

# 100 Gb/s per Lane for Electrical Interfaces and PHYs

CFI Consensus Building

Beth Kochuparambil (Cisco Systems)  
John D'Ambrosia (Futurewei, Subsidiary  
of Huawei)

Kent Lusted (Intel)  
Adam Healey (Broadcom)  
David Ofelt (Juniper)

# Objective

- Build consensus of starting a study group investigating a “100 Gb/s per Lane for Electrical Interfaces and PHYs” project
- We do **not** need to:
  - Fully explore the problem
  - Debate strengths and weaknesses of solutions
  - Choose a solution
  - Create a PAR or 5 Criteria
  - Create a standard
- Anyone in the room may vote or speak

# Introductions for today's presentation

John D'Ambrosia, FutureWei Technologies  
& Beth Kochuparambil, Cisco Systems, Inc.

Intro

David Ofelt, Juniper Networks

Market Drivers

Adam Healey, Broadcom

Technical Feasibility

Beth Kochuparambil, Cisco Systems, Inc.

Why Now? & Close

Panel, incl. Kent Lusted, Intel

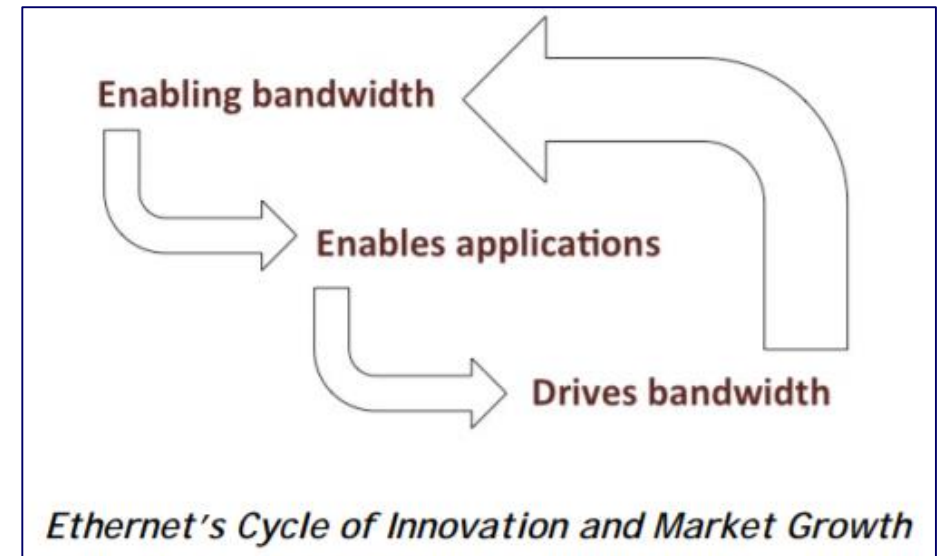
Q&A

# Motivation for 100 Gb/s per Lane

With next steps in Ethernet, comes the needed next step in interfaces.













- Faceplate density
- Chip breakout
- System throughput

They are all tied together!



\*Web-scale data centers and cloud based service are presented as leading applications

# Electrical interfaces come in many shapes and sizes.



IEEE 802.3 November 7, 2017, Consensus Building

Copper lives on...  
we may need to  
open our minds to  
what a channel is!

# Tonight's Meeting

- To present the  
market ***NEED***,  
technical **Feasibility**,  
and *Why Now??*  
of 100Gb/s per lane of electrical signaling.
- To gain consensus towards Thursday's motion to form a study group.
- We are NOT discussing specific implementations or objectives – these are just some of the reasons that we need a study group!

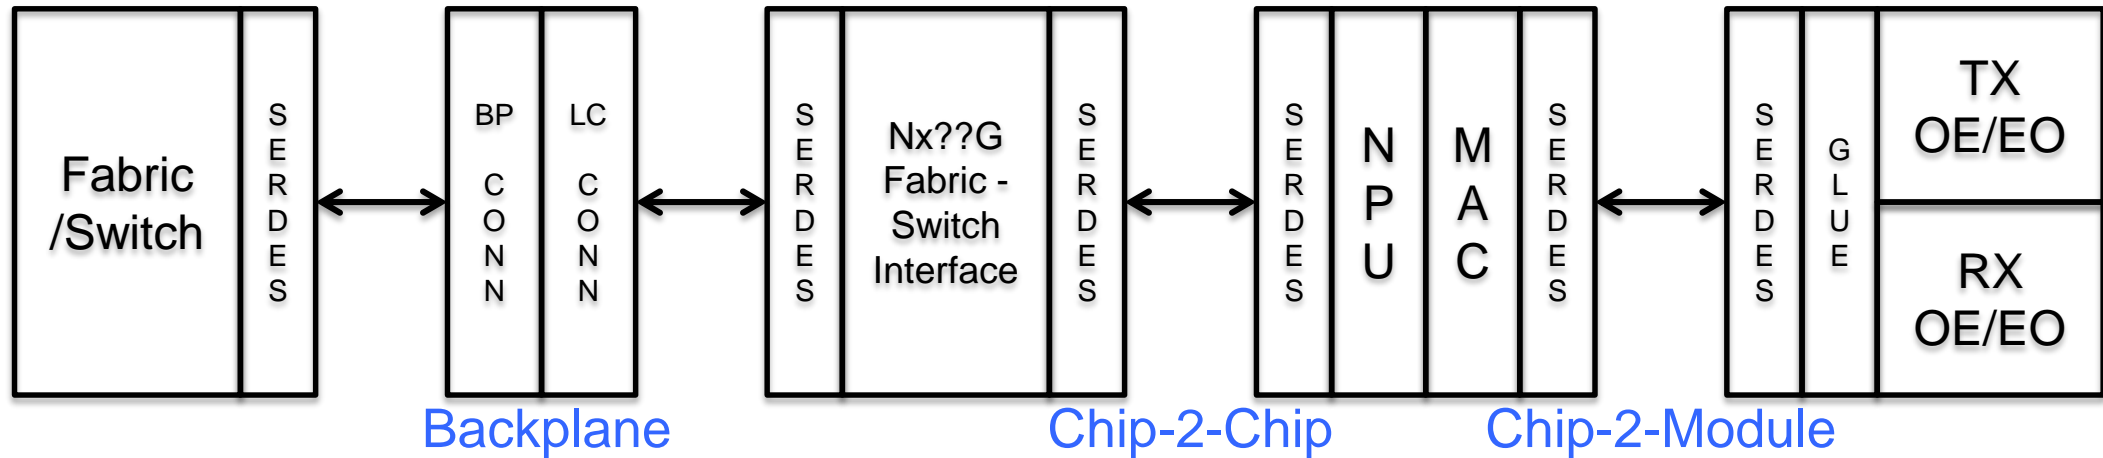


# Market Drivers


for 100 Gb/s per lane for Electrical Interfaces

Go *Faster*  
to  
Go **Denser**  
to  
Continue to **Grow**.

# What Are We Talking About?




Consider how many instances of the interface can exist in the system...





# Why Copper Cable??



Need to study highly cost sensitive and very short reach market.

## Interconnection Volume

- Four sections per colo & multiple colos ( $\geq 4$ ) per data center
- Volumes below are per section (except DCR to Metro)

A End	Z End	Volume	Reach (max)	Medium	Cost Sensitivity	Market Space
Server ‡	TOR	10k – 100k	3 m	Copper	Extreme	LAN
TOR	LEAF	1k – 10k	20 m	Fiber (AOC)	High	
LEAF	SPINE	1k – 10k	400 m	SMF	High	Campus
SPINE	DCR	100 – 1000	1,000 m	SMF	Medium	
DCR	Metro	100 – 300	10 - 80 km	SMF	Low	WAN

‡ Server-TOR links may be served by breakout cables

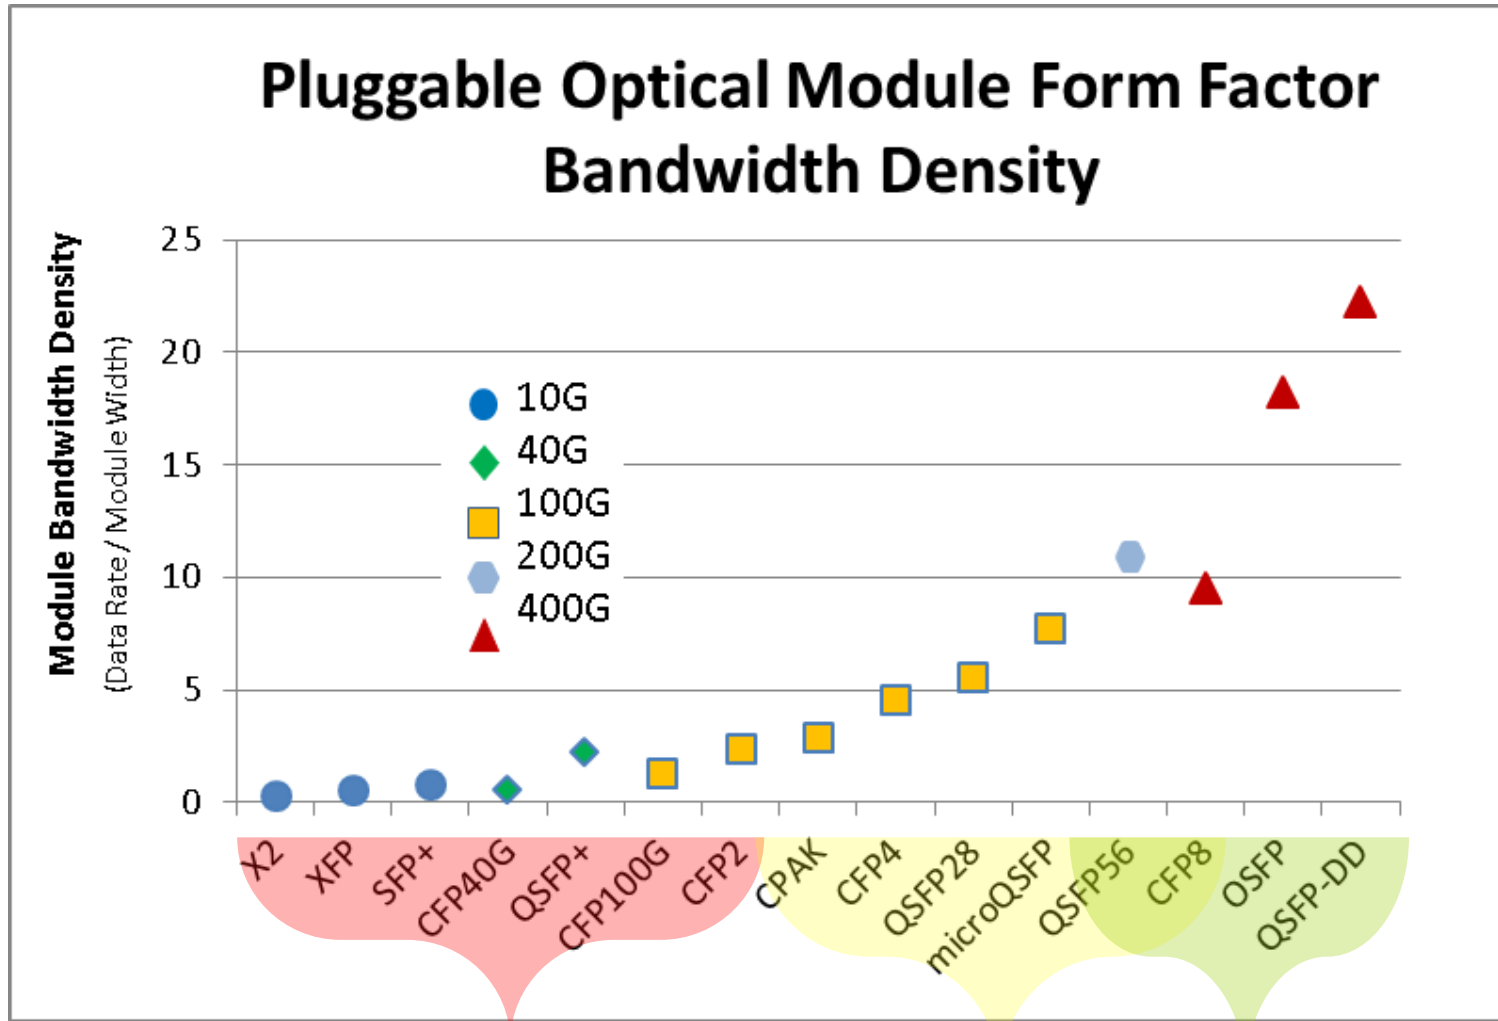
9

IEEE 802.3 400G Study Group - November 2013

Source: Brad Booth, Microsoft [http://www.ieee802.org/3/400GSG/public/13\\_11/booth\\_400\\_01a\\_1113.pdf](http://www.ieee802.org/3/400GSG/public/13_11/booth_400_01a_1113.pdf)

\*Note that data is from 2013, however data center architecture hasn't drastically changed in recent years

# Faceplate Evolution



Primary Electrical Interface:

IEEE 802.3 November 7, 2017, Consensus Building

Increased faceplate density

\*\*\*\*\*

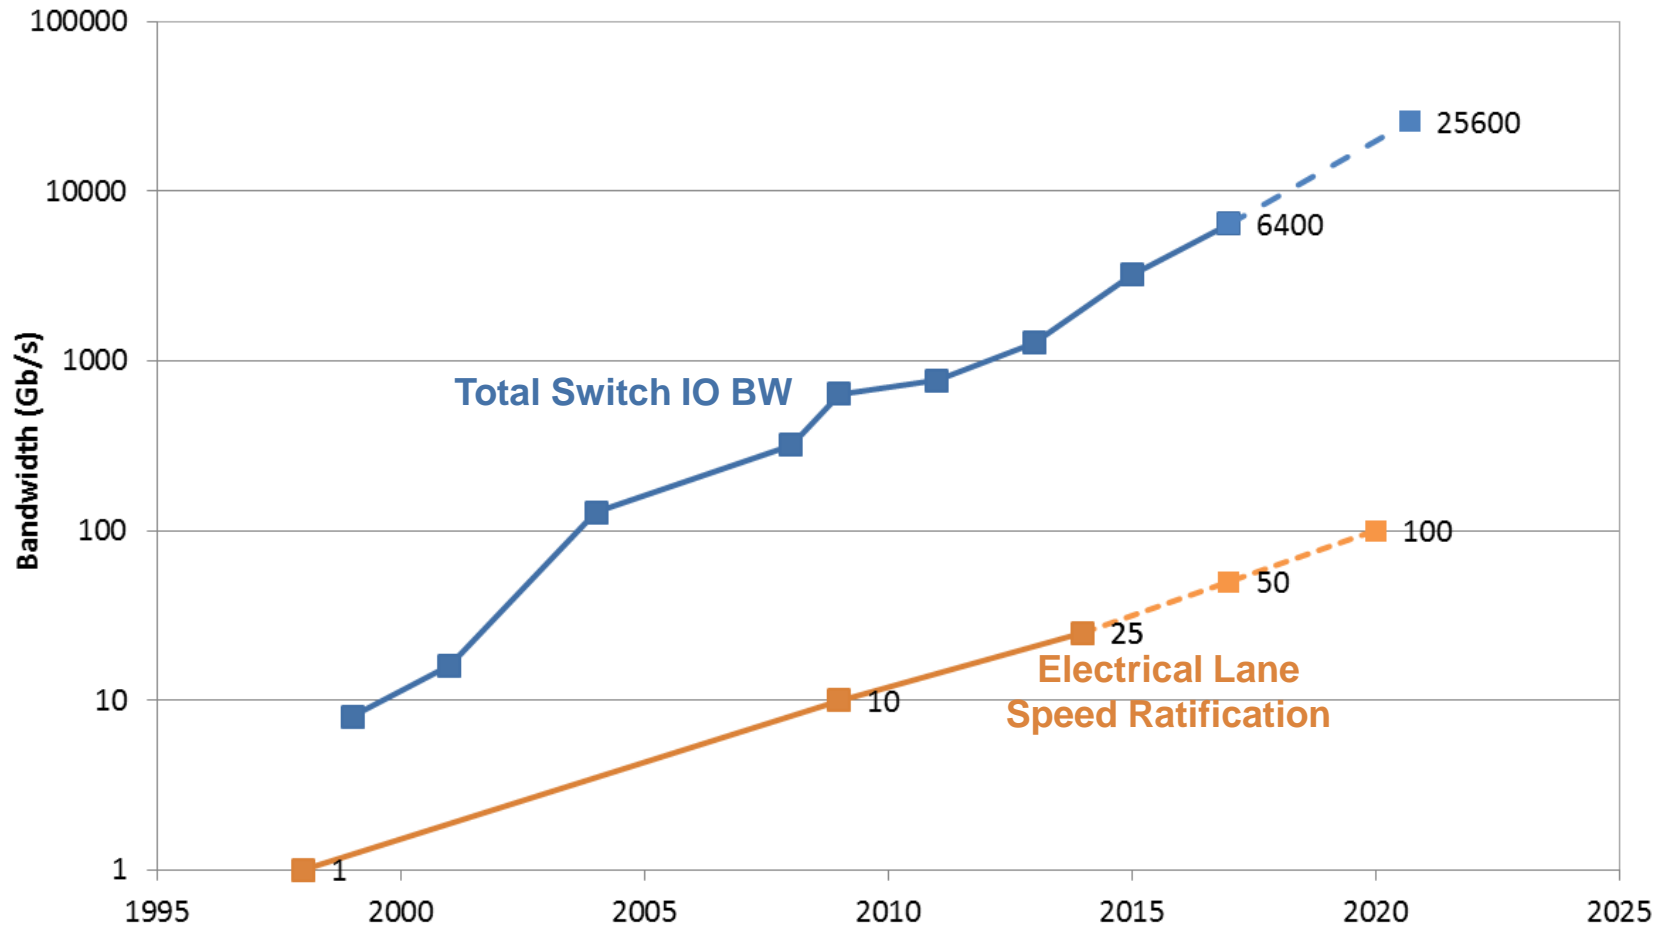
Requires smaller form factors

\*\*\*\*\*

Which is enabled by faster AUIs (per lane speed)

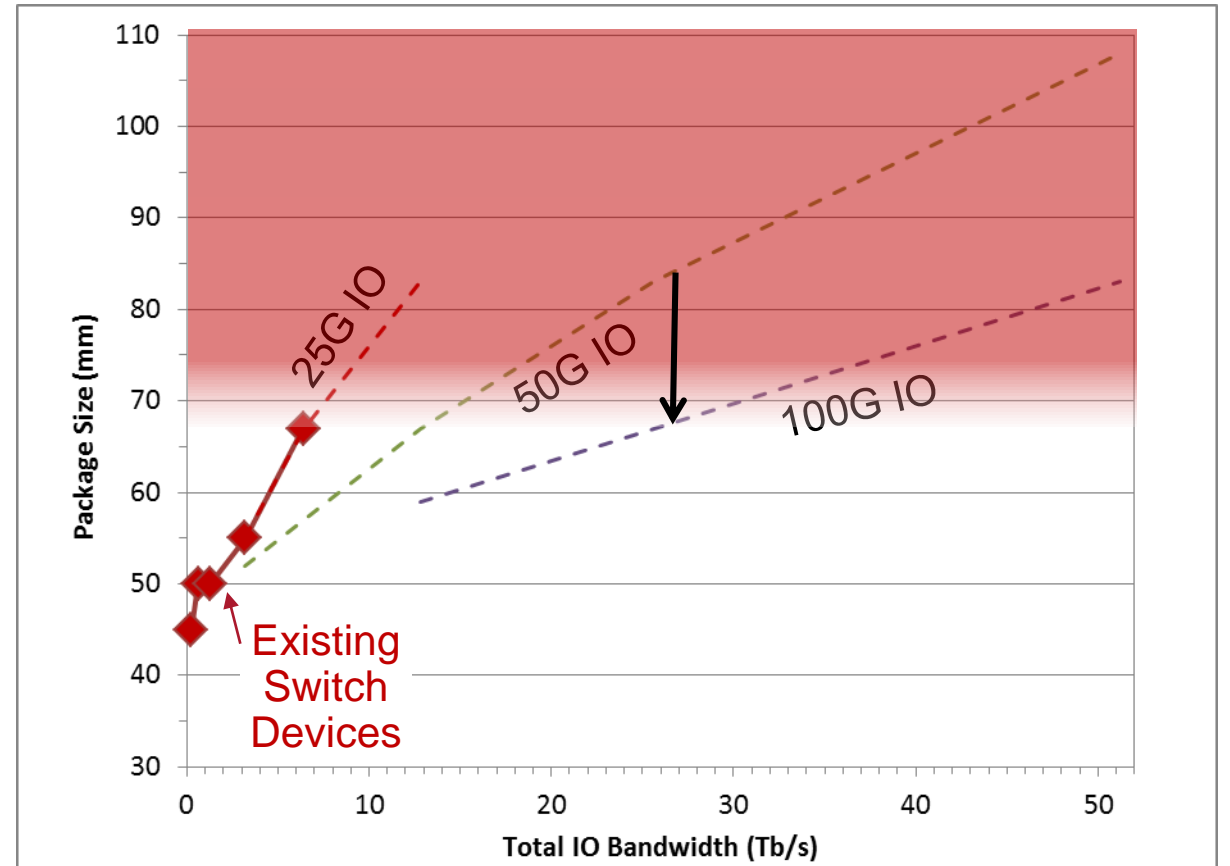

Source of graph:  
<http://www.qsfp-dd.com/wp-content/uploads/2017/03/QSFP-DD-whitepaper-15.pdf>

# Historical Perspective Shows What's Coming



- Historical curve fit to highest rate switch products introduced to market (blue squares)
- Single ASIC IO capacity doubling every ~ 2 years

# IO Escape forcing transition to higher lane speeds




- ~ 70mm package is a current BGA practical maximum (due to coplanarity / warpage)
- BGA devices with > 14Tb/s of aggregate bandwidth are forced to transition to lane rates beyond 50G

# Backplane is easily system bottleneck

Finite space created by line card size and card pitch

Example: 52-55 diff. pairs per inch



# The Current Ethernet Family (100 Gb/s and Above)

	Signaling (Gb/s)	Electrical Interface	Backplane	Twin-ax	MMF	500m SMF	2km SMF	10km SMF	40km SMF
100GBASE-	10	CAUI-10		CR10	SR10		<u>10X10</u>		
	25	CAUI-4 / 100GAUI-4	KR4	CR4	SR4	<u>PSM4</u>	<u>CWDM4 CLR4</u>	LR4	ER4
	50	100GAUI-2	KR2	CR2	SR2		-		
	100	?	?	?		DR			
200GBASE-	25	200GAUI-8							
	50	200GAUI-4	KR4	CR4	SR4	DR4	FR4	LR4	
	100	?	?	?					
400GBASE-	25	400GAUI-16			SR16				
	50	400GAUI-8					FR8	LR8	
	100	?	?	?		DR4			

Includes Ethernet standards in development

Underlined – indicates industry MSA or proprietary solutions

Blue – indicates the areas of interest for this CFI



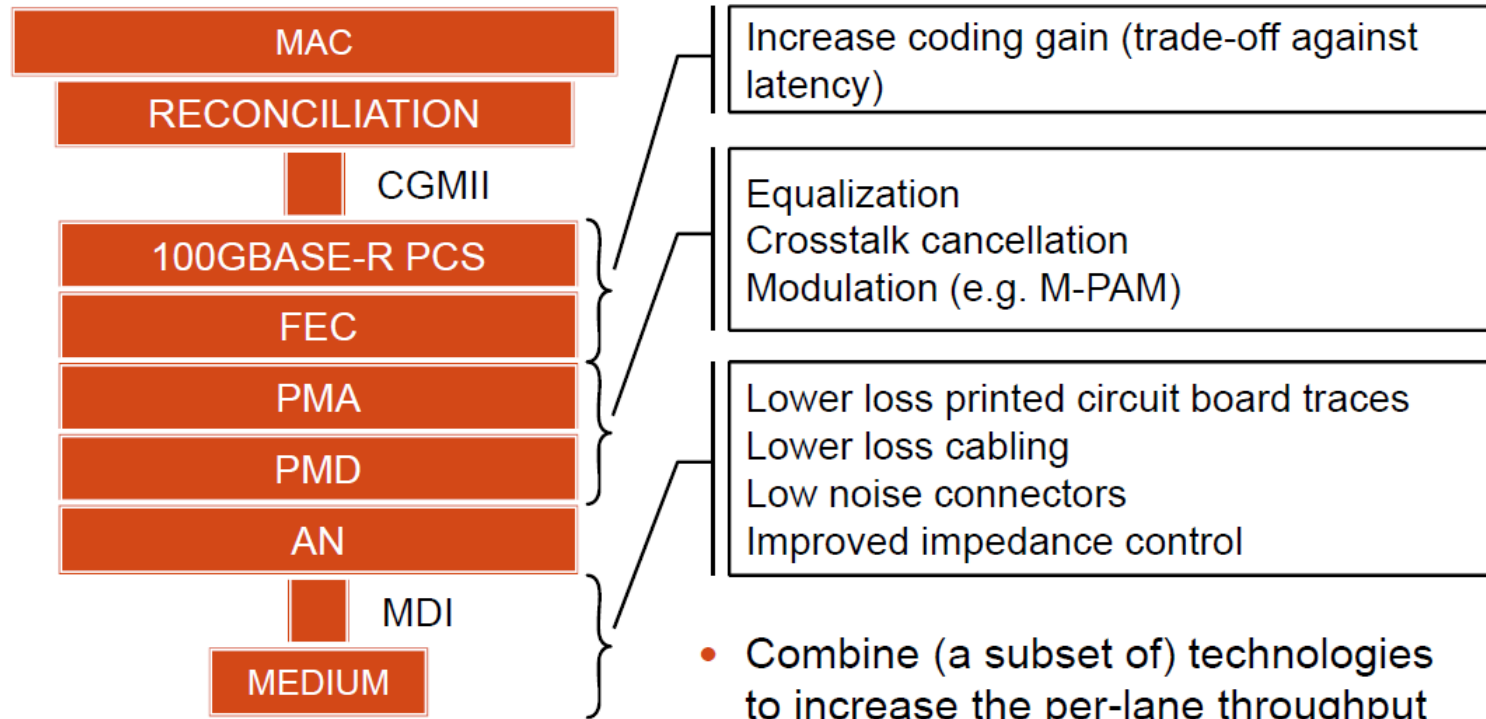
# Technical Feasibility

for 100 Gb/s per lane for Electrical Interfaces

Go *Faster*  
to  
Go **Denser**  
to  
Continue to **Grow**.

# It's time to open the toolbox again...

## Potential enablers for more Gb/s/lane



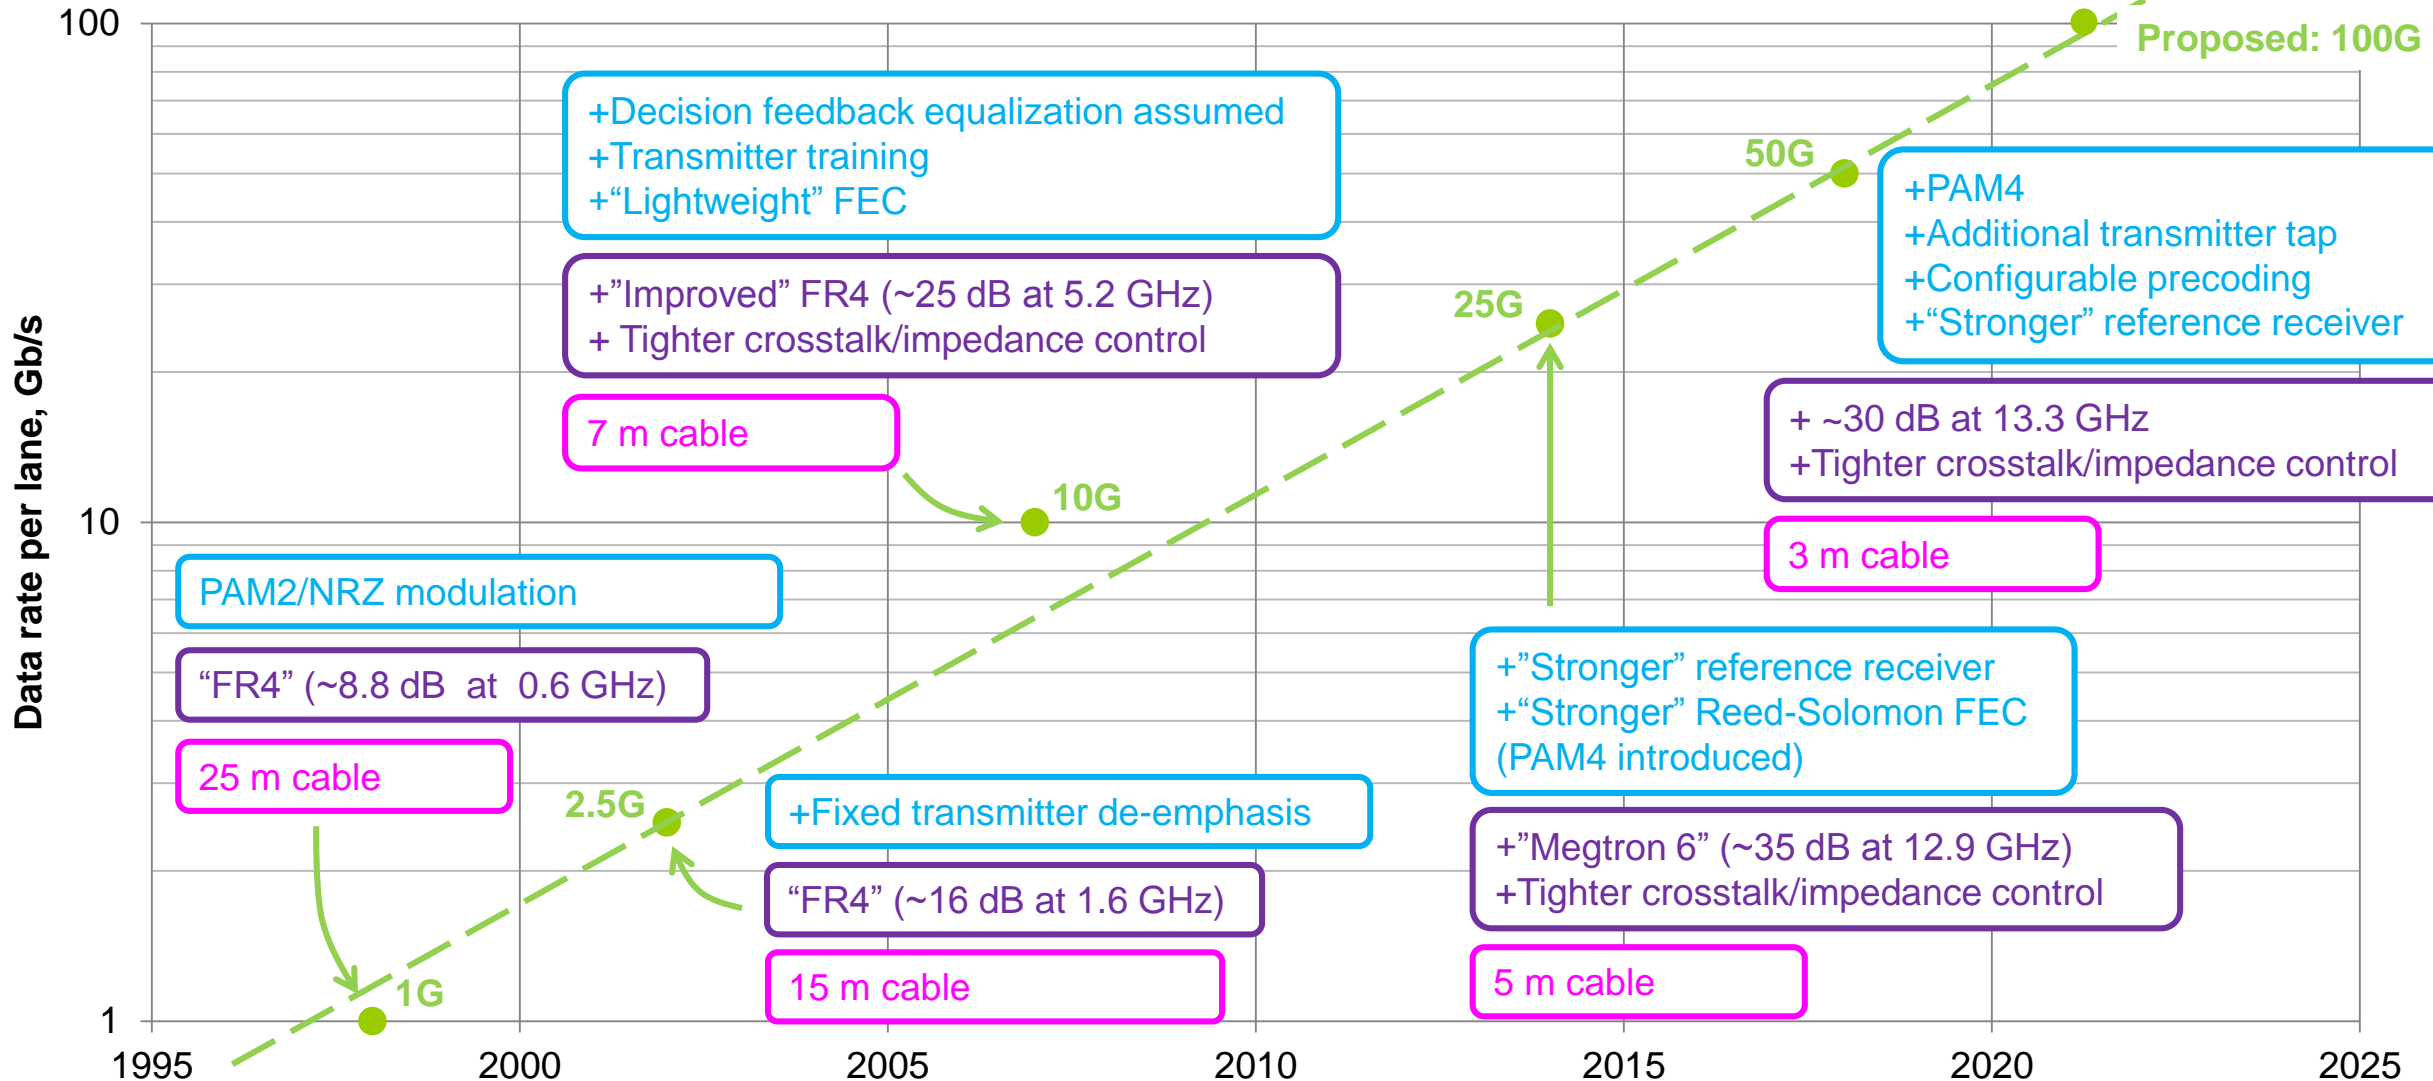
From the [100GbE Electrical Backplane / Cu Cabling Call-For-Interest](#) consensus building presentation, November 2010

November 9, 2010



# Solutions for each generation

SerDes    Backplane    Cable



SerDes Speed per lane

IEEE 802.3 November 7, 2017, Consensus Building

\* Dates are approximate

# Different constraints for different applications

## Chip-to-module

- What is the insertion loss range that supports useful applications?
- Can we “Coexist” with defined PHYs, including FEC & PCS?
- Consider improved PCB materials, PCB vs. cable, improvements in impedance/noise control

## Chip-to-chip and “backplane”

- What are useful topologies and reaches? What material and power implications are acceptable?
- Consider improved PCB materials, PCB vs. cable, improvements in impedance/noise control?

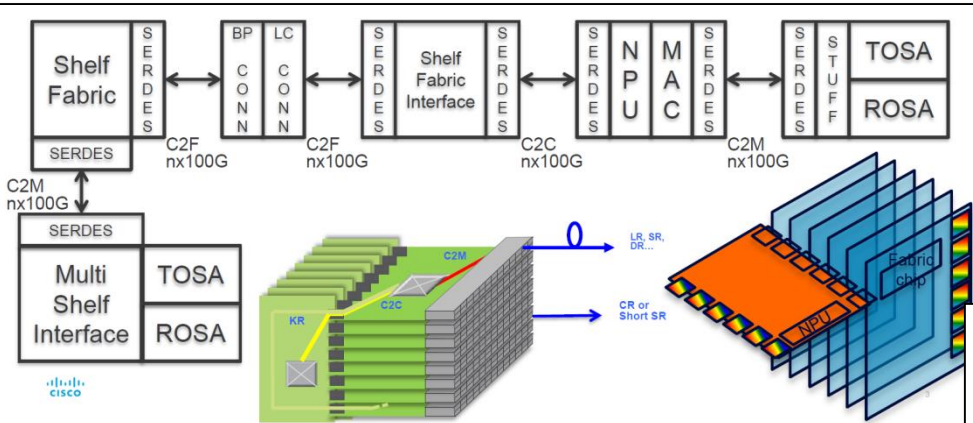
## Cable

- What is the minimum useful reach?
- Consider “middle-of-rack” topologies?
- Can we “Coexist” with defined PHYs, including FEC & PCS?


**Apply signal processing to meet the needs of each application**

# The discussion is already underway

From the proceedings of the IEEE 802.3 New Ethernet Applications ad hoc




**System considerations**  
**Channel options**  
**Higher-speed SerDes**

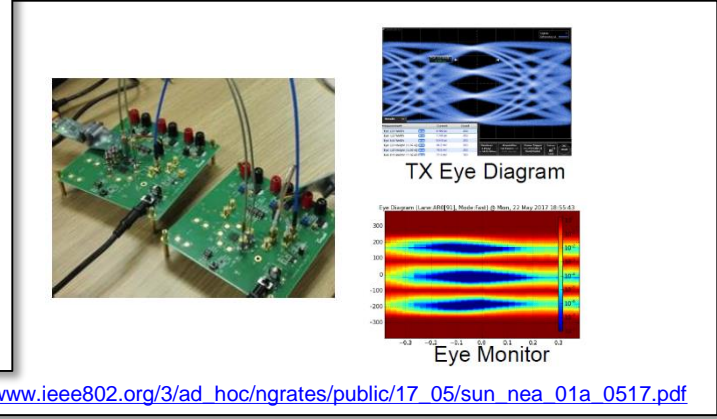



[http://www.ieee802.org/3/ad\\_hoc/ngrates/public/17\\_05/palkert\\_nea\\_02\\_0517.pdf](http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/palkert_nea_02_0517.pdf)

### Reduce Channel Length




### Use Lower Loss Channel



[http://www.ieee802.org/3/ad\\_hoc/ngrates/public/17\\_05/sun\\_nea\\_01a\\_0517.pdf](http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/sun_nea_01a_0517.pdf)

**OIF has CEI 112G projects are already underway**



[http://www.ieee802.org/3/ad\\_hoc/ngrates/public/17\\_05/mellitz\\_nea\\_01a\\_0517.pdf](http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/mellitz_nea_01a_0517.pdf)

# Technical feasibility summary

- Rich signal integrity and signal processing toolbox that can be applied to the problem of “100 Gb/s per lane electrical signaling”
- We must be mindful of the different needs for different applications
- We have done this many times before
- The discussion is already underway

# Why Now???

100 Gb/s per lane of Electrical Interfaces

Go *Faster*  
to  
Go **Denser**  
to  
Continue to **Grow**.

# The Road Map of Port Rates

First 100G project ratified


25G/lane electricals ratified

OIF est. first 100G/lane project

First 100G/lane optics project start

First 200G/400G project est. ratification

50G/lane electricals, projected ratification



- 100G/200G/400G technology is already out
- 100G per lane optics has begun, more to come
- OIF is already working on this
- 100G per lane is coming...

→ IEEE needs to study and frame it **NOW** so the industry can plan

Gb/s

# Supporters (individuals from 45 companies)

- Andy Zambell, Amphenol
- Erdem Matoglu, Amphenol
- Amir Bar-Niv, Aquantia
- Ramin Farjad, Aquantia
- Adam Healey, Broadcom
- Henry Chen, Broadcom
- Raj Hegde, Broadcom
- Rob Stone, Broadcom
- Arthur Marris, Cadence
- Mike Dudek, Cavium
- Marek Hajduczenia, Charter
- Ryan Tucker, Charter
- Gary Nicholl, Cisco Systems
- Jane Lim, Cisco Systems
- Joel Goergen, Cisco Systems
- Mark Nowell, Cisco Systems
- Alex Umnov, Corning
- Steve Swanson, Corning
- Haoli Qian, Credo Semiconductor
- Phil Sun, Credo Semiconductor
- Jeff Twombly, Credo Semiconductor
- David Piehler, Dell
- Jon Lewis, Dell EMC
- Martin Zielinski, Emerson
- Vipul Bhatt, Finisar
- Ali Ghiasi - Ghiasi Quantum LLC
- John Ewen, Global Foundries
- Jacky Chang, HPE
- Kenghua Chuang, HPE
- Andre Szczepanek, HSZ Consulting
- Yasuo Hidaka, Independent
- Ted Sprague, Infinera
- Kapil Shrikhande, Innovium
- Adee Ran, Intel
- Howard Heck, Intel
- Kent Lusted, Intel
- Mike Li, Intel
- Tom Issenhuth, Issenhuth Consulting
- Jerry Pepper, IXIA
- Rick Rabinovich, IXIA
- Thananya Baldwin, IXIA
- David Ofelt, Juniper
- Jeff Maki, Juniper
- Steve Sekel, Keysight Technologies
- Dale Murray, Lightcounting
- Matt Brown, Macom
- David Malicoat, Malicoat Networking Solutions
- Jacov Brener, Marvell
- Liav Ben-Artzi, Marvell
- Venu Balasubramonian, Marvell
- Piers Dawe, Mellanox
- Brad Booth, Microsoft
- Scott Sommers, Molex
- Tom Palkert, Molex/Macom
- Mabud Choudhury, OFS
- Rick Pimpinella, Panduit
- Chris DiMinico, PHY-SI
- Bharat Tailor, Samtec
- Mark Kimber, Samtec
- Rich Mellitz, Samtec
- Toshiaki Sakai, Socionext
- David Estes, Spirent
- Ed Nakamoto, Spirent
- Lokesh Kabra, Synopsys
- Michael Lynch, Synopsys
- Rita Horner, Synopsys
- Pavel Zivny, Tektronix
- Megha Shanbhag, TTM Technologies
- Nathan Tracy, TTM Technologies
- Paul Brooks, Viavi Solutions
- Mark Gustlin, Xilinx
- Takeshi Nishimura, Yamaichi Electronics USA
- Yang Zhiwei, ZTE

# Summary

- Higher bandwidth requires density improvements in faceplate, chip breakout, and backplane connections.
- 100 Gb/s per electrical lane is the next logical step
- We've moved to the "unknown" before and the industry flourished.
- Technical details need to be rebalanced for the next speed.
- 100Gb/s per lane is coming, we need to frame the discussion
- Let's form a Study Group!!



# Thank You!



## Contributors/Q&A panel

- John D'Ambrosia, Futurewei
- David Ofelt, Juniper
- Adam Healey, Broadcom
- Kent Lusted, Intel

Go *Faster*  
to  
Go **Denser**  
to  
Continue to **Grow**.

# Straw Poll: Study Group Formation


Should a study group be formed for “100Gb/s per Lane for Electrical Interfaces and PHYs”?

Y:

N:

A:

Room count:



# Straw Poll: Study Group Participation

I would participate in a “100Gb/s per lane for Electrical Interfaces and PHYs” study group in IEEE 802.3.

Tally:

My company would support participation in a “100Gb/s per lane for Electrical Interfaces and PHYs” study group.

Tally:

# Next Steps

- Make a motion to the 802.3 working group at Thursday's closing meeting to request the formation of a study group.
- If the motion passes, David Law will make the request to 802 EC on Friday.
- If approved:
  - Teleconference(s) in December to start the discussion (will post on NEA Adhoc reflector)
  - First study group meeting would be during Jan. 2018 IEEE 802.3 interim meeting in Geneva

