

SAN and FICON Long Distance Connectivity

David Lytle, BCAF
Global Solutions Specialist
Brocade Communications
[*dlytle@brocade.com*](mailto:dlytle@brocade.com)

Uli Schlegel
Director, Global Business Development
Datacenter Solutions
ADVA Optical Networking
[*uschlegel@advaoptical.com*](mailto:uschlegel@advaoptical.com)

Session
16001





Agenda



- Who are Brocade and ADVA
- Fundamentals of SAN and FICON FC Long Distance Connectivity
 - Mainframe Channel Cards and Long Distance Connectivity
 - MAN / WAN / SONET / SDH
 - Direct-attached storage
 - Switch-attached storage but FICON non-cascaded
 - Switch-attached storage and FICON cascaded
- Brocade and ADVA Products
- WDM options and benefits





Brocade Communications Inc., Today

- Founded in 1995, currently has about 4,100 employees worldwide
- Serves a wide range of industries and customers in more than 160 countries
- An industry leader in providing reliable, high-performance network solutions
- Brocade provides our users:
 - Unmatched simplicity to overcome today's complexity
 - Non-stop networking to maximize business uptime
 - Optimized applications to increase business agility and gain a competitive advantage
 - Investment protection to provide a smooth transition to new technologies while leveraging existing infrastructure
- 90 percent of the Global 1000 rely on Brocade solutions
- 38.7 million SAN switch ports shipped, 200,000+ SANs in production, 50,000+ Brocade directors installed worldwide

ADVA Optical Networking Today

Mission

Our MISSION is to be the trusted partner for innovative connectivity solutions that ADVANCE next-generation networks for cloud and mobile services.



Key Facts

Our NUMBERS
>1400 employees
€311* million revenue
20 years of innovation

Our CUSTOMERS
Hundreds of carriers
Thousands of enterprises

Our QUALITY
TL 9000, ISO 14001
Award-winning supply chain

*2013

We bring differentiation, quality and ease-of-use to next-generation networks



FUNDAMENTALS OF FIBRE CHANNEL LONG DISTANCE CONNECTIVITY



FICON: Analysis of Mainframe Channel Cards

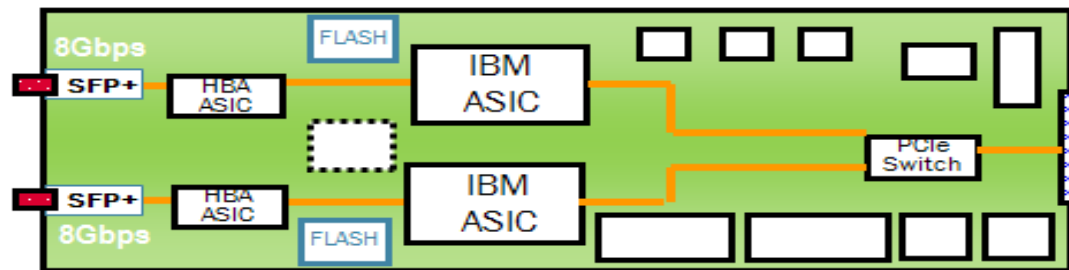
Standard PCIe card



FICON Express8S – 2 ports
800MBps+800MBps=1600MBps

FICON Express8S

- zEC12, zBC12, z196, z114
- 2, 4 or 8 Gbps link rate
- FICON never creates full frames
- Buffer Credits:
 - 2Gbps – 107 BCs per port
 - 4Gbps – 200 BCs per port
 - 8Gbps – 40 BCs per port
 - Out to 5km assuming 1K frames
 - **No “Long Distance” capability**



- For FICON, zHPF, and FCP environments
 - CHPID types: FC and FCP (2 PCHIDs/CHPIDs)
- Auto-negotiates to 2, 4, or 8Gbps
- Increased performance versus FICON Express8
- 10KM LX - 9 micron SM fiber
 - Unrepeated distance - 10 km which 6.2 miles
- SX - 50 or 62.5 micron multimode fiber
 - Distance variable with link data rate and fiber type
- 2 channels of LX or SX (no mix)
- Small form factor pluggable (SFP) optics
 - Concurrent repair/replace action for each SFP



Metropolitan-Area/Regional-Area Networks

- A MAN or RAN covers a North American metropolitan area, or a small to medium-sized country in Europe or Asia
- Provides an optical ring/mesh topologies with adequate back-up and protection
- Main technologies:
 - SONET/SDH
 - OTN
 - Gigabit
 - 10-Gigabit Ethernet
 - CWDM, DWDM
- Several LANs could be connected to a single MAN
- The graphic shows how a ring topology might be beneficial.





Wide-Area Networks (WAN)

- Long haul intra-city and intra-country connections
- Typically government-regulated or in the public network environment
 - WANS originated in telephony
- Main technologies: SONET/SDH, OTN, WDM
 - Voice circuits vs. data packets



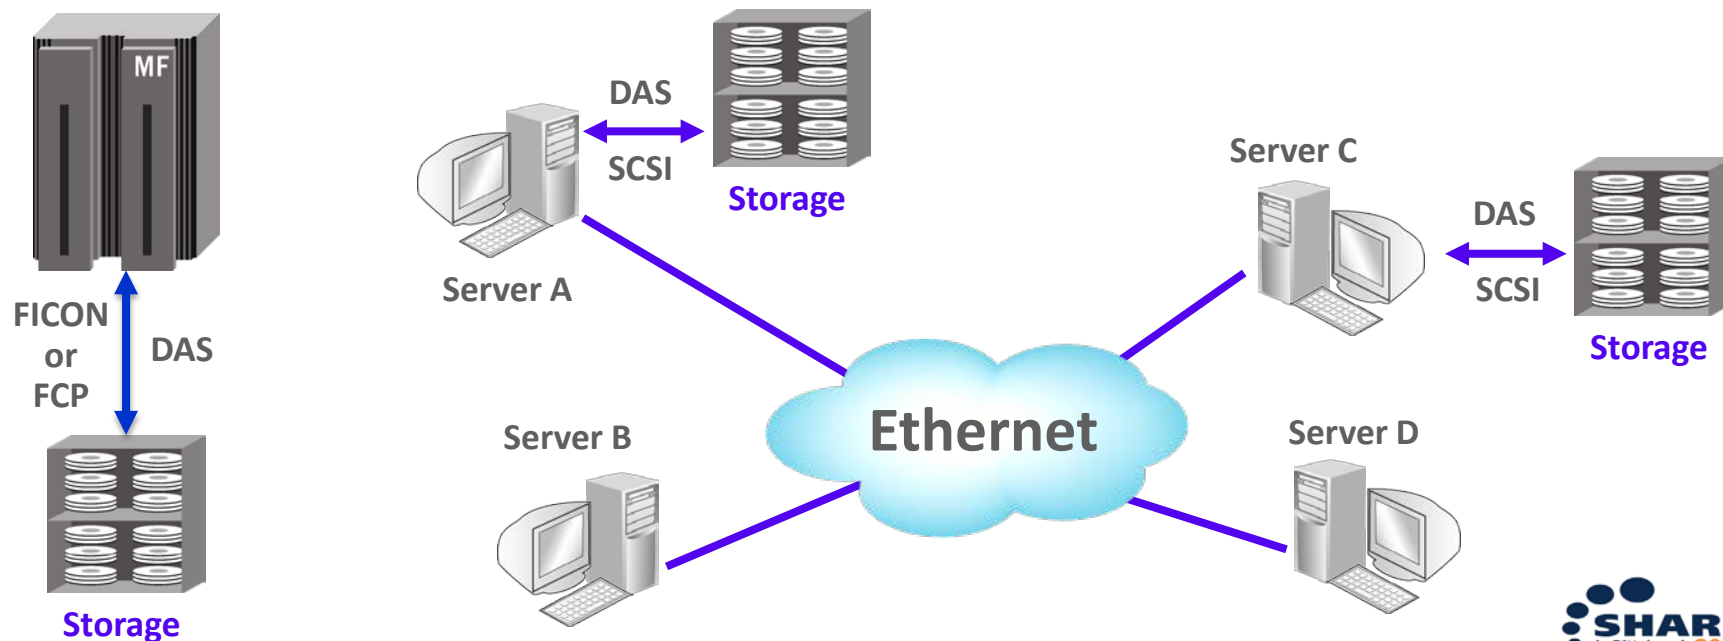
SONET/SDH - OTN

- SONET (Synchronous Optical Networks) is the Time Division Multiplexing (TDM) optical network standard for North America (called SDH in the rest of the world)
 - It is the de-facto standard for fiber backhaul networks
 - WDM is usually the underlying transport structure
 - OC-1 (with a frame format of SONET STS-1) using optical fiber has:
 - A transmission speed of up to 51.84 Mbps
 - OC-12 (with a frame format of SONET STM-4) using optical fiber has:
 - A transmission speed of up to 622.08 Mbps
 - OC ranges from OC-1 up to OC-768 which has:
 - A transmission speed of up to 39.813 Gbps
- OTN (Optical Transport Network) is the new standard for optical data transmission
 - Is the evolutionary successor of SONET/SDH technology
 - OTU-1 with a transmission speed of 2.67Gbps
 - OTU-2 with a transmission speed of 10.7Gbps
 - OTU-3 with a transmission speed of 43.0Gbps
 - OTU-2 with a transmission speed of 111.8Gbps




Storage: Direct attached

- Data storage connected directly to a server or workshop through Host Bus Adapter (HBA), there is no network between storage and host servers





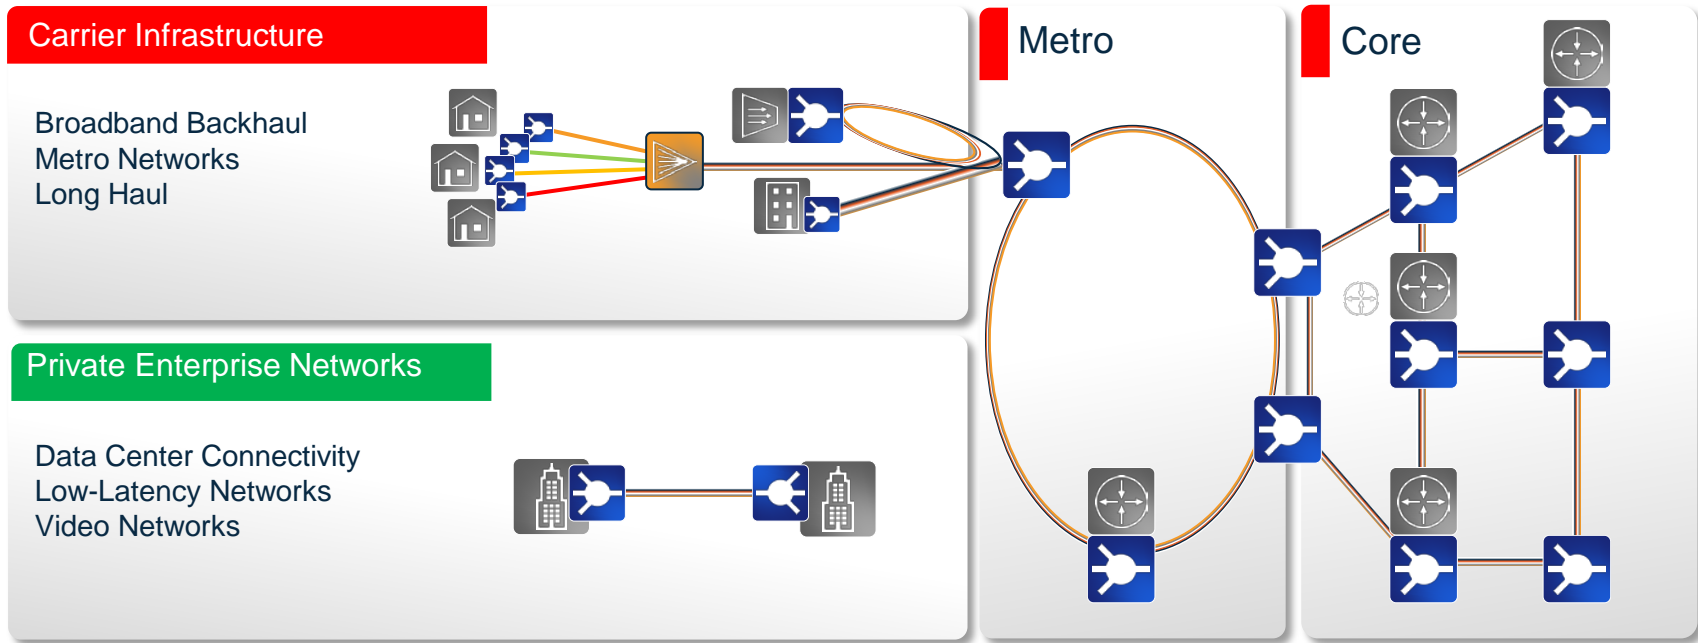
Storage: Direct attached <10km




- Scalability ? ☹️
- Management ? ☹️
- Link Utilization ? ☹️
- Long Distance ? ☹️

— Coupling
— FICON/FC
— PPRC

WDM transport – The big picture



Wavelength Division Multiplexing (WDM)



- Multiple channels of information carried over the same fibre, each using an individual wavelength
- Attractive multiplexing technique
 - High aggregate bit rate without high speed electronics or modulation
 - Low dispersion penalty for aggregate bit rate
 - Very useful for upgrades to installed fibres
 - Commonly used for distances up to 3000km
- Loss, crosstalk and non-linear effects are potential problems

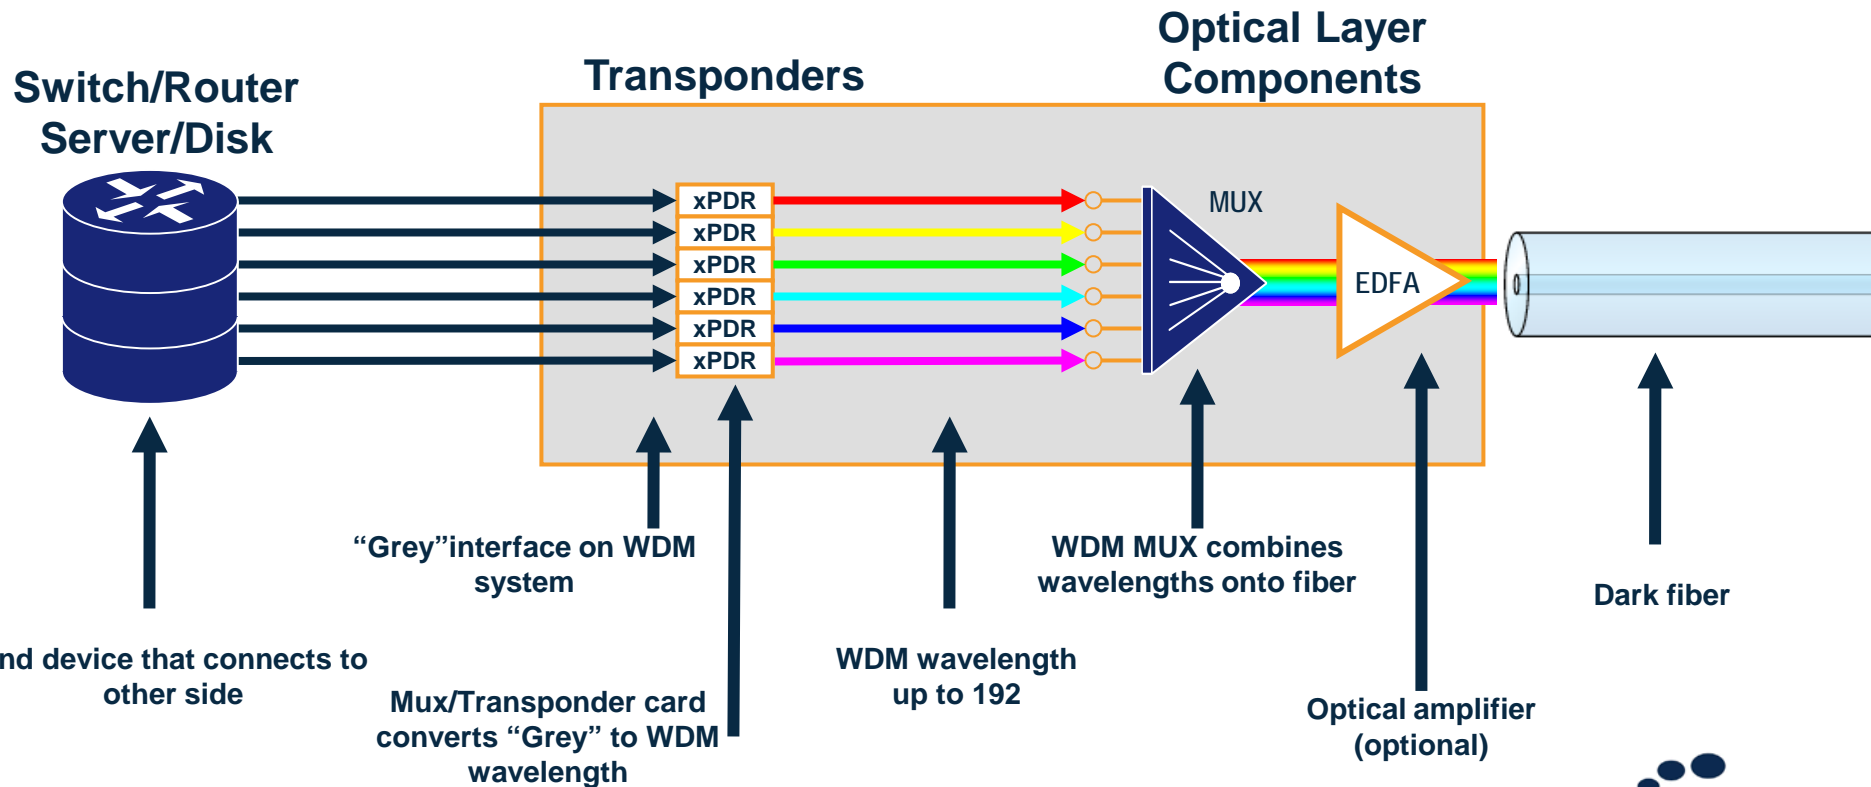


Wavelength Division Multiplexing


- Two WDM flavors available
 - CWDM (Coarse WDM)
 - Up to 16 optical lambdas max
 - Cheaper than DWDM
 - No amplification, optical switching (80km max)
 - DWDM (Dense WDM)
 - Up to 192 optical lambdas
 - Up to thousands of kilometers
 - Lambda switching, 100G, advanced features
 - DWDM is qualified for System Z only



Basic WDM scheme



Storage: Direct attached with WDM <10km

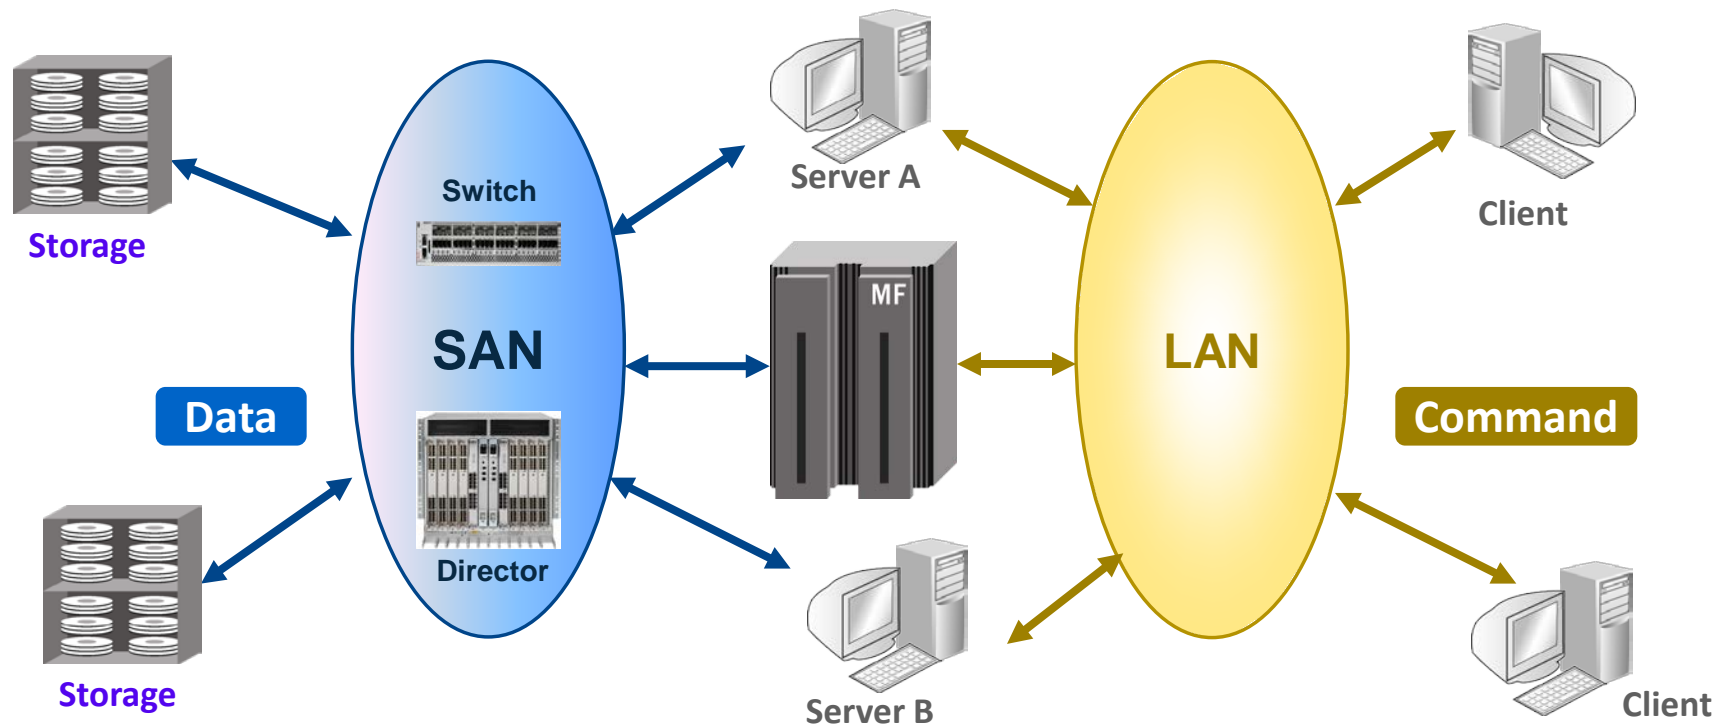


- Scalability ? ☺
- Management ? ☹
- Link Utilization ? ☹
- Long Distance ? ☹


— Coupling
— FICON/FC
— PPRC

Storage: Switch attached

- A dedicated data storage network which can be accessed by multiple servers



Storage: Switched but non-cascaded <10km



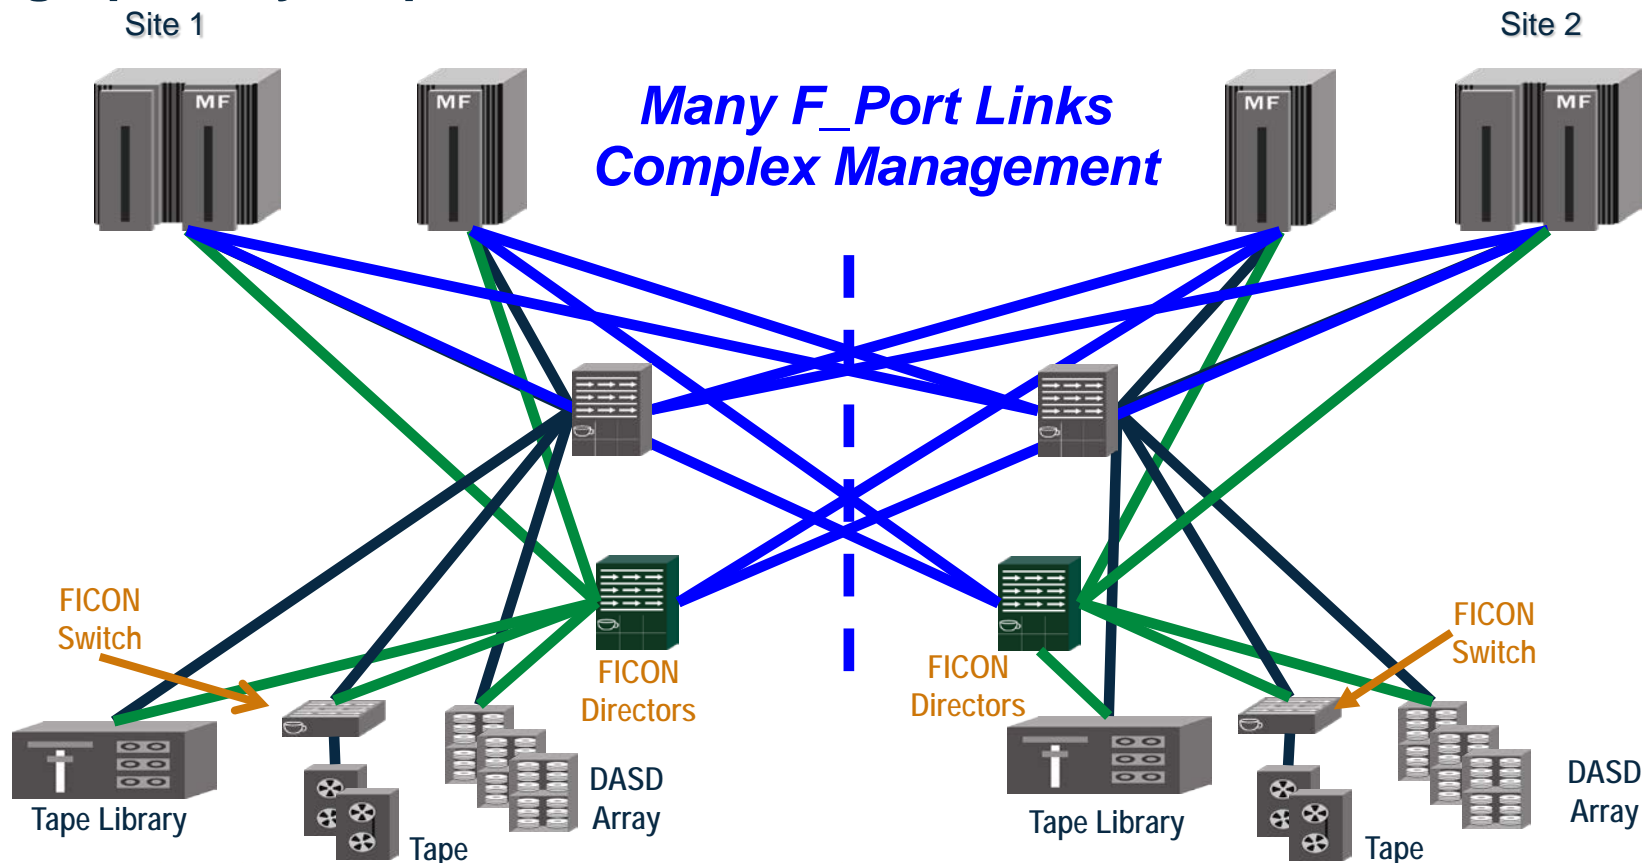
- Scalability ? ☹️
- Management ? ☹️
- Link Utilization ? ☹️
- Long Distance ? ☹️

— Coupling
— FICON/FC
— PPRC




Non-Cascaded FICON Environment

Geographically Dispersed Data Center Connections





Storage: Switched non-cascaded with WDM



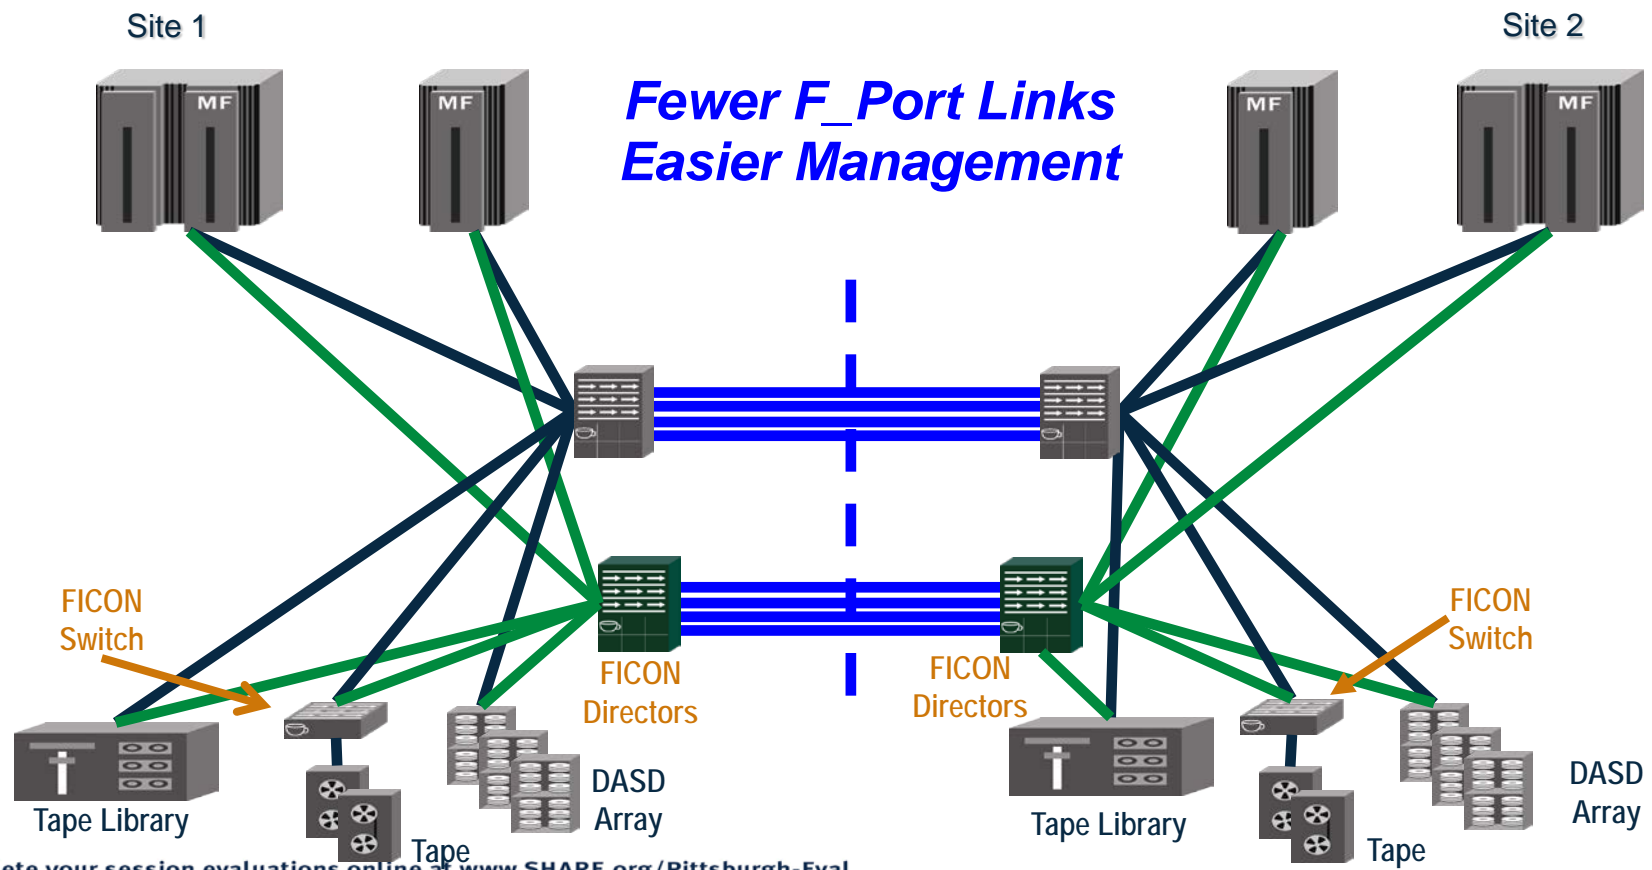
- Scalability ? ☺
- Management ? ☺
- Link Utilization ? ☺
- Long Distance ? ☹

— Coupling
 — FICON/FC
 — PPRC



Cascaded FICON Environment


Geographically Dispersed Data Center Connections



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



Storage: Switched – cascaded >10km




- Scalability ? ☹️
- Management ? ☹️
- Link Utilization ? 😊
- Long Distance ? ☹️

- Coupling
- FICON/FC
- PPRC
- ISL's



Storage: Switched – cascaded with WDM >10km




- Scalability ? ☺
- Management ? ☺
- Link Utilization ? ☺
- Long Distance? ☺

- Coupling
- FICON/FC
- PPRC
- ISL's

Brocade Gen 5 6510 Switch

Front view



- 48x16/10/8/4/2 Gbps Fibre Channel ports
- System Ethernet port (RJ45) for out-of-band management
- System RS232 console port (RJ45)
- USB port for firmware upgrades and system log downloads
- Switch ID pull-out tab containing serial number and MAC address
- Small footprint (1U and less than 18 inches deep) for flexible deployments

Brocade Gen 5 8510 Director

Up to 384 ports with 1:1 subscription at 8 Gbps

- FICON: 256 ports 1.0:1 at 16 Gbps (z/OS limitation)
- FCP: 384 ports 1.0:1 at 16 Gbps using local switching

Aggregate Bandwidth:

- 4 Tbps per chassis for central switching
 - $256 \text{ user ports} * 16 \text{ Gb} = 4 \text{ Tbps/chassis}$

512 Gbps data rate bandwidth per connectivity s

12-slot card cage:

- 8 port and/or special purpose blades (e.g. FCIP)
- 2 control processor blades
- 2 core routing blades

Buffer Credits

- 8,000 per 16-port group on 32-port blades

16Gbps blades are interchangeable between the 8510-4 and 8510-8



Brocade-branded Optics



8510-4

8510-8

ADVA FSP 3000


- State of the art 1G-100G transport platform
- Up to 192 optical channel/19.2Tbps per fiber pair
- Qualified with all SAN/Storage vendors and applications
- Physical Layer inspection via optical line monitoring and build in OTDR
- Dedicated modules especially for datacenter connectivity
- Ultra low latency (down to 5ns) per link
- Less than 1W per Gbit transport
- Full support of all Brocade features over distance



Brocade and ADVA for System z


The best of two worlds

- Long distance trunking over WDM System
- 1000+ 16G FC; 2000 8G FC ports per WDM System
- Physical Layer inspection and WDM failover for System z connectivity
- Encryption on Brocade and or ADVA WDM possible
- WDM Encryption fro Ethernet/FC/FICON/Coupling up to 100G



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

A real world example



- 10 x 8G FC for PPRC
- 10 x 8G FICON over distance (each Host site to each disk side)
- 4 x Parallel Sysplex InfiniBand (PSIFB) for Coupling
- Only one link used for calculations



Model comparison

Different models (one link) vs long distance fiber need

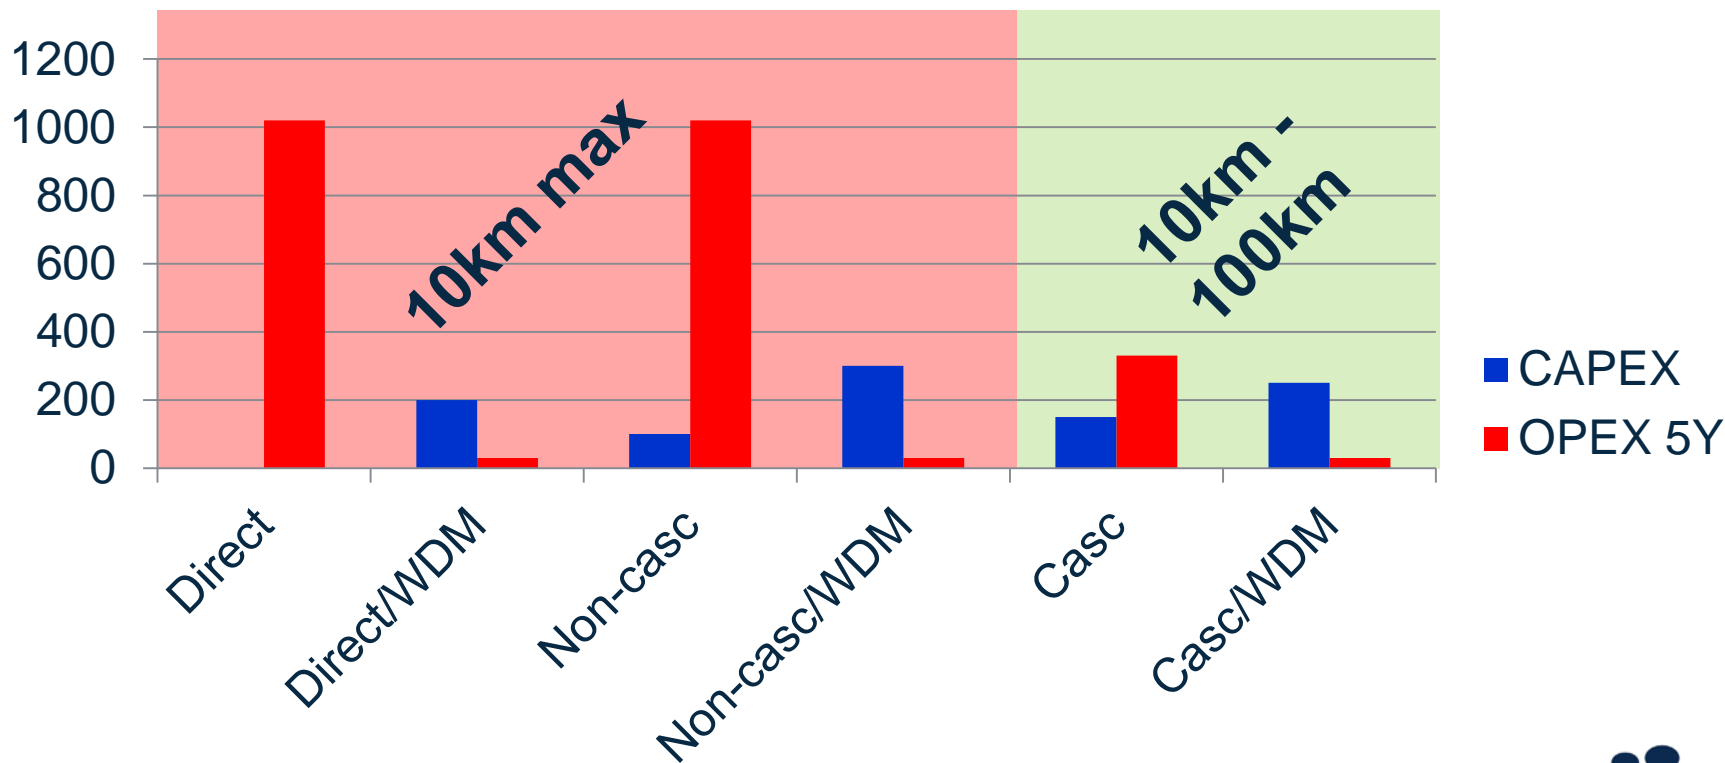
- Direct : 34 x Singlemode cables
- Direct w/ WDM: 1 Singlemode cable, 34 WDM links
- Switched: 34 x Singlemode cables
- Switched w/ WDM: 1 Singlemode cable, 34 WDM links

- Cascaded 10 x Singlemode long dist. Cables (16G ISL's) + 4 WDM links
- Cascaded w/ WDM: 1 Singlemode cable, 14 WDM links

- Fiber Cost estimated USD 600 per fiber pair, per km and per year
- With cascaded directors, the long distance connection bandwidth was reduced from 30 x 8G to 10 x 16G due to better link utilization using 16G ISL's and trunking



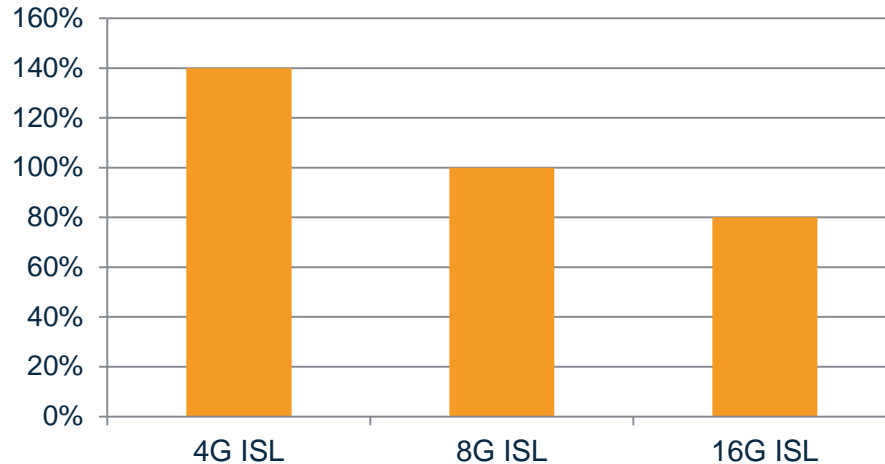
CAPEX/OPEX in 1000USD





WDM Connectivity cost versus ISL Speed

WDM Capex



- 10 x 8G FC/FICON/ISL connectivity over a 50km WDM link
- WDM HW only
- + additional savings on FC director ports
- + better bandwidth utilization on 16G ISL's
- + more upgrade capacity on WDM and FC director

Conclusion

- WDM should not be seen separately, it is part of the SAN/MF architecture to achieve best performance
- Cascaded directors is the #1 choice for longer distances, better performance/utilization
- ISL's should be used at the highest speed possible regardless the local connection speed
- See more on IBM Redbook SG248047:
 - ***System z End-to-End Extended Distance Guide***

Your
Eval
↓

Our `Reaction!

- 5 = "Aw shucks. Thanks!"
- 4 = "Mighty kind of you!"
- 3 = "Glad you enjoyed this!"
- 2 = "A Few Good Nuggets!"
- 1 = "You Got a Nice Nap!"

David Lytle, BCAF
Principal Engineer
dlytle@brocade.com

Uli Schlegel
Director BD
uschlegel@advaoptical.com



Session
16001

