Next Generation 100 Gigabit Optical Ethernet

Call For Interest

Dan Dove – HP Networking
Kapil Shrikhande – Force10 Networks
Pete Anslow – Ciena
Jonathan King – Finisar
Ryan Latchman – Mindspeed Technologies

Objective for this meeting

- To measure the interest in starting a study group for *Next Generation 100Gb Optical Ethernet*

- We don't need to
  - Fully explore the problem
  - Debate strengths and weaknesses of solutions
  - Choose any one solution
  - Create PAR or five criteria
  - Create a standard or specification

- Anyone in the room may speak / vote
- RESPECT… give it, get it

Motivation of this Presentation

- There is a market need for higher density front panel interfaces to support bandwidth growth demand and to reduce cost and power of 100G optics:
  - Module electrical Interface density can be increased
  - MMF port density can be increased
  - SMF port density can be increased
- Define potential project opportunities
- Explain “Why IEEE Project Now?”

Acknowledgement & Support (Contributors and Supporters)

Jon Anderson - Opnext
Thananya Baldwin – Ixia
Michael Bennett - Lawrence Berkeley Lab
Vipul Bhatt - Lightwire
Sudeep Bhoja – Broadcom
Matt Brown - APM
Steve Carlson – High Speed Design
Martin Carroll - Verizon
Frank Chang - Vitesse
Chris Cole - Finisar
Kai Cui - Huawei
John D'Ambrosia - Force 10 Networks
Piers Dawe - IPtronics
Wael Diab - Broadcom
Mike Dudek - Qlogic
Harry Forbes - Nexans
Freddy Hongyan FU - Huawei
Ali Ghiasi - Broadcom
Mark Gustlin - Cisco
Hiroshi Hamano - Fujitsu Labs Ltd
Ziad Hatab – Vitesse
Huangxi - Huawei
Osamu Ishida – NTT
Hiro Iwadate - Sumitomo
Myles Kimmitt - Emulex
Scott Kipp – Brocade
Satoshi Kodama, NTT
Paul Kolesar - CommScope
Masashi Kono - Hitachi
David Lewis – JDSU
Robert Lingle, Jr. – OFS
Kent Lusted - Intel
Phil McClay - TEC
Jeff Maki - Juniper
John McDonough - NEC America
Gary Nicholl - Cisco
Ronald Nordin - Panduit
Mark Nowell – Cisco
David Ofelt – Juniper
George Oulundsen – OFS
Mike Peng Li – Altera
Petar Pepeljugoski - IBM
Jerry Pepper – Ixia
John Petrilla – Avago Technologies
Rick Pimpinella – Panduit
Iain Robertson - Texas Instruments
Sam Sambasivan - AT&T
Oren Sela - Mellanox
Koichiro Seto - Hitachi Cable Ltd
Megha Shanbhag - TEC
Song Shang - Semtech Corporation
Siddharth Sheth - Inphi
Ted Sprague - Infinera
Peter Stassar - Huawei
Andre Szczepanek - Inphi
Nathan Tracy - TEC
Matt Traverso - Cisco
Francois Tremblay - Gennum
Steve Trowbridge, Alcatel-Lucent
Paul Vanderlaan - Nexans
Tim Warland - APM
Zengli – Huawei

Agenda

- Team members
- Review of 802.3ba
- Market Opportunity
- Proposed Project
- Areas of Study
- Organizational Structure
- Straw Polls

Team Members

- Dan Dove – HP Networking
- Kapil Shrikhande – Force 10 Networks
- Pete Anslow – Ciena
- Jonathan King – Finisar
- Ryan Latchman – Mindspeed Technologies

Review of 802.3ba 100G Optical PHYs

Name	Description
100GBASE-SR10	100 Gb/s PHY using 100GBASE-R encoding over ten lanes of multi-mode fiber, with reach up to at least 100 m (see Clause 86)
100GBASE-LR4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 88)
100GBASE-ER4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 88)

Technology Used for 802.3ba (1 of 3)

- SMF solution based on 25Gb/s optics
  - 10x10G electrical interface to module
    - 25Gb/s signaling was too challenging at the time
    - Connector technology for 25Gb/s was not available at the time
  - 4x25G optical wavelengths
    - Lowest long term laser cost and power alternative
    - Enables investment in optimum long term technology
  - Gearbox converts 10x10 interface to internal 4x25G electrical lanes

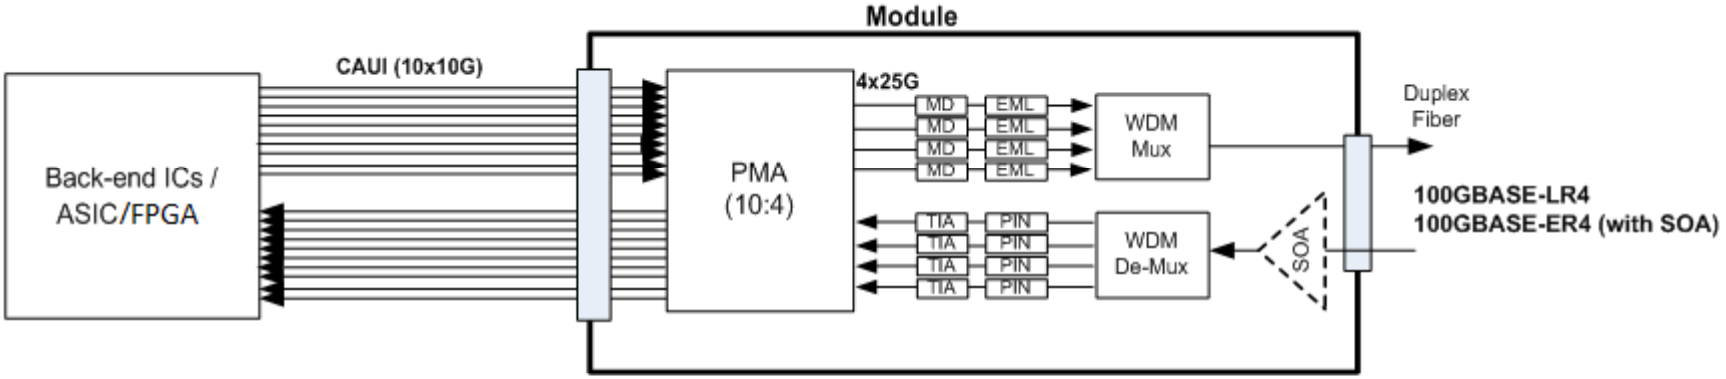

Technology Used for 802.3ba (2 of 3)

- MMF solution based on 10Gb/s VCSELs
  - 10x10G optical channels -> Two 12-fiber optical cables per link
    - 25Gb/s signaling was too challenging at the time
    - Connector technology for 25Gb/s was not available at the time
    - 25G/s VCSELs were not available at the time

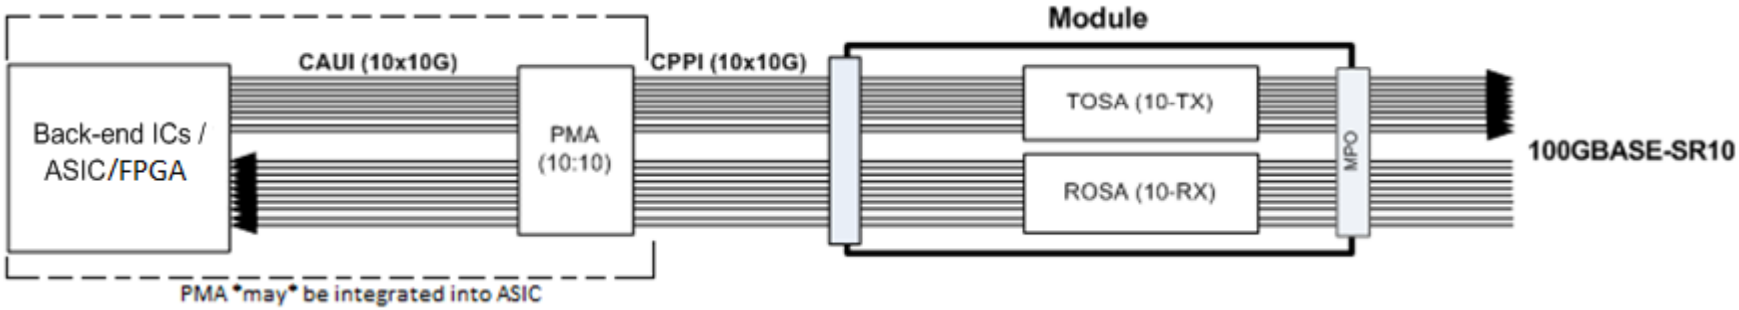

Note: figure shows implementation using un-retimed module. Retimed module also possible

Technology Used for 802.3ba (3 of 3)

- Electrical interfaces based on 10 Gb/s
  - 10x10G electrical channels, differential -> **40 high speed pins** just for data
  - ASIC/FPGA I/O challenged to achieve high enough density
    - 32 CXP connectors possible on a faceplate would require 320 - 10G channels
  - 25Gb/s signaling too challenging at the time
  - Connector technology for 25Gb/s not available at the time
  - Low power 25Gb/s I/O not available at the time

CAUI Interface Loss Summary

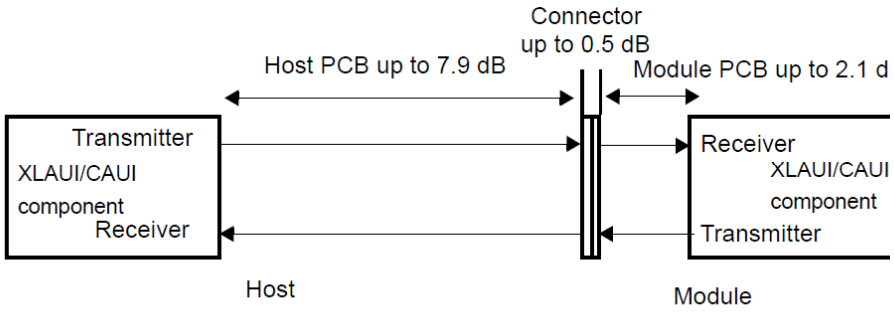

Figure 83B–2—Chip-module loss budget at 5.15625 GHz

CPPI Interface Loss Summary

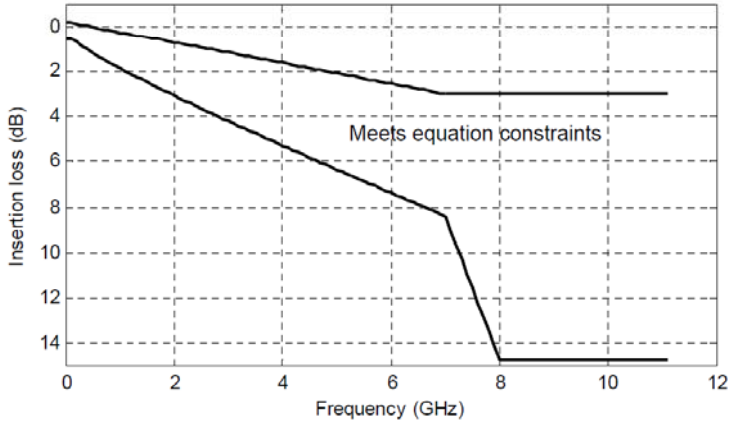

Figure 86A–11—Recommended insertion loss limits of host PCB, connector and HCB

100G SMF Port Density Opportunity

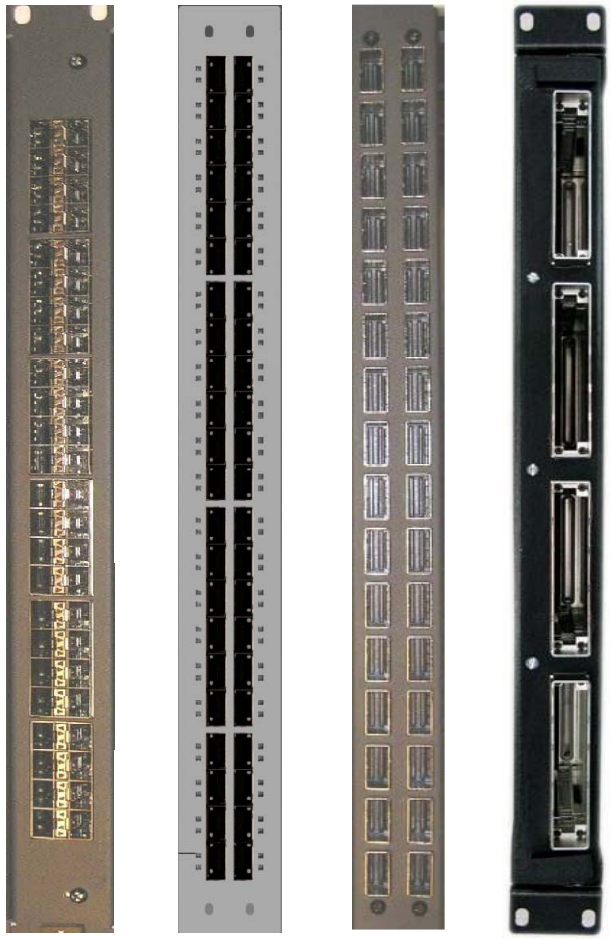

a b c d

Source: 100GbE Electrical Backplane/Cu Cable CFI

Line card illustrations

a. 48 ports SFP+ @ 10GbE = 480 Gb/s
b. 44 ports QSFP @ 40GbE = 1.76 Tb/s
c. 32 ports CXP@ 100GbE= 3.2 Tb/s (MMF only)
d. 4 ports CFP @ 100GbE= 400 Gb/s

Perceived opportunity @ 100G:

- CFP provides lower bandwidth density than10Gb/s
- 100Gb/s SMF solutions limited by CFP size
  - First Gen ICs & Optics require large module size to fit in all components and dissipate power
  - 10x10G interface requires 40 high speed data pins

Multimode Cable Cost/Density Opportunity

Example 100GBASE-SR10 end-to-end channel

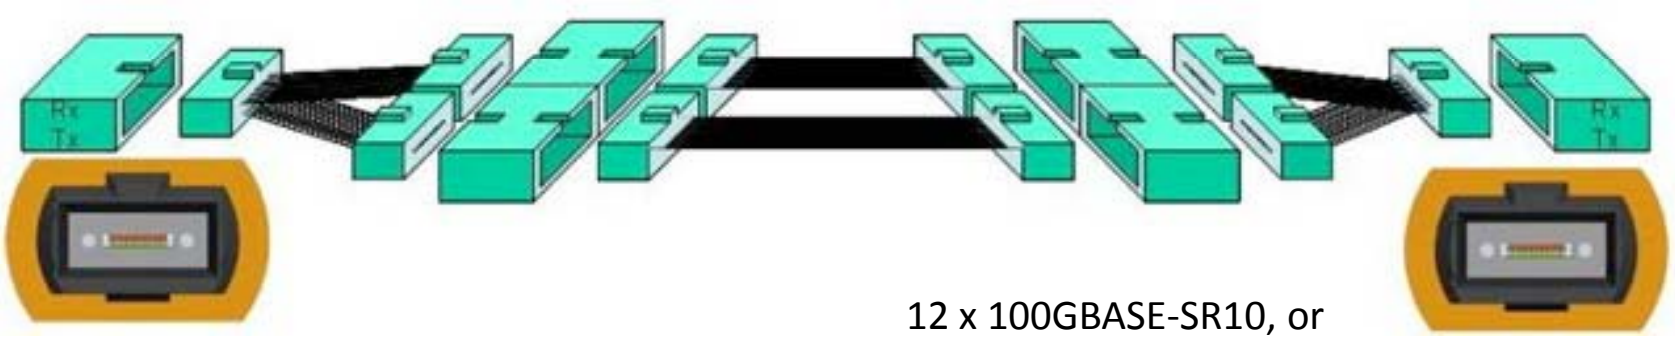

Using 12-fiber cables (common), one 100GBASE-SR10 link makes 2 appearances on the panel using 2 MPO connectors

Going from 10Gb/s to 25Gb/s cuts infrastructure needs by half.

12 x 100GBASE-SR10, or
24 x "100GBASE-SR4"

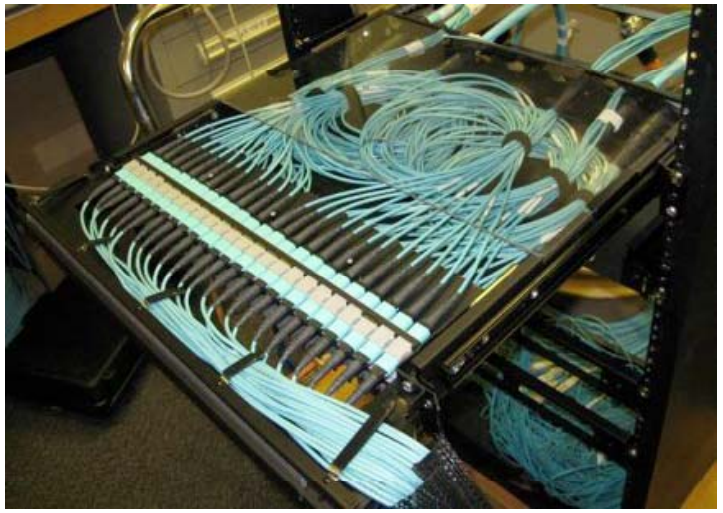

All pictures in the slide are courtesy of CommScope

100G MMF Fiber Density/Cost Opportunity

- 20 fibers/link can be reduced to 8 fibers/link
- Significantly reduced lane count: 10 -> 4 pairs per duplex link
  - Lower power/bit, lower cost/bit, lower infrastructure costs
  - Reducing number of lasers/PDs will reduce port power/cost
- A narrower interface permits a higher density short reach optical interface for multimode fiber, which will offer multiple potential benefits compared with 100GBASE-SR10
- Direct scaling from 40G to 100G using the same cable infrastructure by increasing 4x10G to 4x25G

Images in this slide are courtesy of Finisar

Market Need for Bandwidth

- Bandwidth demand growing and driving density
  - Historically this has been seen with 10G
  - Market trends leading to 100G repeating that evolution

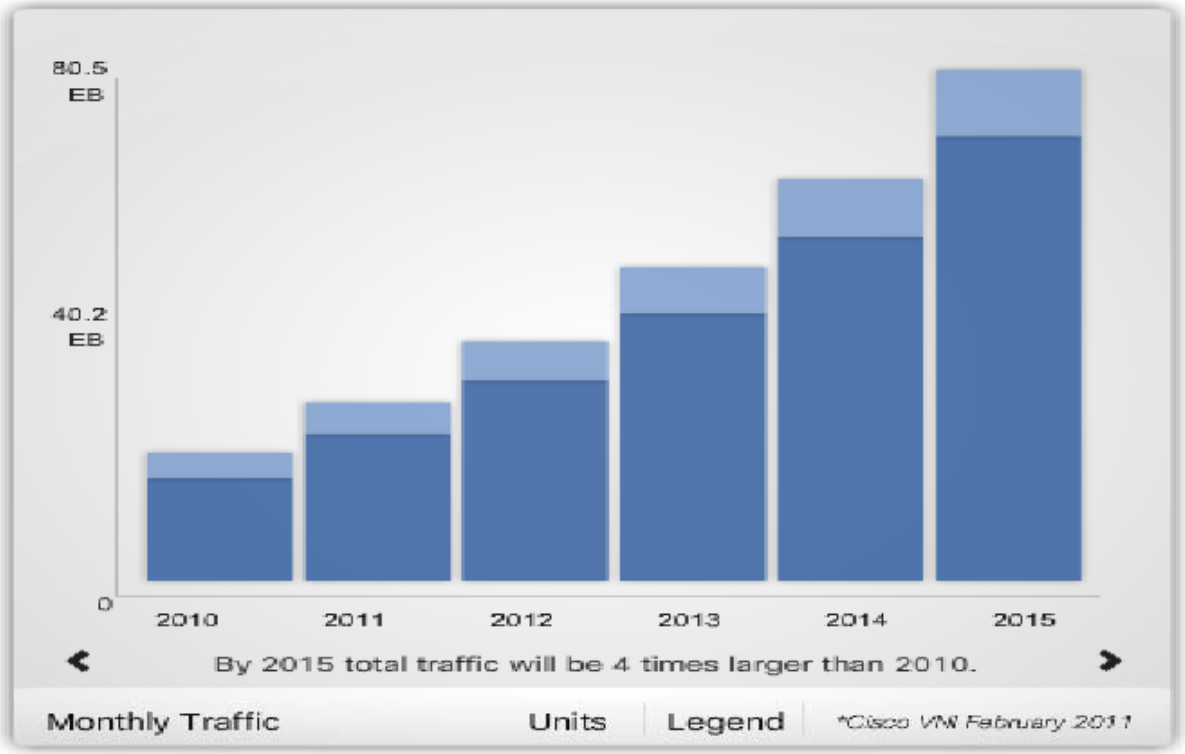

Business
Consumer

Global IP Traffic, 2010-2015
By Segment (ExaBytes per Month)

Courtesy: Cisco Systems

http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html#

BW Growth: Example AMS-IX Data

Peak traffic in bits/sec from 2005-2010 shows exponential growth

Courtesy AMS-IX

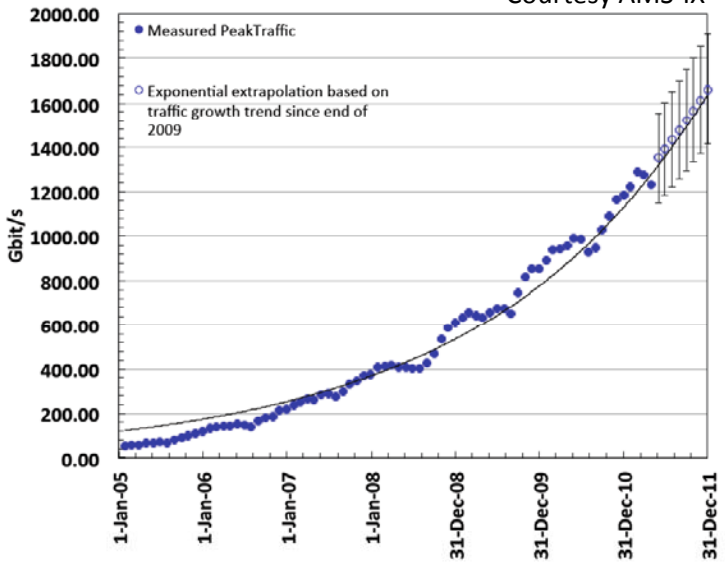

PBytes per month, May 2007 to Apr 2011 shows ~ 45% AGR

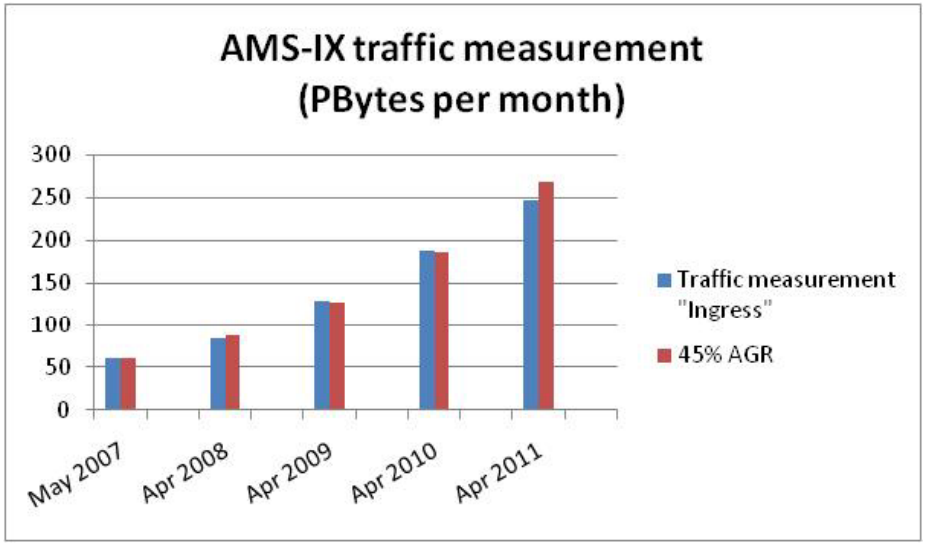

Plotted from raw data available at http://www.ams-ix.net/statistics/

Peak traffic in bits/sec, Jan 2010 to Apr 2011

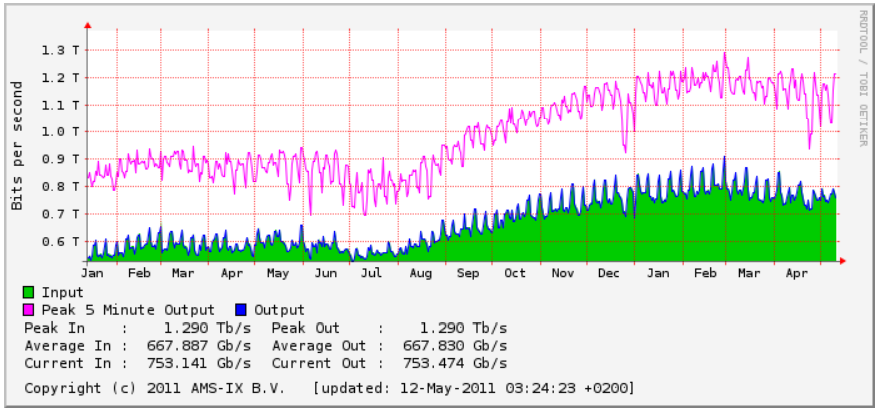

Courtesy AMS-IX

Market Opportunity for Size Reduction

- 40 high speed connector signal pins (62 including signal GND pins) and 40 PCB RF traces are reduced to 16 high speed connector pins (26 including signal GND pins) and 16 PCB RF traces
- Reducing optical module size allows higher switch density and improved rack utilization
- Reductions in space will provide opportunity to expand market faster and more efficiently

Market Opportunity for Power Reductions

- Power inside optical module is driven by optical components and electrical interface ICs types, and the number of optical lanes
- Reducing the number of optical lanes from 10 to 4 reduces the number of optical components from 10 to 4 and results in large reduction in module power
- Replacing the Gearbox IC by a Re-timing IC leads to additional power savings
- If technically feasible and necessary, removal of re-timing IC leads to further power savings
- Reductions in size will provide opportunity to expand market faster and more efficiently

Market Opportunity for Cost Reduction

- For MMF, reducing the number of lasers from 10 to 4, and number of required parallel MMFs from 20 to 8 leads to cost reduction. (SMF already uses 4 lasers)
- Reducing the size and power of modules allows higher port density
- Increase in port density results in better amortization of switch infrastructure (fans, PC, chassis, etc.) which reduces cost per port
- Reductions in cost will provide opportunity to expand market faster and more efficiently

10G Ethernet Single-mode Optics: Evolutionary Path

10GbE evolution → improvement in power, density, and cost due to:

- Higher electrical lane speed enabled reduction in number of electrical lanes
- Improved technology enabled lower power electrical and optical components

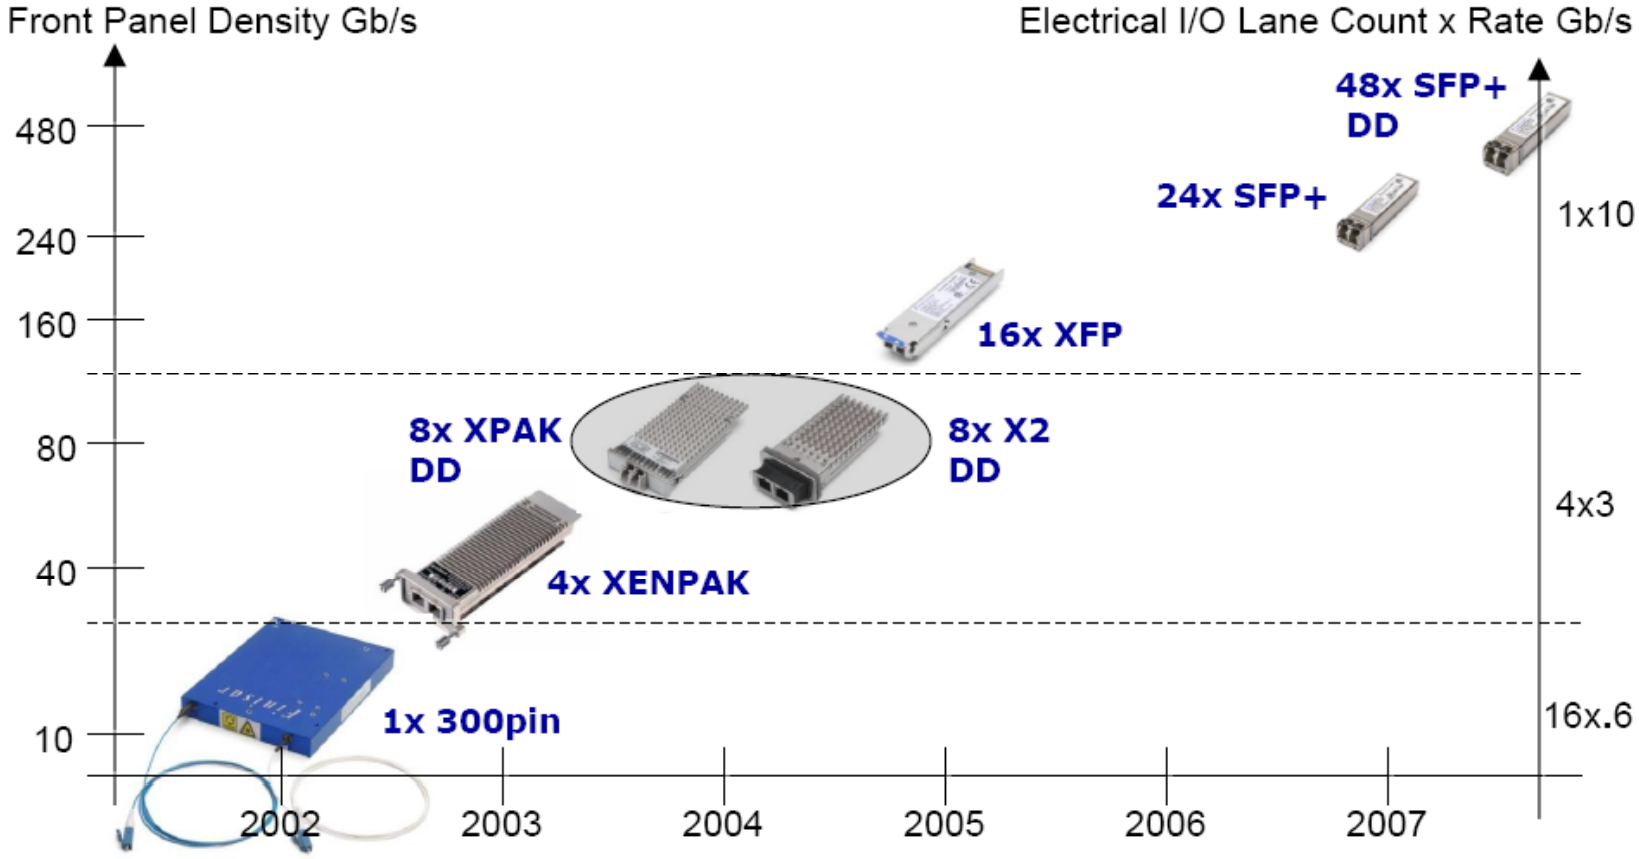

100G Ethernet Single-mode Optics: Possible Evolutionary Path

100GbE evolution → improvement in power, density, and cost due to:

- Higher electrical lane speed enabling reduction in number of electrical lanes
- Improved technology enabling lower power electrical and optical component

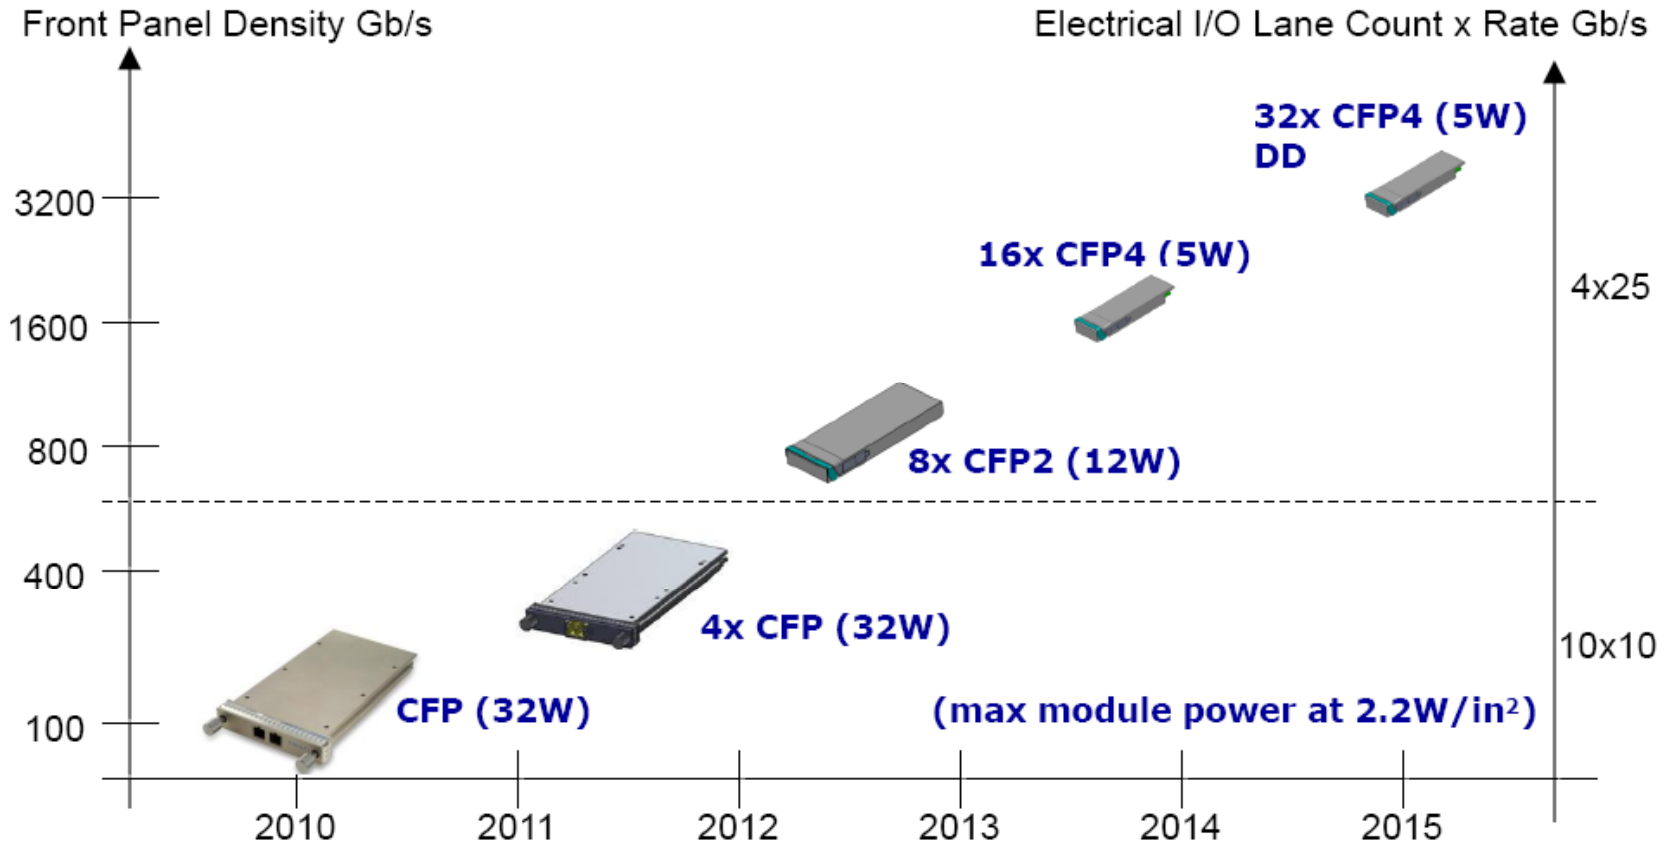

100G Ethernet Single Mode Optics: Potential Evolutionary Path using 4x25G interfaces

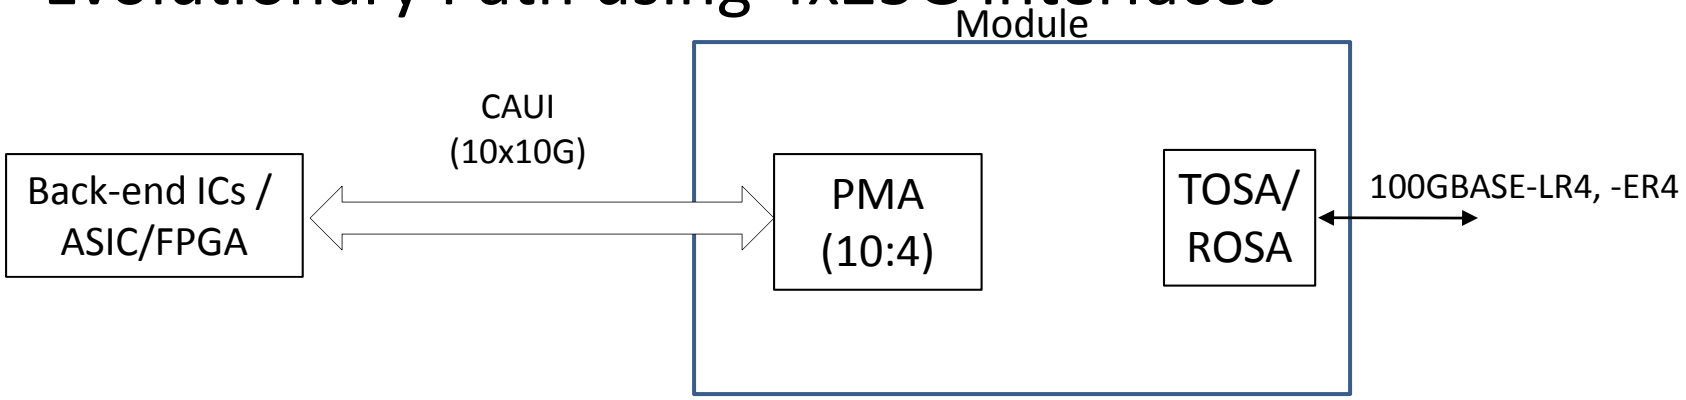

100G Ethernet Single Mode Optics: Potential Evolutionary Path using 4x25G interfaces

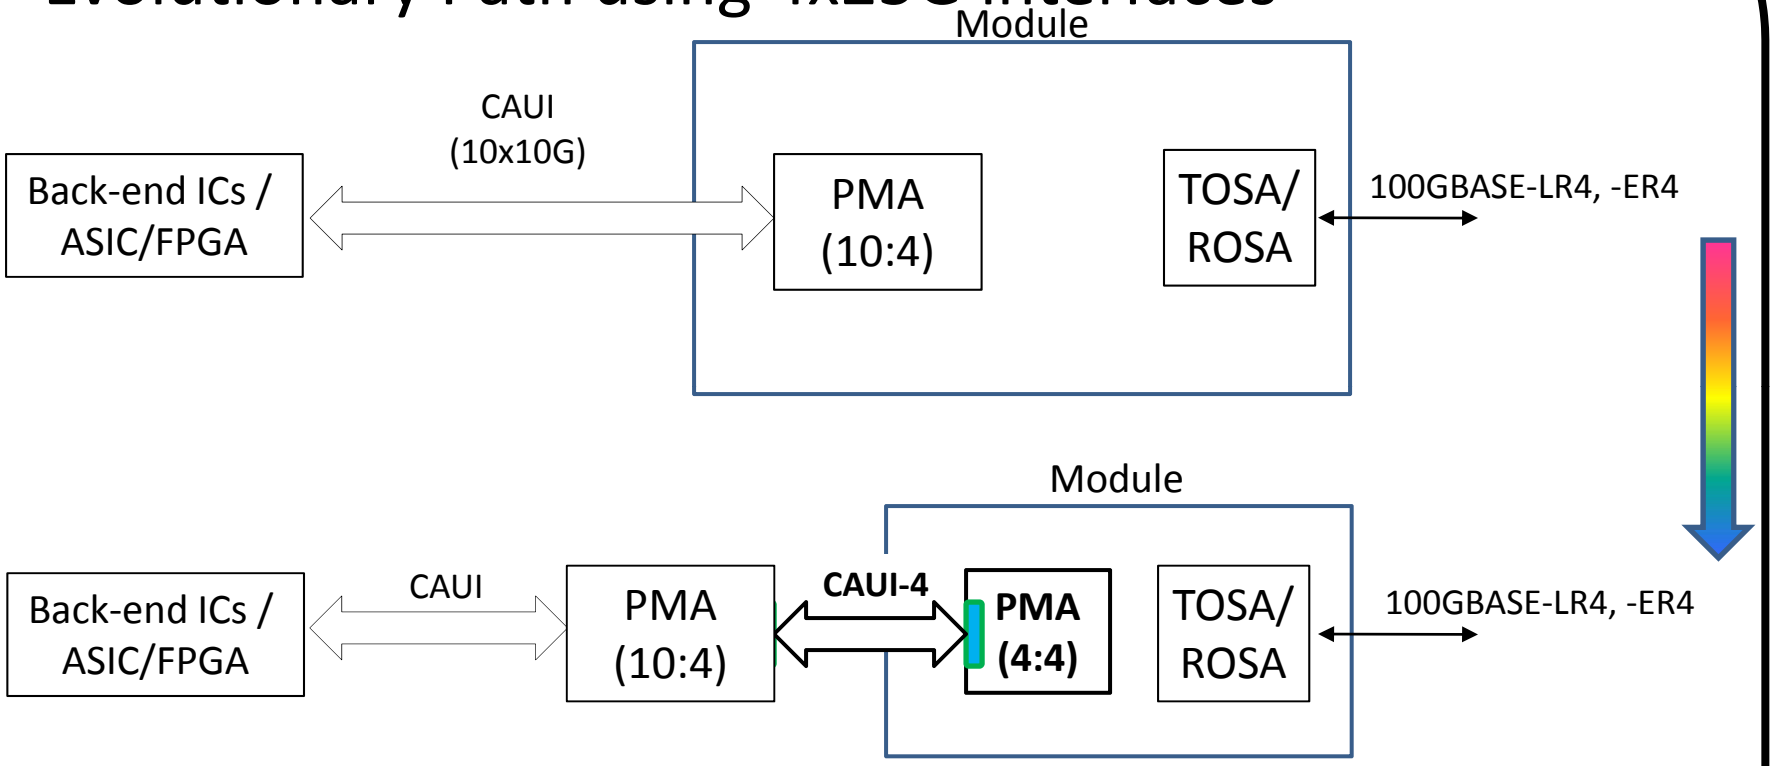

100G Ethernet Single Mode Optics: Potential Evolutionary Path using 4x25G interfaces

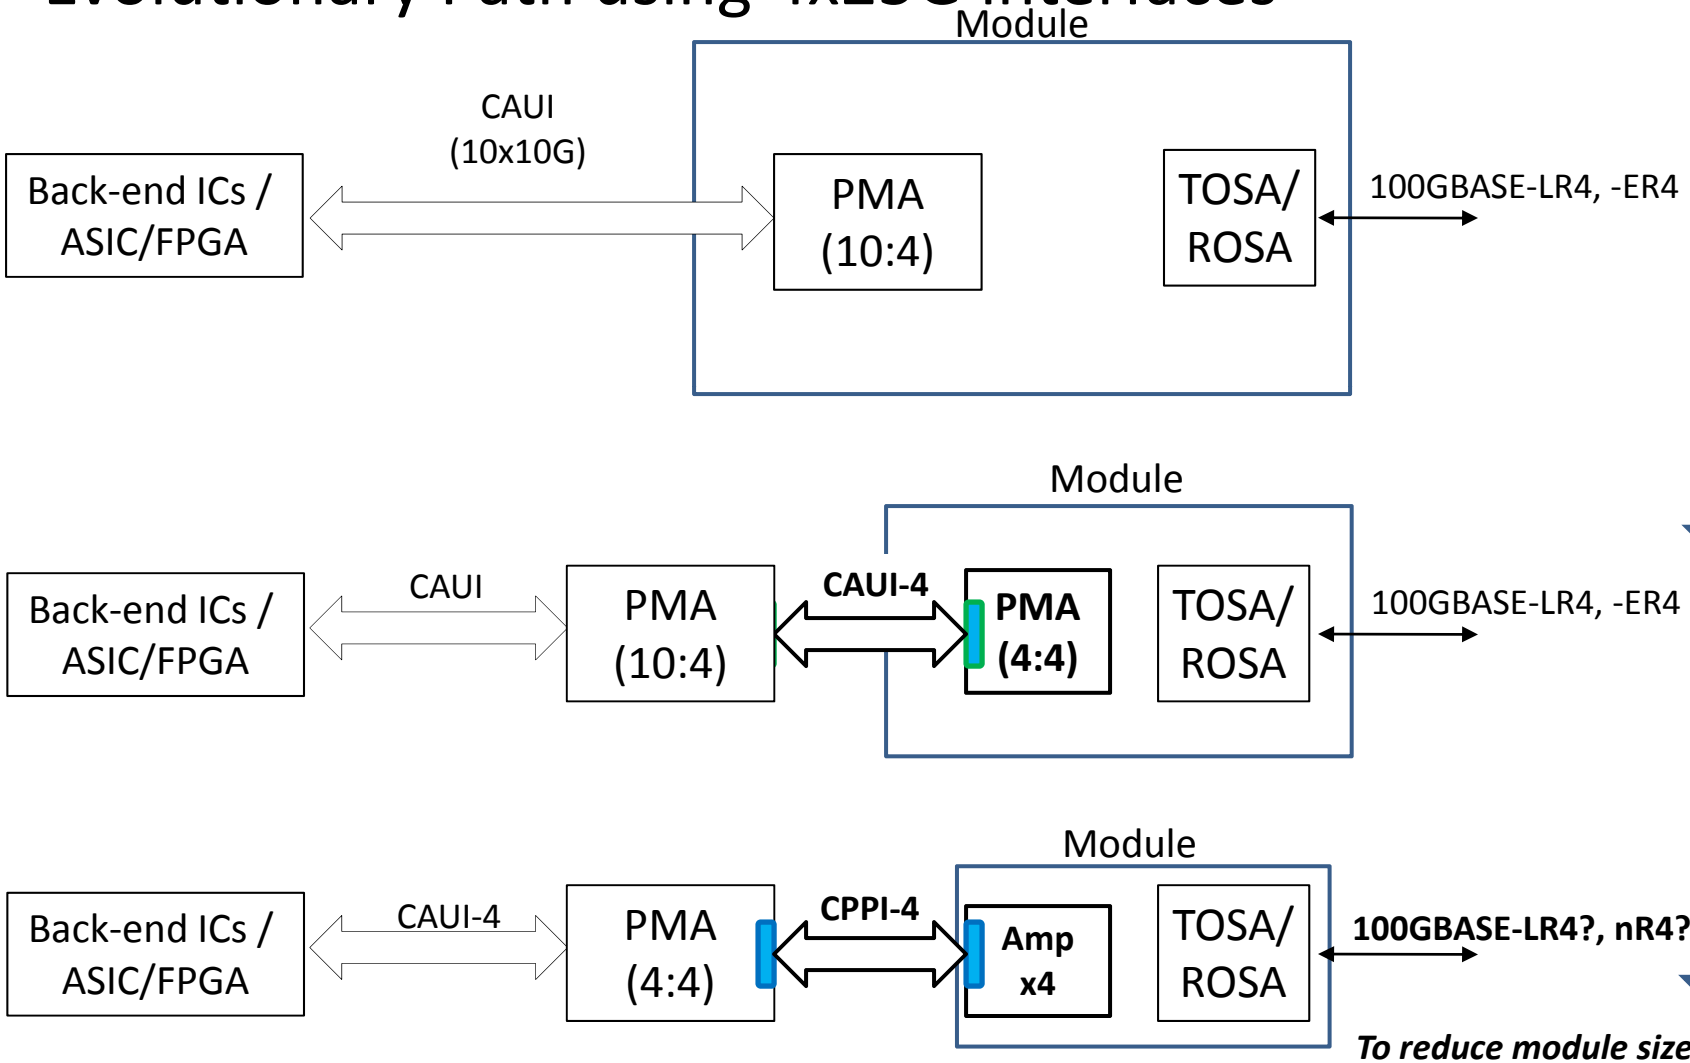

Potential areas for SMF study

- CAUI-4 (retimed 4 lane) and/or CPPI-4 (un-retimed or equivalent 4 lane) electrical interfaces
- Study alternate PMD technologies to determine if there is significant opportunity for additional size, power and cost reduction
- Reach on G.652 fiber
- Optical transmitter and receiver performance
- Link power and jitter budgets
  - Start with signaling rate scaled budgets
  - Use of retiming in module, EDC, and/or host card retiming near connector
  - Impact of power, cost, size requirements
- Performance of 4 x 25 GBd electrical connector
- Use of FEC and impact on budgets
  - valuable if it has no/v. low overhead (e.g. KR or similar) and v. low latency However, resides in the host ASIC/FPGA so only available for new designs

Potential areas for MMF study

- 4 x 25 GBd, 4 parallel fibers, each direction
- NRZ modulation format
- CAUI-4 (retimed 4 lane) and/or CPPI-4 (un-retimed or equivalent 4 lane) electrical interfaces
- Reach on OM3, OM4
- Optical transmitter and receiver performance
- Link power and jitter budgets
  - Start with signaling rate scaled budgets
  - Use of retiming in module, EDC, and/or host card retiming near connector
  - Impact of power, cost, size requirements
- Performance of 4 x 25 GBd electrical connector
- Use of FEC and impact on budgets
  - valuable if it has no/v. low overhead (e.g. KR or similar) and v. low latency. However, resides in the host ASIC/FPGA so only available for new designs

Potential areas for Cu study (electrical interface)

- 4 x 25 GBd
- NRZ modulation format
- CAUI-4 (retimed 4 lane)
- CPPI-4 (un-retimed or equivalent 4 lane)
  - Impact upon optical transmitter and receiver performance
- Performance of 4 x 25 GBd electrical connector and channel
- Equalization & de-emphasis requirements
- Electrical interface jitter budget, amplitude requirements

Why an IEEE Project Now

- Bandwidth demands are going up exponentially
- The only way to meet this demand is to increase density
- Density is being increased in the backplane and passive Cu already
- Density needs to be increased on the front panel
  - New electrical interfaces: CAUI-4, CPPI-4
  - New MMF interface: -SR4
  - Determine if there is sufficient market demand and technical maturity of feasible optics alternatives to justify a new duplex SMF interface in addition to existing -LR4
- Allow coordination with Backplane & Copper Cables while there is an opportunity to influence their direction

25G Lane Rate Standards Activity is Ramping Up

Organization	Project	Notes
Optical Internetworking Forum (OIF) [1]	CEI-28G-SR, -VSR	• 19.9 to 28.05 GBd/lane • Chip-to-chip, chip-to-module
	CEI-25G-LR	• 19.9 to 25.8 GBd/lane • Backplane
Infiniband Trade Association (IBTA)	EDR	• 25.78125 GBd/lane • Passive copper, active cables including optical fiber
INCITS T11 Fibre Channel	32GFC	• 28.05 GBd (single lane) • Chip-to-module, multimode and single-mode fiber
IEEE 802.3	100 Gb/s Backplane and Copper Study Group	• ~25 Gb/s/lane • Backplane, copper cable

[1] Refer to the proceedings of the IEEE 802.3 Ethernet Working Group OIF CEI-28G-VSR liaison response ad hoc (http://www.ieee802.org/3/ad_hoc/OIF_VSR_liaison/index.html).

Organizational Structure of Project

- Division of activities
  - Multimode
  - Single-mode
  - Copper interface

- Proposed schedule
  - Tuesday, July 19th 3:30 PM to 5:00 PM – CFI consensus building meeting
  - Thursday, July 21st – 802.3 Vote to form SG
  - Friday, July 22nd – EC vote whether to approve SG
  - Monday, August 8th – Website & Reflector set up (pending EC approval)
  - Week of September 12th – First SG meeting (pending EC approval)

Straw Polls - CFI

- Should a study group be formed for “Next Generation 100GbE Optical Interfaces”?

Y___ N___

Straw Polls - Participation

- Number of CFI Attendees

Number ___

- I would participate in the Next Generation 100GbE Optical Interfaces Study Group in IEEE 802.3.

Number ___

- My company would support participation in the Next Generation 100GbE Optical Interfaces Study Group in IEEE 802.3.

Number ___

Straw Polls - Participation

- If a study group is formed, I plan to attend and participate in the September meeting.

Number ___

- If a study group is formed, I would like to attend both this SG and the 100G Backplane & Copper Cables meetings

Number ___

THANK YOU!