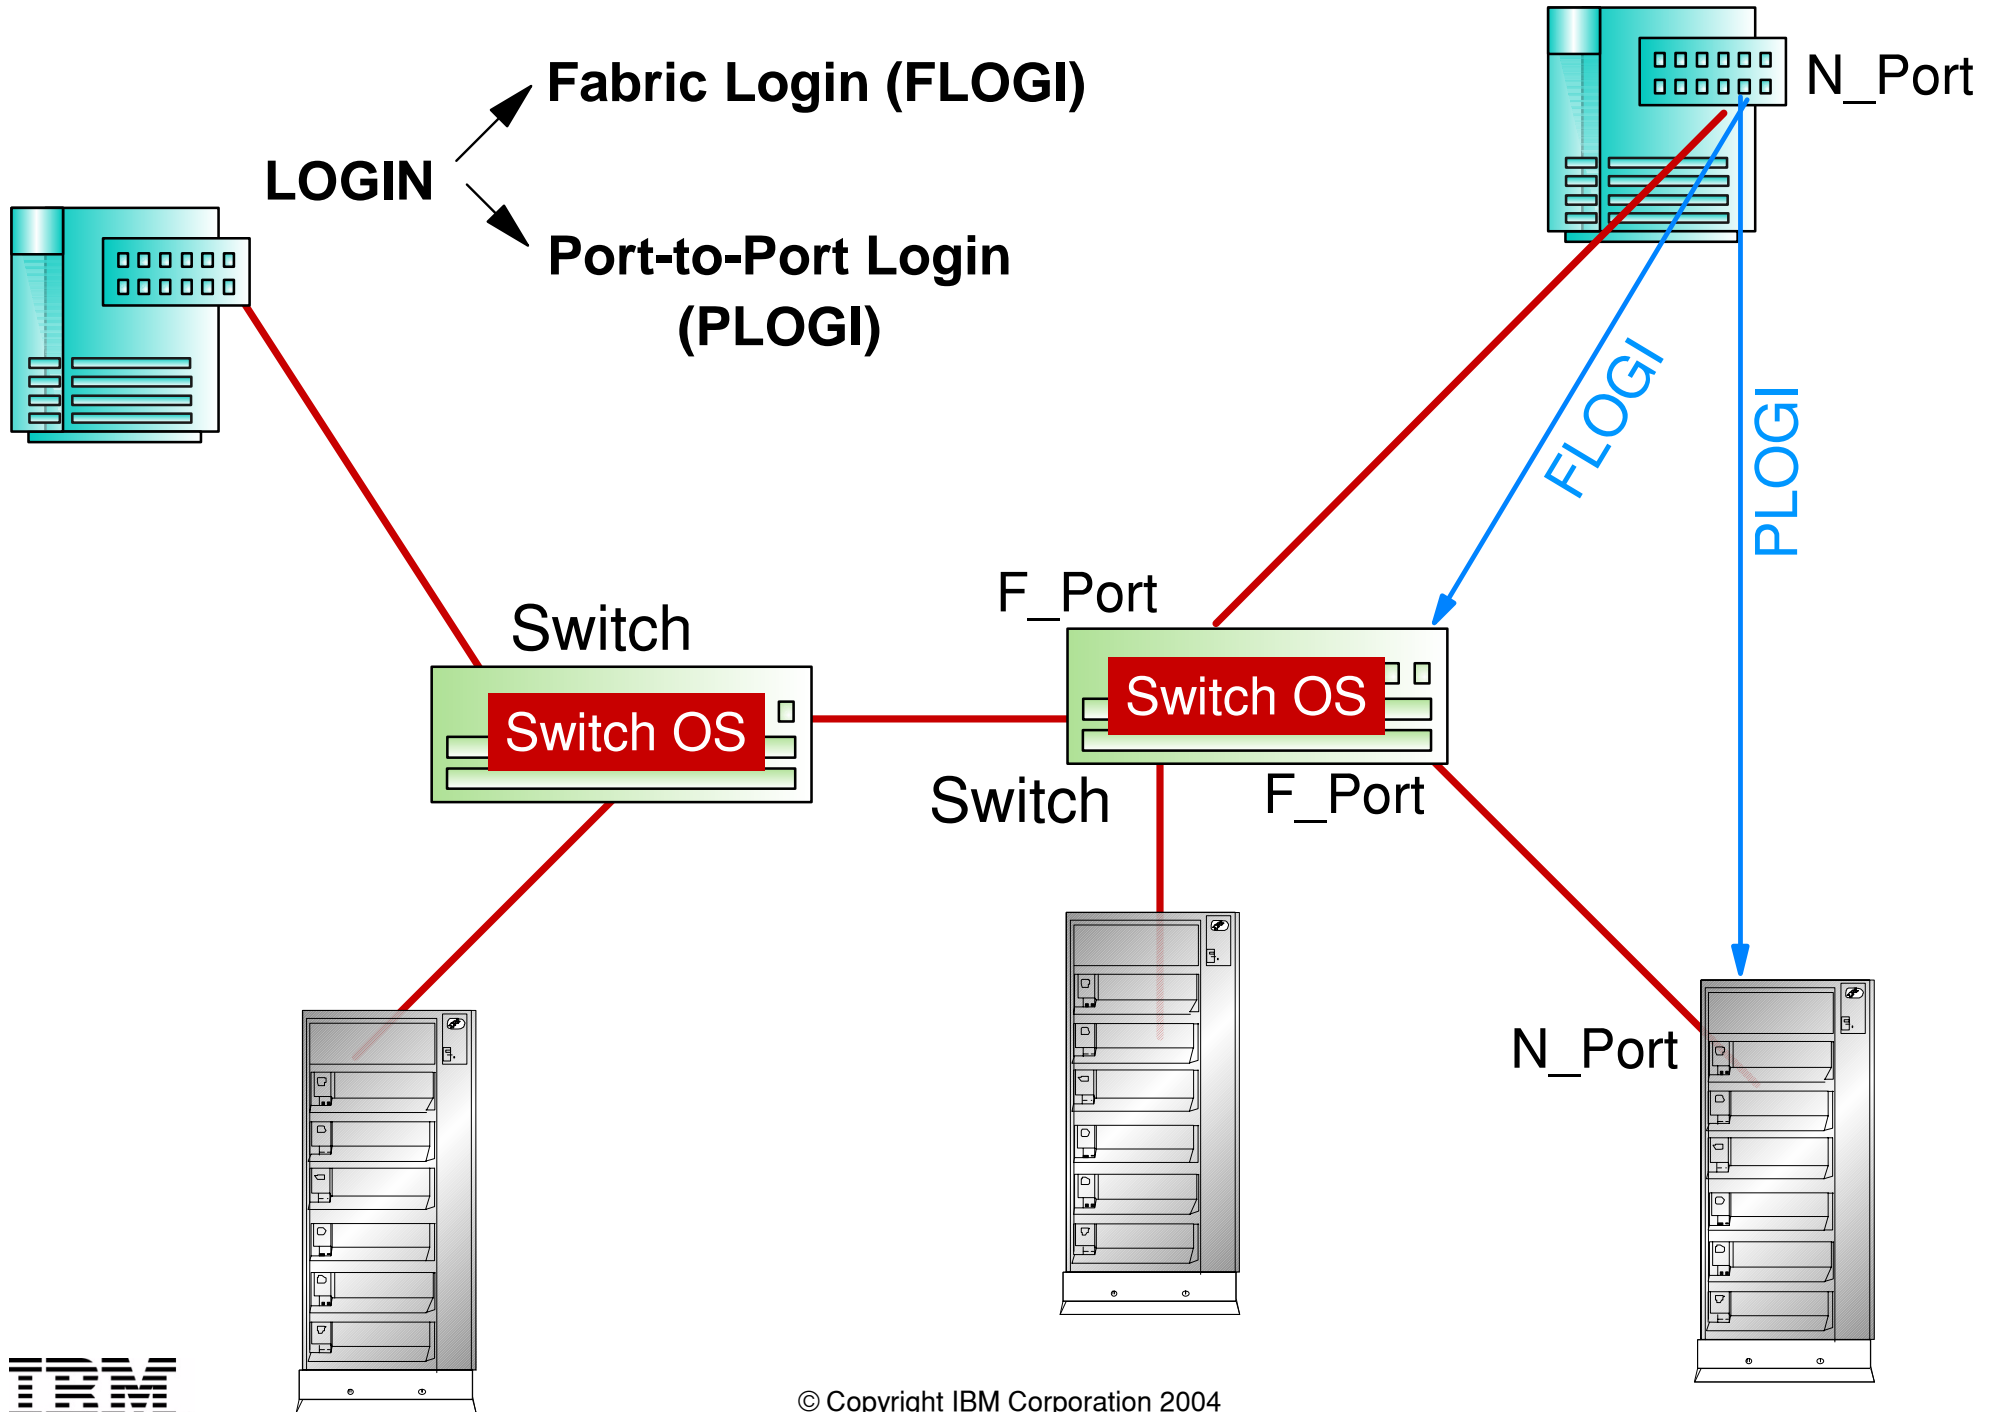
<http://standards.ieee.org/regauth/oui/index.html>

Format:

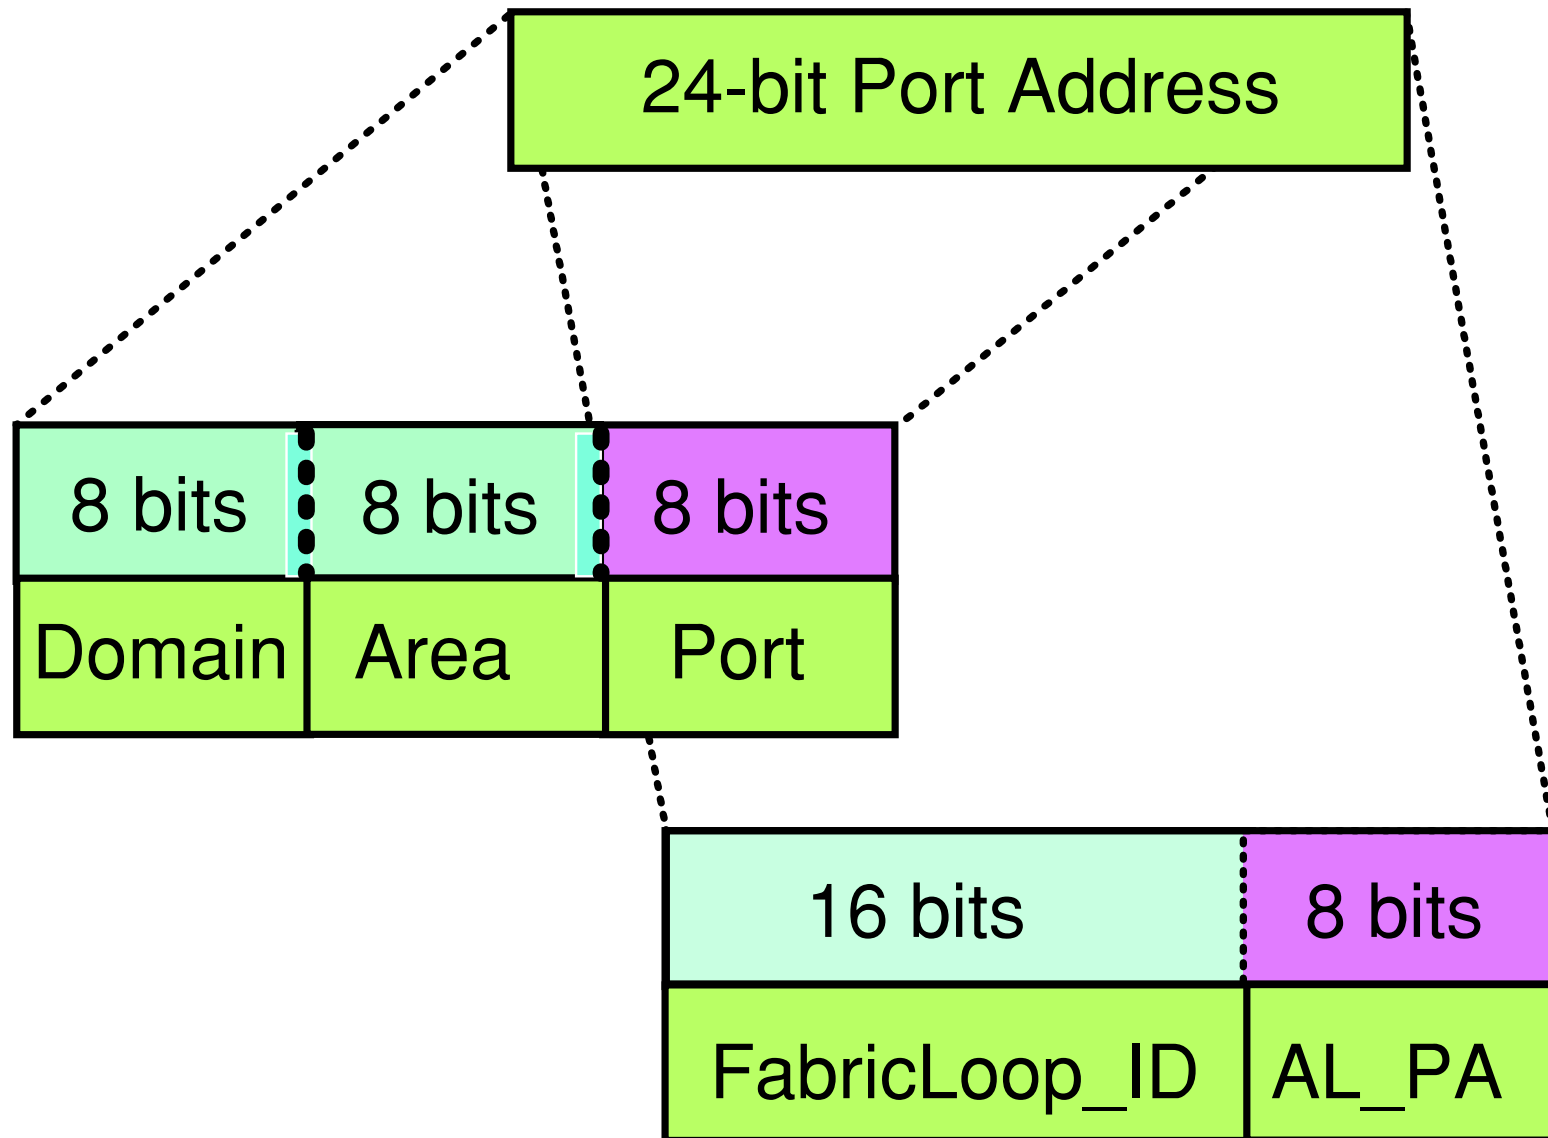
▲ 1000 - Standard

▲ 2000 - Extended


# Determining the Operating Environment



# Fabric Addressing (Port\_ID or FC\_ID)




# Point-to-Point Topology





# Arbitrated Loop with Hub

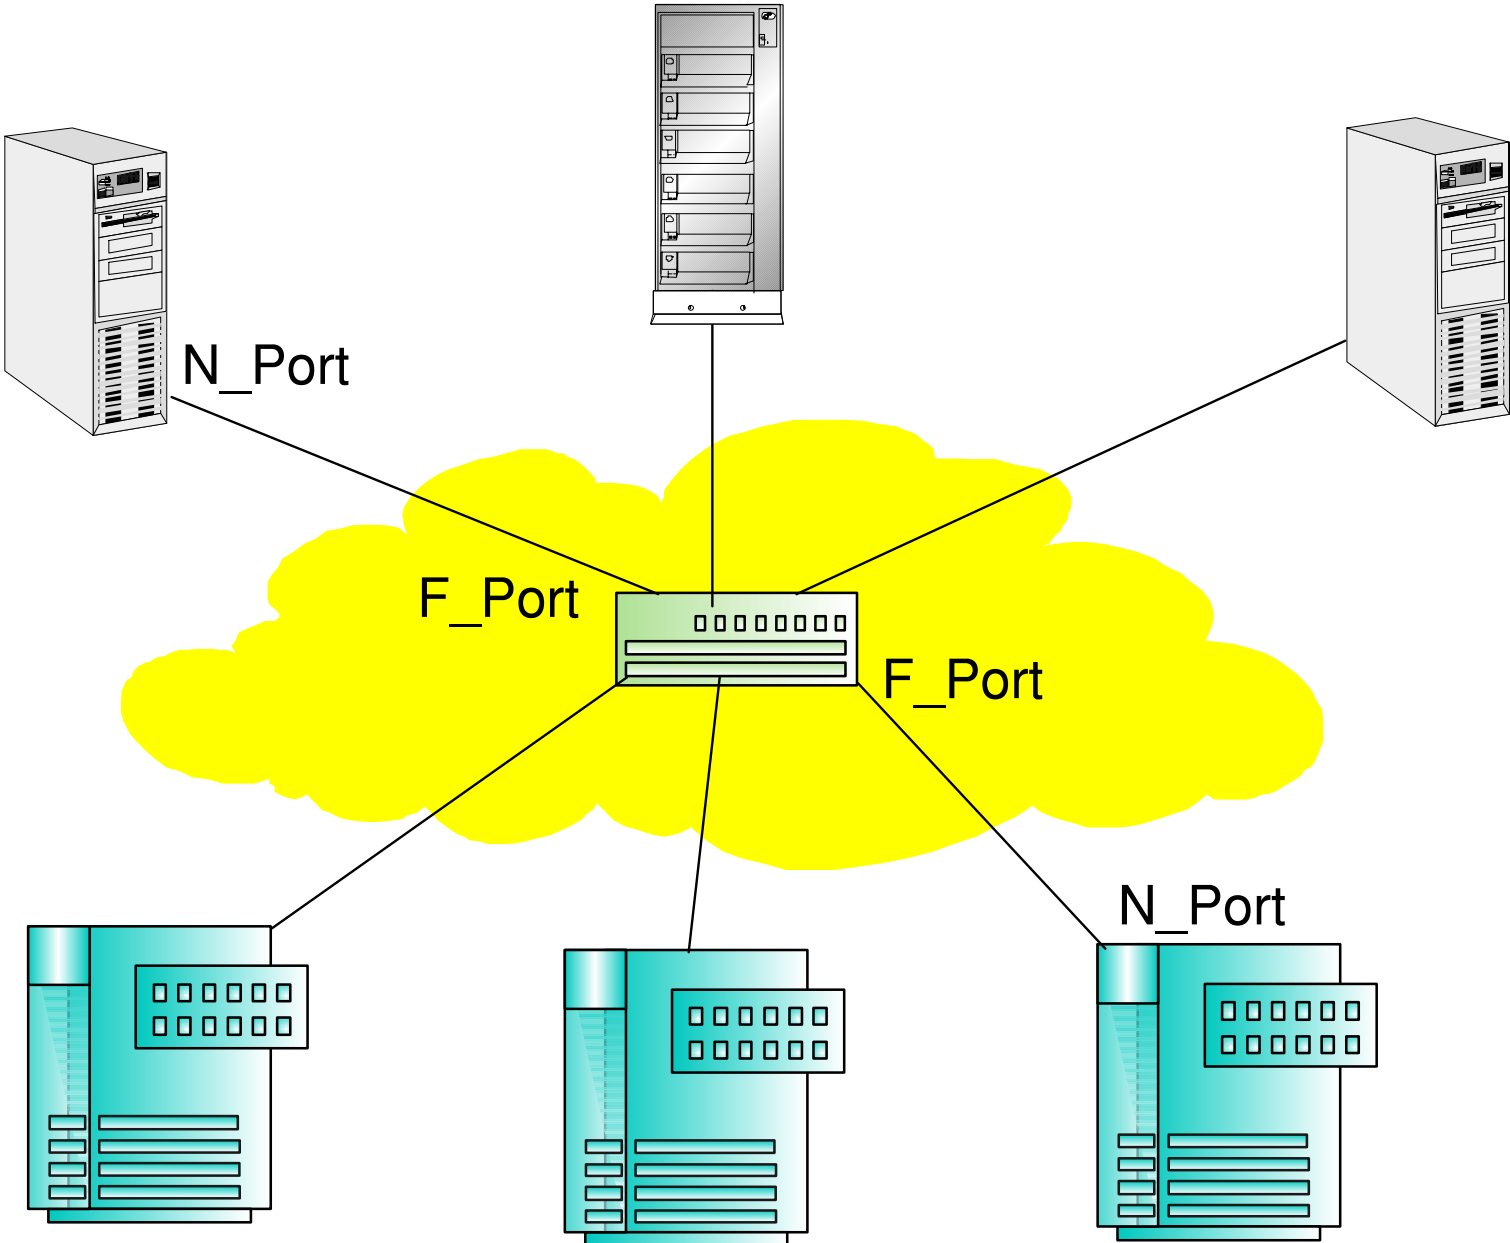


Shared Bandwidth  
© Copyright IBM Corporation 2004



PBC-Port Bypass Circuitry

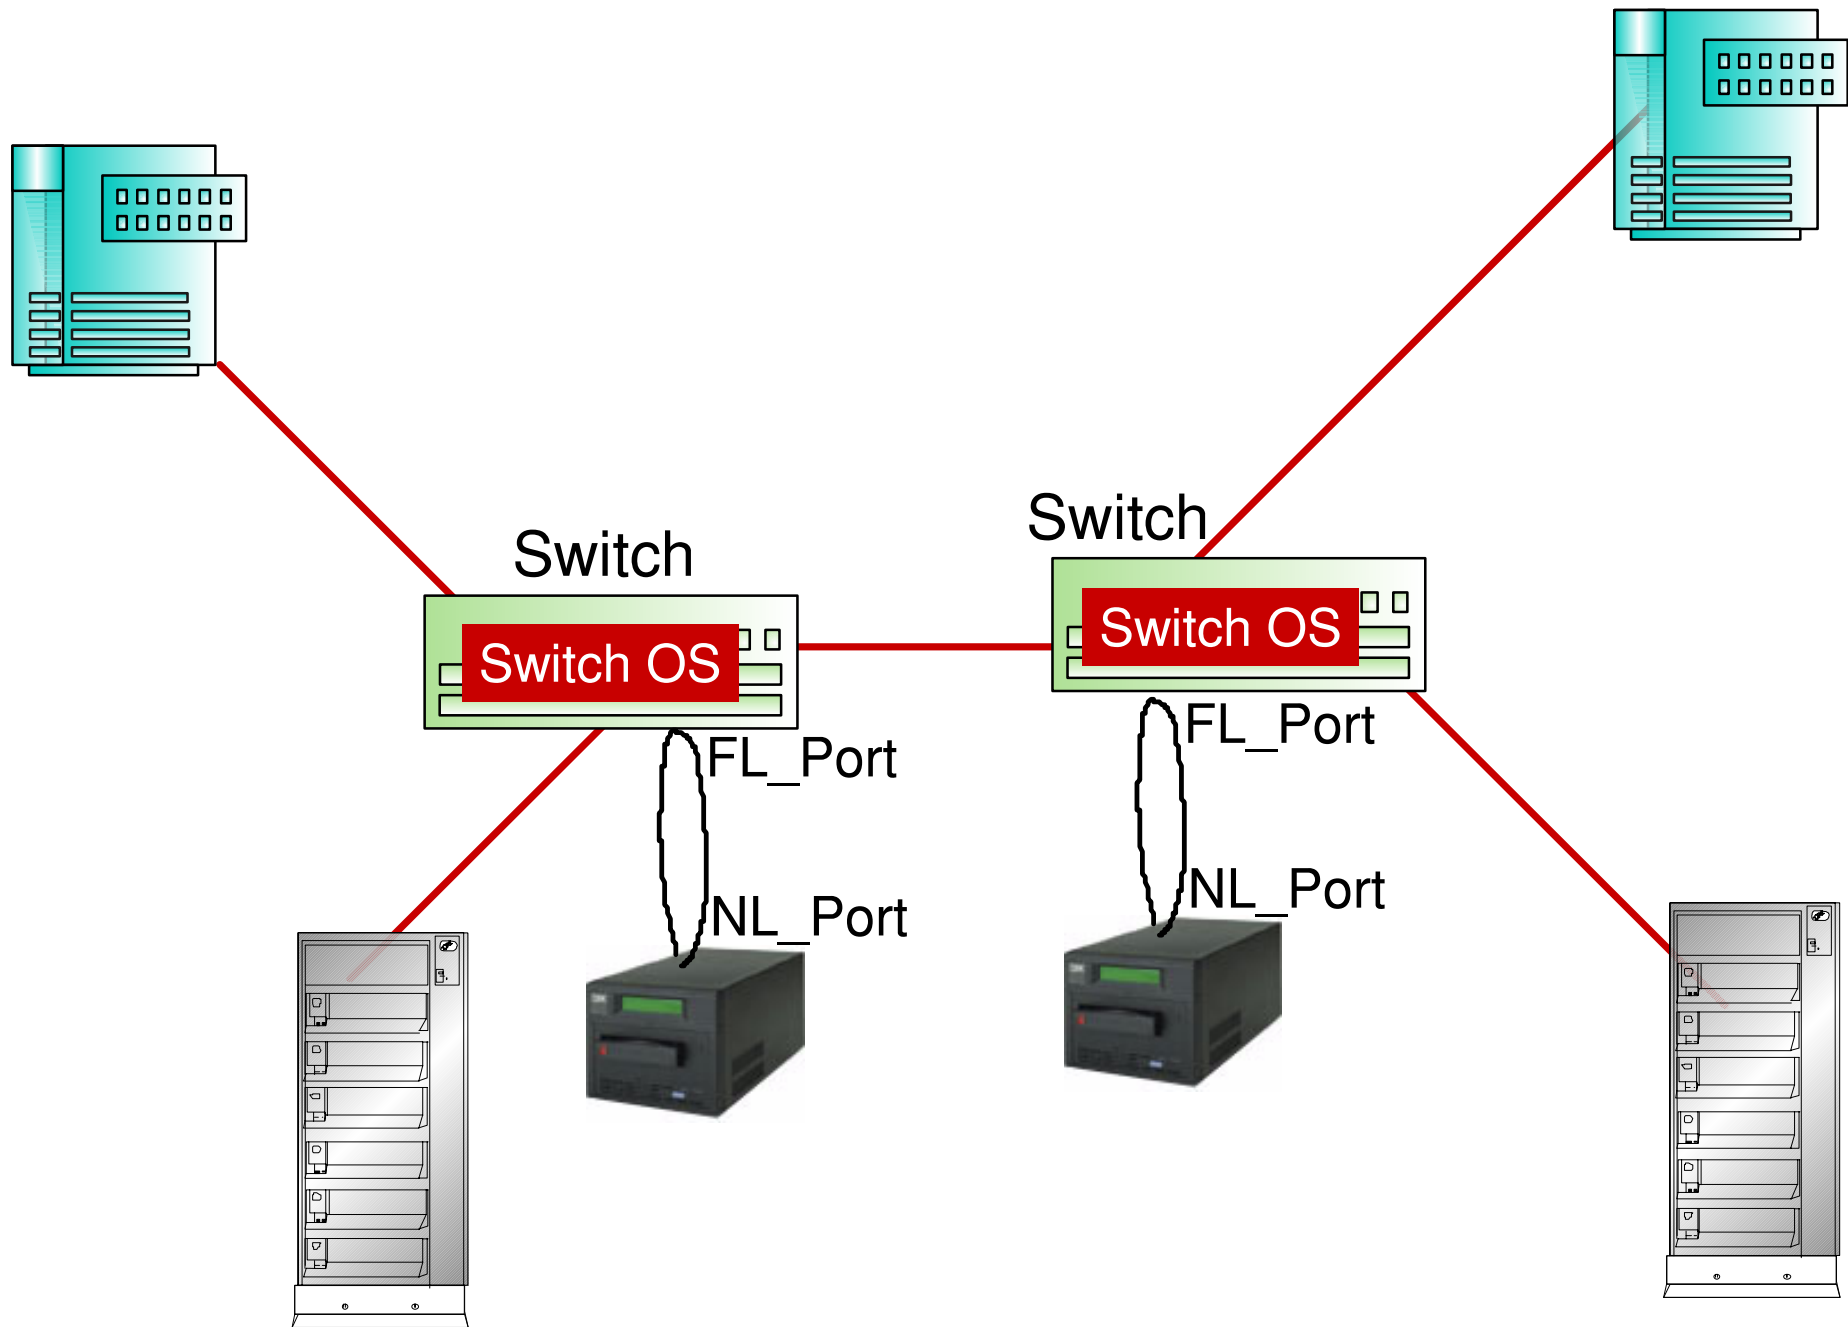
# Switched Fabric




Scalable Bandwidth



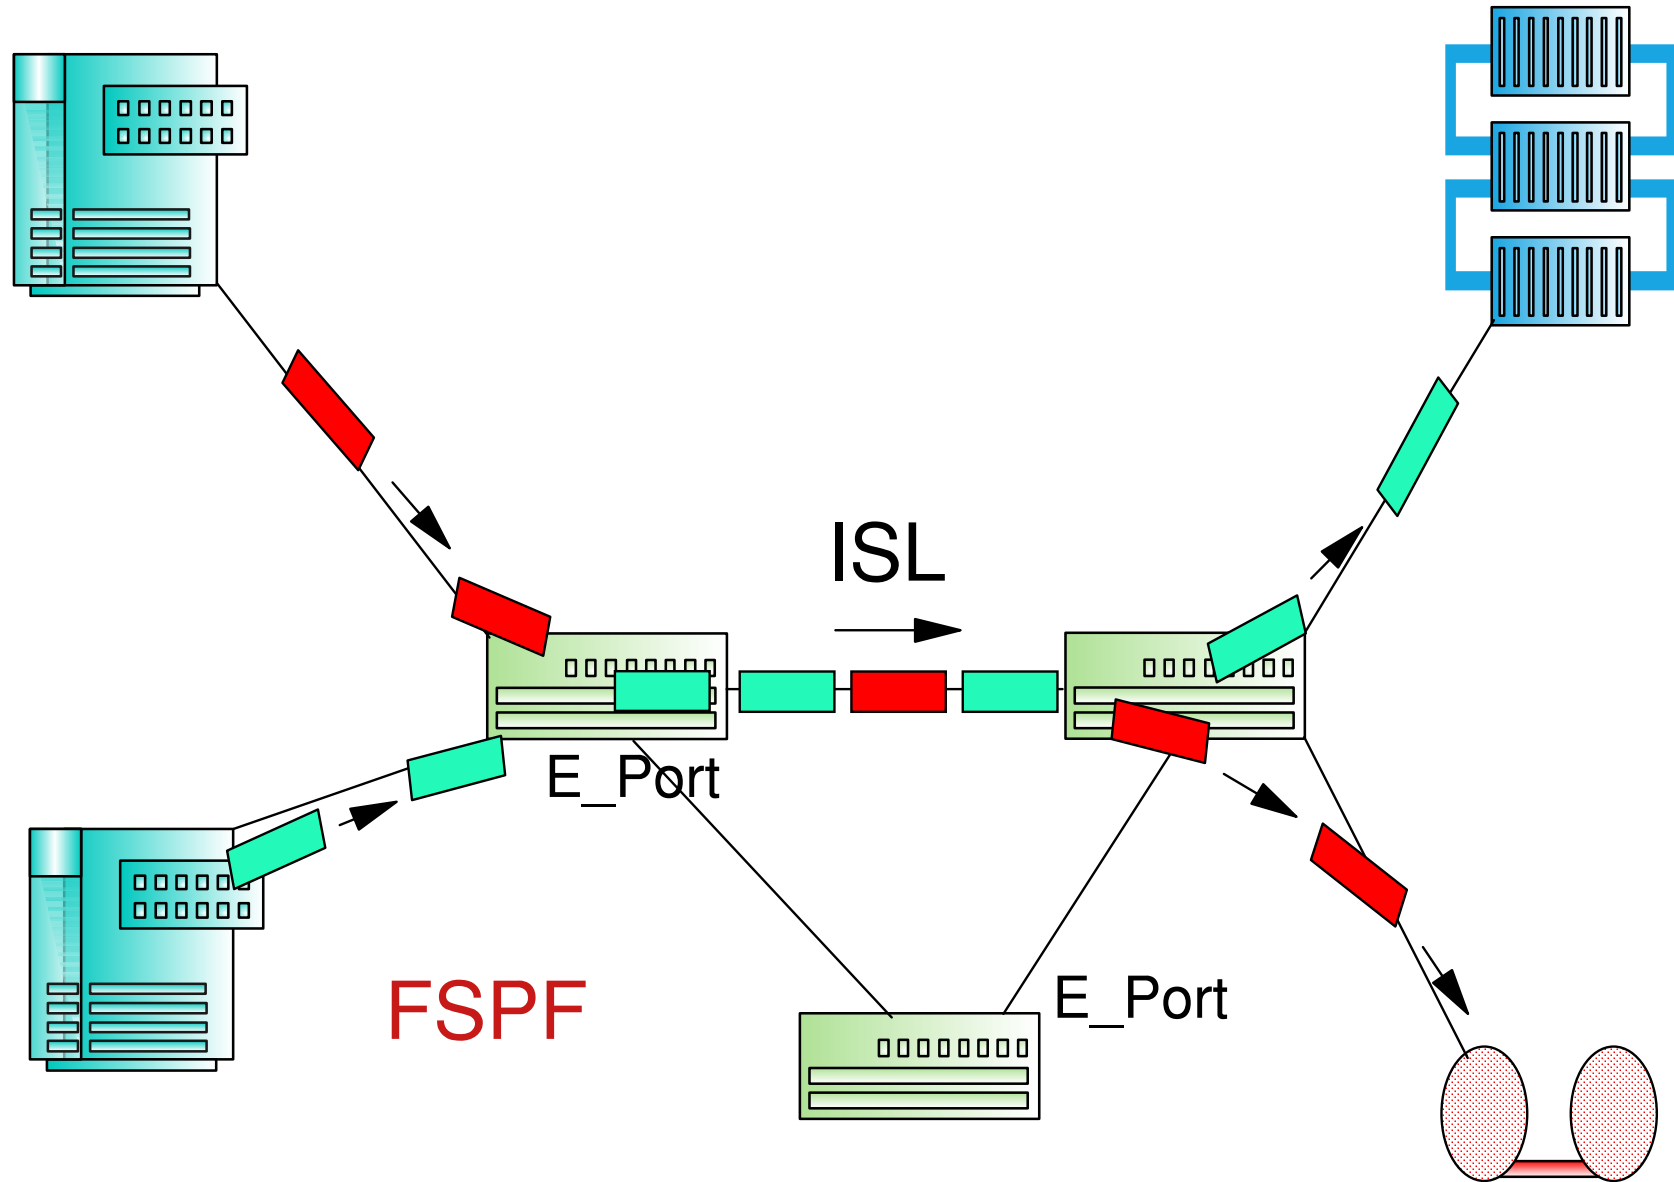
# Arbitrated Loop Devices - Switch Attached



# Switched Fabric - Dual Fabrics



# Frame Routing - Multi-Switch Environment



FSPF - Fabric Shortest Path First

ISL-InterSwitch Link



# Review: Ports

---

★ N\_Port (node)

★ F\_Port (fabric)

★ NL\_Port (node loop)

★ FL\_Port (fabric loop)

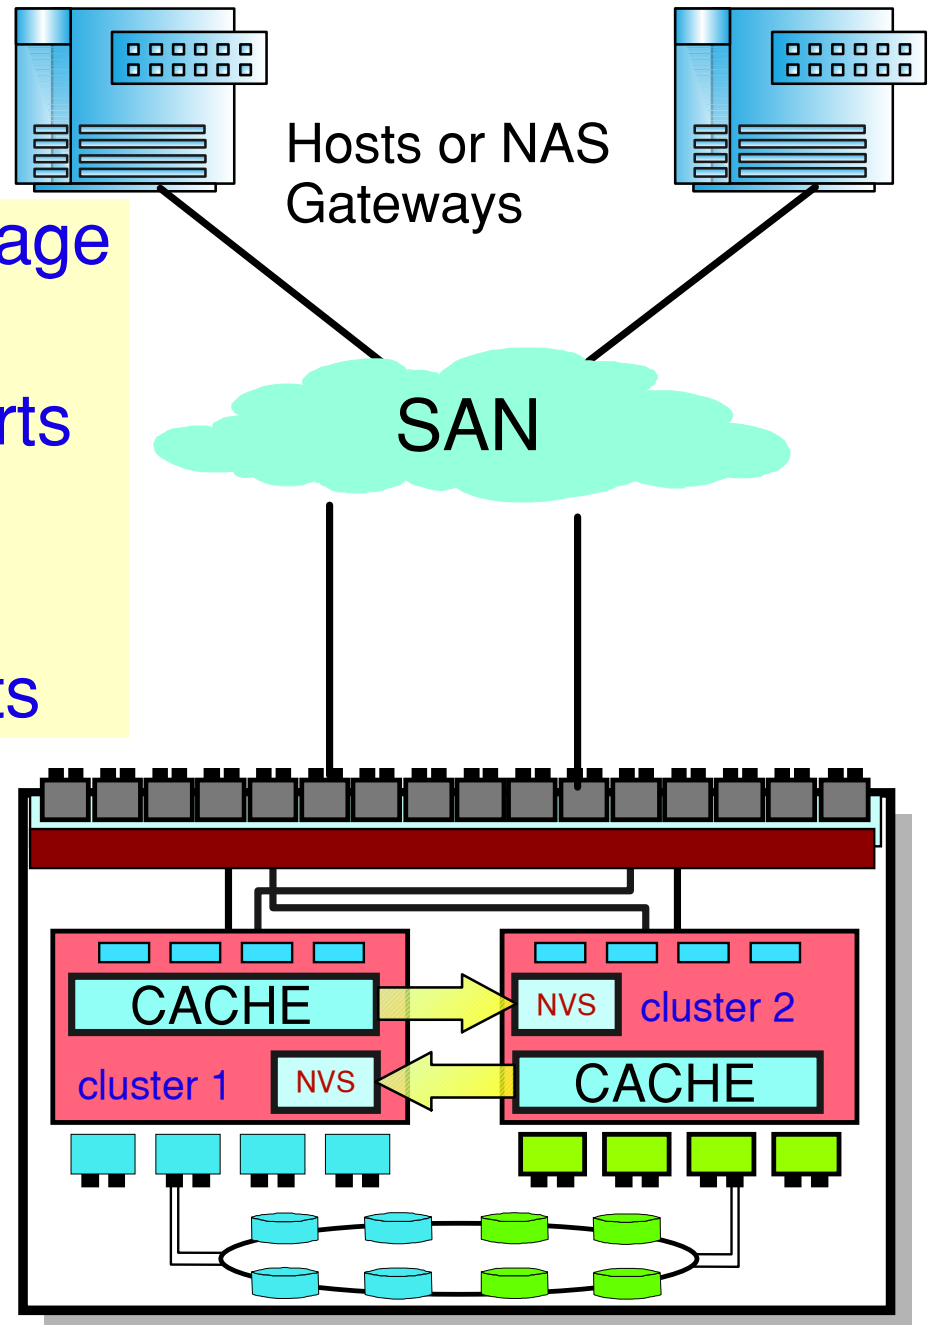
★ E\_Port (expansion)

★ G\_Port (generic)

★ U\_Port (universal)

# ESS: Host LUN Assignment Steps

1. Create Open Systems Storage
2. Modify Host Systems
3. Configure Host Adapter Ports
4. Configure Disk Groups
5. Add Volumes
6. Modify Volume Assignments



IBM TotalStorage  
ESS Specialist



# ESS: Define Host System Port

Specialist - wavec1.storage.sanjose.ibm.com

Enterprise Storage Server Specialist

## Modify Host Systems

**Host Attributes**

Nickname: SINESTRO\_QL3

Host Type: Microsoft Windows

Host Attachment: Fibre Channel attached

Hostname/IP Address: SINESTRO

World-Wide Port-Name: 210000E08B0EA350  
(select from list of known WWPNs)

Fibre-Channel Ports: All installed ports  
Rev 1, Card 2, Port A

**Host Systems List**

SanBoot_LP_W2003	Microsoft Windows
SanBoot_LP_W2K	Microsoft Windows
SanBoot_QL_W2003	Microsoft Windows
SanBoot_QL_W2K	Microsoft Windows
SEAWEED1	IBM RS/6000
SEAWEED2	IBM RS/6000
simonsez_fs2	IBM RS/6000
simonsez_fs3	IBM RS/6000
SINESTRO_QL1	Microsoft Windows
SINESTRO_QL2	Microsoft Windows
SINESTRO_QL3	Microsoft Windows
SINESTRO_QL4	Microsoft Windows
Stone1_QL1	Microsoft Windows
Stone1_QL2	Microsoft Windows
Stone3_QL1	Microsoft Windows
Stone3_QL2	Microsoft Windows
Stone4_LP1	Microsoft Windows
Stone4_LP2	Microsoft Windows
Stone5_QL1	Microsoft Windows

Buttons: Add >>, Modify >>, << Remove


Buttons: Perform Configuration Update, Cancel Configuration Update

Warning: Applet Window

Windows Host



SAN



LUNs for Windows Host



# ESS: Host System and Assigned Volumes

The screenshot displays the 'Enterprise Storage Server Specialist' interface. On the left is a navigation pane with options: TotalStorage Solutions, Status, Problem Notification, Communications, Storage Allocation, Users, and Licensed Internal Code. The main area is titled 'Open System Storage' and contains two tables.

**Host Systems Table:**

Nidname	Host Type	Attachment	WWPN	Hostname/IP Address
SINESTRO_QL1	Microsoft Windows	FC	210000E08B002818	SINESTRO
SINESTRO_QL2	Microsoft Windows	FC	210000E08B00EC0A	SINESTRO
SINESTRO_QL3	Microsoft Windows	FC	210000E08B0EA350	SINESTRO
SINESTRO_QL4	Microsoft Windows	FC	210000E08B0EE24E	SINESTRO
Stone1_QL1	Microsoft Windows	FC	210000E08B0A9CA8	Stone1

A yellow arrow points to the 'Host server port WWN' column of this table.

**Assigned Volumes Table:**

Volume	Vol Type	Size	Storage Type	Location	LSS	Shared
501-12028	Open System	01.0 GB	RAID-5 Array	Device Adapter Pair 3 Cluster 2, Loop A Array 1, Vol 001	LSS: 0x015	Yes
502-12028	Open System	01.0 GB	RAID-5 Array	Device Adapter Pair 3 Cluster 2, Loop A Array 1, Vol 002	LSS: 0x015	Yes

The 'Assigned Volumes' table is circled in green, and a note '(Total: 2 volumes)' is shown to its right. Below the tables are buttons for 'Modify Host Systems', 'Configure Host Adapter Ports', 'Configure Disk Groups', 'Add Volumes', and 'Modify Volume Assignments'. The status bar at the bottom indicates 'Java Applet Window'.

# ESS: Volume Assignments ID and LUNs

Specialist - wavec1

Enterprise Storage Server Specialist

## Modify Volume Assignments

Volume Assignments Refresh Status Print Table Perform Sort

no sort no sort no sort no sort no sort no sort second first

Volume	Location	LSS	Volume Type	Size	Storage Type	Host Port	Host Nicknames
501-12028	Device Adapter Pair 3 Cluster 2, Loop A Array 1, Vol 001	0x15	Open System	001.0 GB	RAID-5 Array	Fibre Channel ID 00, LUN 0000	SINESTRO_QL3
502-12028	Device Adapter Pair 3 Cluster 2, Loop A Array 1, Vol 002	0x15	Open System	001.0 GB	RAID-5 Array	Fibre Channel ID 00, LUN 0001	SINESTRO_QL3
501-12028	Device Adapter Pair 3 Cluster 2, Loop A Array 1, Vol 001	0x15	Open System	001.0 GB	RAID-5 Array	Fibre Channel ID 00, LUN 0000	SINESTRO_QL4

Action

- Assign selected volume(s) to target hosts
- Use same ID/Lun in source and target
- Unassign selected volume(s) from target hosts

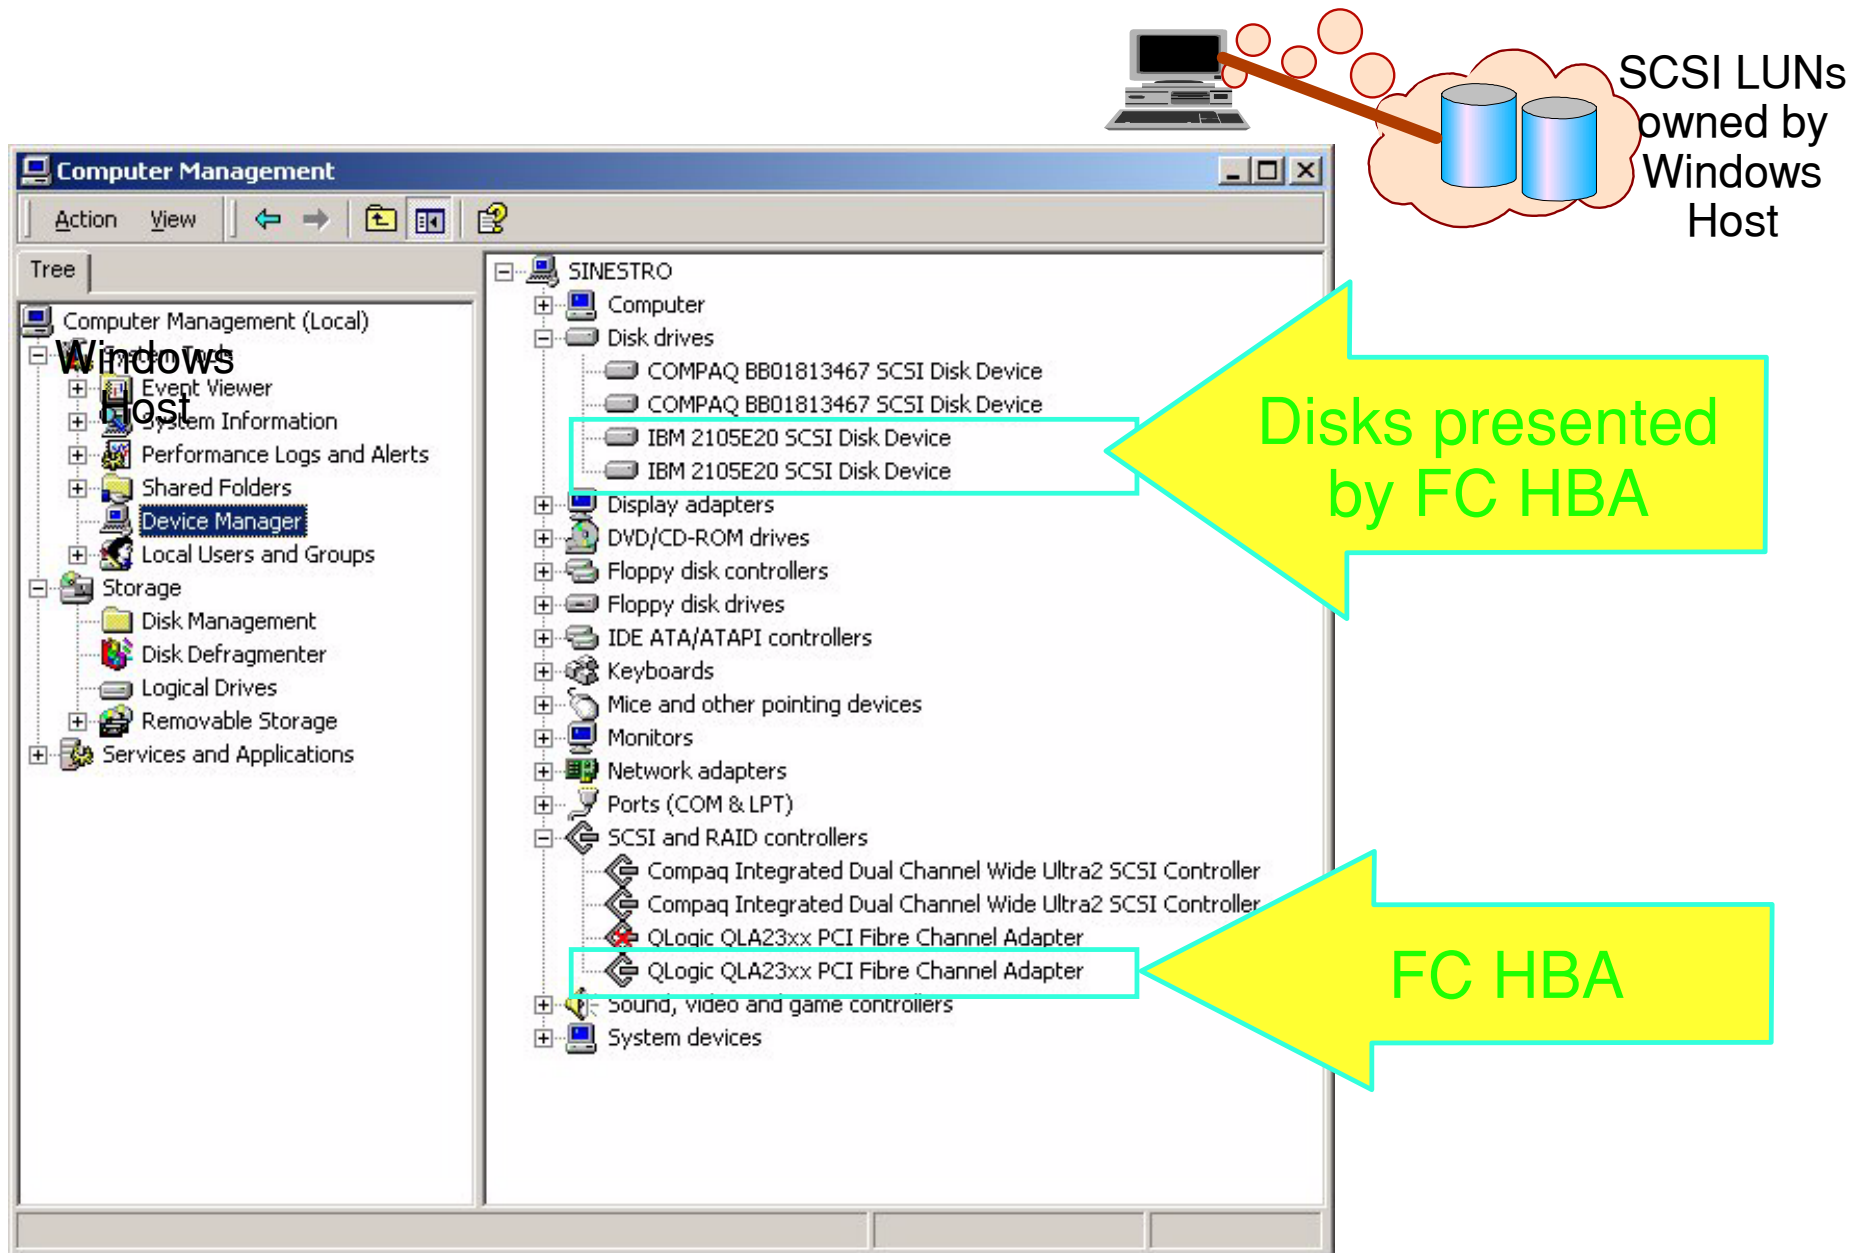
Perform Configuration Update

Target

Target ID and LUN numbers

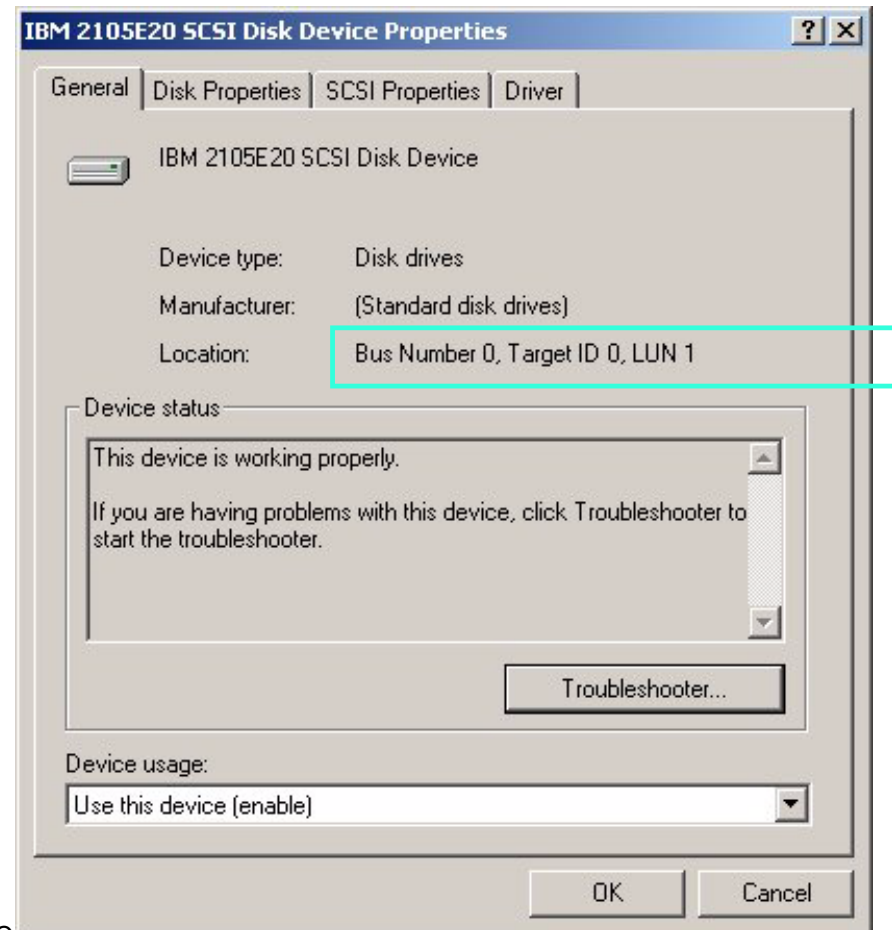
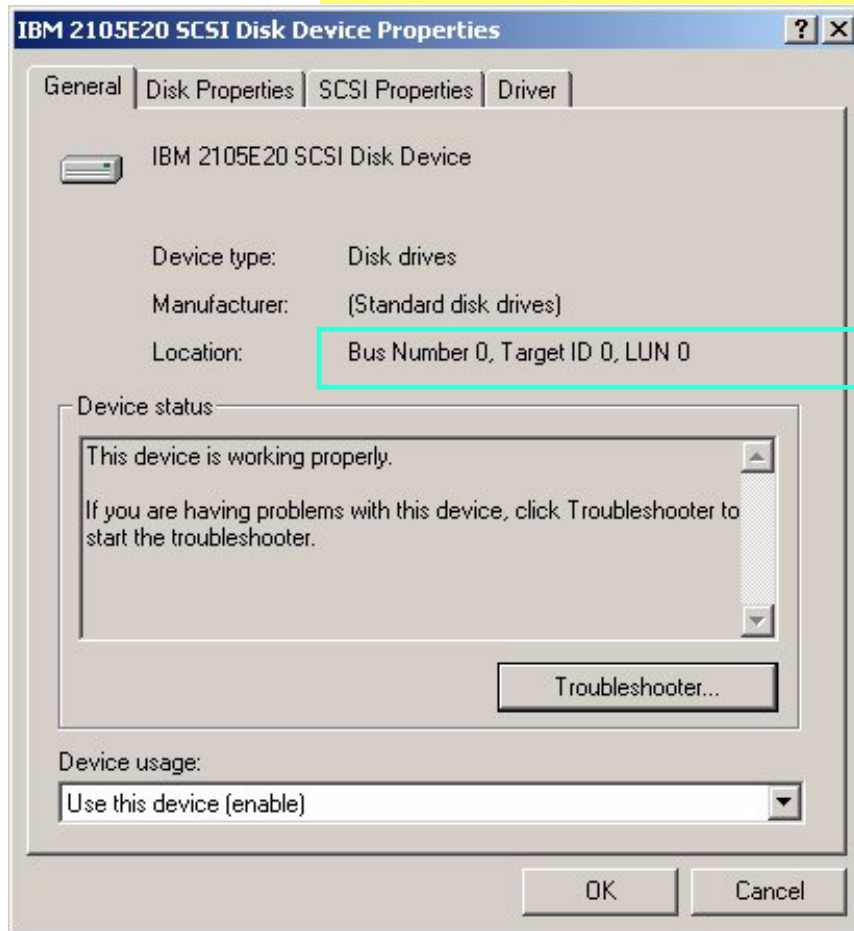
Java Applet Window

# Windows Host: SCSI Disk Devices



# SCSI Target ID and LUN Numbers

LUNs from the ESS  
presented to host by Fibre Channel HBA  
as SCSI Target ID and LUNs



# Windows Host: Disk Management

The screenshot displays the Windows Computer Management console. The left pane shows the 'Storage' tree with 'Disk Management' selected. The main pane shows a list of disks and their partitions. A 'Disk 2 Properties' dialog box is open, showing details for Disk 2.

Volume	Layout	Type	File Sys
	Partition	Basic	FAT
	Partition	Basic	FAT
MirrorCat (F:)	Mirror	Dynamic	NTFS
NT (E:)	Partition	Basic	NTFS
Win2k (C:)	Partition	Basic	NTFS

Disk	Layout	Type	File Sys
Disk 0	39 MB FA Healthy (	Win2k (C:) 16.84 GB NTFS Healthy (System)	78 M Una
Disk 1	39 MB I Health;	NT (E:) 4.00 GB NTFS Healthy (Active)	12.92 GB Unallocated
Disk 2	MirrorCat (F:) 949 MB NTFS Healthy		
Disk 3	MirrorCat (F:) 949 MB NTFS Healthy		


  

Volume	Capacity
MirrorCat (F:)	948 MB

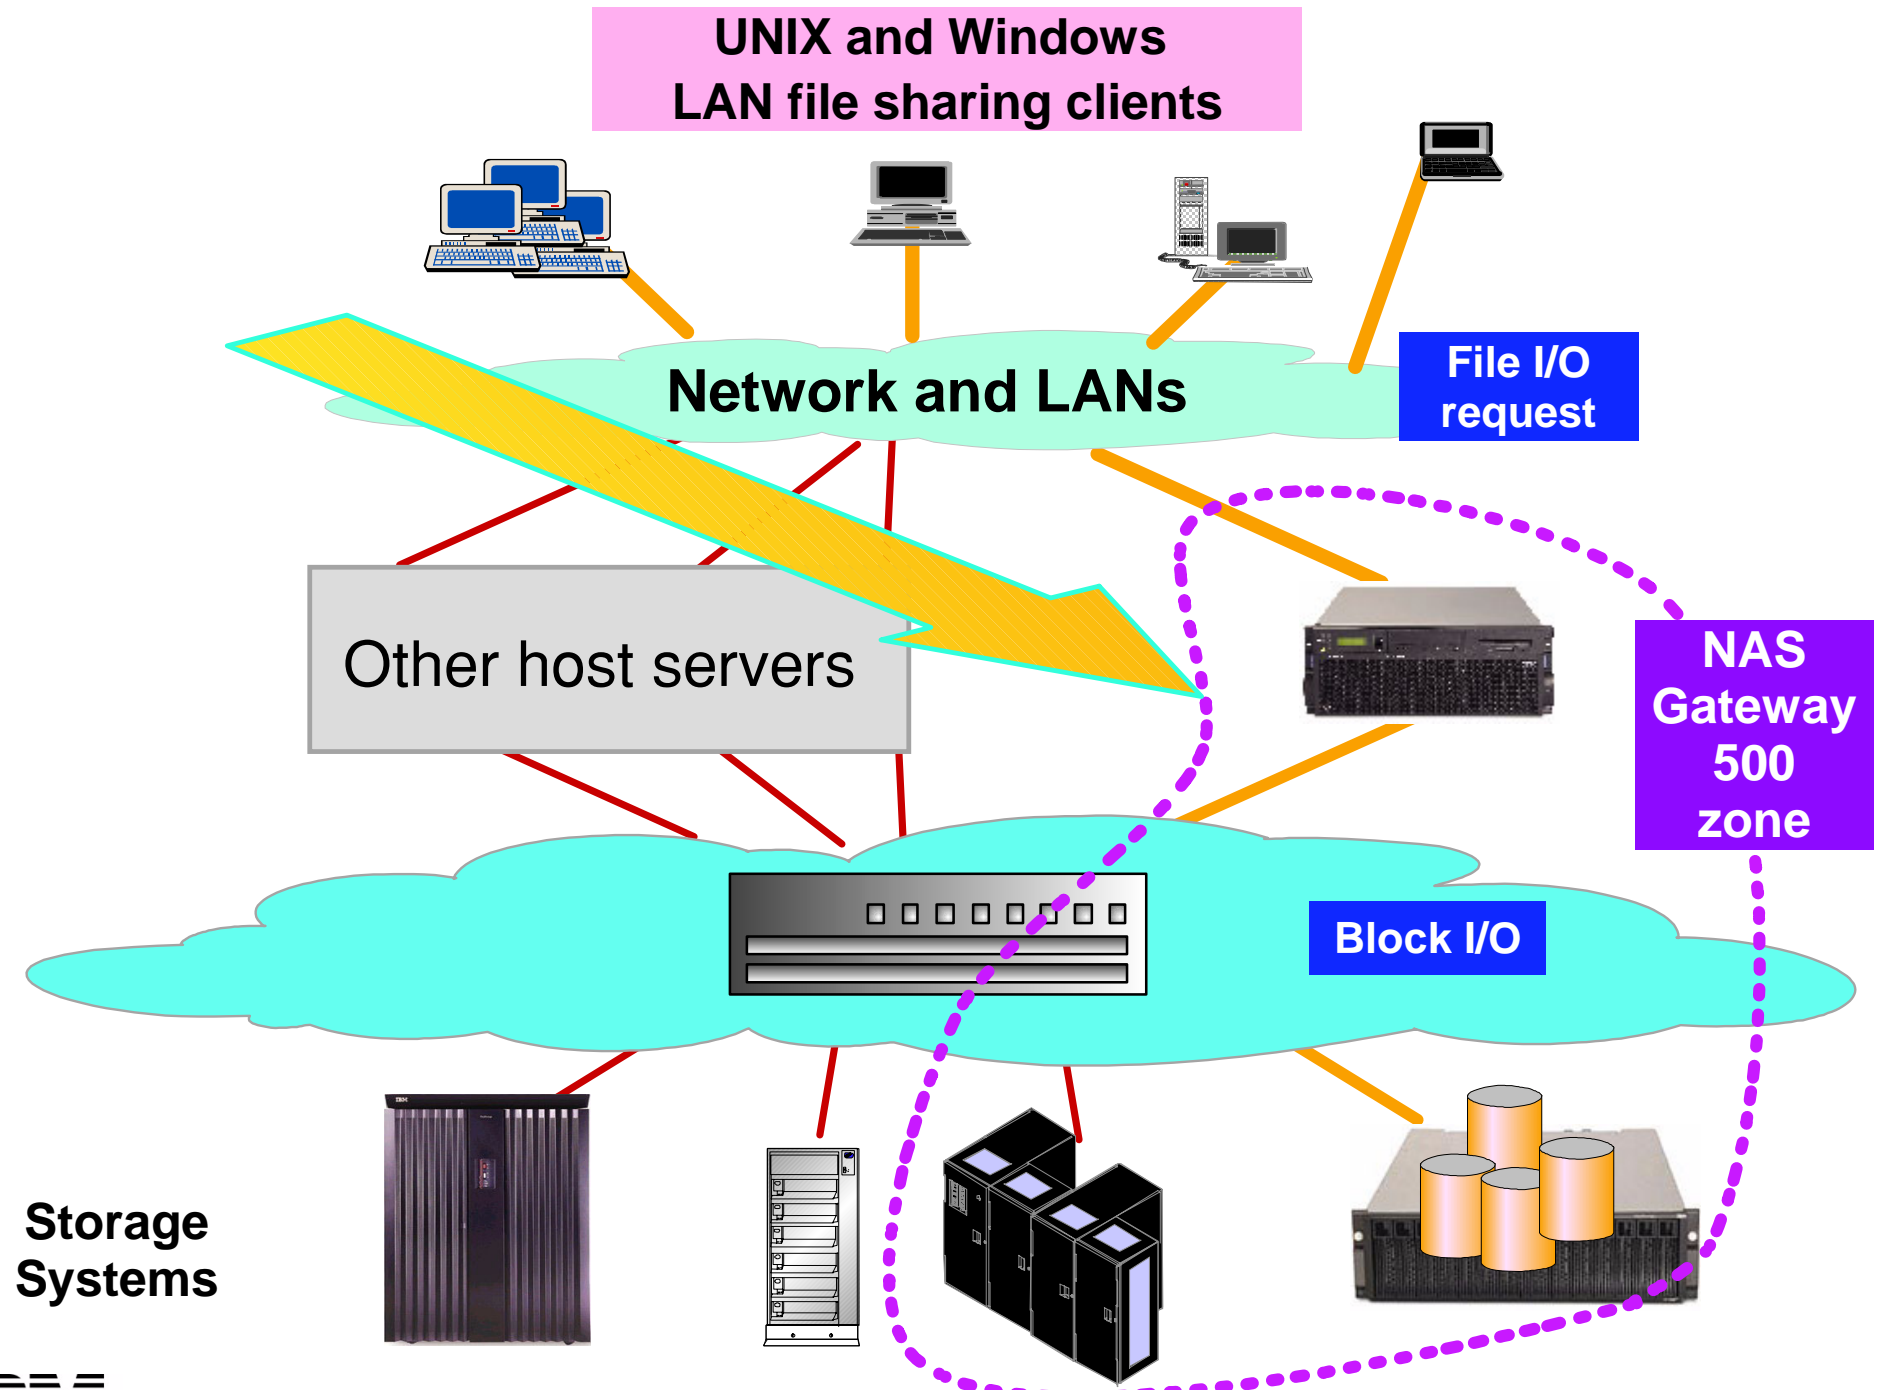
Device Type: **SCSI (Port:4, Target ID: 0, LUN:0)**

Hardware Vendor: IBM 2105E20 SCSI Disk Device

Adapter Name: QLogic QLA23xx PCI Fibre Channel Adapter

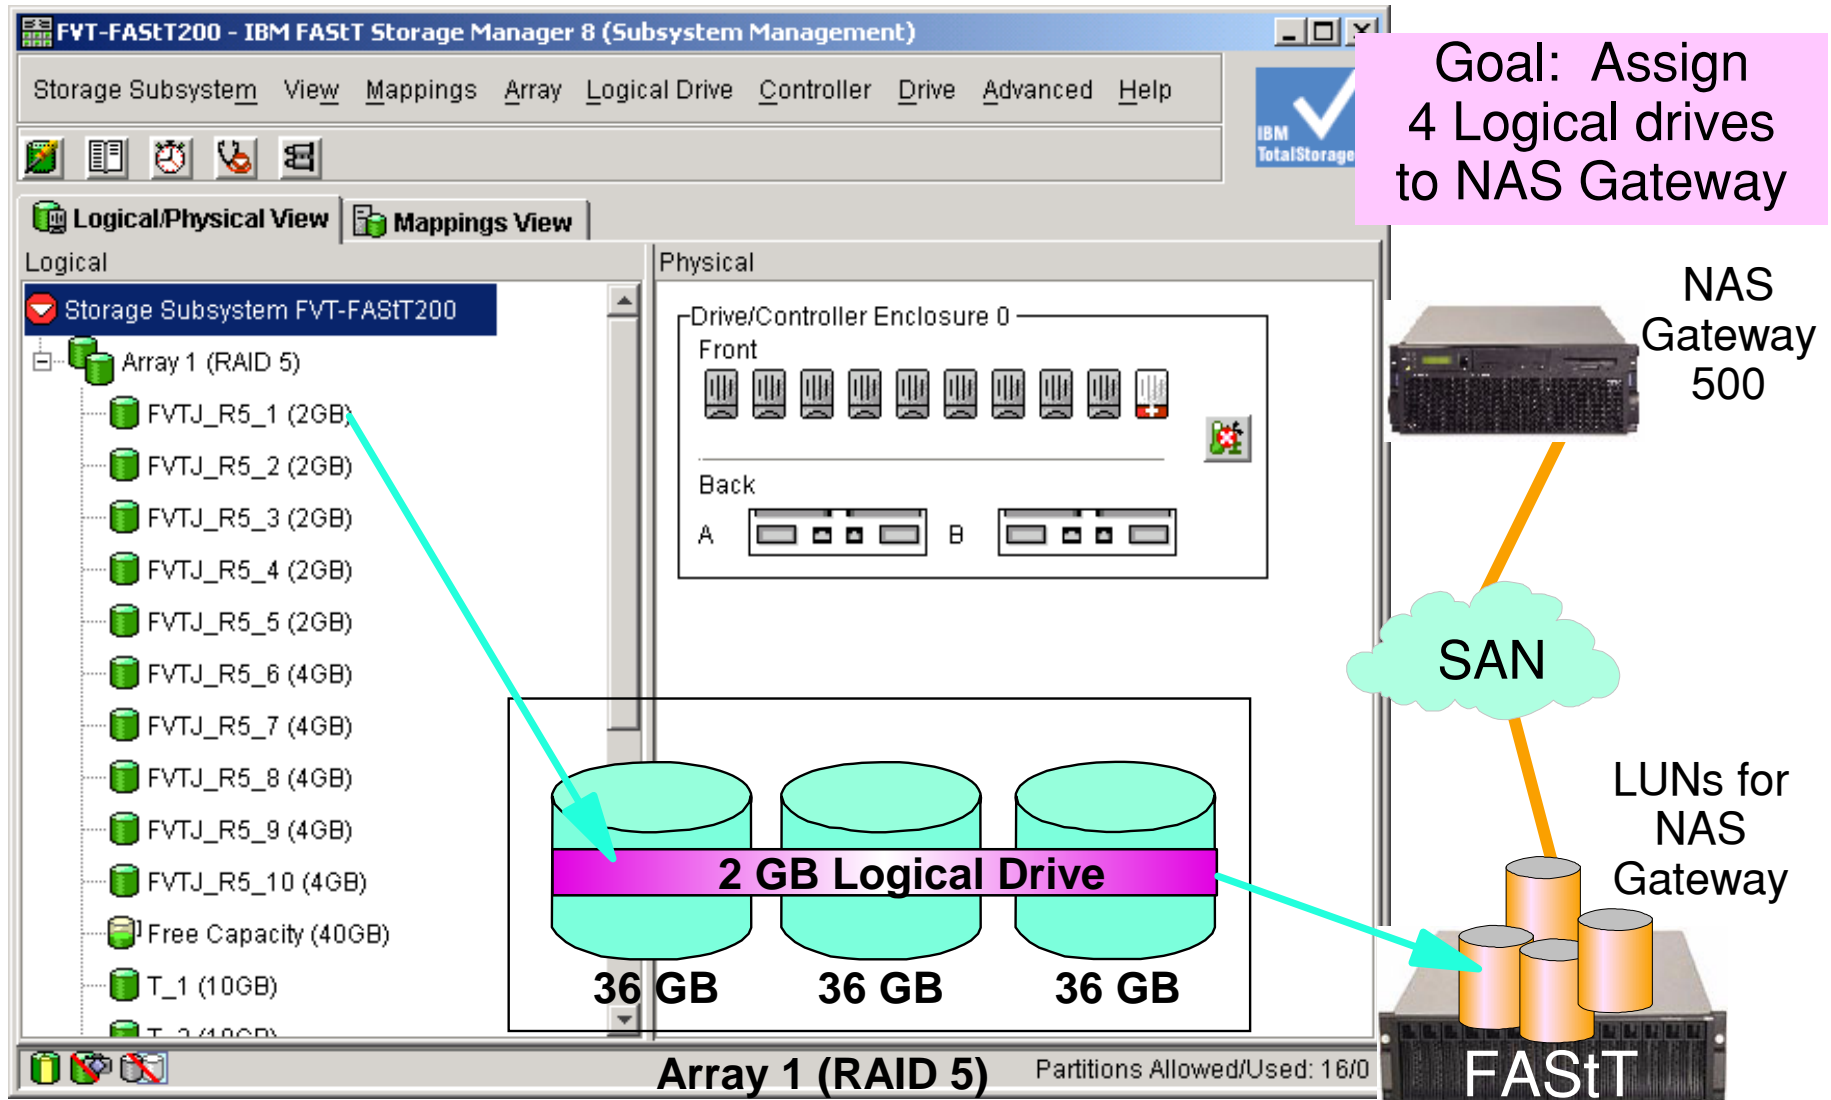


# NAS Gateway - Facilitate File Access to SAN



# FAStT Storage Manager

## Logical and Physical Views:



# FAStT SM: Define Host Group and Host

NAS Gateway  
is a host to  
storage systems

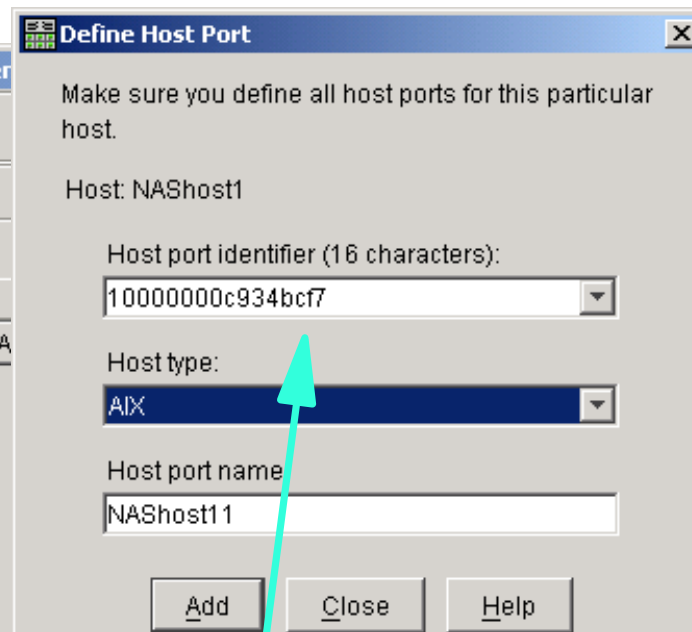
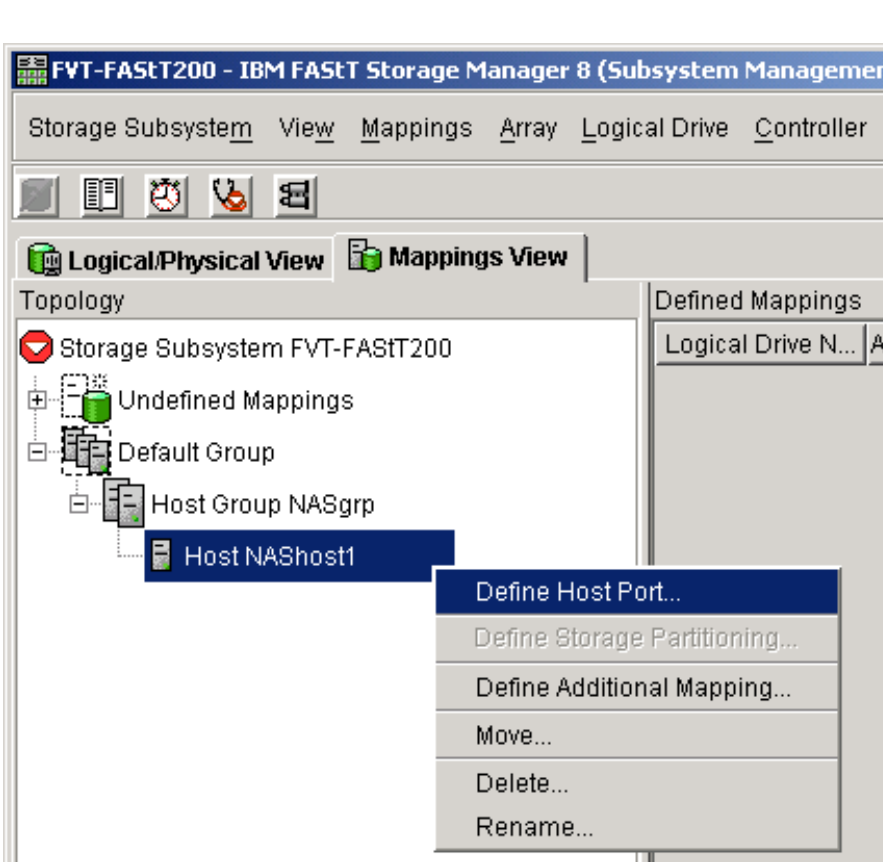
The screenshot displays the IBM FASTT Storage Manager 8 (Subsystem Management) interface. The main window shows the 'Logical/Physical View' and 'Mappings View' tabs. The 'Topology' pane on the left shows a tree structure with 'Storage Subsystem FVT-FAST200' expanded to show 'Undefined Mappings' and 'Default Group'. A context menu is open over 'Default Group', with 'Define Host Group...' selected. A red arrow points from this menu item to the 'Define Host Group' dialog box. The dialog box contains the text: 'After closing this dialog, make sure you define the hosts associated with each host group. Highlight the newly-defined host group and then select Mappings>>Define>>Host.' Below this text is a 'Host group name:' label and a text input field containing 'NASgrp'. At the bottom of the dialog are 'Add', 'Close', and 'Help' buttons.

A second red arrow points from the 'Define Host...' option in the context menu (which is highlighted in a red box) to the 'Define Host' dialog box. The dialog box contains the text: 'After closing this dialog, make sure you define all host ports associated with each host. Highlight the newly-defined host and then select Mappings>>Define>>Host Port.' Below this text is a 'Host name:' label and a text input field containing 'NAShost1'. At the bottom of the dialog are 'Add', 'Close', and 'Help' buttons.

A horizontal dashed red line is drawn across the middle of the image, separating the two dialog boxes.




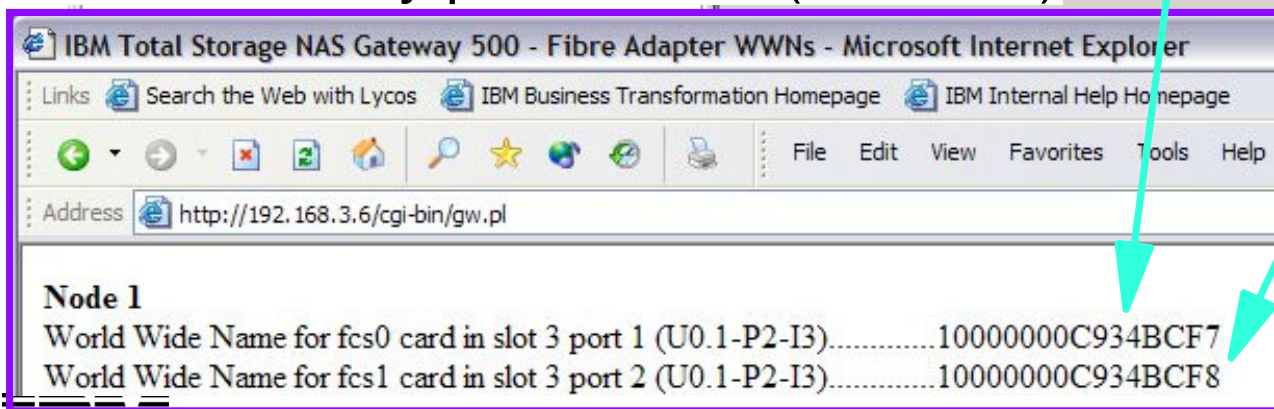
# FAStT SM: Define Host Ports



Port1

Port2

## NAS Gateway port WWNs (WWPNs)



# FASTT SM: Storage Partitioning

Storage Partitioning ↔ LUN Masking

Assign LUNs to host group NASgrp

Storage Partitioning Wizard - Select Logical Drives/LUNs

Select the logical drives and assign the logical unit numbers (LUNs) that the host group will use to access the logical drives in this partition.

Host group: NASgrp  
Hosts: NAShost1

Logical Drive/LUN assignment

Select logical drive: Assign LUN: (0 to 31)

Select logical drive:	Assign LUN: (0 to 31)
Access	4
FVTJ_R5_10 4.000GB	
FVTJ_R5_3 2.000GB	
FVTJ_R5_4 2.000GB	
FVTJ_R5_5 2.000GB	
FVTJ_R5_8 4.000GB	
FVTJ_R5_9 4.000GB	

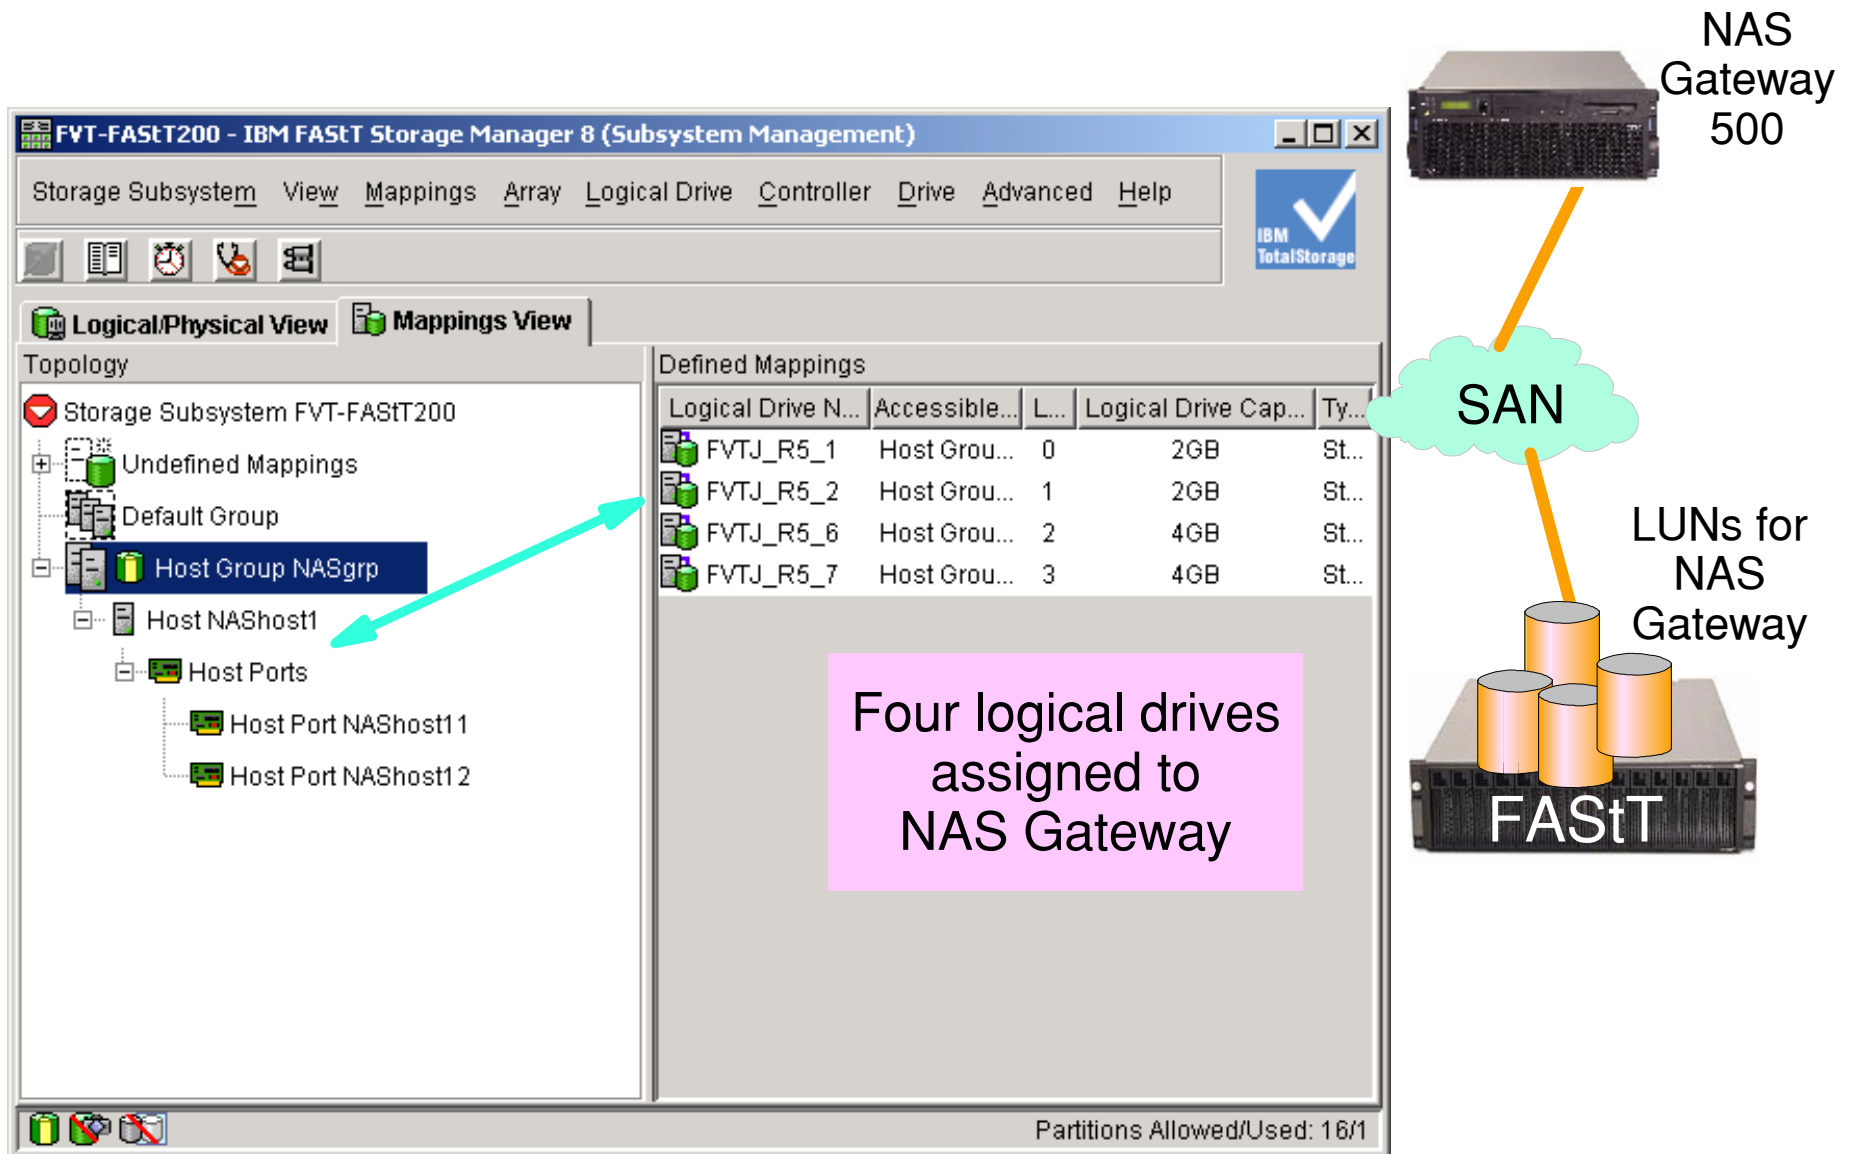
Add -> <- Remove

Logical Drives to include in partition:

Logical Drive	LUN
FVTJ_R5_1	0
FVTJ_R5_2	1
FVTJ_R5_6	2
FVTJ_R5_7	3

< Back Finish Cancel Help

# FAStT: LUNs Mapped to Host Group



# Volume Wizard - Create NAS Volumes

The screenshot shows the IBM FASTT Storage Manager 8 interface. The 'Logical/Physical View' pane on the left displays a tree structure under 'Storage Subsystem FVT-FASTT200', including 'Undefined Mappings', 'Default Group', 'Host Group NASgrp', 'Host NASHost1', and 'Host Ports'. The 'Defined Mappings' table on the right lists four mappings:

Logical Drive N...	Accessible...	L...	Logical Drive Cap...	Ty...
FVTJ_R5_1	Host Grou...	0	2GB	St...
FVTJ_R5_2	Host Grou...	1	2GB	St...
FVTJ_R5_6	Host Grou...	2	4GB	St...
FVTJ_R5_7	Host Grou...	3	4GB	St...

Red arrows indicate the flow of information: one points from the 'Host Group NASgrp' in the topology to the 'Defined Mappings' table, and another points from the 'Host Group NASgrp' to the 'Volume selection' dialog box.

LUNs in FASTT assigned to NAS Gateway

The screenshot shows the 'Volume selection' dialog box in the IBM NAS Gateway 500. It prompts the user to 'Select one or more physical volumes used to create the NAS volume:'. The dialog features two list boxes: 'Physical volumes' containing 'hdisk2', 'hdisk3', 'hdisk4', and 'hdisk5'; and 'Selections', which is currently empty. Between the list boxes are 'Add' and 'Remove' buttons. At the bottom, there are 'Help', 'Next', and 'Exit' buttons.



# Summary

---

- Fibre Channel topologies included:
  - Point-to-point
  - Arbitrate loop
  - Switch fabric
- Fibre Channel ULP enables multiple protocols to share a common physical transport.
- A fabric comprises one or more interconnected switches that implement Fibre Channel services to link ports and route data.
- Fibre Channel is the enabling technology for SANs
- Fibre Channel HBAs provides device discovery and shields host OS from Fibre Channel SAN awareness

# References

---

- Course SN700 - Introduction to Storage Networking
- Course SN710 - Planning and Implementing a SAN
- Course SN820 - SAN Volume Controller Planning and Implementation
- [www.ibm.com/services/learning](http://www.ibm.com/services/learning) - Training information
- SG24-6419 Designing and Optimizing an IBM SAN
- [www.ibm.com/san](http://www.ibm.com/san) IBM Storage Area Network
- [www.fibrechannel.org](http://www.fibrechannel.org) Fibre Channel Industry Association
- [www.t11.org](http://www.t11.org) Device Interfaces and Drafts of FC Standards
- [www.snia.org](http://www.snia.org) Storage Network Industry Association

