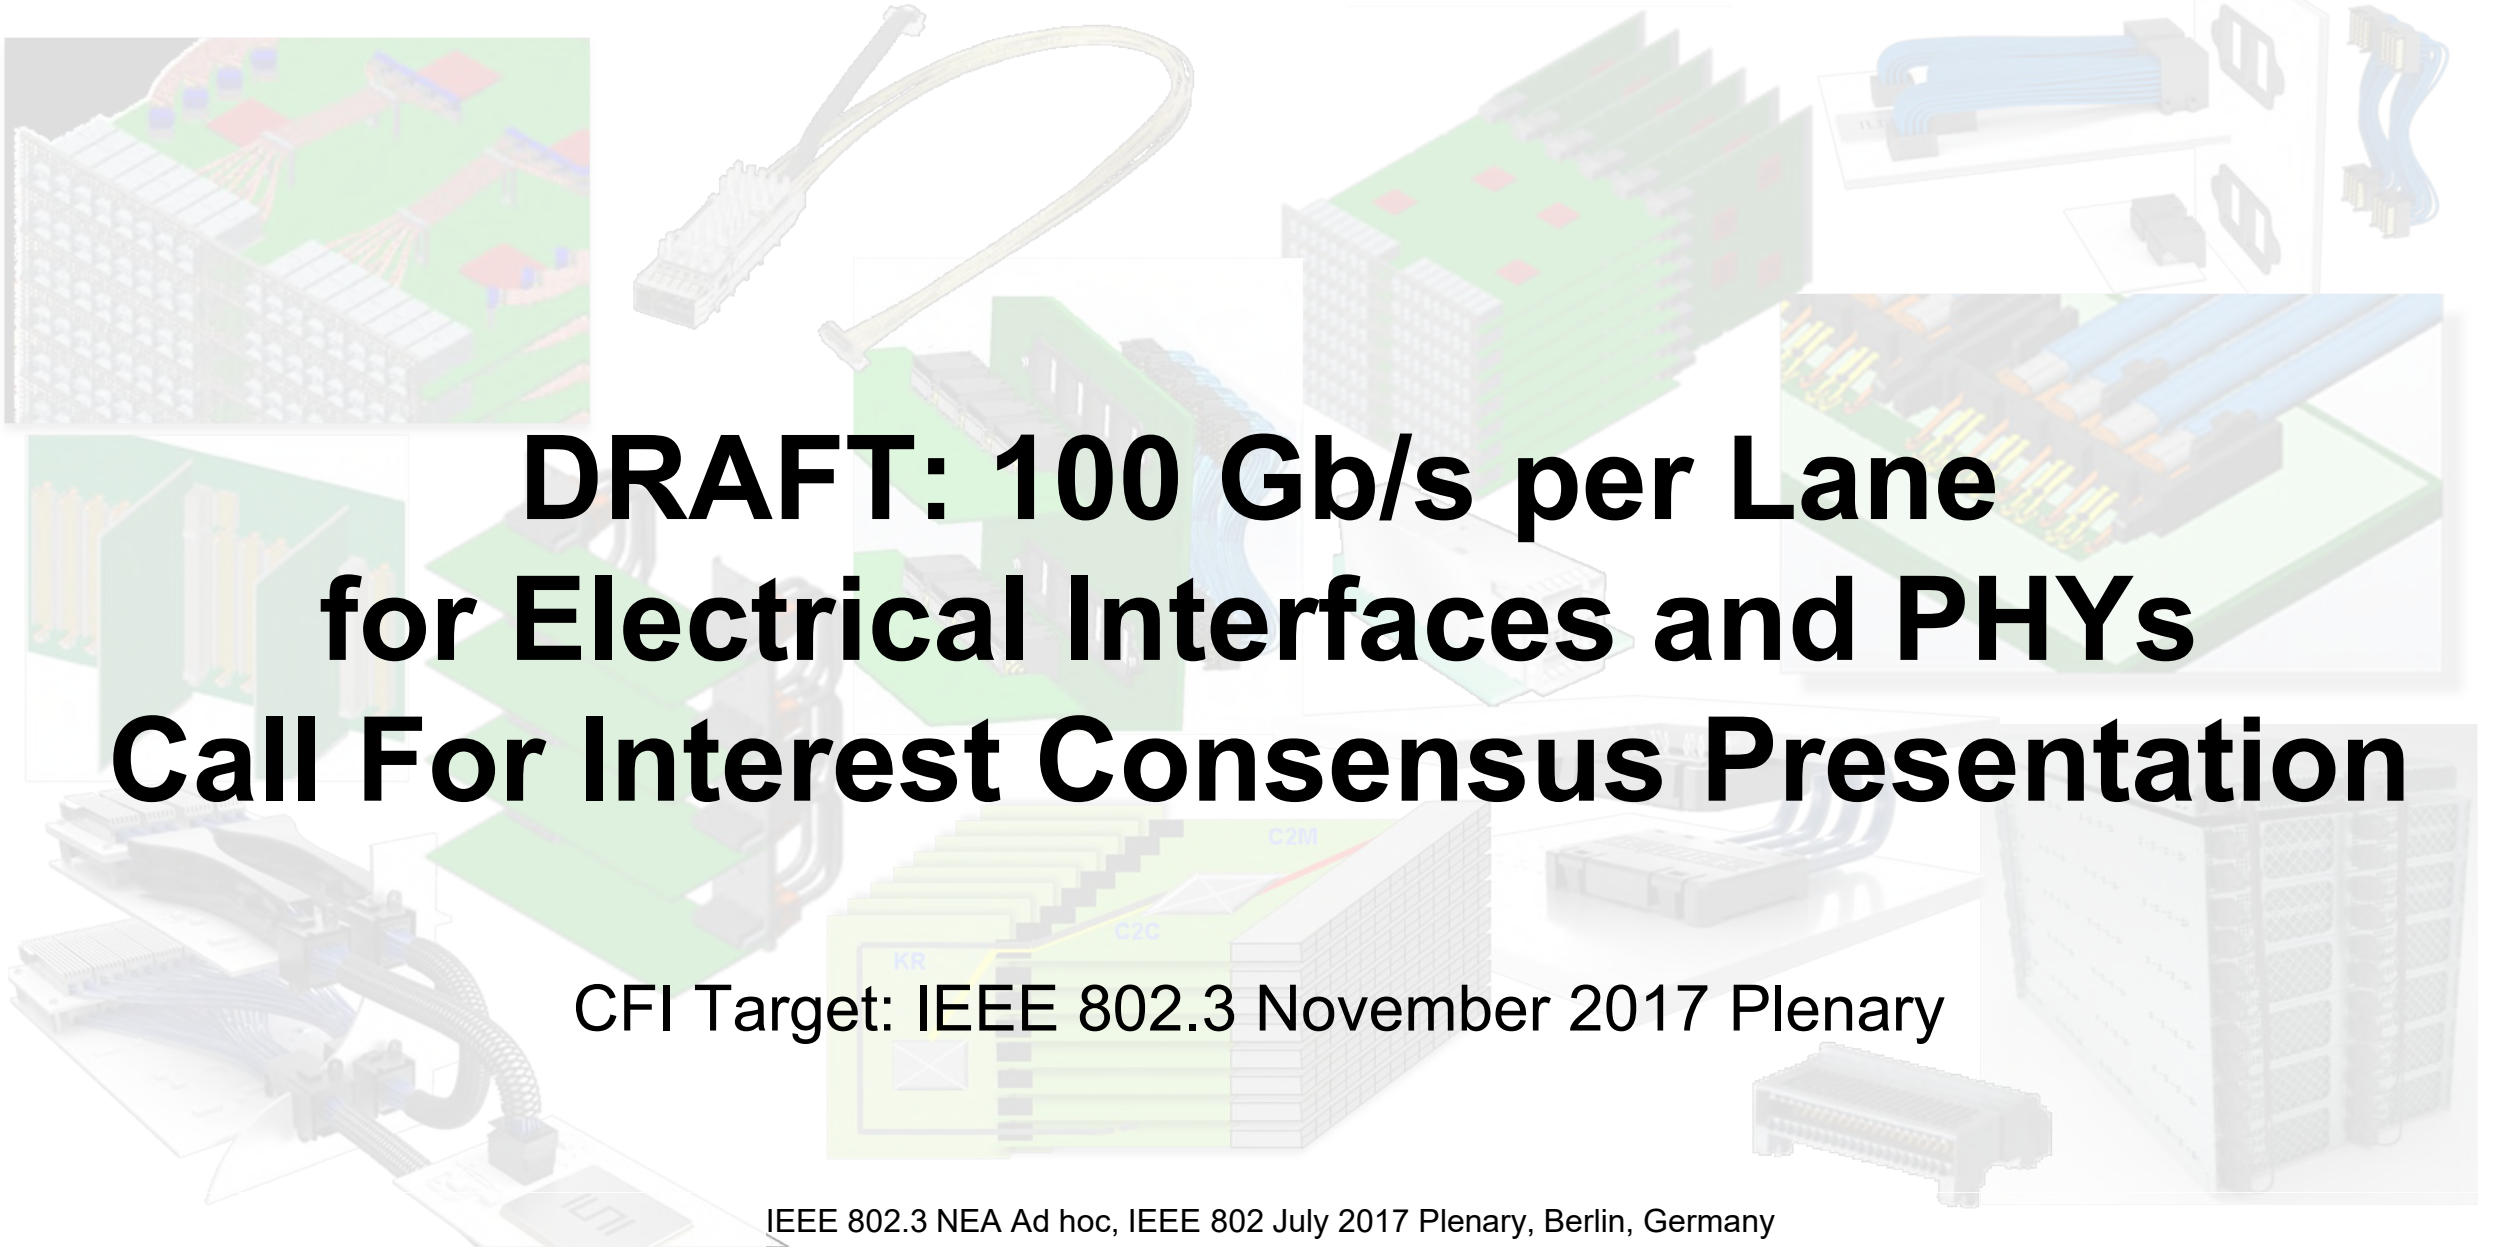

DRAFT: 100 Gb/s per Lane for Electrical Interfaces and PHYs Call For Interest Consensus Presentation

CFI Target: IEEE 802.3 November 2017 Plenary

Objective

- To gauge the interest and 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

Higher speed copper transmission

What are the limits of copper?

- Higher speed copper was predicted to be dead a decade ago– but this has also been the case for the last 30 years
- Copper keeps on pushing the frequency limits
- Copper vs optical – gap is closing as speeds increase
- Although optics offer reach and density, electrical still offers lower cost and power
- Equalization technology and modulation techniques continue to be improved
  - PAM4, ENRZ, Duobinary, etc.
- The economics at stake are huge, "Do you really want to bet against copper?"

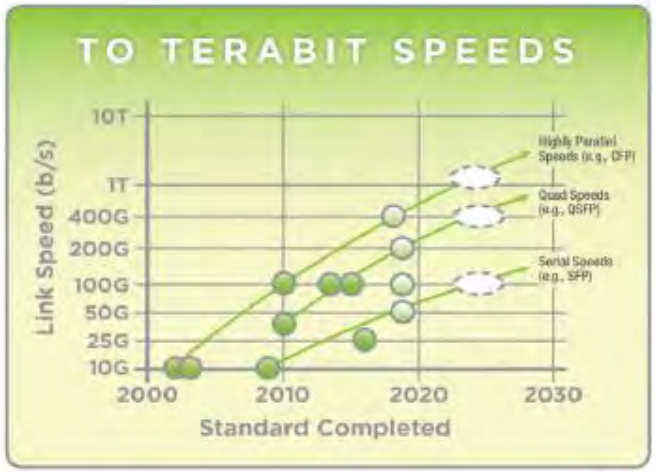

DATA COMMUNICATIONS

TE connectivity

Source: Nathan Tracy, TE Connectivity, **IEEE 802.3 NEA Ad Hoc. May 2017**
http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/tracy_nea_01_0517.pdf

Motivation for 100 Gbps per Lane

- Parallel interfaces limit abilities
- Cost optimization (motivation for much of the industry)
- Maximize system efficiency
- 100Gbps per lane could be leveraged for many parts in the Ethernet family

Web-scale data centers and cloud based services are presented as leading applications

Electrical interfaces come in many shapes and sizes.

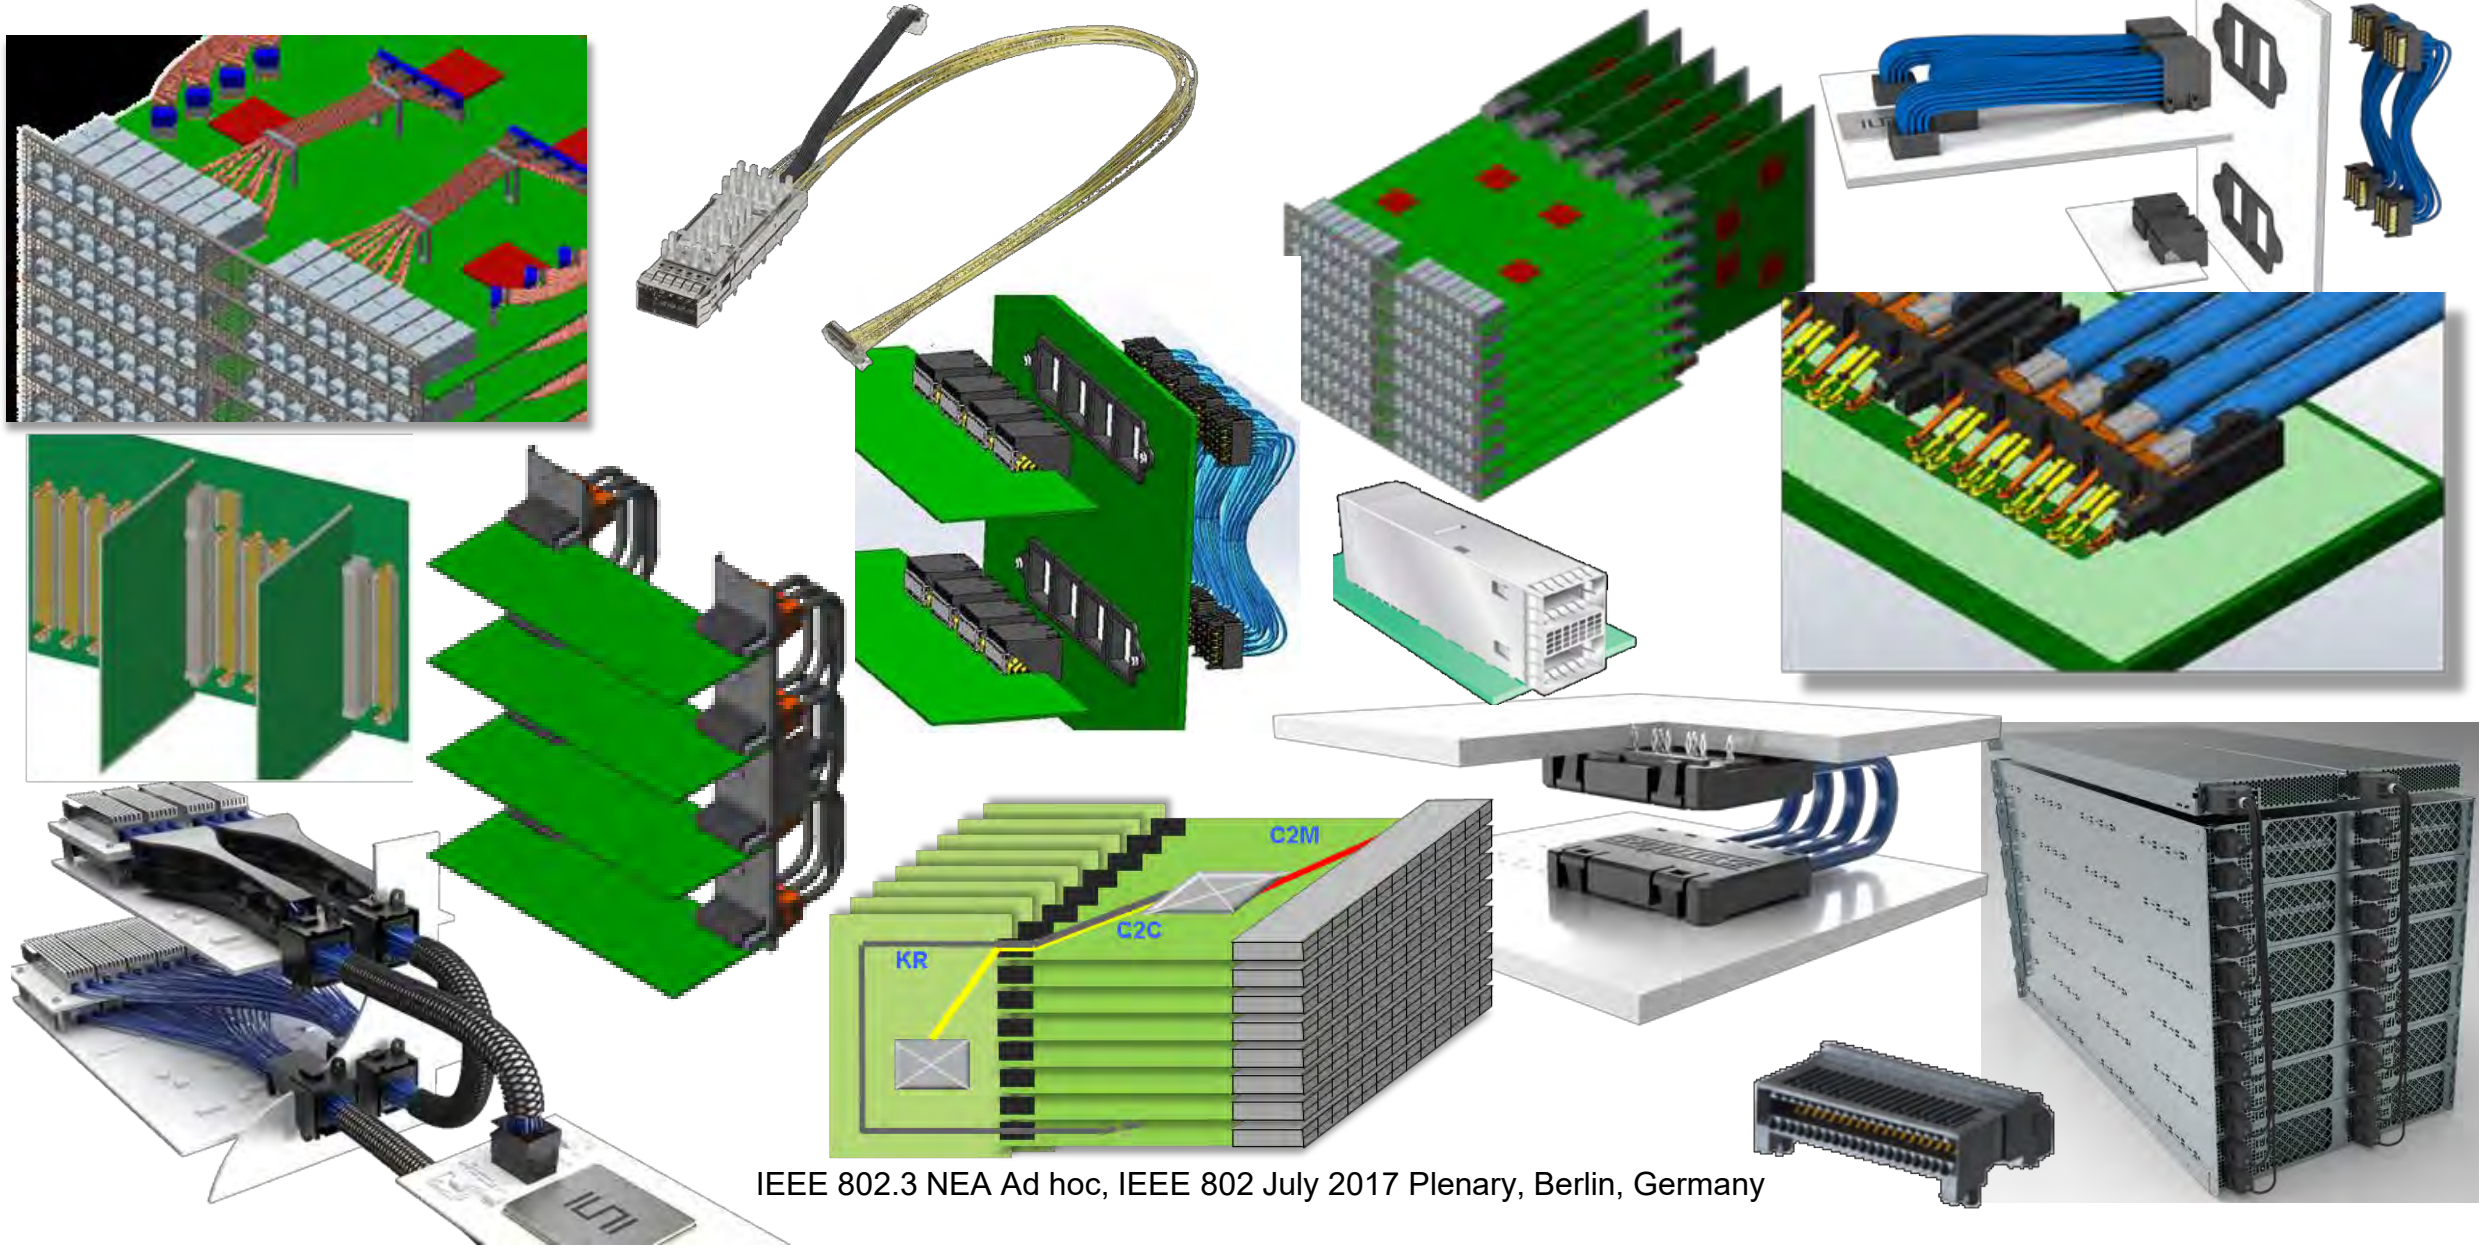

Tonight's Meeting

- To present the

  market ***NEED***,

  technical **Feasibility**,

  and *impact* on relevant PHYs and data rates

  of 100Gbps/lane-based electrical signaling.

- To gain consensus towards Call-for-Interest 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

100 Gb/s per Electrical Lane Ethernet Market Drivers -

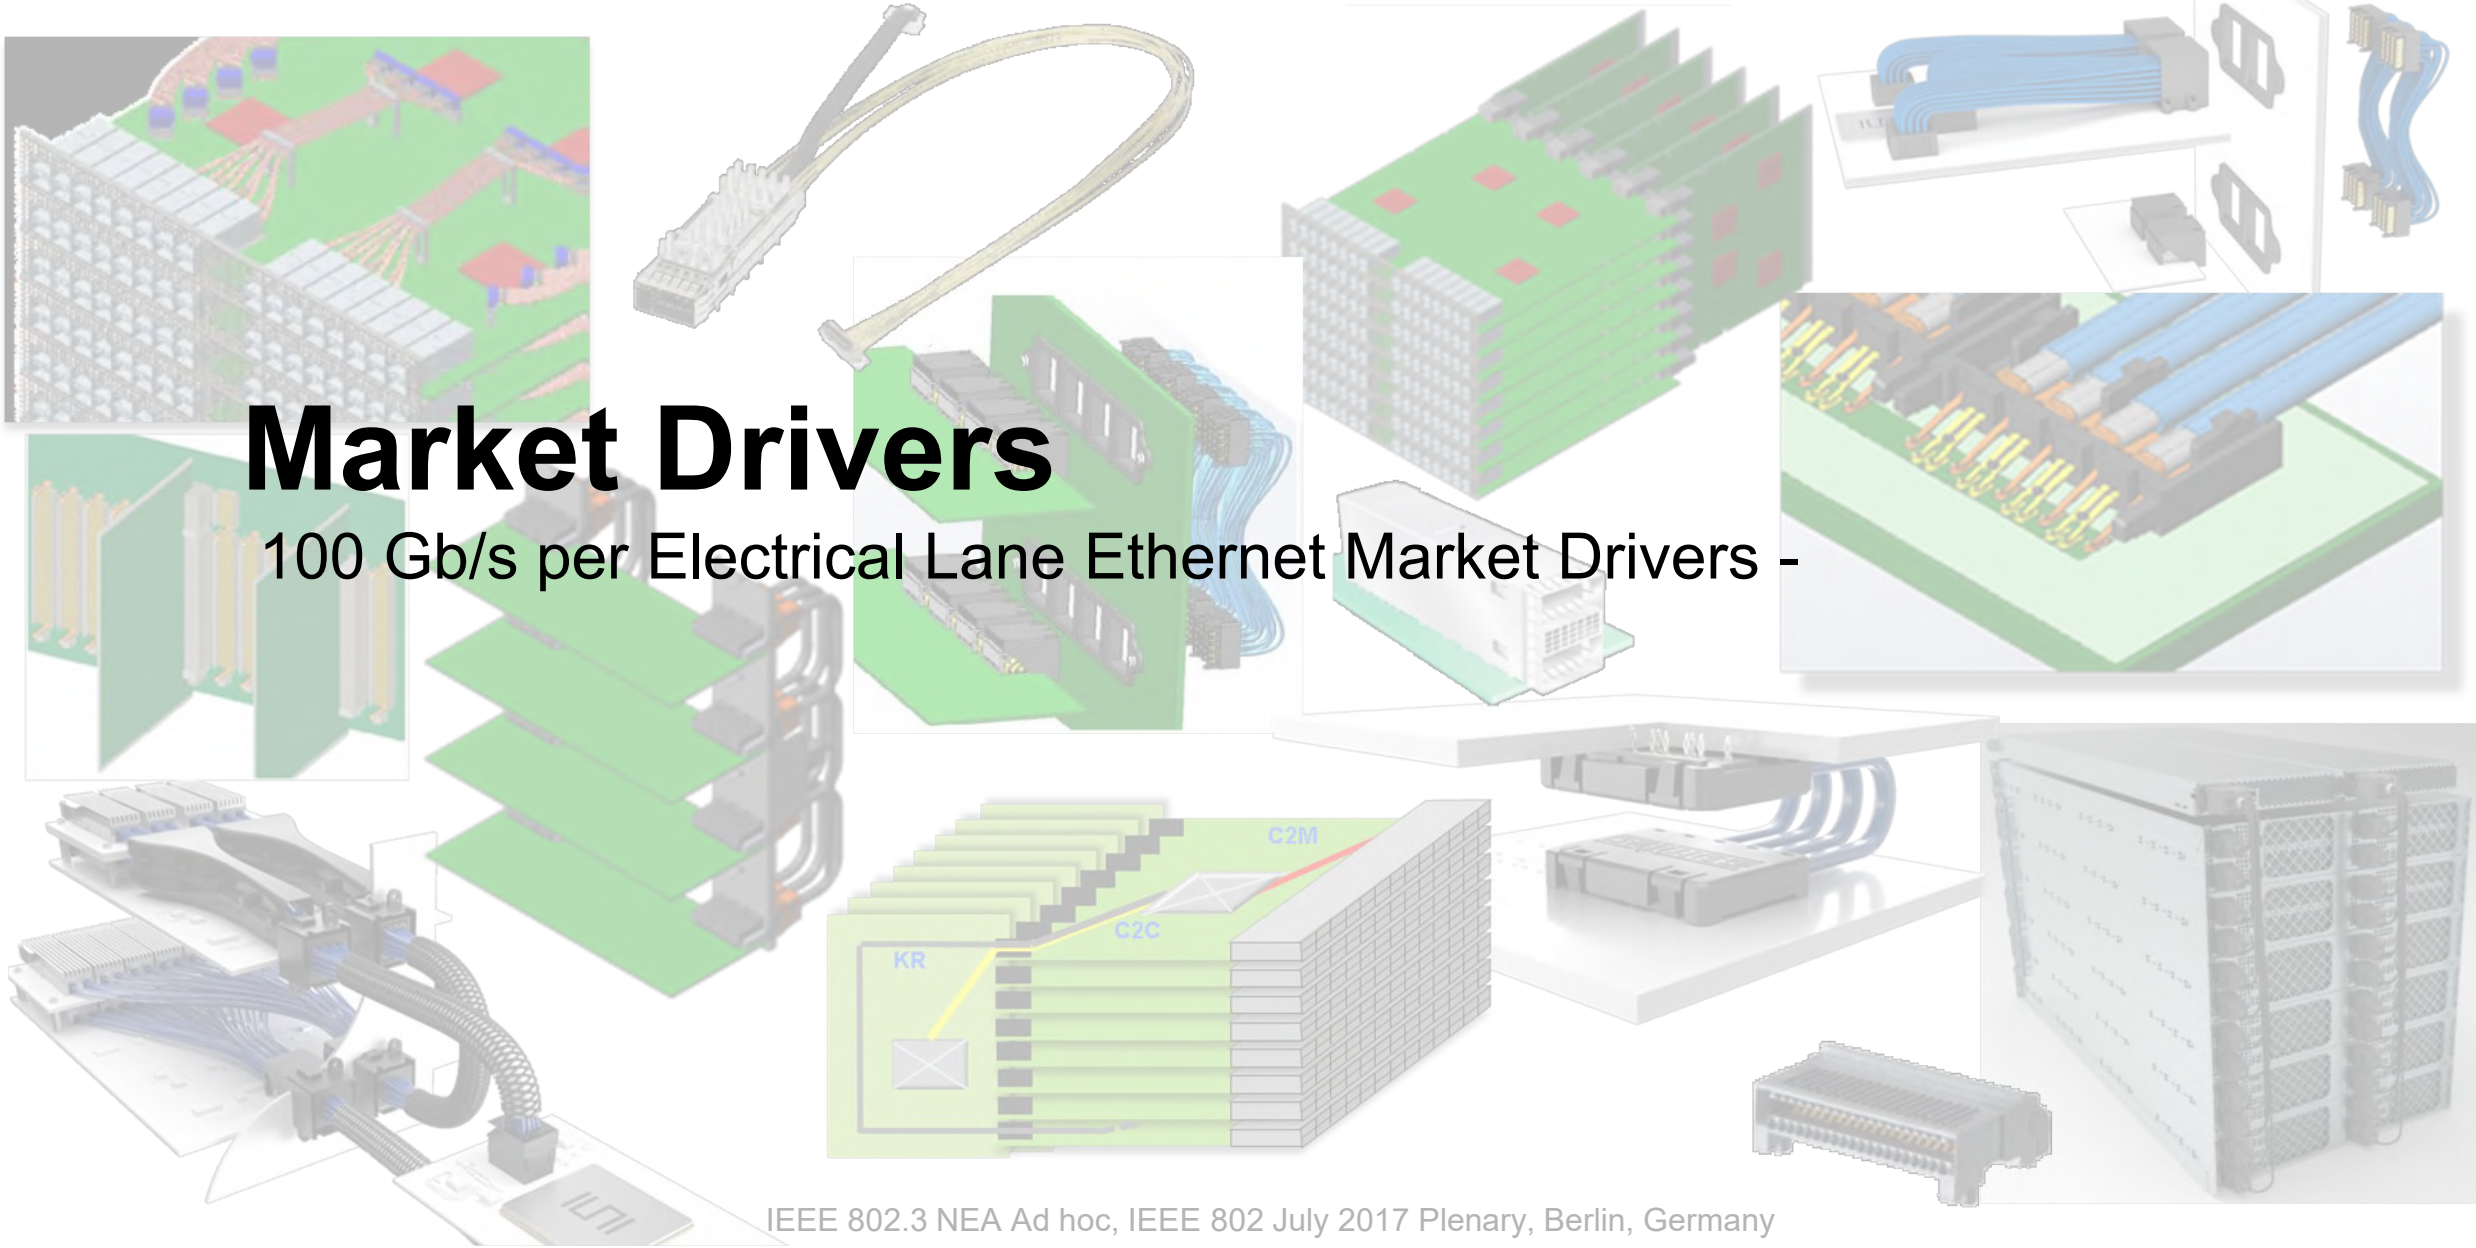

High Speed Front End Interconnects
The Road Map of Port Rates

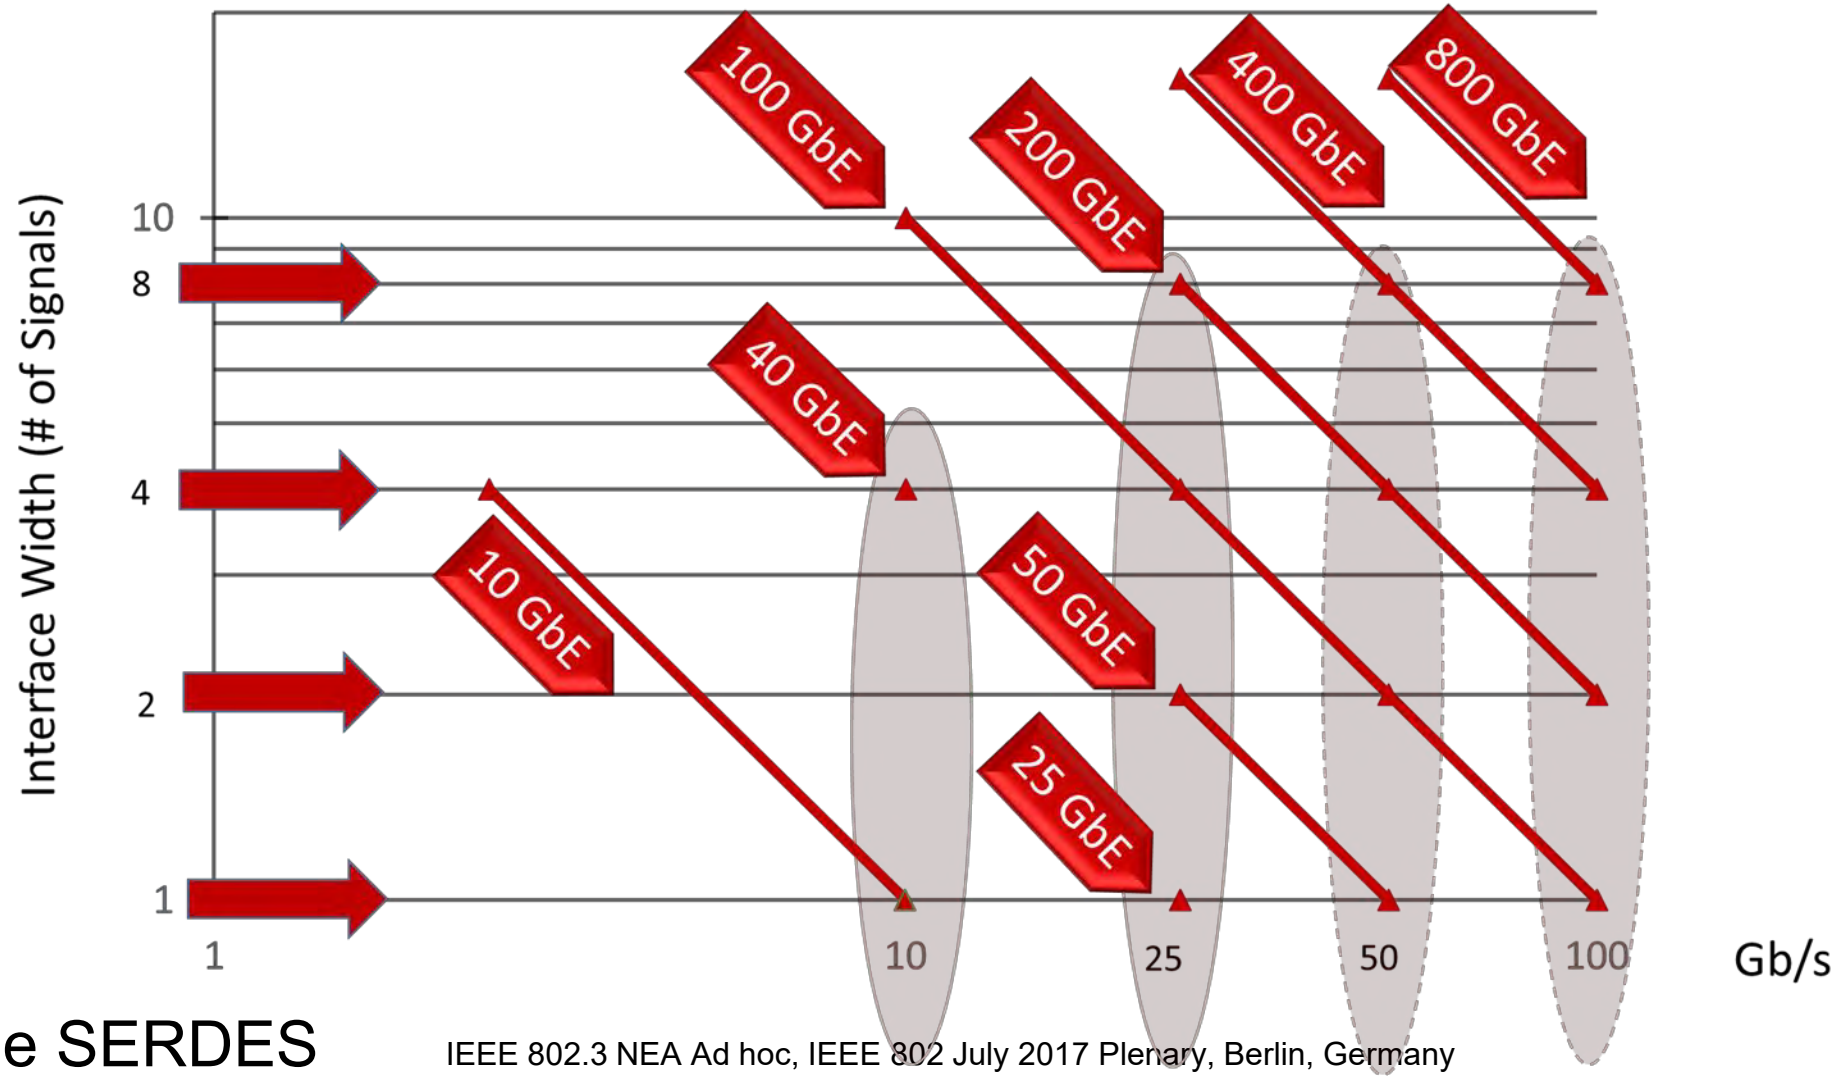

Follow the SERDES

The 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		10X10		
	25	CAUI-4 / 100GAUI-4	KR4	CR4	SR4	PSM4	CWDM4 CLR4	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

What Are We Talking About?

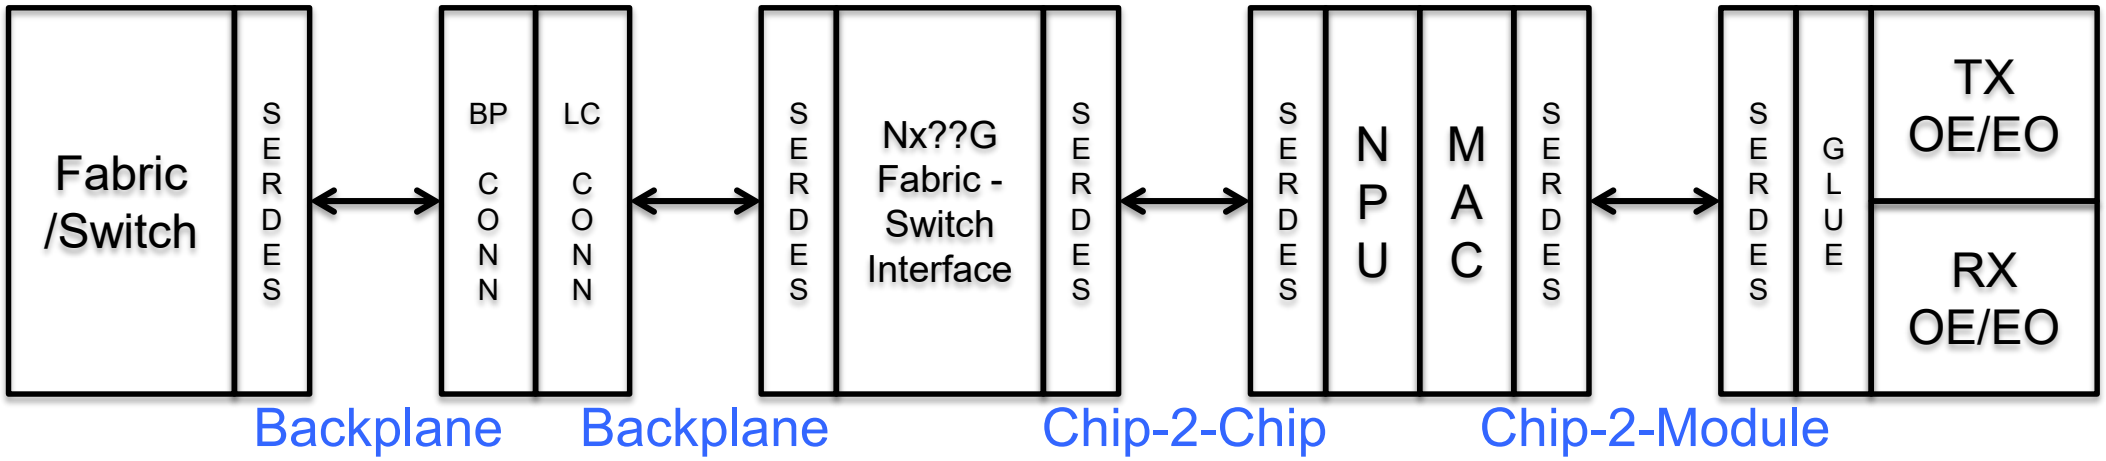

Consider how many instances of the interface can exist in the system…

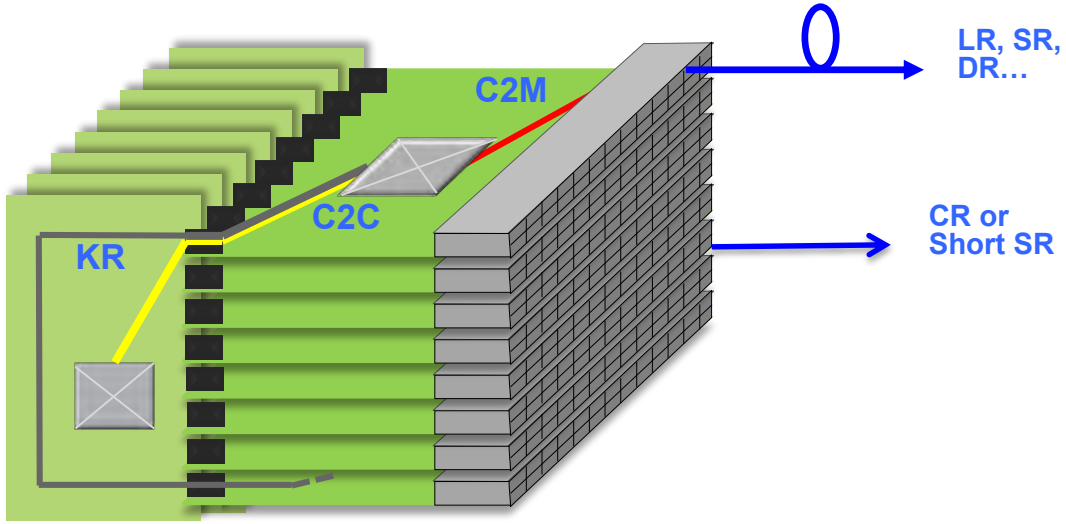

Interface Width: High Level Math

- System Connections
  - Consider a system has 10s of thousands of interconnects (say 15k)
  - For the same throughput:
    - 25Gbps/lane “by 4” ➔ 60k interfaces
    - 100Gbps/lane “by 1” ➔ 15k interfaces
- Pin Count of a 10Tb chip

Speed/lane	# Lanes in 1Tbps	# Balls Needed
10 Gbps	1000	4000
25 Gbps	400	1600
100 Gbps	100	400

Enabling higher throughput

Yikes!!

* Single lane = 2 differential pairs (4 balls)

IO Escape forcing transition to higher lane speeds

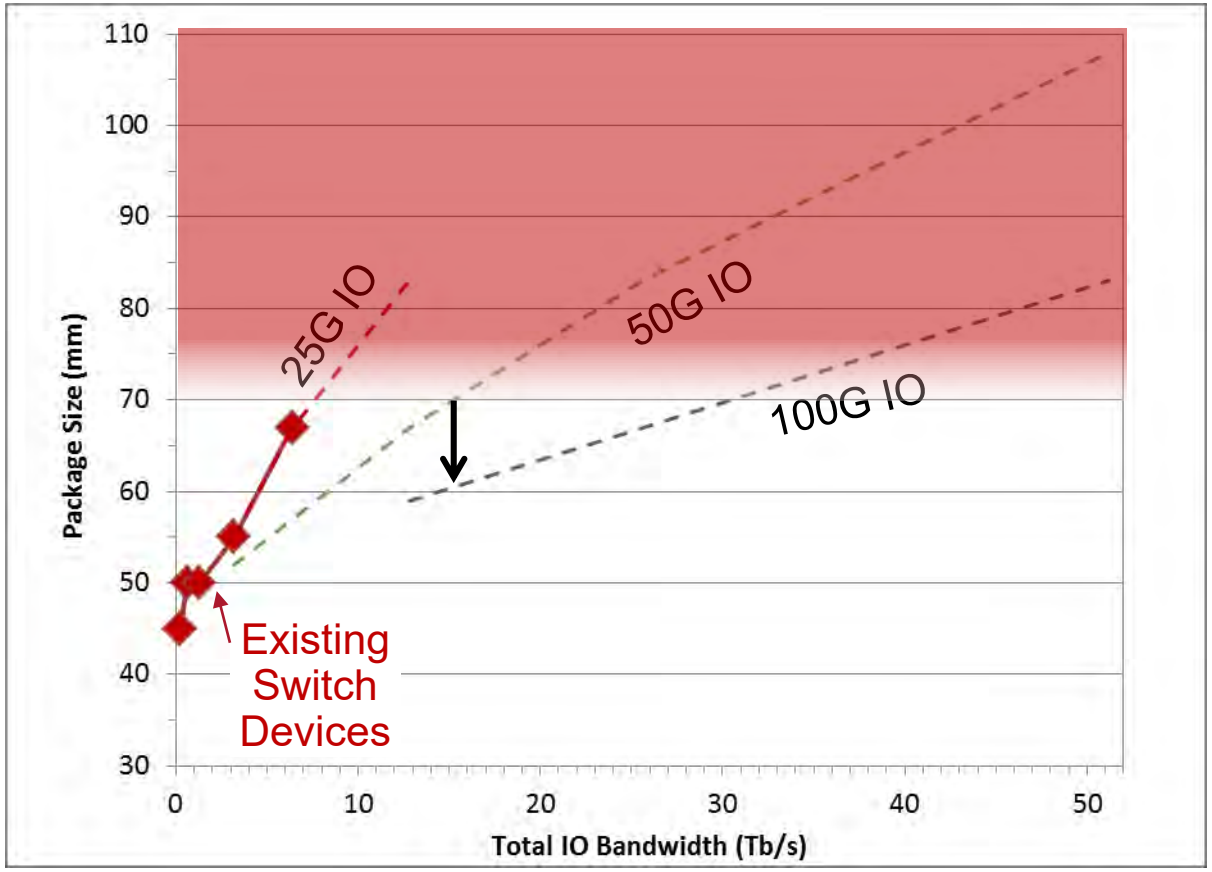

- ~ 70mm package is a current BGA practical maximum (due to coplanarity / warpage)
- This will force BGA devices with > 14Tb/s of aggregate bandwidth to transition to lane rates of higher greater than 50G (possibly 100G?)

IO Escape forcing transition to higher lane speeds

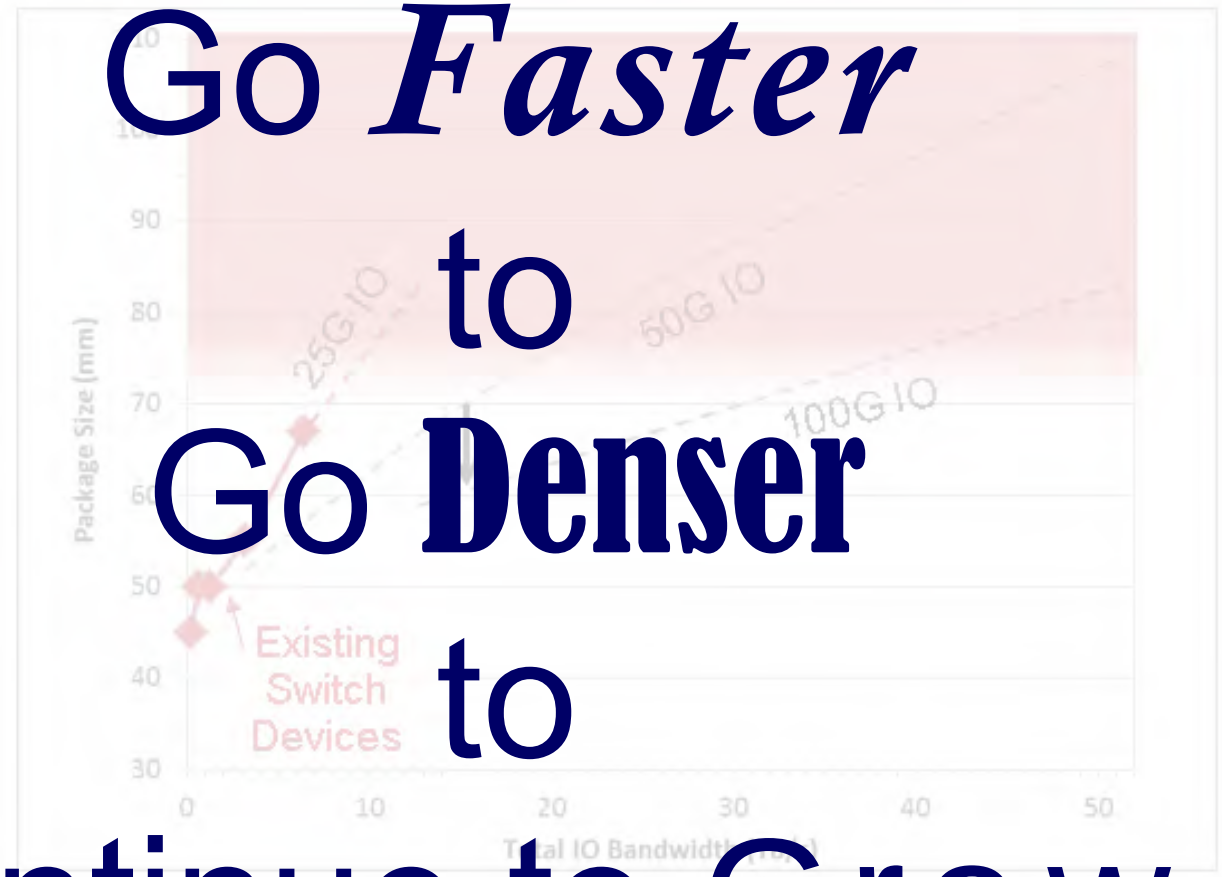

Go *Faster*

to

Go **Denser**

to

Continue to Grow.

- ~ 70mm package is a current BGA practical maximum (due to coplanarity / warpage)
- This will force BGA devices with > 14Tb/s of aggregate bandwidth to transition to lane rates of higher greater than 50G (possibly 100G?)

Technical Feasibility

100 Gb/s per Electrical Lane Ethernet Technical Feasibility -

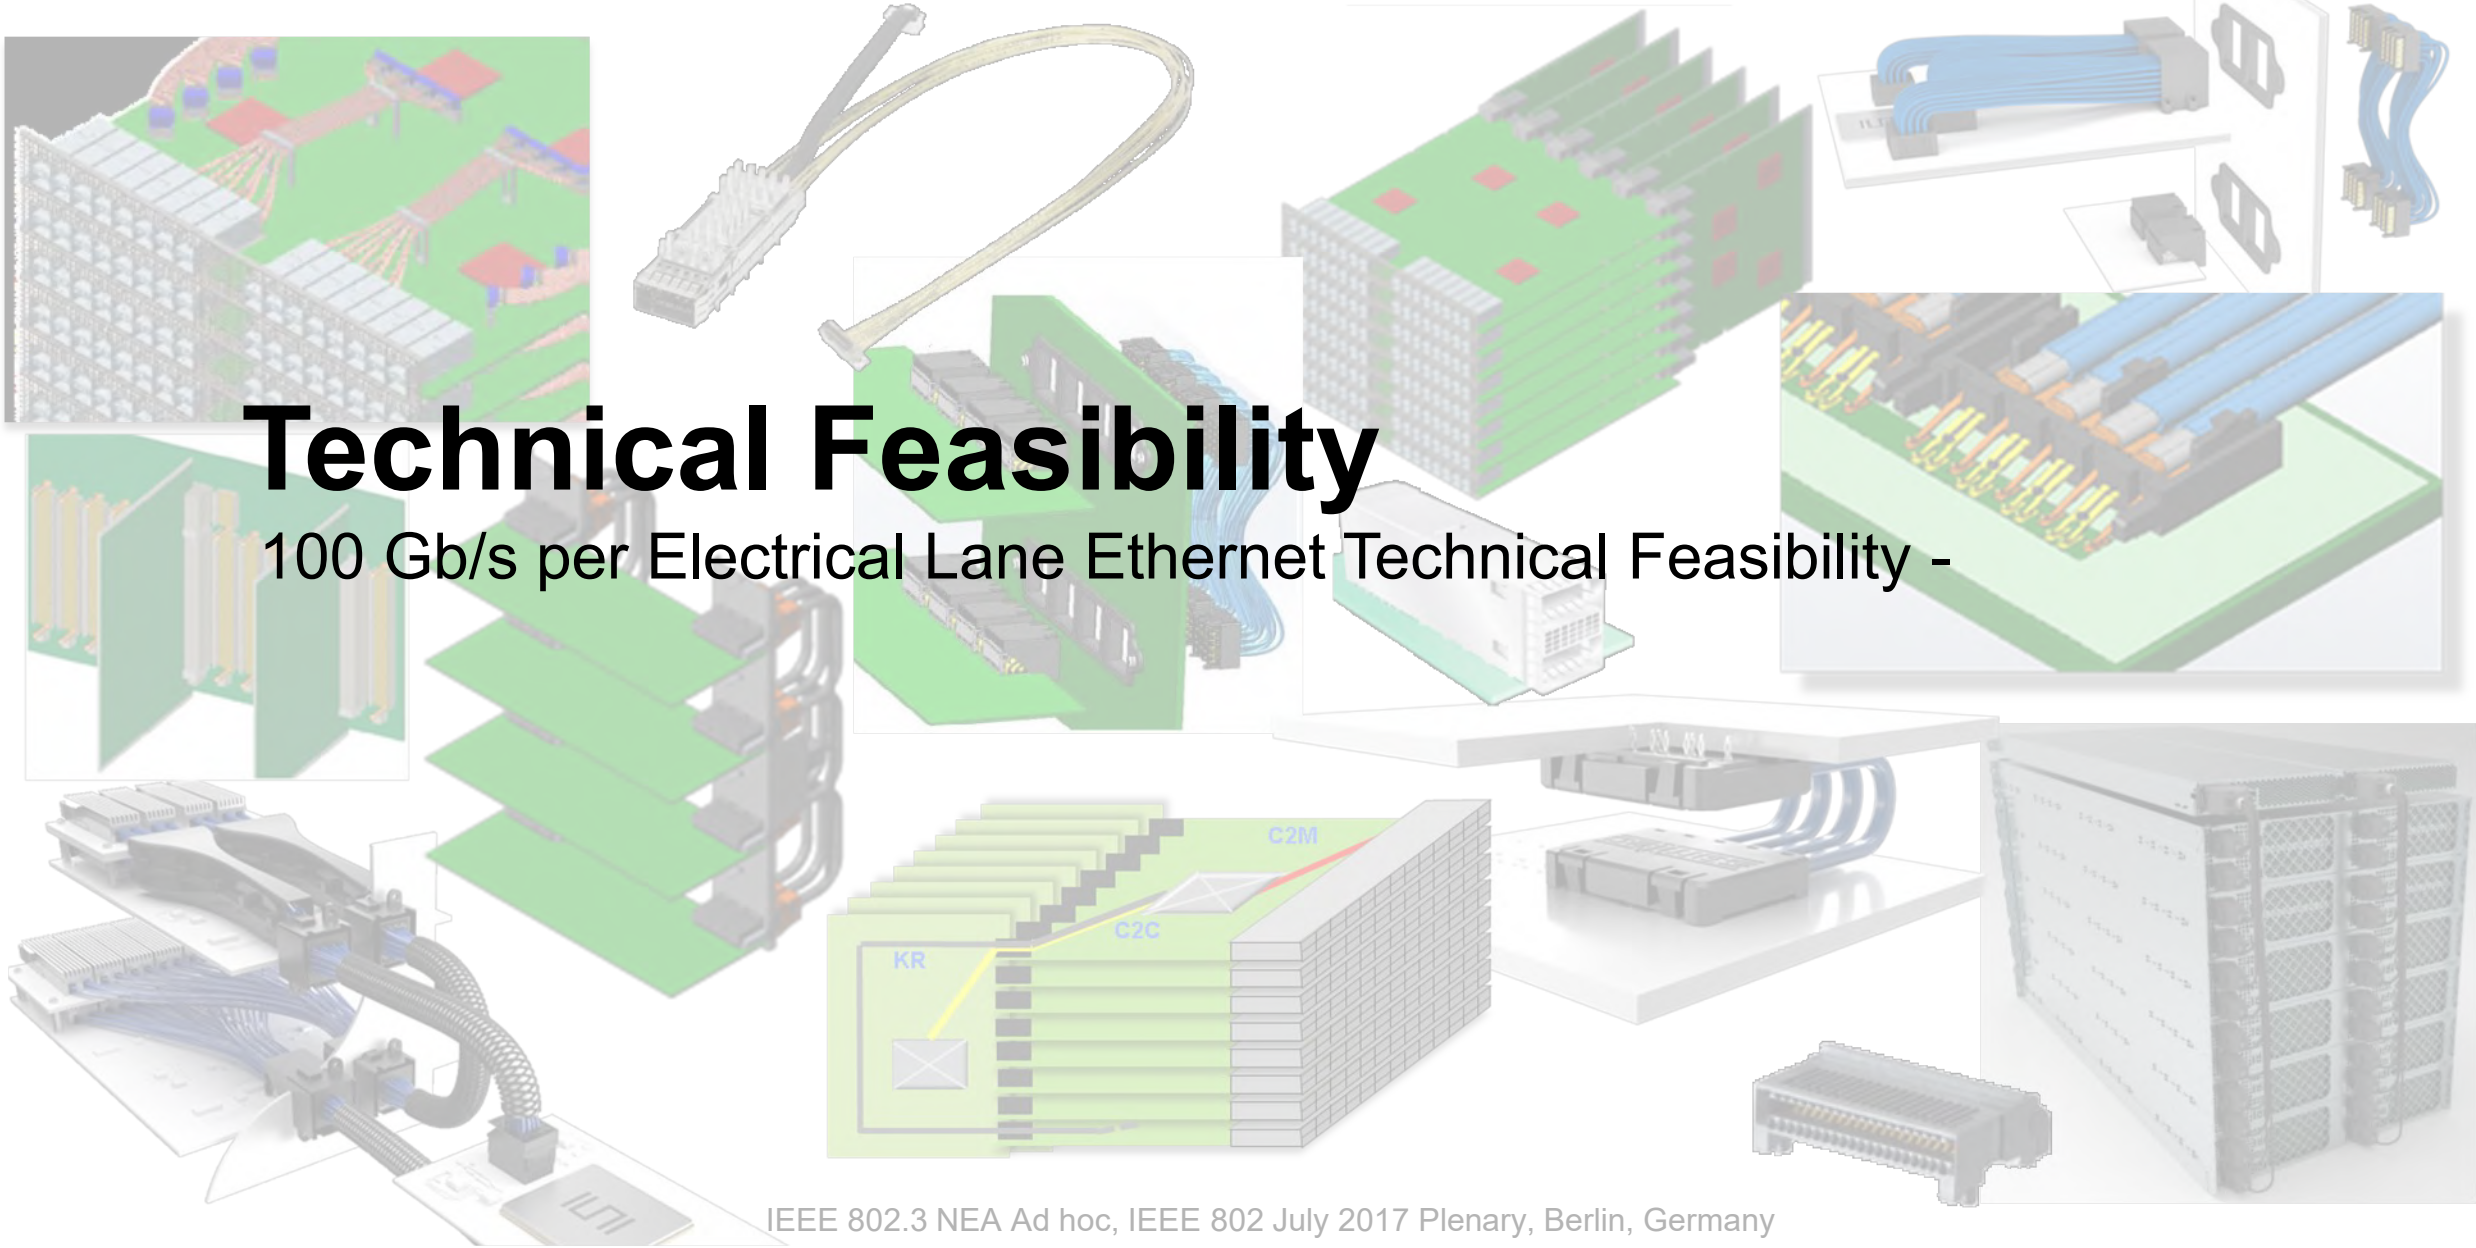

The Story Hasn't Changed

- From 25G Backplane CFI

  http://www.ieee802.org/3/cfi/1110_1/CFI_01_1110.pdf

- "Knobs" for tuning are the same; some have extended their scope, but it's still the same list.

- Now is the time to rebalance for the next speed!

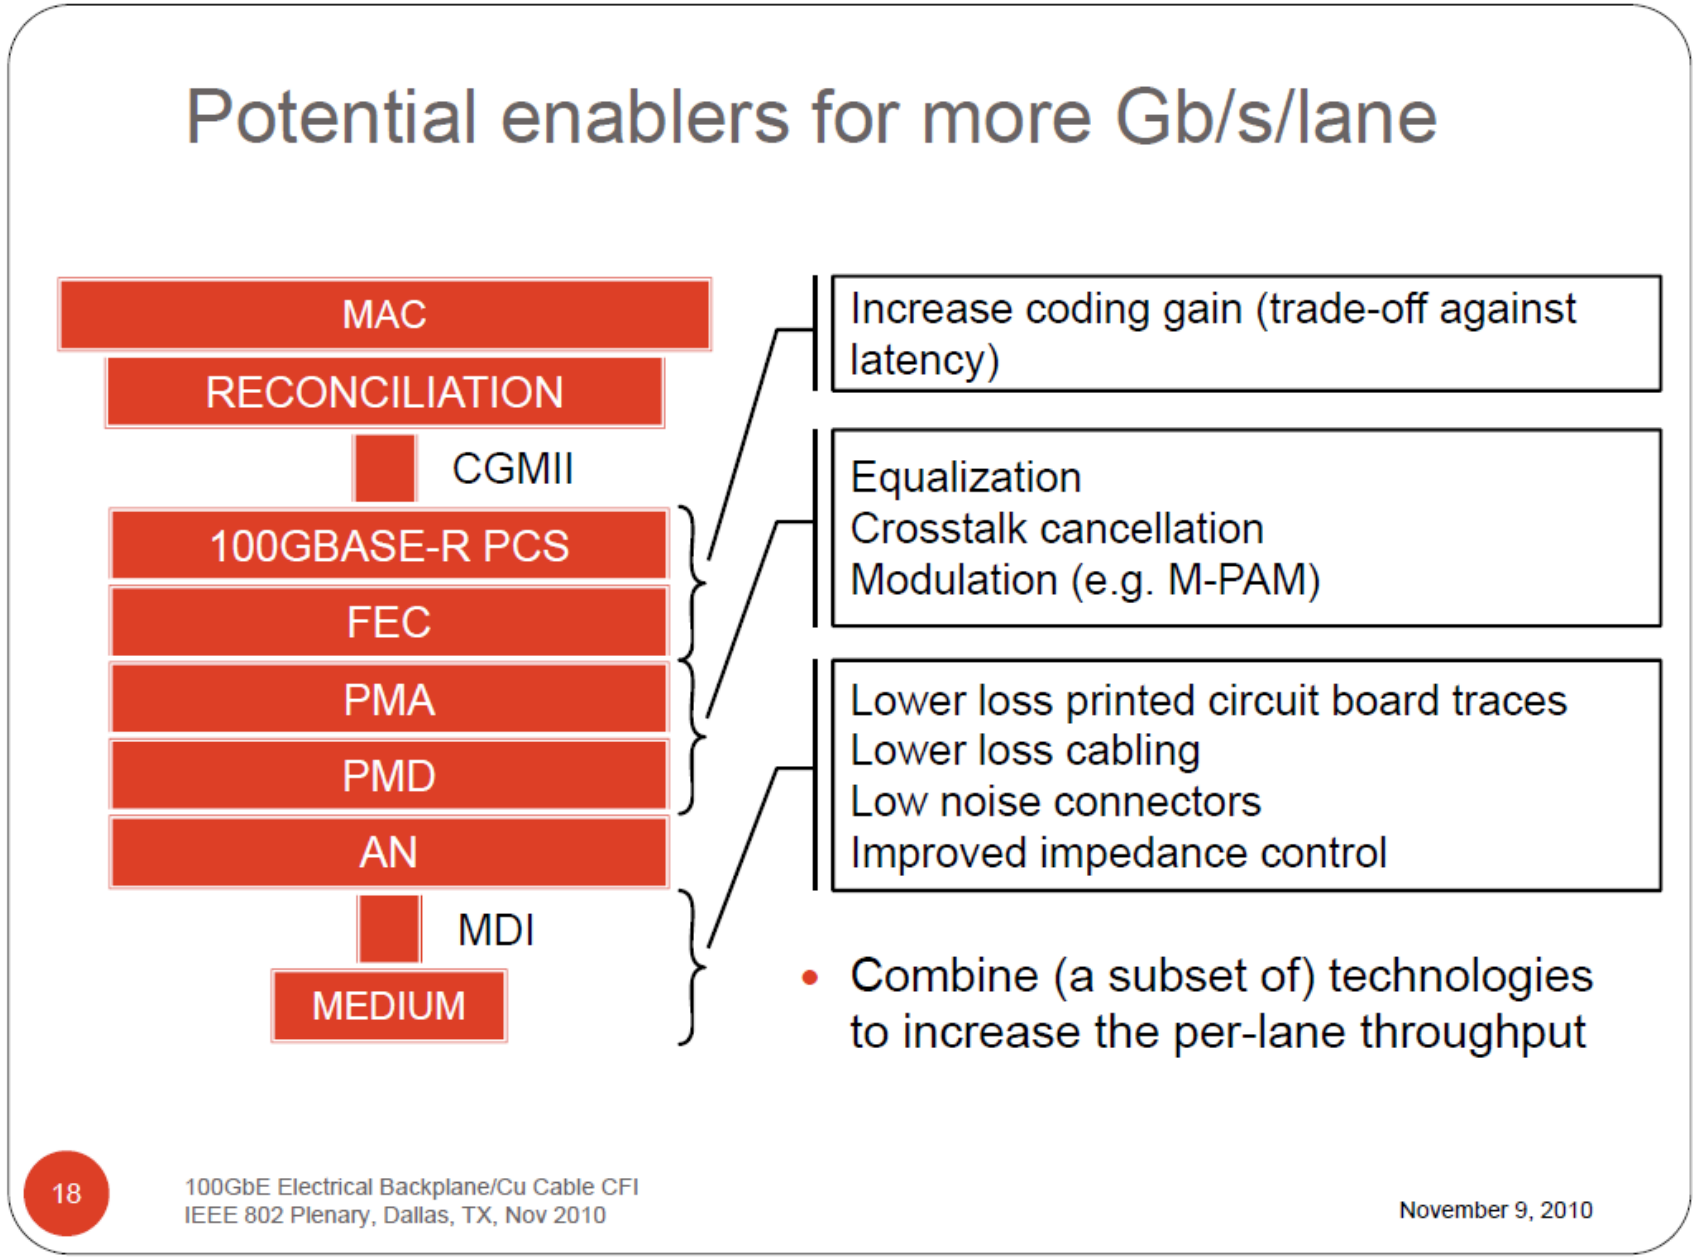

Not-your-Grandmother's Channel

- 25G NRZ channels may be too high loss,
  but we have options to study.

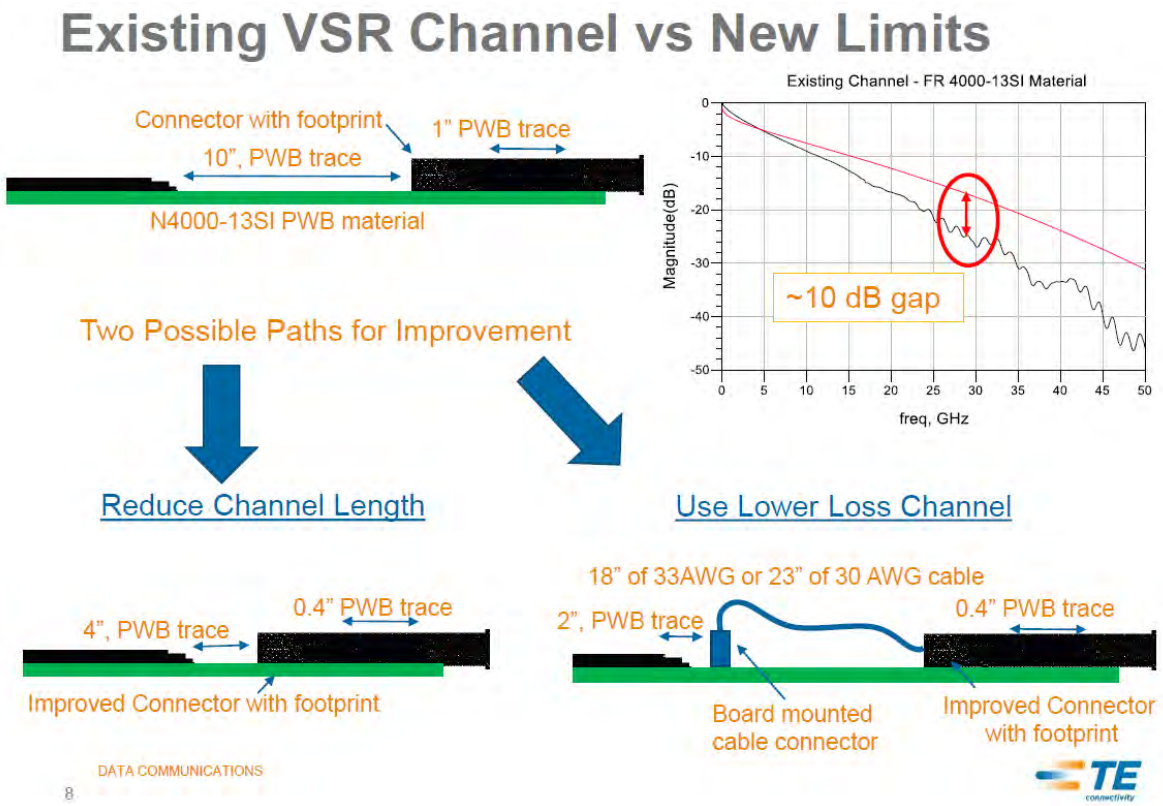

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/tracy_nea_01_0517.pdf

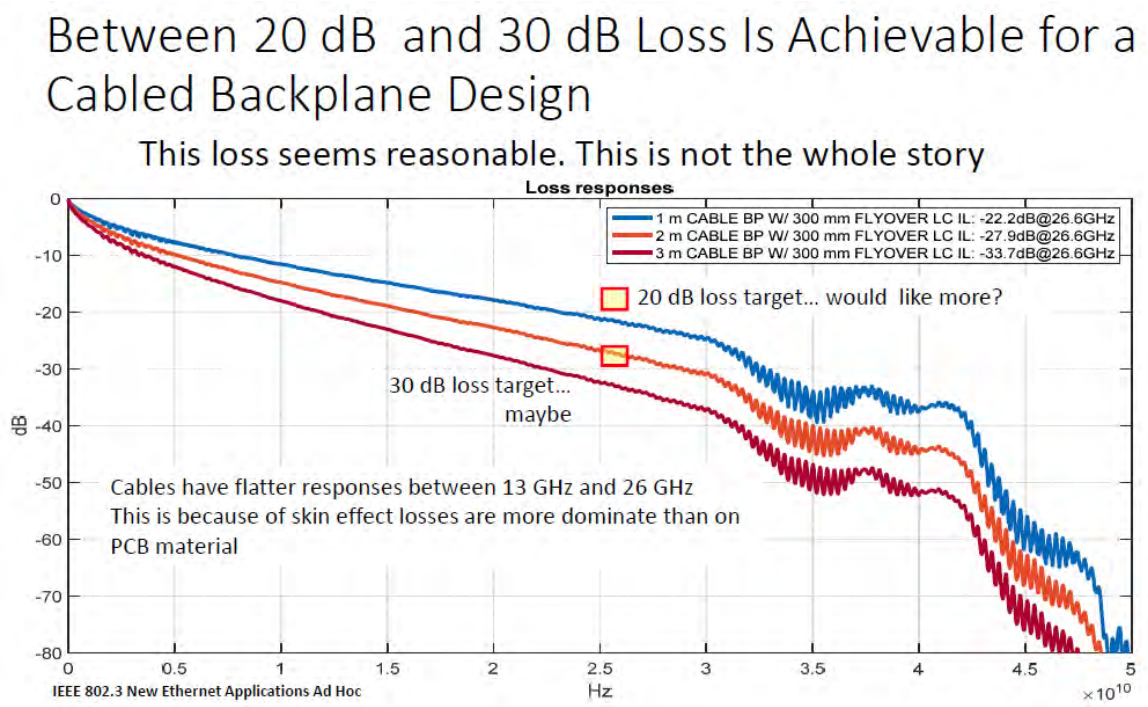

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/mellitz_nea_01a_0517.pdf

Proof of Concepts Shown

PHY Simulation Shows Promising Results

- ❑ BER is about 6.2E-8.
- ❑ TX FIR: 3 pre cursors, 25 post cursors, tail taps are very small.
- ❑ RX: CTLE + DFE.

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/mellitz_nea_01a_0517.pdf

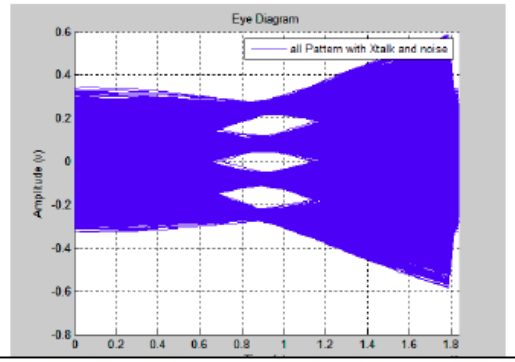

Of course follow-on work will happen.

100G Short Reach Design Results

- A 100Gb/s PAM4 SERDES for short reach has been developed and demoed.
- With 28nm process node, TX eye is clean. Multiple tap TX FIR has been applied for TX eye measurement.

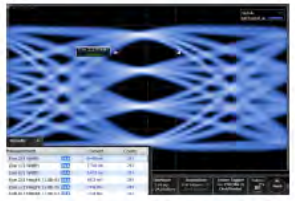
TX Eye Diagram

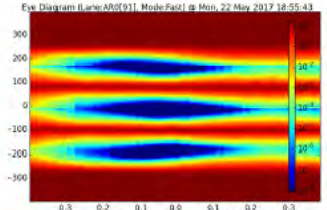
Eye Monitor

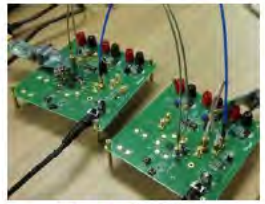
Test Setup

Narva: Proof of Concept Vehicle for 100Gbps SERDES

MACOM

Narva

MACOM

- > Narva
  - > DSP
    - > Tx and Rx Equalization
    - > PRBS31 generator and checker.
  - > Data Converters ADC and DAC.
- > Enables >100Gbps on a single differential pair using PAM-4

- > DAC output @53.125GBaud

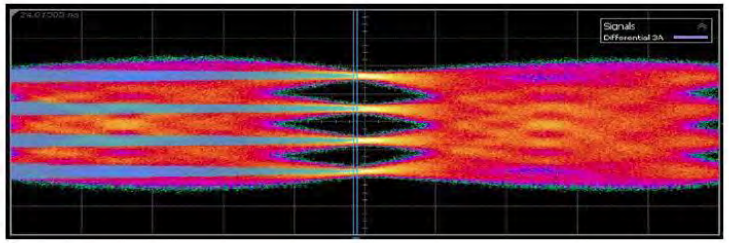

CredÖ

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/sun_nea_01a_0517.pdf

Further Study Needed – Optimize and Tune

PHY Simulation Shows Promising Results

- BER is about 6.2E-8.
- TX FIR: 3 pre cursors, 25 post cursors, tail taps are very small.
- RX: CTLE + DFE.

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/mellitz_nea_01a_0517.pdf

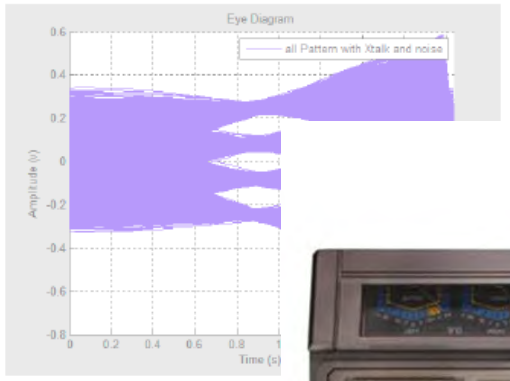

IEEE 802.3 New Ethernet Applications Ad Hoc

Narva

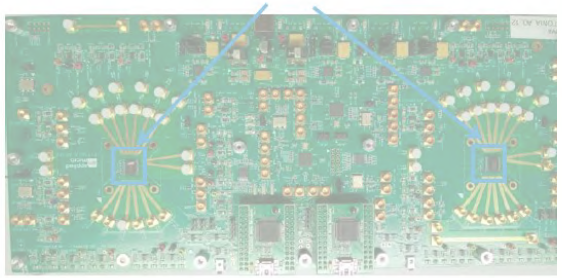

MACOM

> DAC output @53.125GBaud

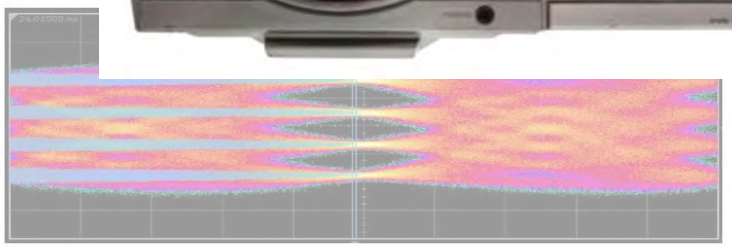

100G Short Reach Design Results

each has been developed and demoed.

ean. Multiple tap TX FIR has been applied for TX eye

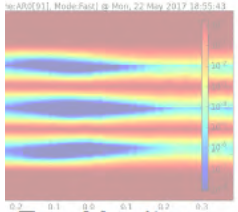

Eye Monitor

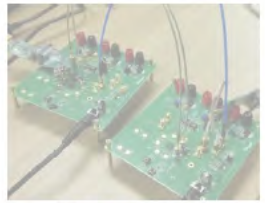

Test Setup

Coding Gain vs. Latency
Equalization
Crosstalk Cancellation
Modulation
Channel Loss
Channel Noise
Impedance Control
Others?

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/palkert_nea_02_0517.pdf

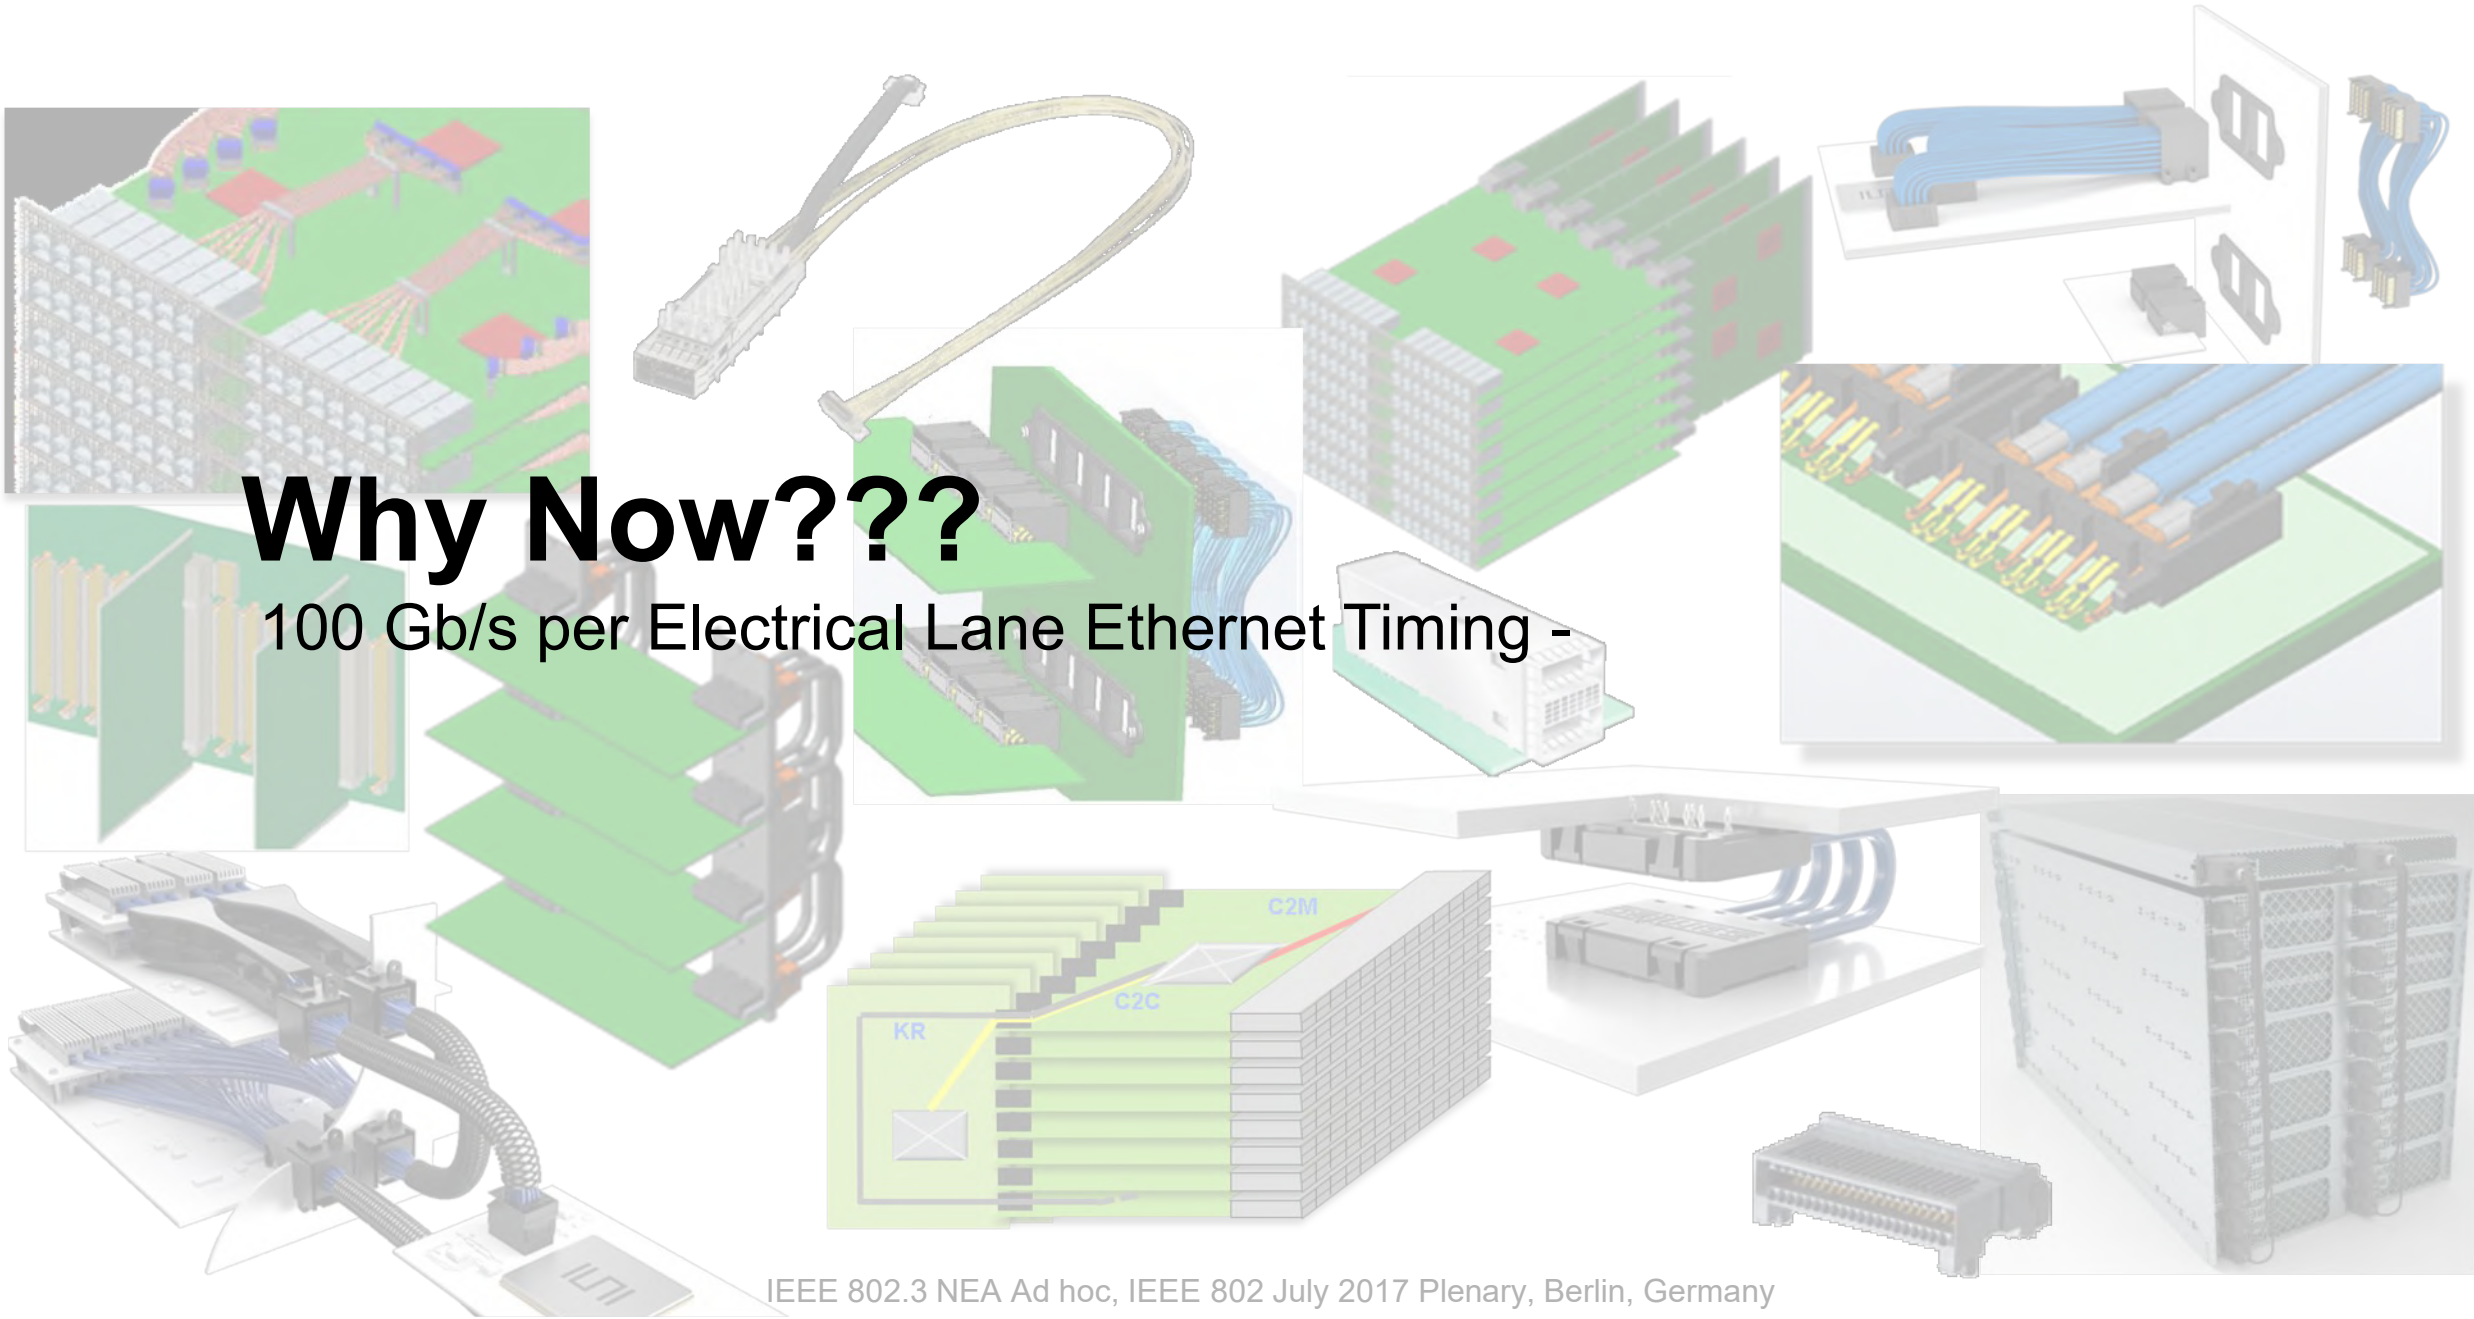

Why Now???

100 Gb/s per Electrical Lane Ethernet Timing -

Historical Perspective – Why 100G Now?

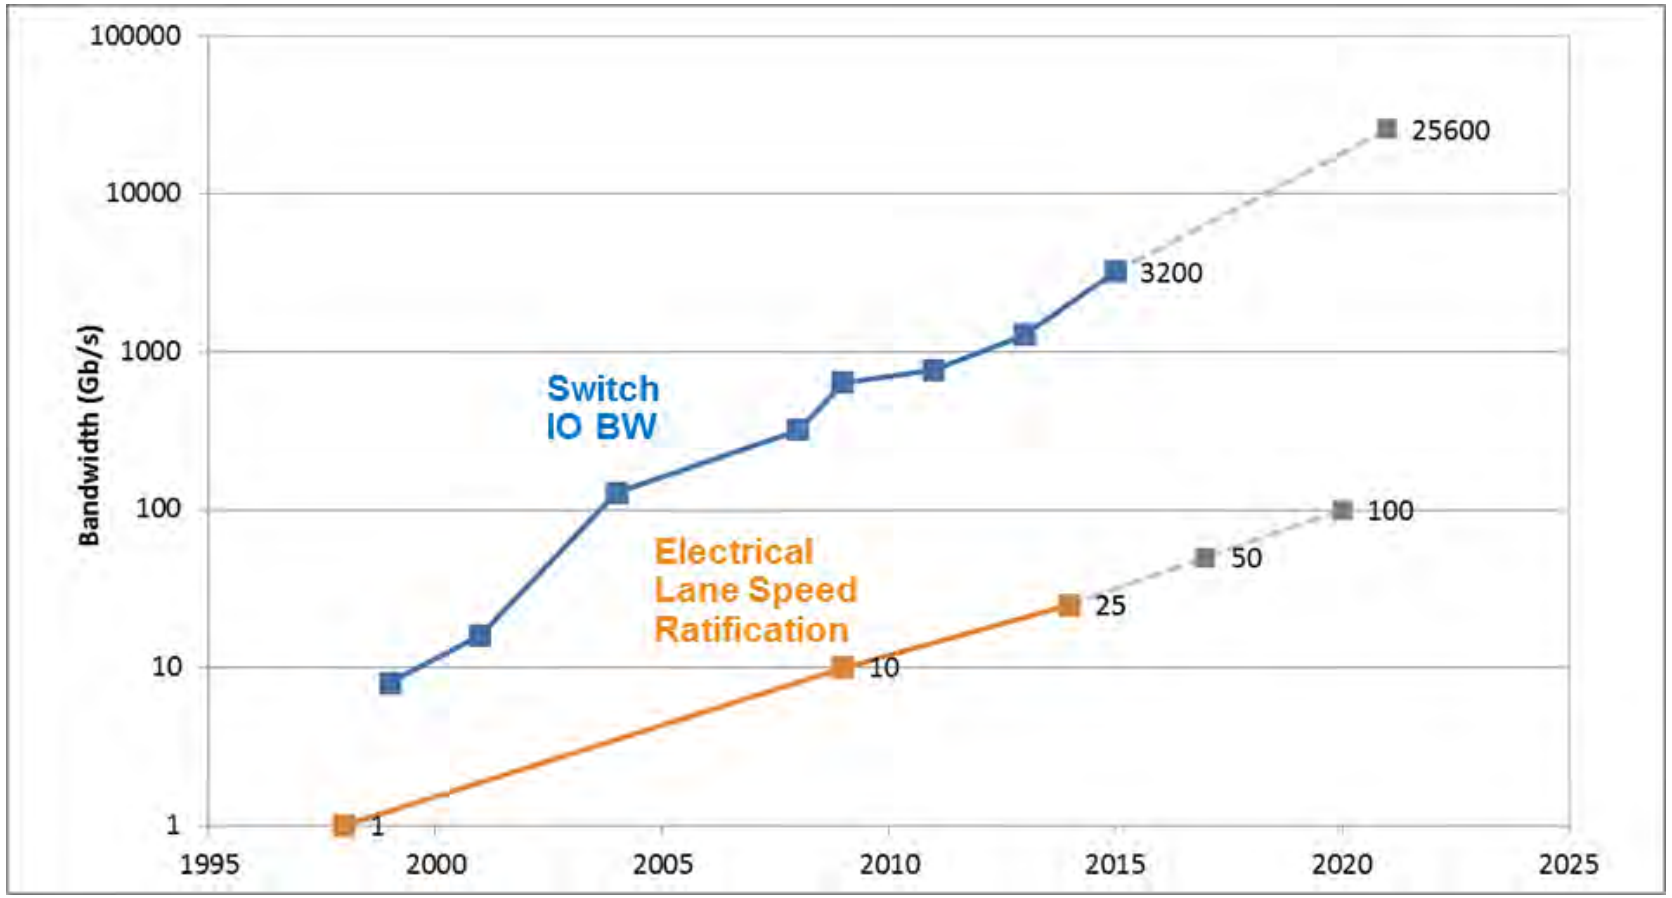

- Historical curve fit to highest rate switch products introduced to market (blue squares)
- Single ASIC IO capacity doubling every ~ 2 years

Enable the Future

	Signaling (Gb/s)	MMF	500m SMF	2km SMF	10km SMF	40km SMF
100GBASE-	10	SR10		10X10		
	25	SR4	PSM4	CWDM4 CLR4	LR4	ER4
	50	SR2				
	100		DR			
200GBASE-	25					
	50	SR4	DR4	FR4	LR4	
	100					
400GBASE-	25	SR16				
	50			FR8	LR8	
	100		DR4			

Includes Ethernet standards in development
Underlined – indicates industry MSA or proprietary solutions

- ***Not*** a statement of objectives, however this effort has the potential to impact many areas of the industry.

The Interest is Here

Straw Poll #2
- Is there interest in developing AUI's based on 100 Gb/s electrical signaling per lane?
- Results
  - Yes – 43
  - No – 2
  - Maybe – 15

Straw Poll #3
- Is there interest in developing Backplane / Copper Cable PHYs based on 100 Gb/s electrical signaling per lane?
  - Yes – 18
  - No – 10
  - Maybe – 20

March 2017**

Straw Polls

May 2017*

Straw Poll #1
- I would support development of a CFI that includes:
  a) new backplane PHY,
  b) new Passive Copper Cable PHY
  c) Chip-to-chip (C2C AUI)
  d) Chip-to-module (C2M AUI)
  e) other
  f) not at this time,
  g) none of the above
- Results
  a) 32
  b) 26
  c) 40
  d) 48
  e) 0
  f) 2
  g) 0

Taken from NEA Ad Hoc unapproved Minutes

* http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/minutes_nea_0517_unapproved.pdf

** http://www.ieee802.org/3/ad_hoc/ngrates/public/17_03/minutes_nea_0317_unapproved.pdf

Areas of Study

1. Chip-to-Chip AUI
2. Chip-to-Module AUI
3. Backplane
4. Potential changes of existing PHYs

NOTE: This is not a laundry list of items to be debated today, but this CFI enables study in these areas.

Summary

- 100 Gb/s is the next step on “Follow the SerDes” and continues existing market trends
  - Switching capacity progression
  - Reduction of interface width
- 100 Gb/s per lane seems feasible with a few tools still available
- Impact of 100 Gb/s Electrical Signaling is broad across the Ethernet Family
  - AUIs for existing PHYs for existing rates
  - For new PHYs for existing rates
  - New AUI’s / PHYs for new rates?
- Let’s form a Study Group!!

Straw Polls for CFI

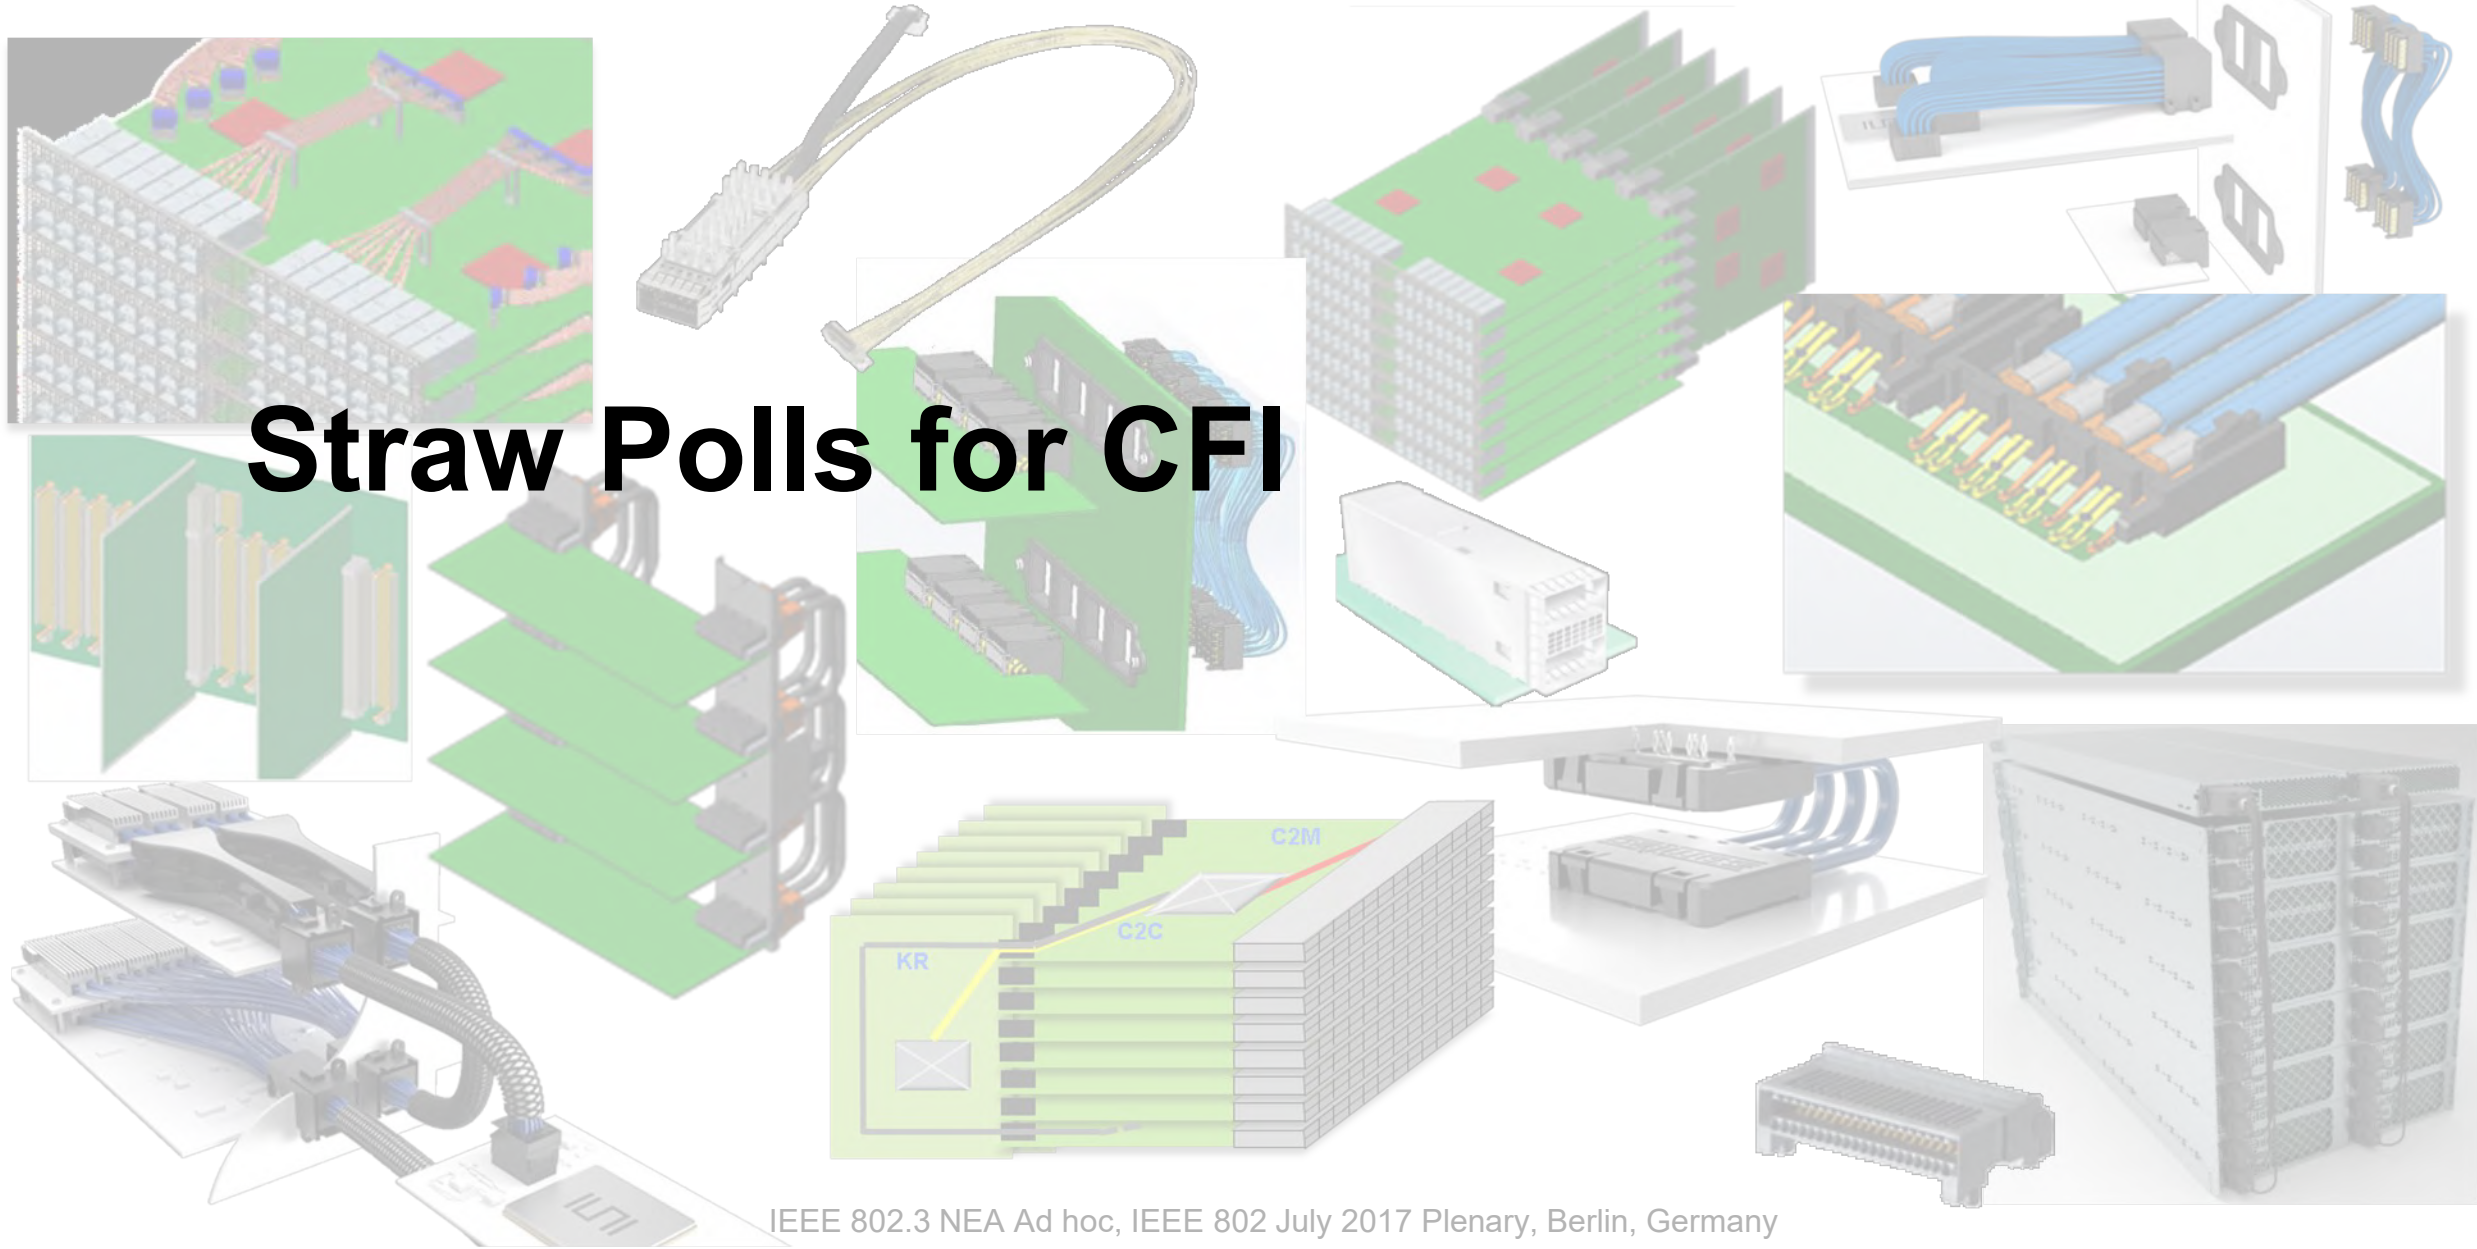

Call-for-Interest Consensus

- Should a study group be formed for "100 Gigabit/s per Lane for Electrical Interfaces and PHYs"?
- Y: N: A:

- Room count:

Participation

- I would participate in a “100 Gigabit/s per Lane for electrical interfaces and PHYs” study group in IEEE 802.3
  - Tally:

- My company would support participation in a “100 Gigabit/s per Lane of Electrical Interfaces and PHYs” study group
  - Tally:

Future Work

- Ask 802.3 at Thursday's closing meeting to form a "100 Gigabit/s per lane for electrical interfaces and PHYs" study group
- Prepare ITU liaison letter for WG approval if Study Group formation is approved by WG.
- If approved:
  - 802 EC informed on Friday of formation of the study group
  - First study group meeting would be during January 2018 IEEE 802.3 interim meeting

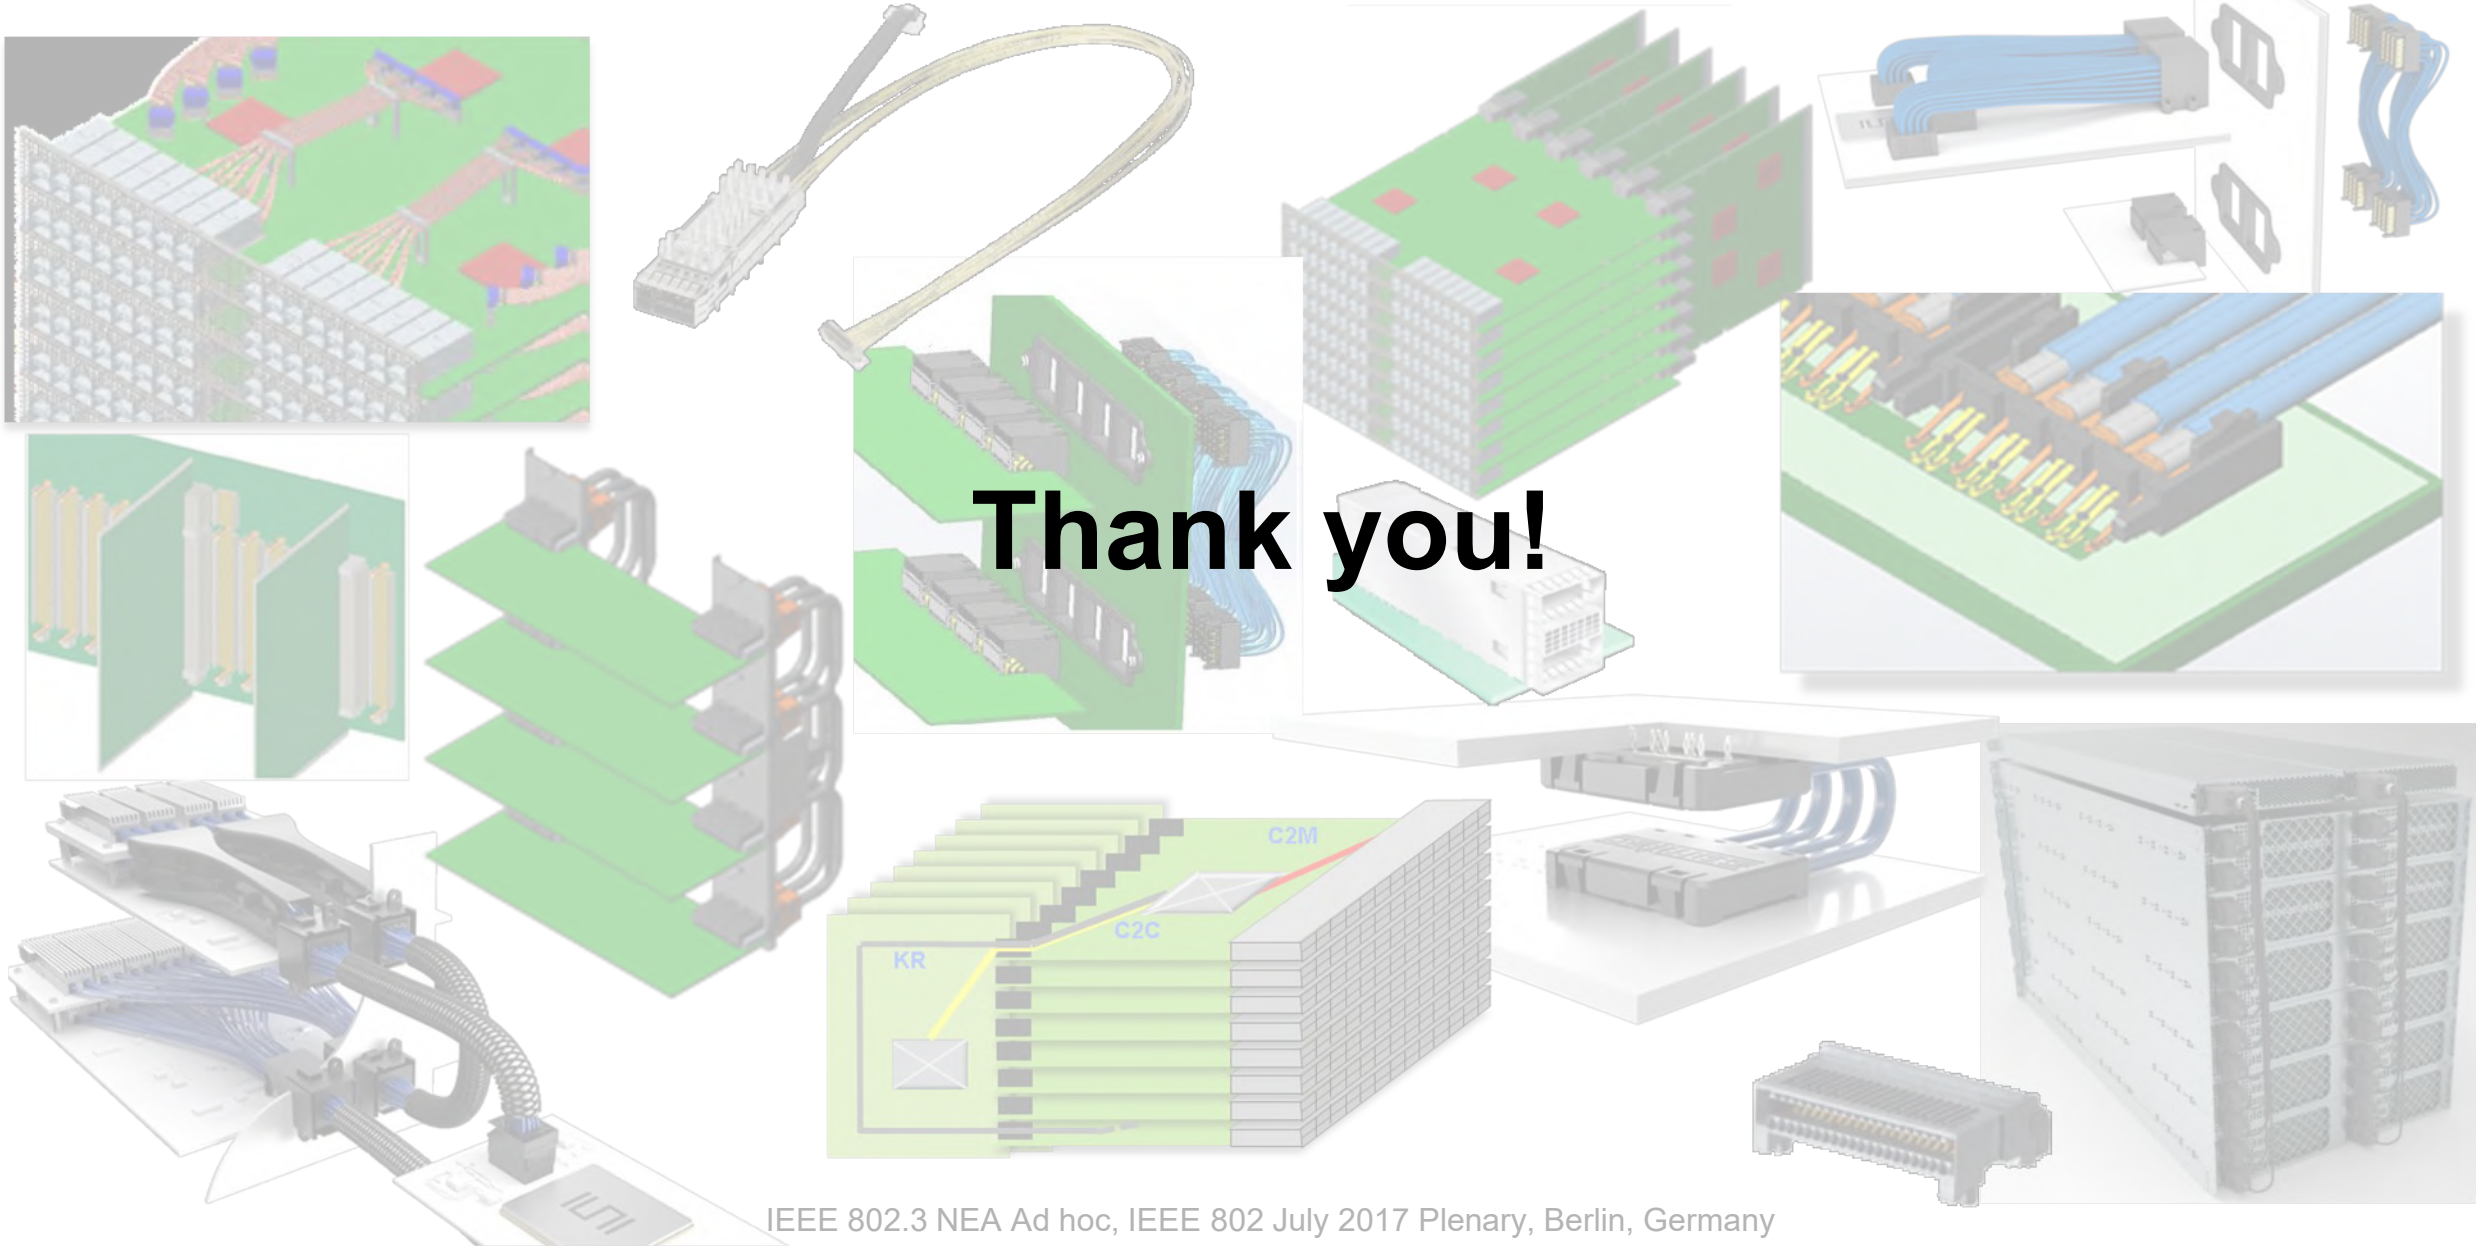

Thank you!