IEEE 802.3
Call For Interest
100 Gb/s Per Lane Optical PHYs for 2 and 10 km for 100 GbE and 400 GbE

Consensus Presentation

Objective for this Meeting

- To *measure the interest* in starting a study group to address:
  - 100 Gb/s Per Lane Optical PHYs for 2 and 10 km for 100 GbE and 400 GbE
- 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

Agenda

- Market Drivers
- Technical Feasibility
- Why Now?
- Q&A Panel
- Straw Polls

Presenters and Panelists

- TBD
-
-
-

The request

Ethernet has a successful track record of reusing and leveraging technology in order to enable new cost-optimized solutions for broad market adoption. Recently, both the IEEE 802.3bs and IEEE 802.3cd Ethernet projects defined optical interfaces based on 100 Gb/s PAM4 per lane optics for 500m reaches on single-mode fiber. This technology enables a lower component count implementation that can lead towards a lower cost solution.

The successful industry adoption of 100 Gb/s is resulting in higher volumes and continuing cost pressures on optical interfaces and the hyper-scale data centers, being aggressive adopters of cost-effective solutions, are looking to enable the next generation of lower cost solutions for 400 Gb/s Ethernet interfaces.

At this time, no IEEE 802.3 Ethernet specifications exist for greater than 500m reaches on single-mode fiber using this advanced technology. This Call For Interest is a request for the formation of a Study Group to explore the potential market requirements and feasibility of extending 100 Gb/s PAM4 per lane optical technology to longer 100 Gb/s and 400 Gb/s Ethernet reaches.

Supporters (your name here too – send me an email)

Justin Abbott, Lumentum
Pete Anslow, Ciena
Vittal Balasubramanian, Innovium
Thananya Baldwin, Keysight
Vipul Bhatt, Finisar
Jose Castro, Panduit
Frank Chang, Source Photonics
David Chen, Applied Optoelectronics
Chris Cole, Finisar
Arash Farhoodfar, Inphi
Ali Ghiasi, Ghiasi Quantum
Paul Goldgeier, ColorChip
Jonathan Ingham, Foxconn Optical Interconnect
Hideki Isono, Fujitsu
Kenneth Jackson, Sumitomo Electric
John Johnson, Broadcom
Mark Kimber, Semtech
Jonathan King, Finisar
Greg Lecheminant, Keysight
David Lewis, Lumentum
Hai-Feng Liu, Intel
Karen Liu, Lightwave Logic
Khushrow Machhi , Broadcom
Jeffery Maki, Juniper
Marco Mazzini, Cisco
Shirao Mizuki, Mitsubishi Electric
Ray Nering, Cisco
Gary Nicholl, Cisco
Tom Palkert, MACOM
Rajiv Pancholy, Broadcom
Jerry Pepper, Keysight
Rick Pimpinella, Panduit
Kees Propstra, Multilane
Rick Rabinovich, Keysight
Scott Schube, Intel
Kapil Shrikhande, Innovium
Scott Sommers, Molex
Phil Sun, Credo
Mike Takefman, Inphi
Pirooz Tooyserkani, Cisco
Matt Traverso, Cisco
Eddie Tsumura, Sumitomo Electric
Jeff Twombly, Credo
Ed Ulrichs, Source Photonics
Mike Wang, HiSense
Brian Welch, Luxtera
Chongjin Xie, Alibaba
Simon Ximen, ColorChip
Ryan Yu, Molex
Hua Zhang, HiSense
Kevin Zhang, IDT
Pavel Zivny, Tektronix

Overview: Motivation

Significant industry interest and progress has been made towards extending the existing IEEE 802.3 Optical PHYs using 100 Gb/s per lane optical technology to longer reaches.

This proposed study group would look to develop 2 km and 10 km SMF PHYs for both 100 GbE and 400 GbE

The motivation is to leverage technology to address the ongoing cost pressures on optical interconnects for a set of known and identified markets including:

- Web-scale data centers
- Service Provider
- Enterprise data centers

Lower cost solutions occur due to reduced lane/component count or through enabling higher density solutions.

Today's Point-to-Point SMF Ethernet

	Lanes	500 m	2 km	10 km	20 km	40 km	Up to 80km
1000BASE-	1		LX	LX10 / LH		EX	ZX
10GBASE-	1			LR		ER	ZR
25GBASE-	1			LR		ER	
40GBASE-	4	PSM4		LR4		ER4	
	1		FR				
50GBASE-	1		FR	LR		ER	
100GBASE-	10		10X10				
	4	PSM4	CWDM4 / CLR4	LR4 / 4WDM-10	4WDM-20	ER4 / 4WDM-40	
	1	DR					"ZR"
200GBASE-	4		FR4	LR4		ER4	
400GBASE-	8		FR8	LR8			
	4	DR4				ER8	"ZR"
	1						

Black Text IEEE Standard
Red Text In Standardization
Blue Text Non-IEEE standard but complies to IEEE electrical interfaces

Focus of this CFI

What Are We Talking About?

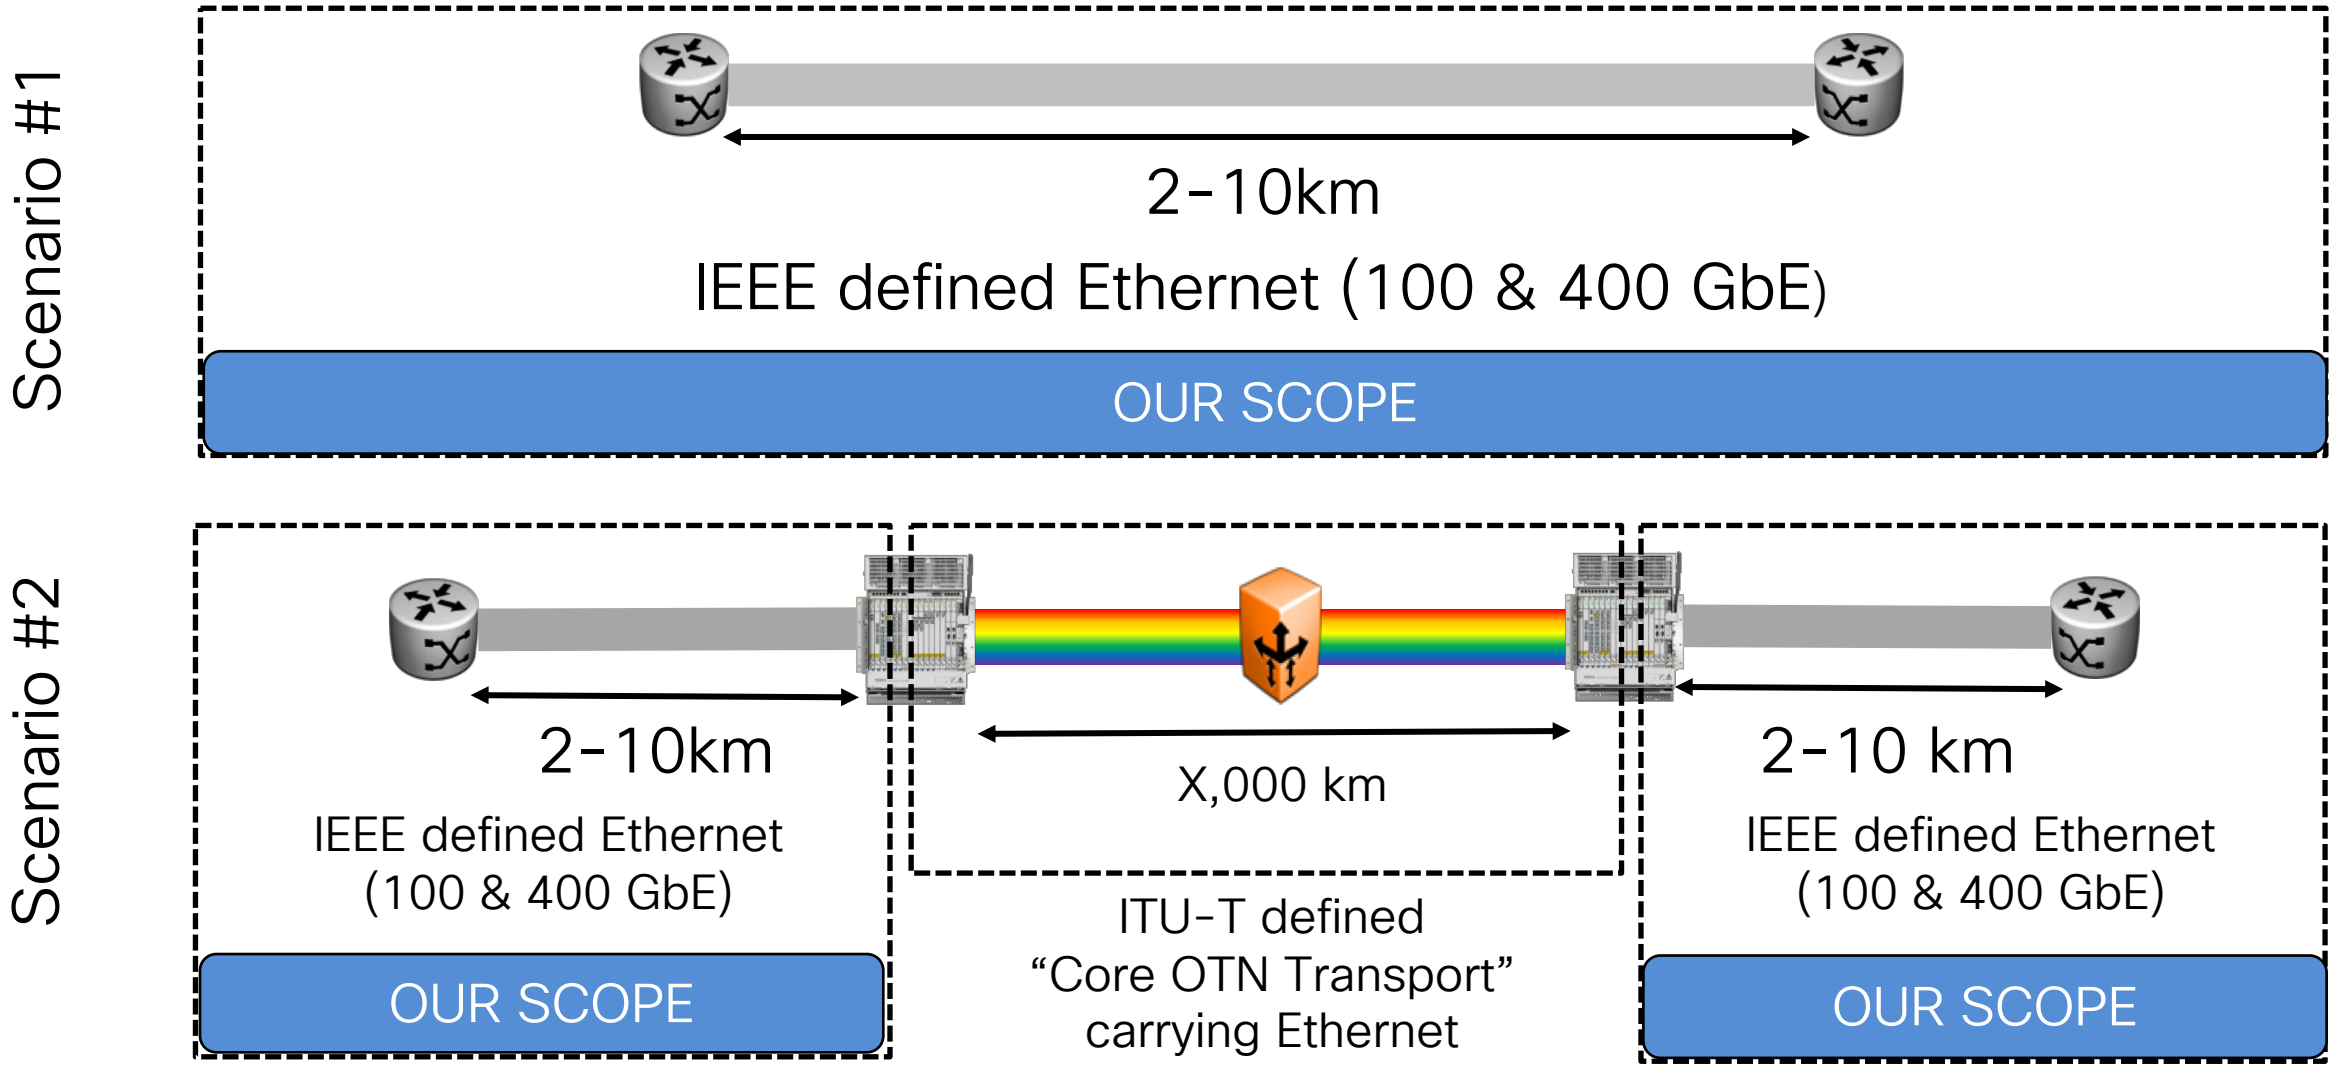

Market Drivers:

longer reach (up to 10 km) 100 Gb/s per lane optical technology

2 & 10 km optics dominate throughout SMF ecosystem

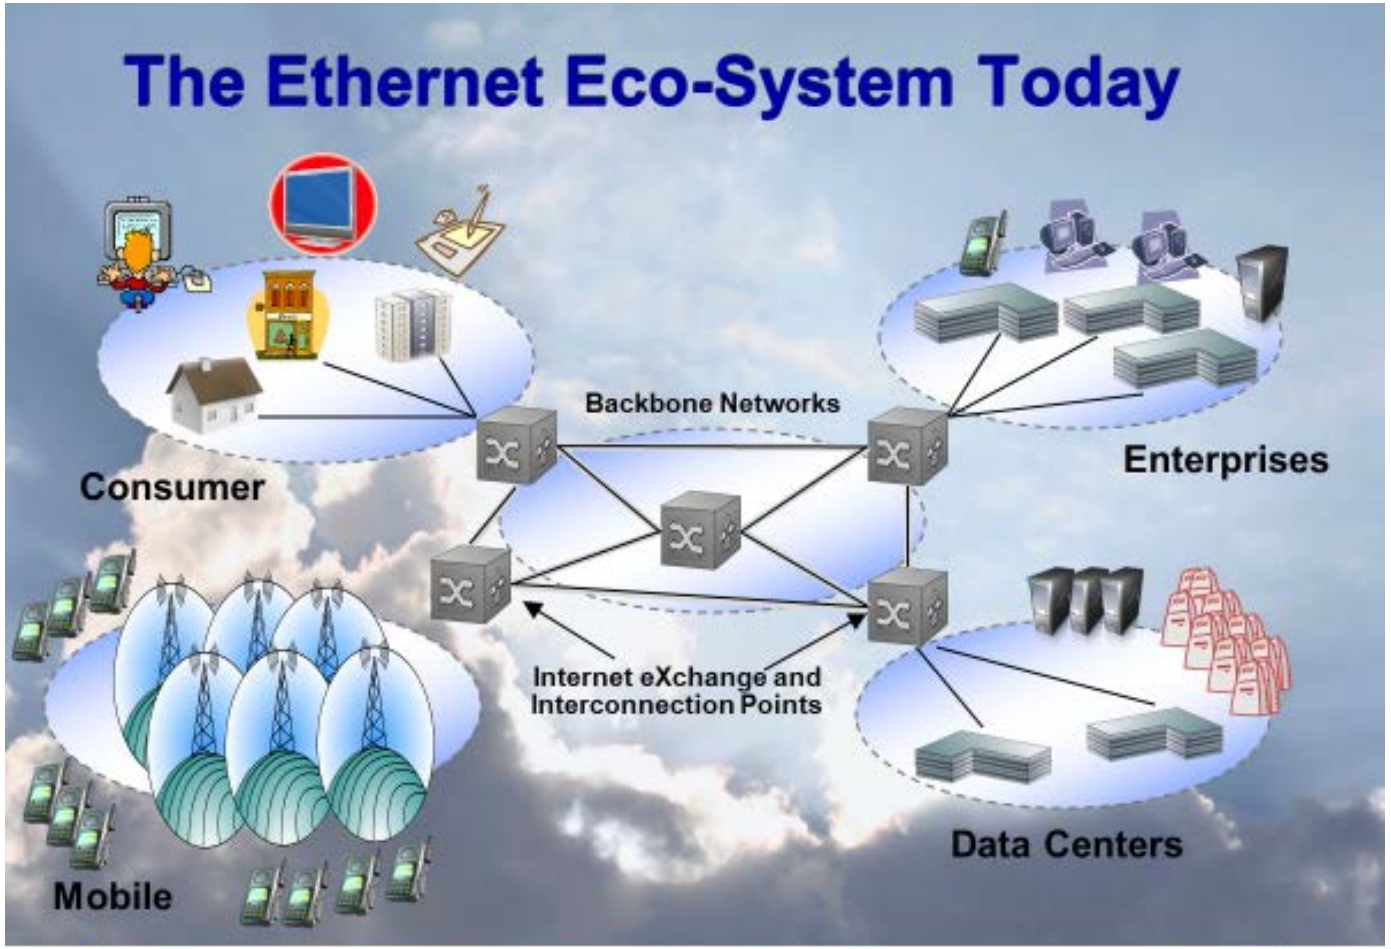

- Equipment interconnect within buildings
  - Web-scale Data Center
  - Service Provider Data Center and Point-of-presence
  - Enterprise Data Center
- Intra-building interconnects in campus environment (up to 10 km)
- Forecast SMF market size (100 GbE and 400 GbE) of 12M Million modules per year in 2023
- This CFI's goal is to add next generation optical technology into that ecosystem

Market Forecast slides

100 GbE Modules by reach

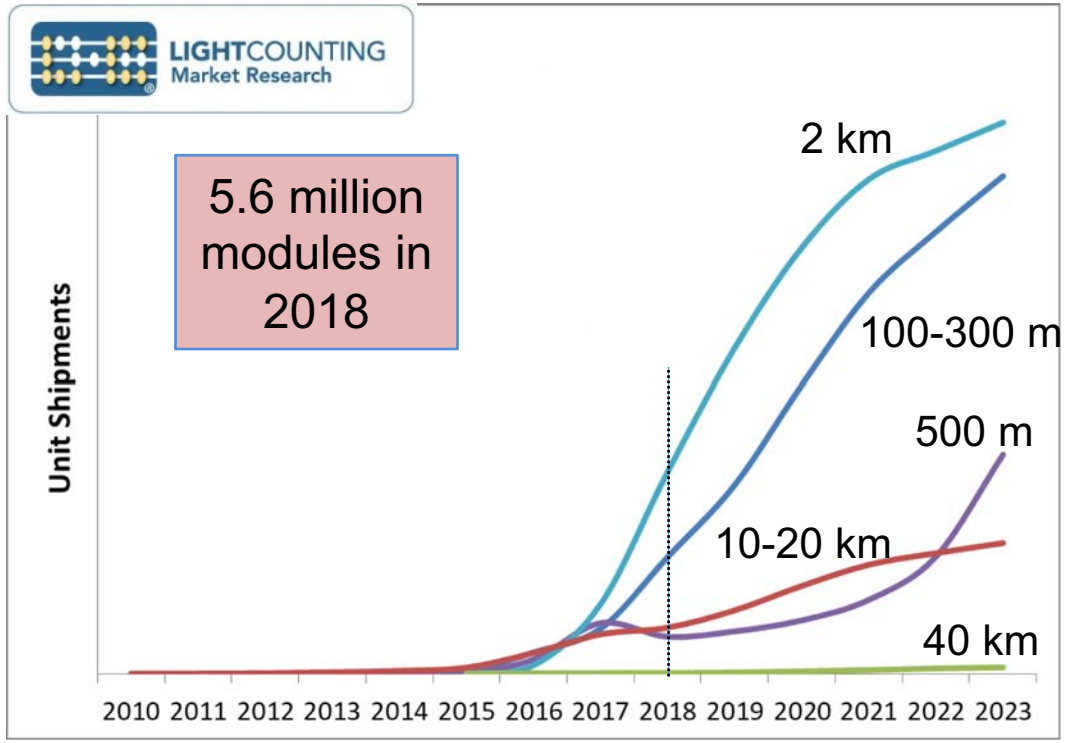

Courtesy Dale Murray, Light Counting

100 GbE & 400 GbE 2 & 10 km SMF Modules

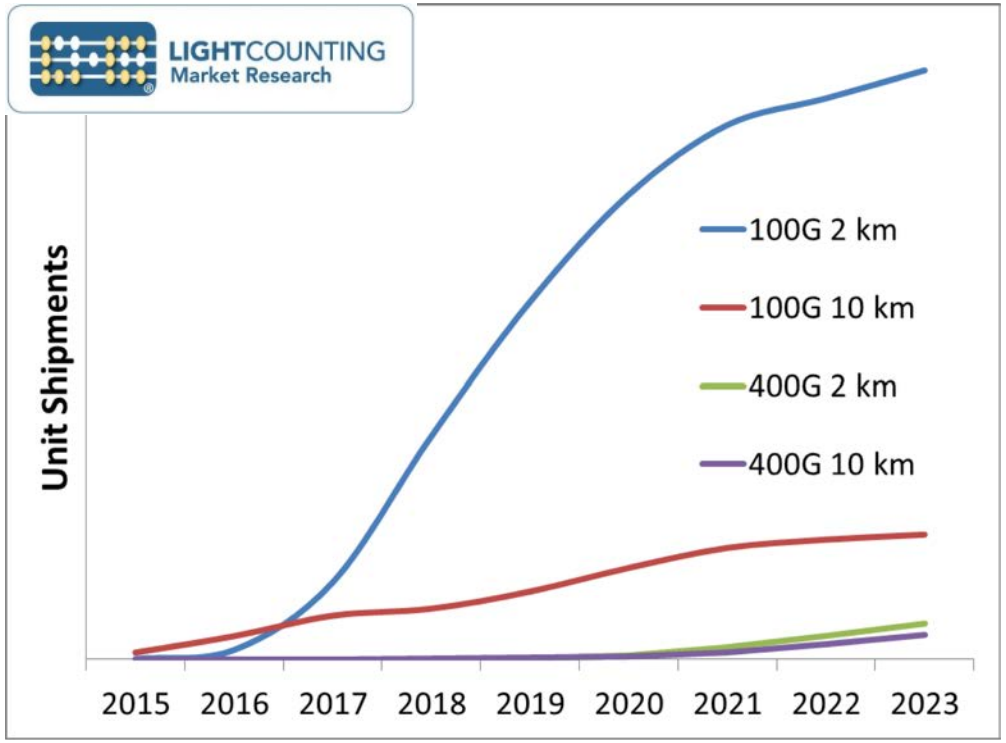

Courtesy Dale Murray, Light Counting

- 100 GbE optics market still in strong growth phase
- 400 GbE at start of its ramp but expected to be fast

} Both market conditions benefit from cost reductions

* At the right time

Ethernet's consistent trend – Narrower/Faster

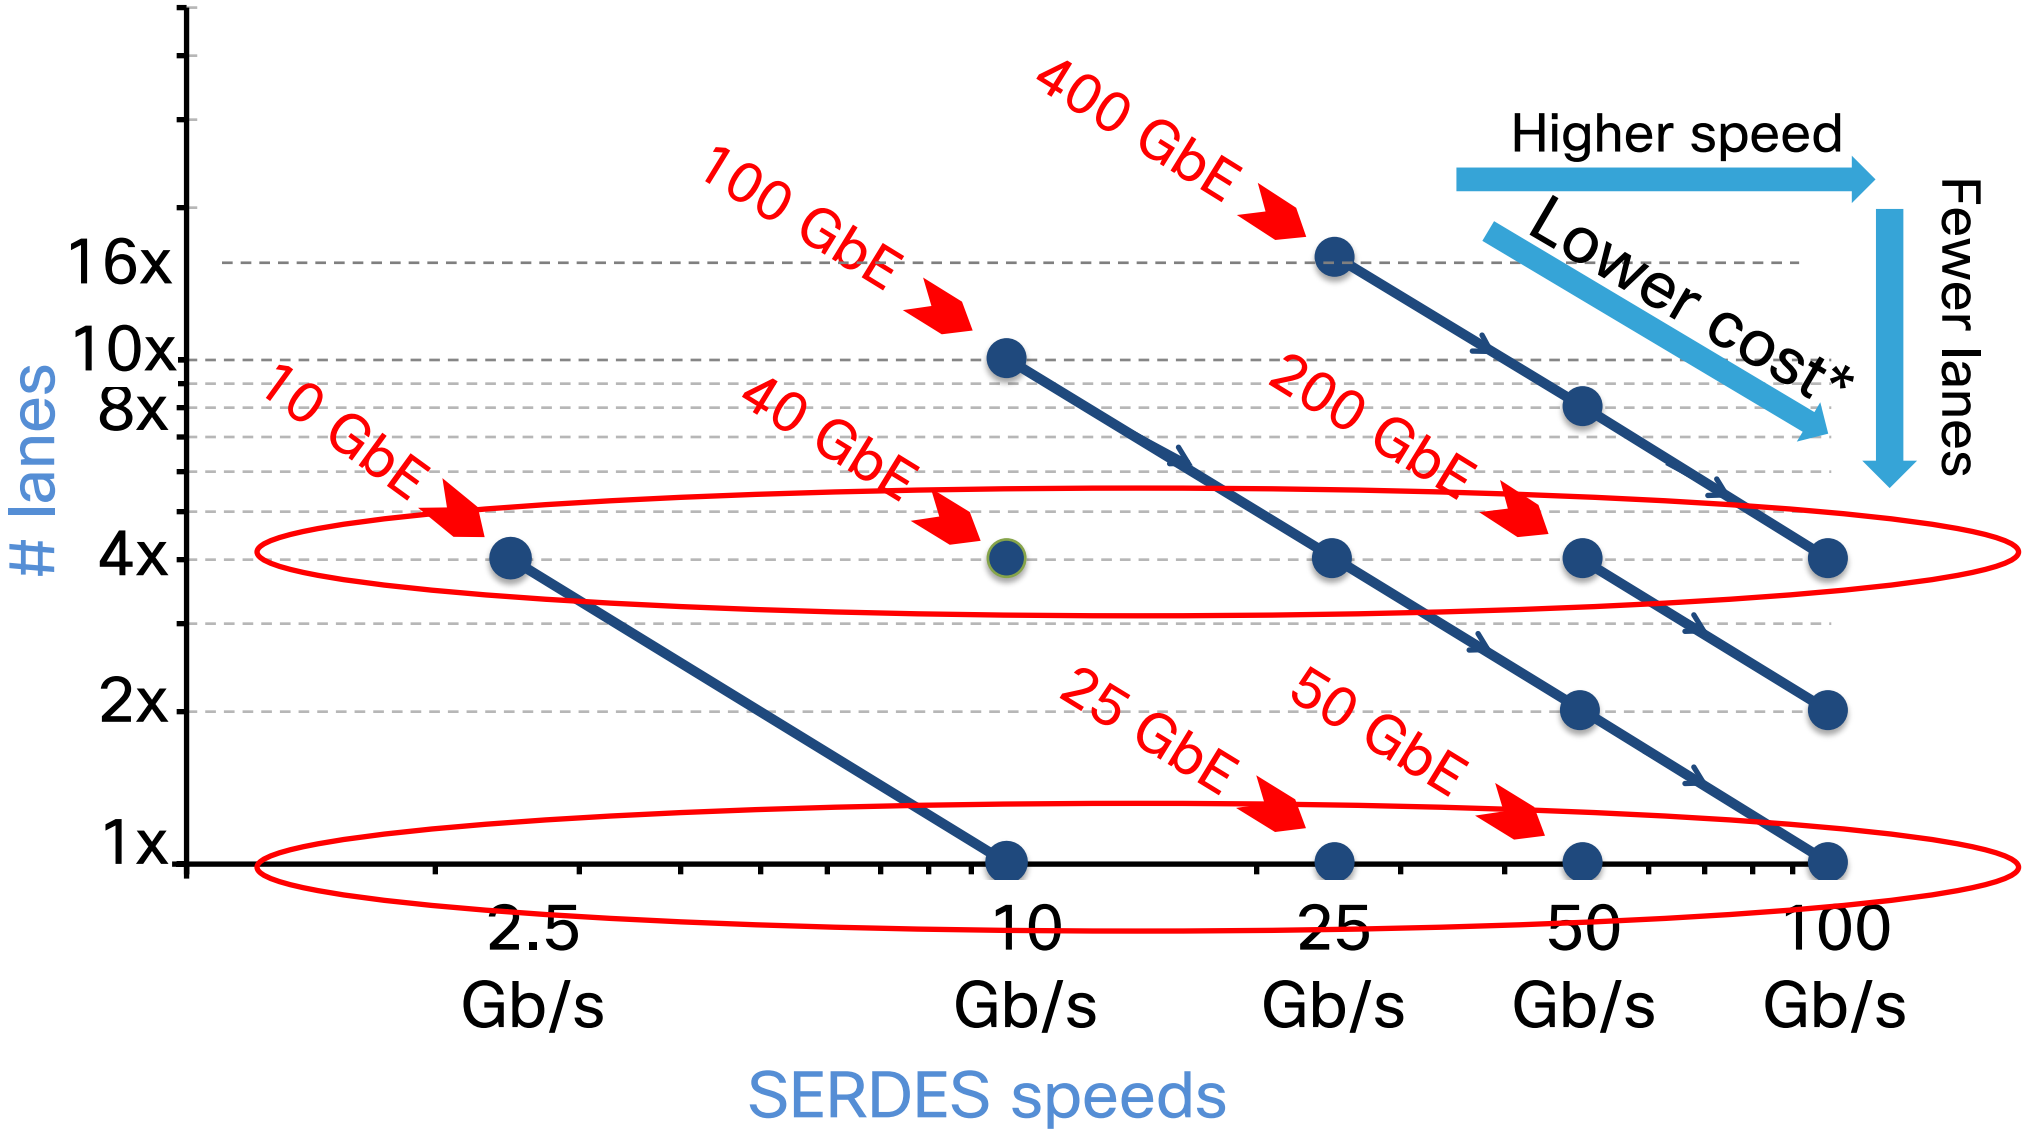

1x and 4x bus widths dominate volume adoption

Reduced complexity leads to lower cost – e.g. 100 GbE

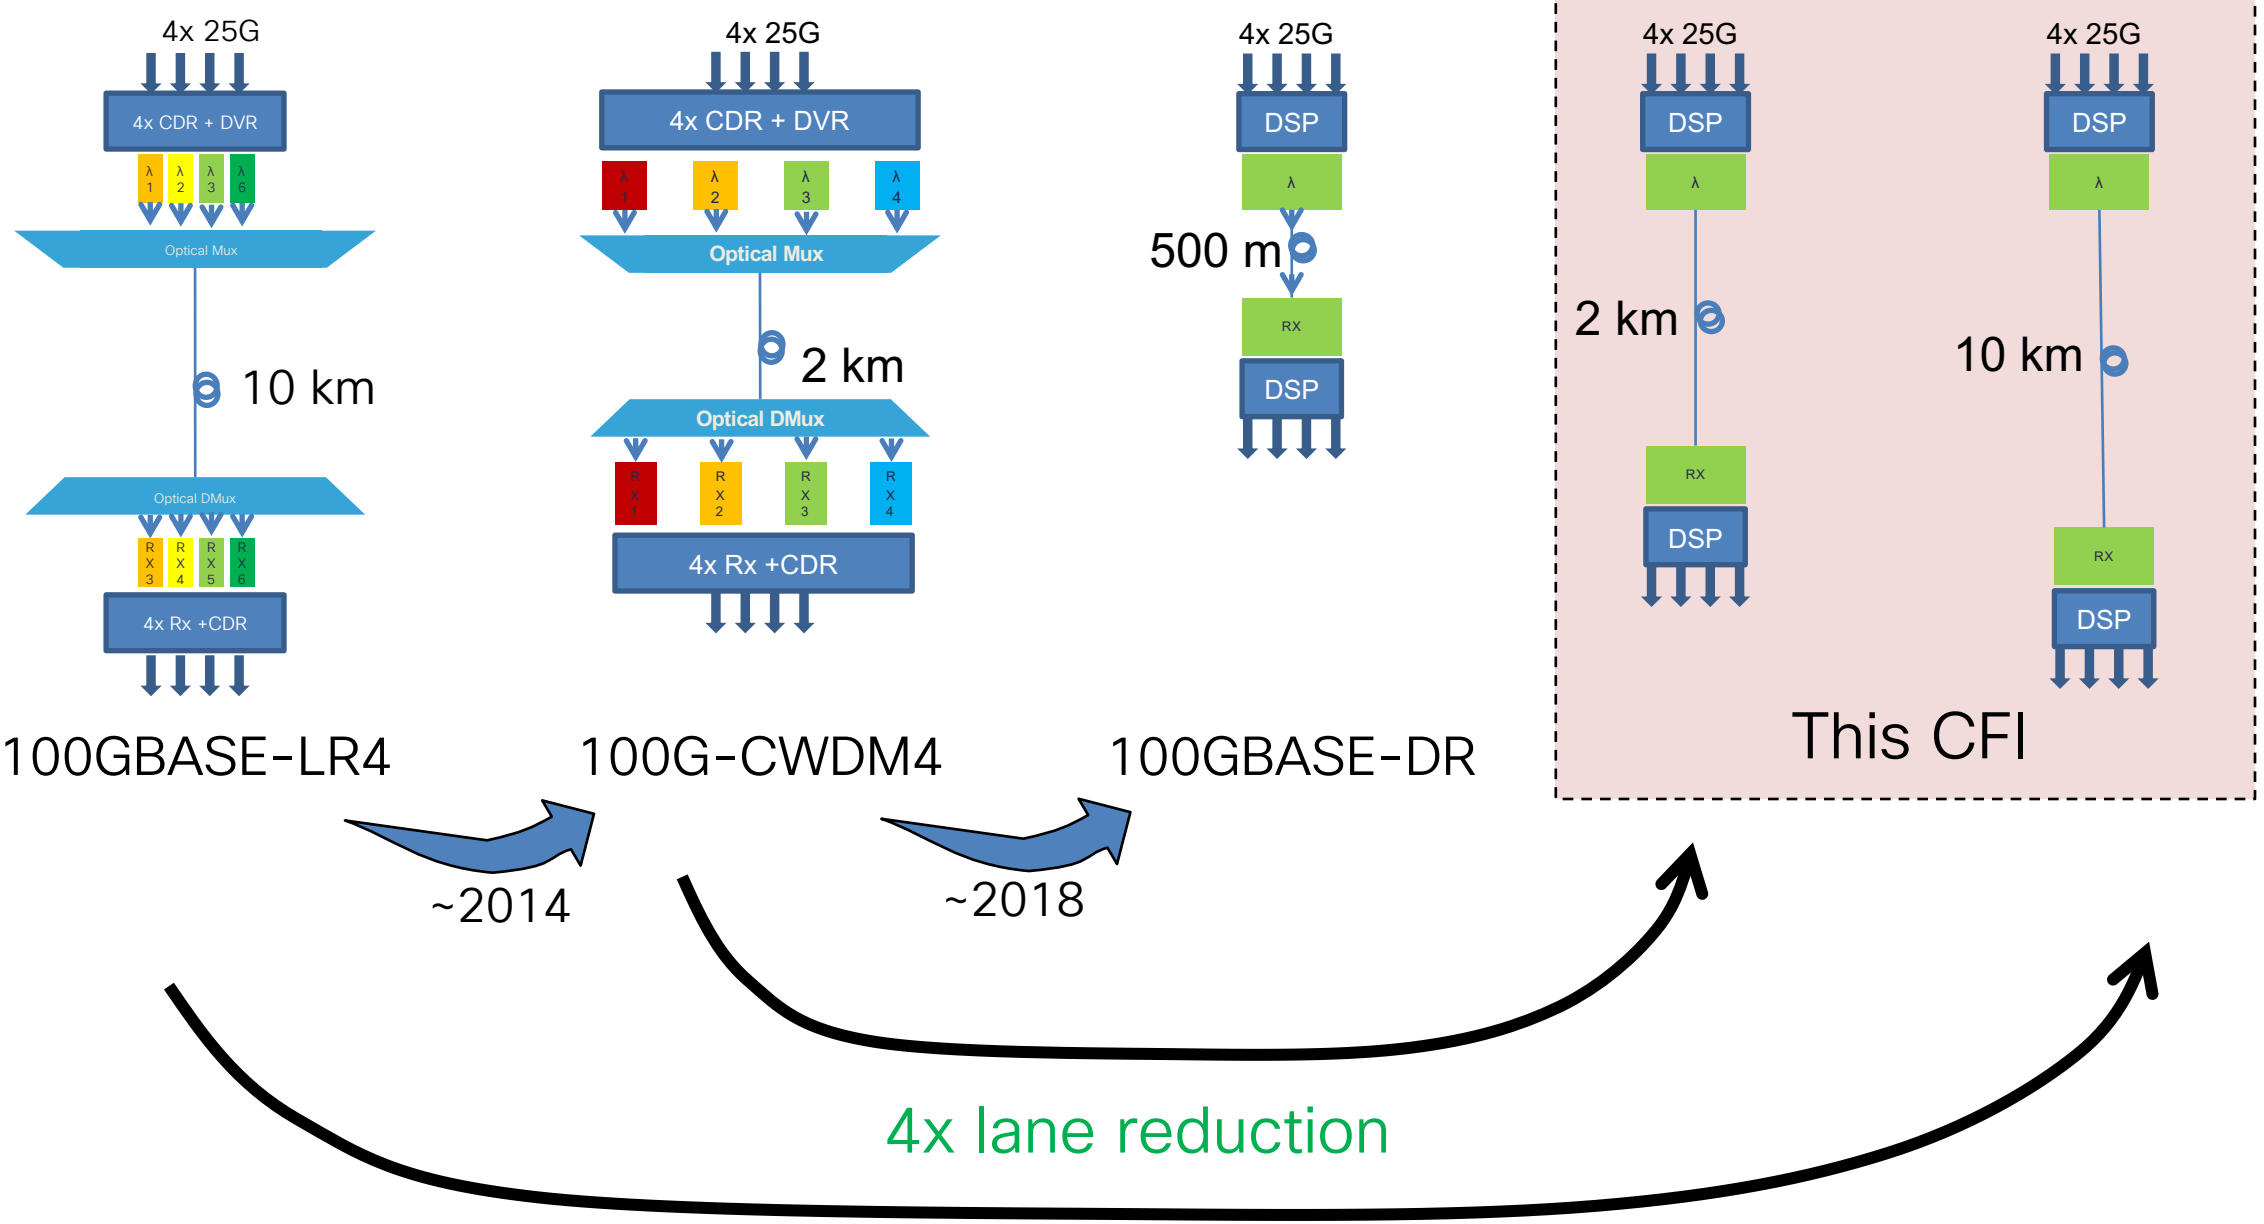

Reduced component count enables denser solutions e.g. 100 GbE

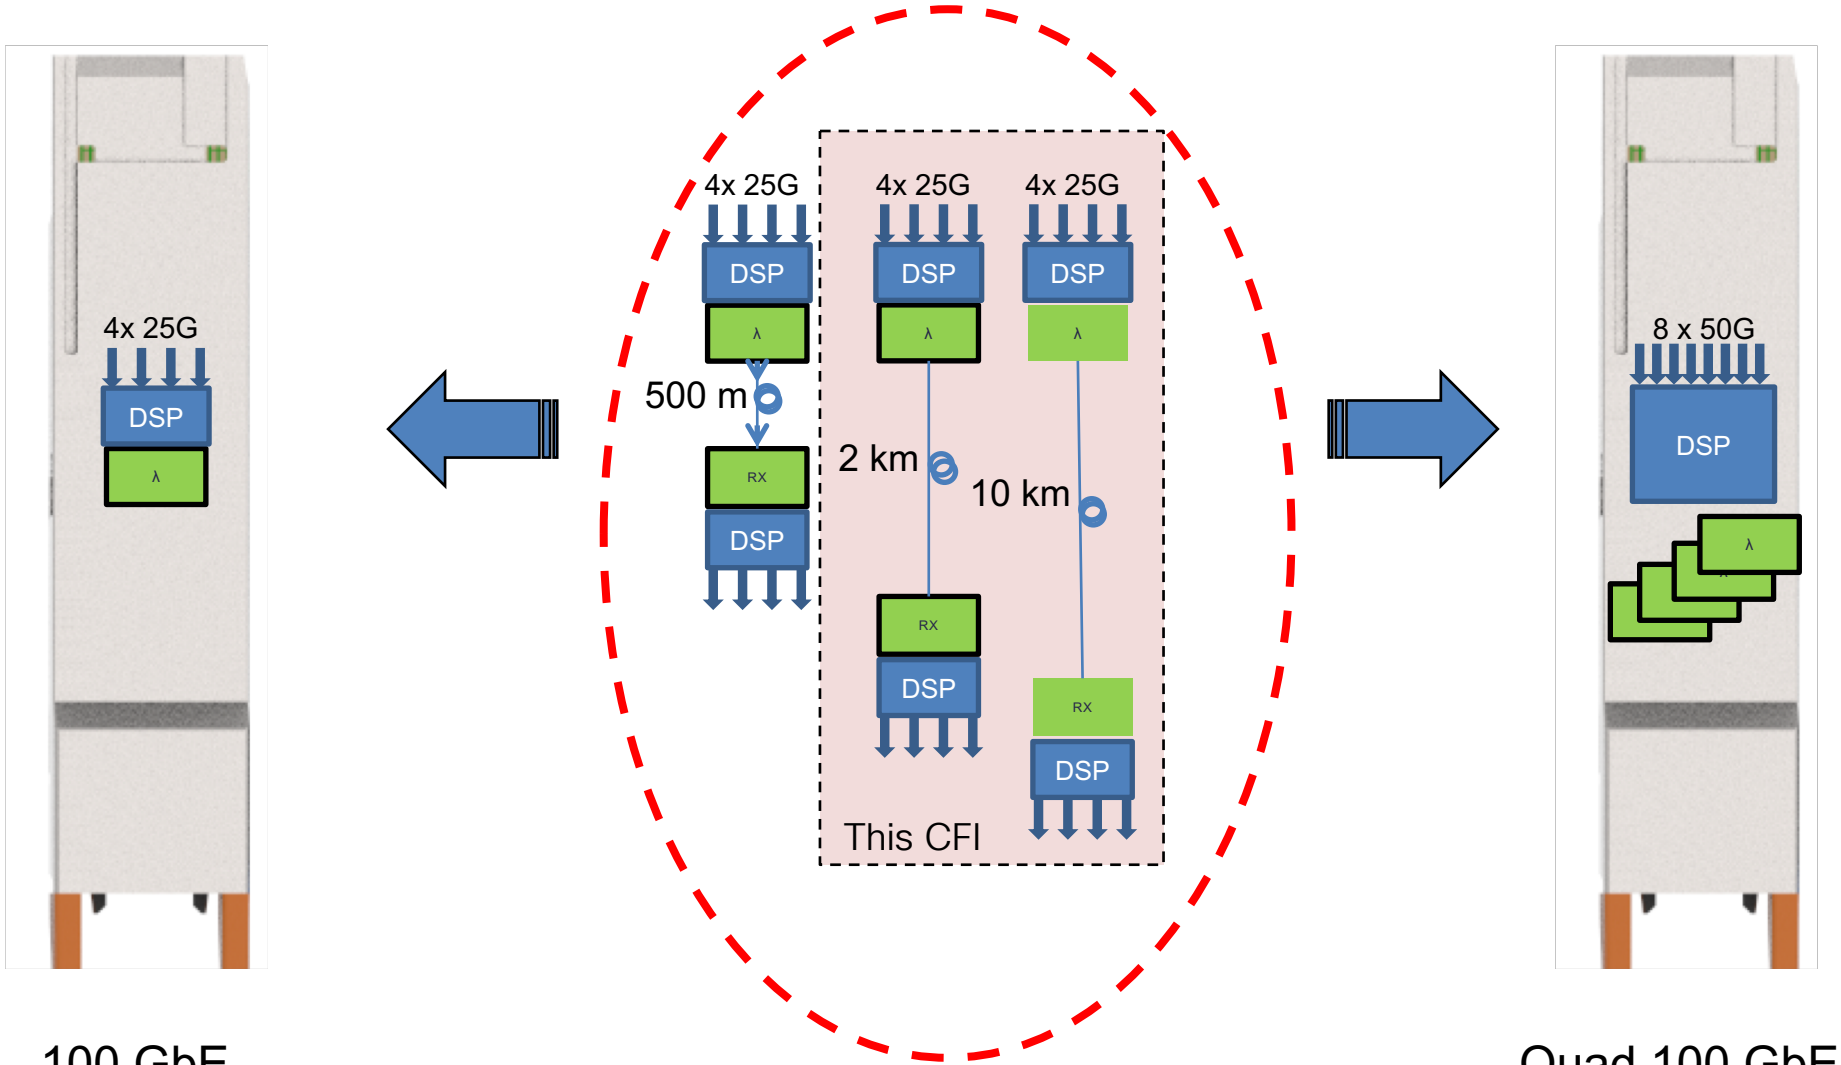

100 GbE
in QSFP28

Quad 100 GbE
in QSFP-DD

400 GbE Duplex SMF Optics – potential complexity reduction

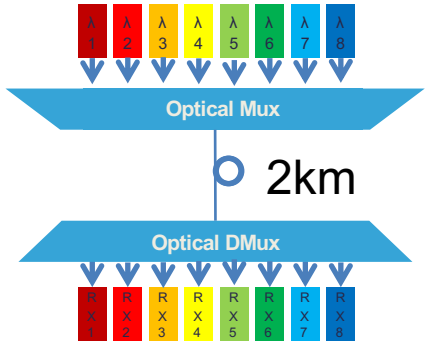

400GBASE-FR8

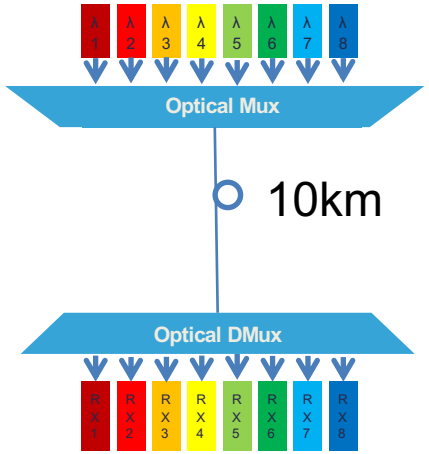

400GBASE-LR8

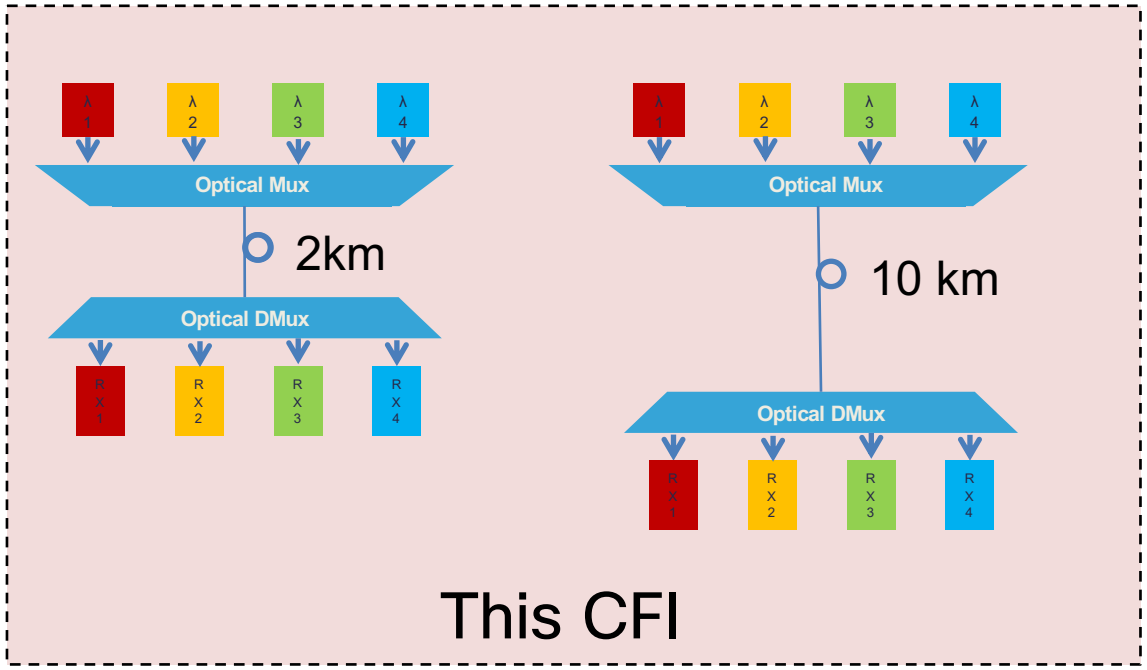

This CFI

2x lane reduction

Moving from 8 lanes to 4 lanes further enables relaxation on wavelength grid to be considered

Market Drivers: Summary

Ethernet has a strong legacy of market success by leveraging newer technology to cost reduce existing solutions

100 GbE SM optics market forecast growth is very strong – lower cost or higher density solutions under demand

400 GbE market adoption about to start in earnest. Significant technology maturity in last few years has led to lower cost solutions than the current Ethernet standardized interfaces being feasible and in demand.

Web-scale, Service Provider and Enterprise Data Centers all identified as potential adopters

Technical Feasibility

IEEE 802.3 Architectural view

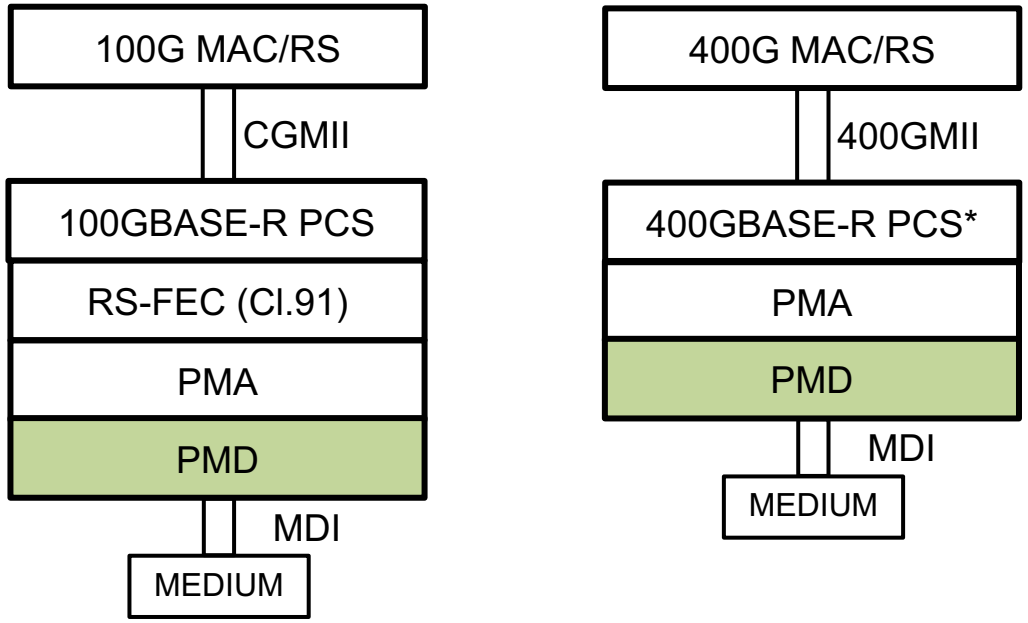

- No architectural changes based on anticipated work and scope of project if approved
- New PMDs to be defined
- No compatibility issues with existing host designs

new

* FEC is part of the 400G PCS sublayer

Industry Progress on 100 Gb/s per lane technology

- 400GBASE-DR4 Completed Dec 2017
- 100GBASE-DR Submitted to RevCom – Oct 2018
- Ethernet Alliance awards “Holy Cup” to 5 companies who were the first to do a public demonstration of 2km SMF 100 Gb/s per lambda in QSFP28 @ ECOC 2018
  - https://twitter.com/EthernetAllianc/status/1044678676799905793
- Multiple public demonstrations of 100 Gb/s per lane technology
  - 100 GbE – 500m, 2km, 10km
  - 400 GbE – 500m, 2km
- See examples on next slides

Matching ASIC IO to Module IO

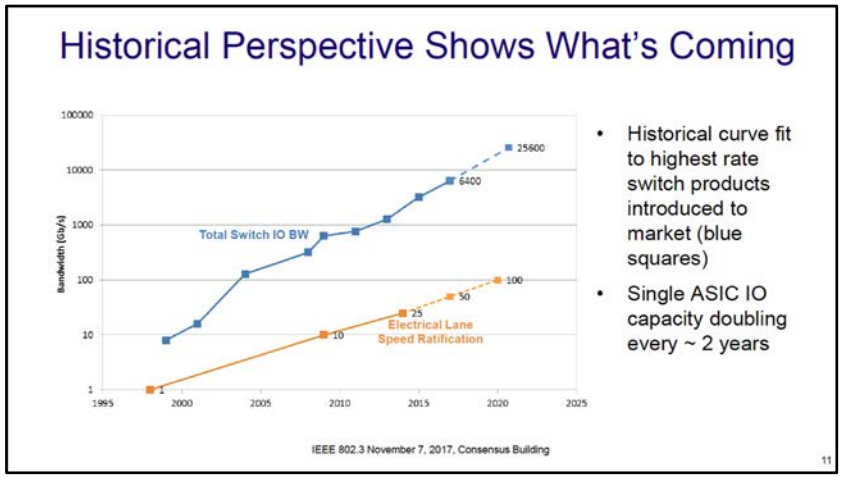

IEEE P802.3ck's CFI:
http://www.ieee802.org/3/cfi/1117_3/CFI_03_1117.pdf

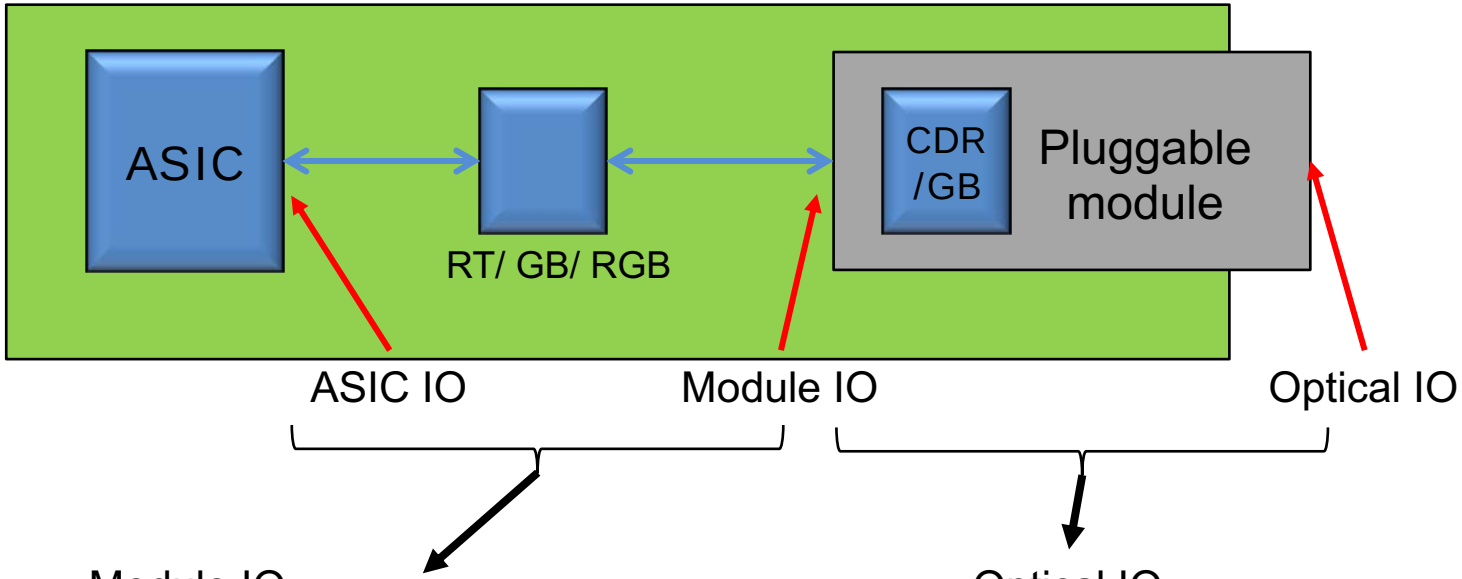

- ASIC IO "needs" to increase
- Module IO "advantage" to match ASIC IO (no mandatory extra host device)
- Optical module simplified when Optical IO matches Module IO

Module IO

ASIC IO	25 Gb/s	50 Gb/s	100 Gb/s
25 Gb/s	RT	GB	GB
50 Gb/s	RGB	RT	GB
100 Gb/s	RGB	RGB	RT

Optional
(vs. Mandatory)

Optical IO

Module IO	25 Gb/s	50 Gb/s	100 Gb/s
25 Gb/s	CDR	CDR/GB	CDR/GB
50 Gb/s	CDR/RGB	CDR	CDR/GB
100 Gb/s	CDR/RGB	CDR/RGB	CDR

Simplest

RT: Retimer GB: Gearbox RGB: Reverse Gearbox

Extending the reach

Link budgets that would extend beyond the current 500m specifications need to deal with:

- Extra fiber loss
- Extra wavelength mux/demux loss (400 GbE 2 & 10km only)
- Extra dispersion penalty

Options to address include:

- Increased launch power
- Increased receiver sensitivity (including PIN or APD)
- Wavelength grid (CWDM vs. LWDM)

Technical feasibility – Transmitters

Various transmitters capable of 100 Gb/s PAM4 have been demonstrated or presented

TDECQ = 1.26 dB (2.2 km @ 5.2 ps/nm)

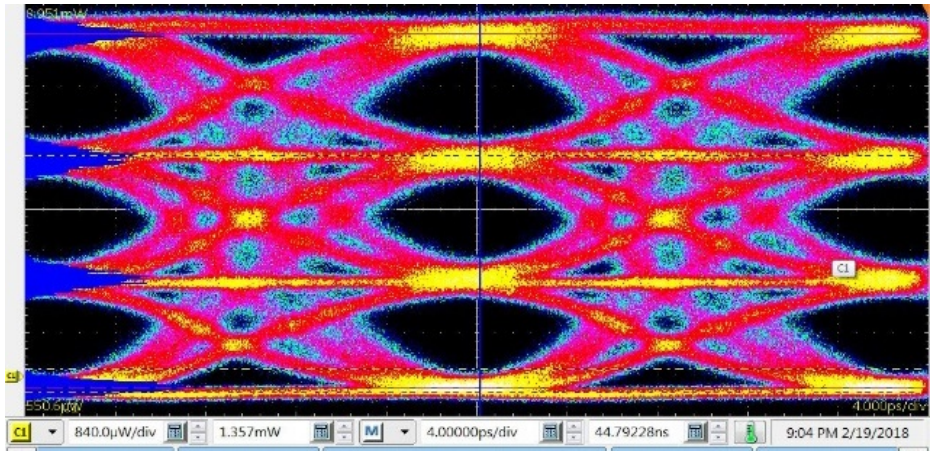

53 GBaud PAM4 (106 Gb/s)

IEEE Pattern PRBS13Q See – mazzini_3cd_01a_0518

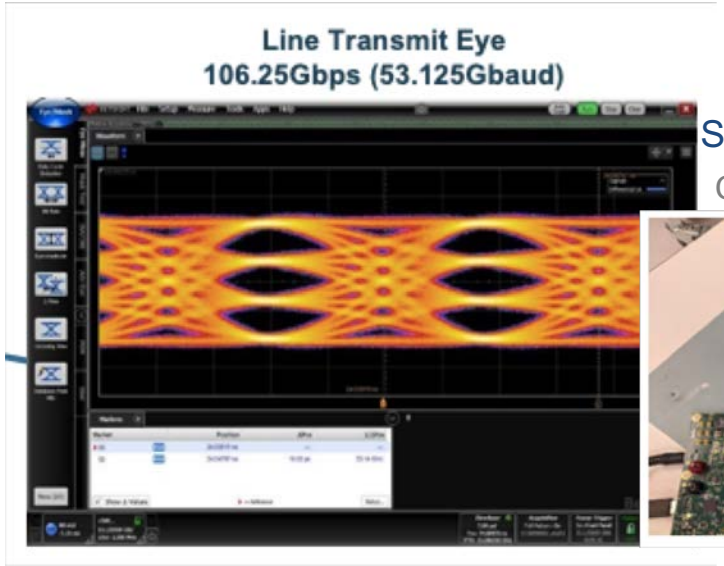

Electrical Loopback
SNR >24dB, BER < 1e-12

Courtesy Inphi

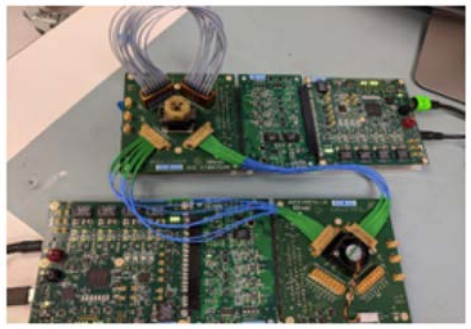

Courtesy Broadcom

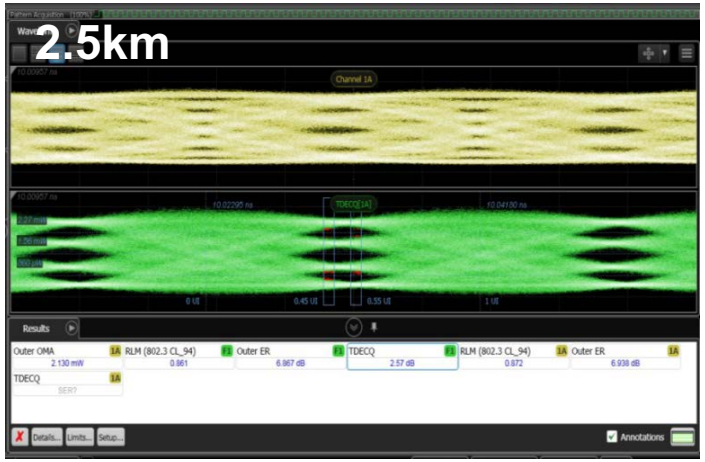

2.5km Penalty = 0.22dB

ER=5.5dB, TDECQ= 2.64dB

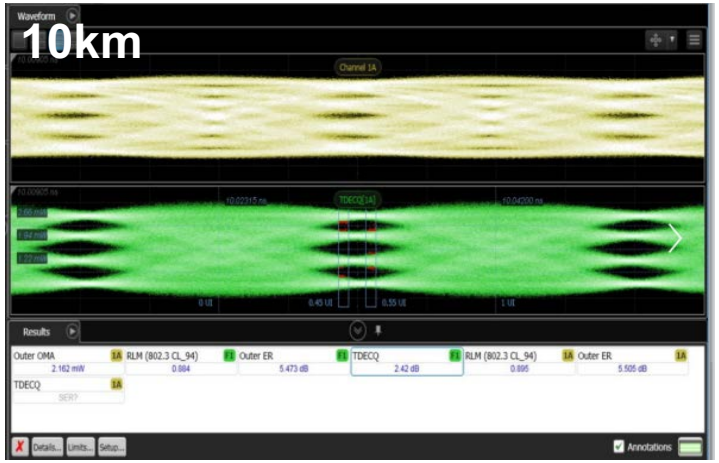

10km Penalty = 0.21dB

19 ps/nm

ER=6.9dB, TDECQ= 2.78dB

Transmitter:

AWG + linear amplifier, Vpp = 1.2V
No emphasis applied at the AWG
SSPRQ pattern at 53 GBd
56GBd EML CoC, λ **= 1330nm**

Technical feasibility – Receivers Both PIN and APD based detectors available

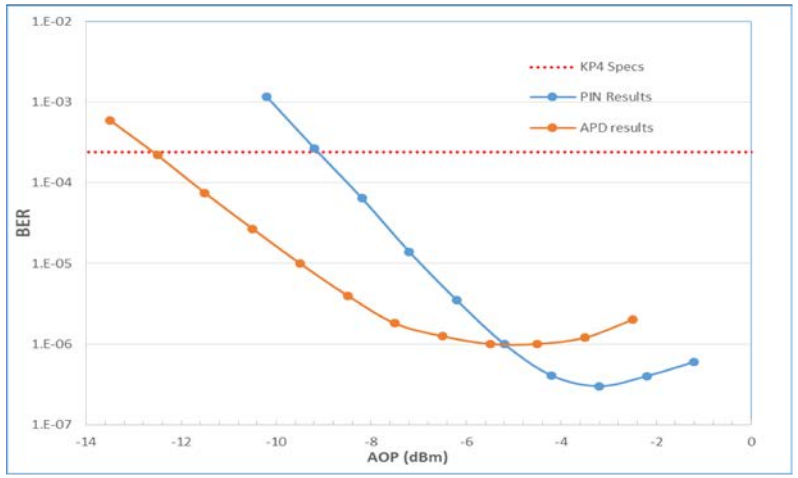

- 53.125GBaud, PAM4 PRBS31Q EML CoC (1304nm), ER ~ 5.1dB, SECQ ~ 1.8dB
- Same TIA, Same DSP

Source: IEEE OI'2018, Santa Fe, 4-6 June 2018 (Inphi/Source)
https://ieee-oi.org/program/

400G - 2km CWDM

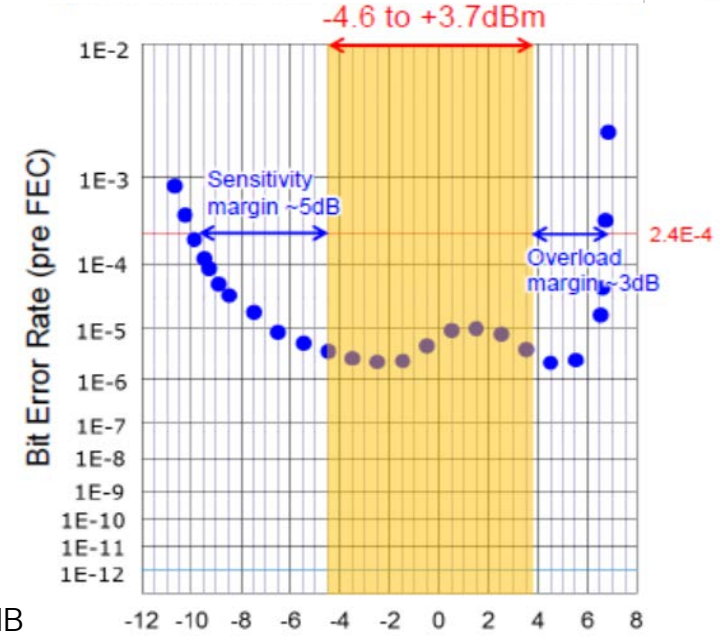

53.125Gbd
PRBS15Q
ER=5.5dB
TDECQ1.6dB

Courtesy: Sumitomo Electric

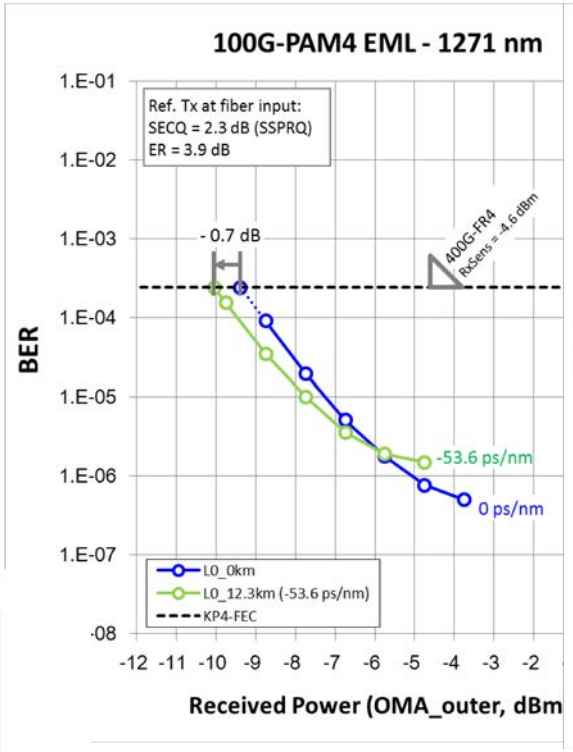

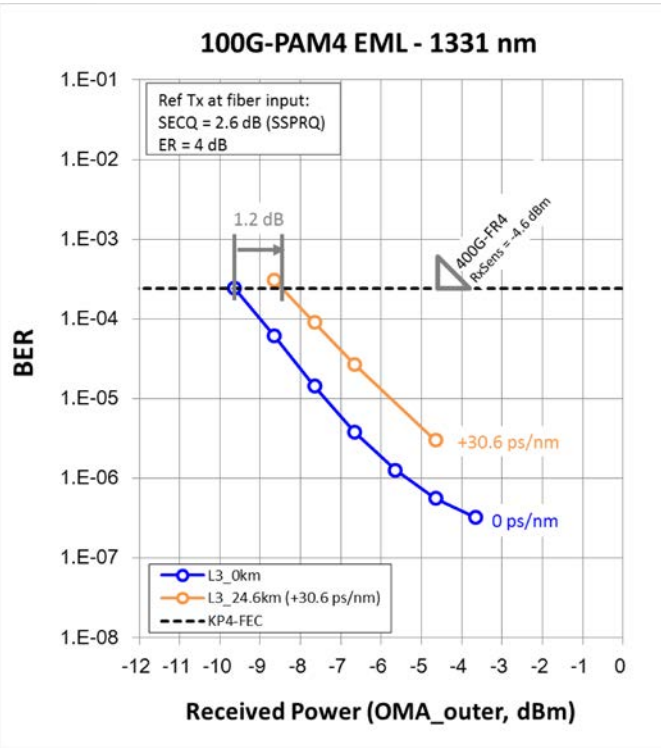

	GVD (ps/nm)
IEEE MIN Spec for L0, 10km	-59.4
Measured (1271 nm, 12.3 km)	-53.6

	GVD (ps/nm)
IEEE MAX Spec for L3, 10km	33.4
Measured (1331 nm, 24.6 km)	30.6

Courtesy: Oclaro

Experimental Configuration:
PRBS15 53.125 Gbaud PAM4
CWDM EML CWDM PIN-PD
DSP (FFE>5 taps)
SM fiber: L0 (12.3 km) & L3 (24.6 km)

Technical feasibility – Test & Measurement

TDECQ Test methodology is solid and numerous test solutions available

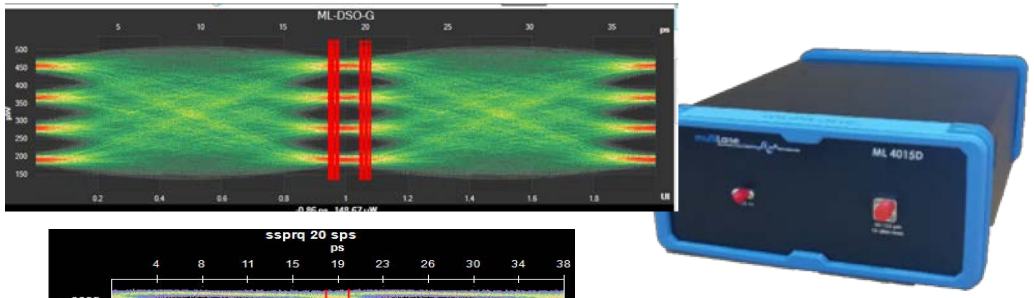

Courtesy Multilane

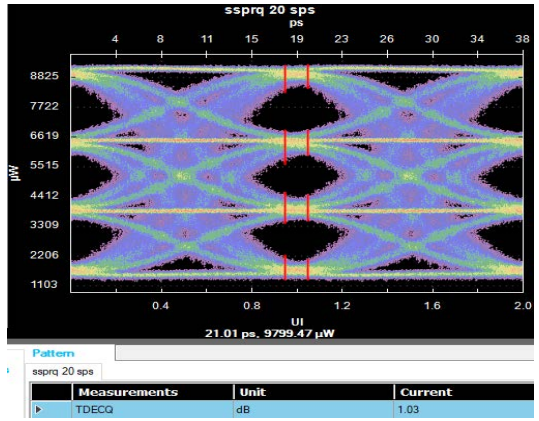

Courtesy Keysight

TDECQ = 1.5015 dB
FFE Taps = 1.0138 0.0059905 0.0016472 -0.02332 0.0019147
OMA Outer = 0.97727mW

Courtesy Tektronix

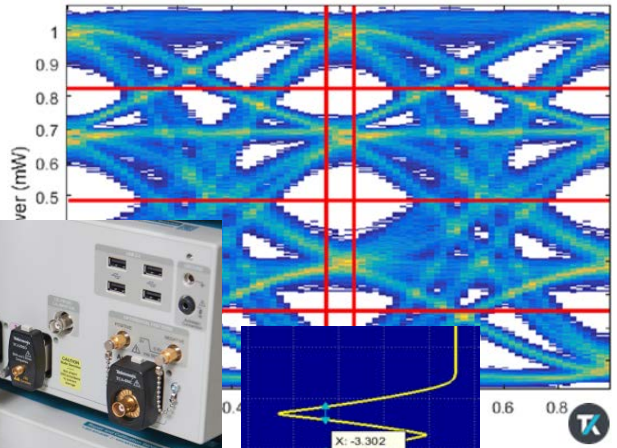

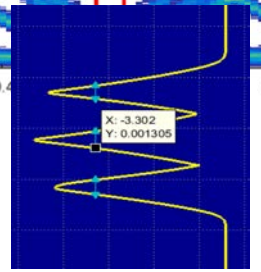

Excellent SECQ vs. Rx Sensitivity correlation demonstrated (50 Gb/s in this published example)

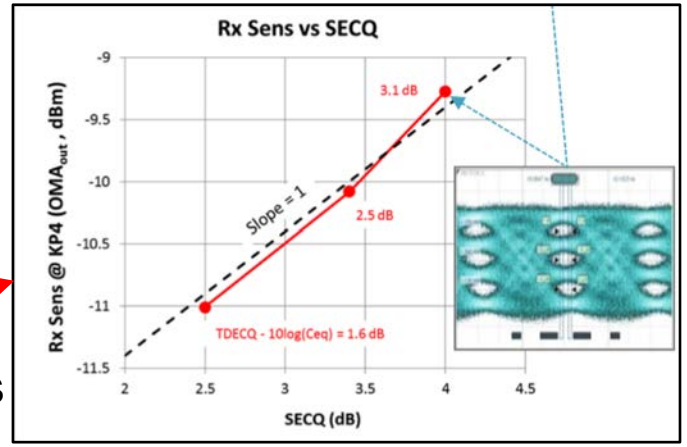

tamura_3cd_01c_0718.pdf

Why Now?

Why Now?

- Technical developments underway already to extend 100 Gb/s per lane technology to longer reaches
- Current IEEE Ethernet solutions not fully aligned with end user demand,
  - especially Web-scale Data Centers looking for solutions based on 100 Gb/s per lane technology
- Technical feasibility demonstrations happening
- Standardization in IEEE 802.3 brings industry convergence and extends Ethernet's solution breadth
- Target markets are:
  - Moving into high volume and therefore cost sensitive (e.g. 100 GbE)
  - Initiating early adoption that cost reduction will accelerate (e.g. 400 GbE)

Straw Polls

Supporters

- Your name could be here – contact Mark!

Straw Poll 1: Call-For-Interest

- Should a Study Group be formed to consider extending 100 Gb/s PAM4 per lane optical technology to longer 100 Gb/s and 400 Gb/s Ethernet reaches up to 10 km?

  Y: N: A:

Room Count:

Participation

- I would participate in the “100G Lambda*” Study Group in IEEE 802.3.

Tally:

- My company would support participation in the “100G Lambda*” Study Group in IEEE 802.3.

  Tally:

* Extending 100 Gb/s PAM4 per lane optical technology to longer 100 Gb/s and 400 Gb/s Ethernet reaches

Supporters (your name here too – send me an email)

Justin Abbott, Lumentum
Pete Anslow, Ciena
Vittal Balasubramanian, Innovium
Thananya Baldwin, Keysight
Vipul Bhatt, Finisar
Jose Castro, Panduit
Frank Chang, Source Photonics
David Chen, Applied Optoelectronics
Chris Cole, Finisar
Arash Farhoodfar, Inphi
Ali Ghiasi, Ghiasi Quantum
Paul Goldgeier, ColorChip
Jonathan Ingham, Foxconn Optical Interconnect
Hideki Isono, Fujitsu
Kenneth Jackson, Sumitomo Electric
John Johnson, Broadcom
Mark Kimber, Semtech
Jonathan King, Finisar
Greg Lecheminant, Keysight
David Lewis, Lumentum
Hai-Feng Liu, Intel
Karen Liu, Lightwave Logic
Khushrow Machhi , Broadcom
Jeffery Maki, Juniper
Marco Mazzini, Cisco
Shirao Mizuki, Mitsubishi Electric
Ray Nering, Cisco
Gary Nicholl, Cisco
Tom Palkert, MACOM
Rajiv Pancholy, Broadcom
Jerry Pepper, Keysight
Rick Pimpinella, Panduit
Kees Propstra, Multilane
Rick Rabinovich, Keysight
Scott Schube, Intel
Kapil Shrikhande, Innovium
Scott Sommers, Molex
Phil Sun, Credo
Mike Takefman, Inphi
Pirooz Tooyserkani, Cisco
Matt Traverso, Cisco
Eddie Tsumura, Sumitomo Electric
Jeff Twombly, Credo
Ed Ulrichs, Source Photonics
Mike Wang, HiSense
Brian Welch, Luxtera
Chongjin Xie, Alibaba
Simon Ximen, ColorChip
Ryan Yu, Molex
Hua Zhang, HiSense
Kevin Zhang, IDT
Pavel Zivny, Tektronix