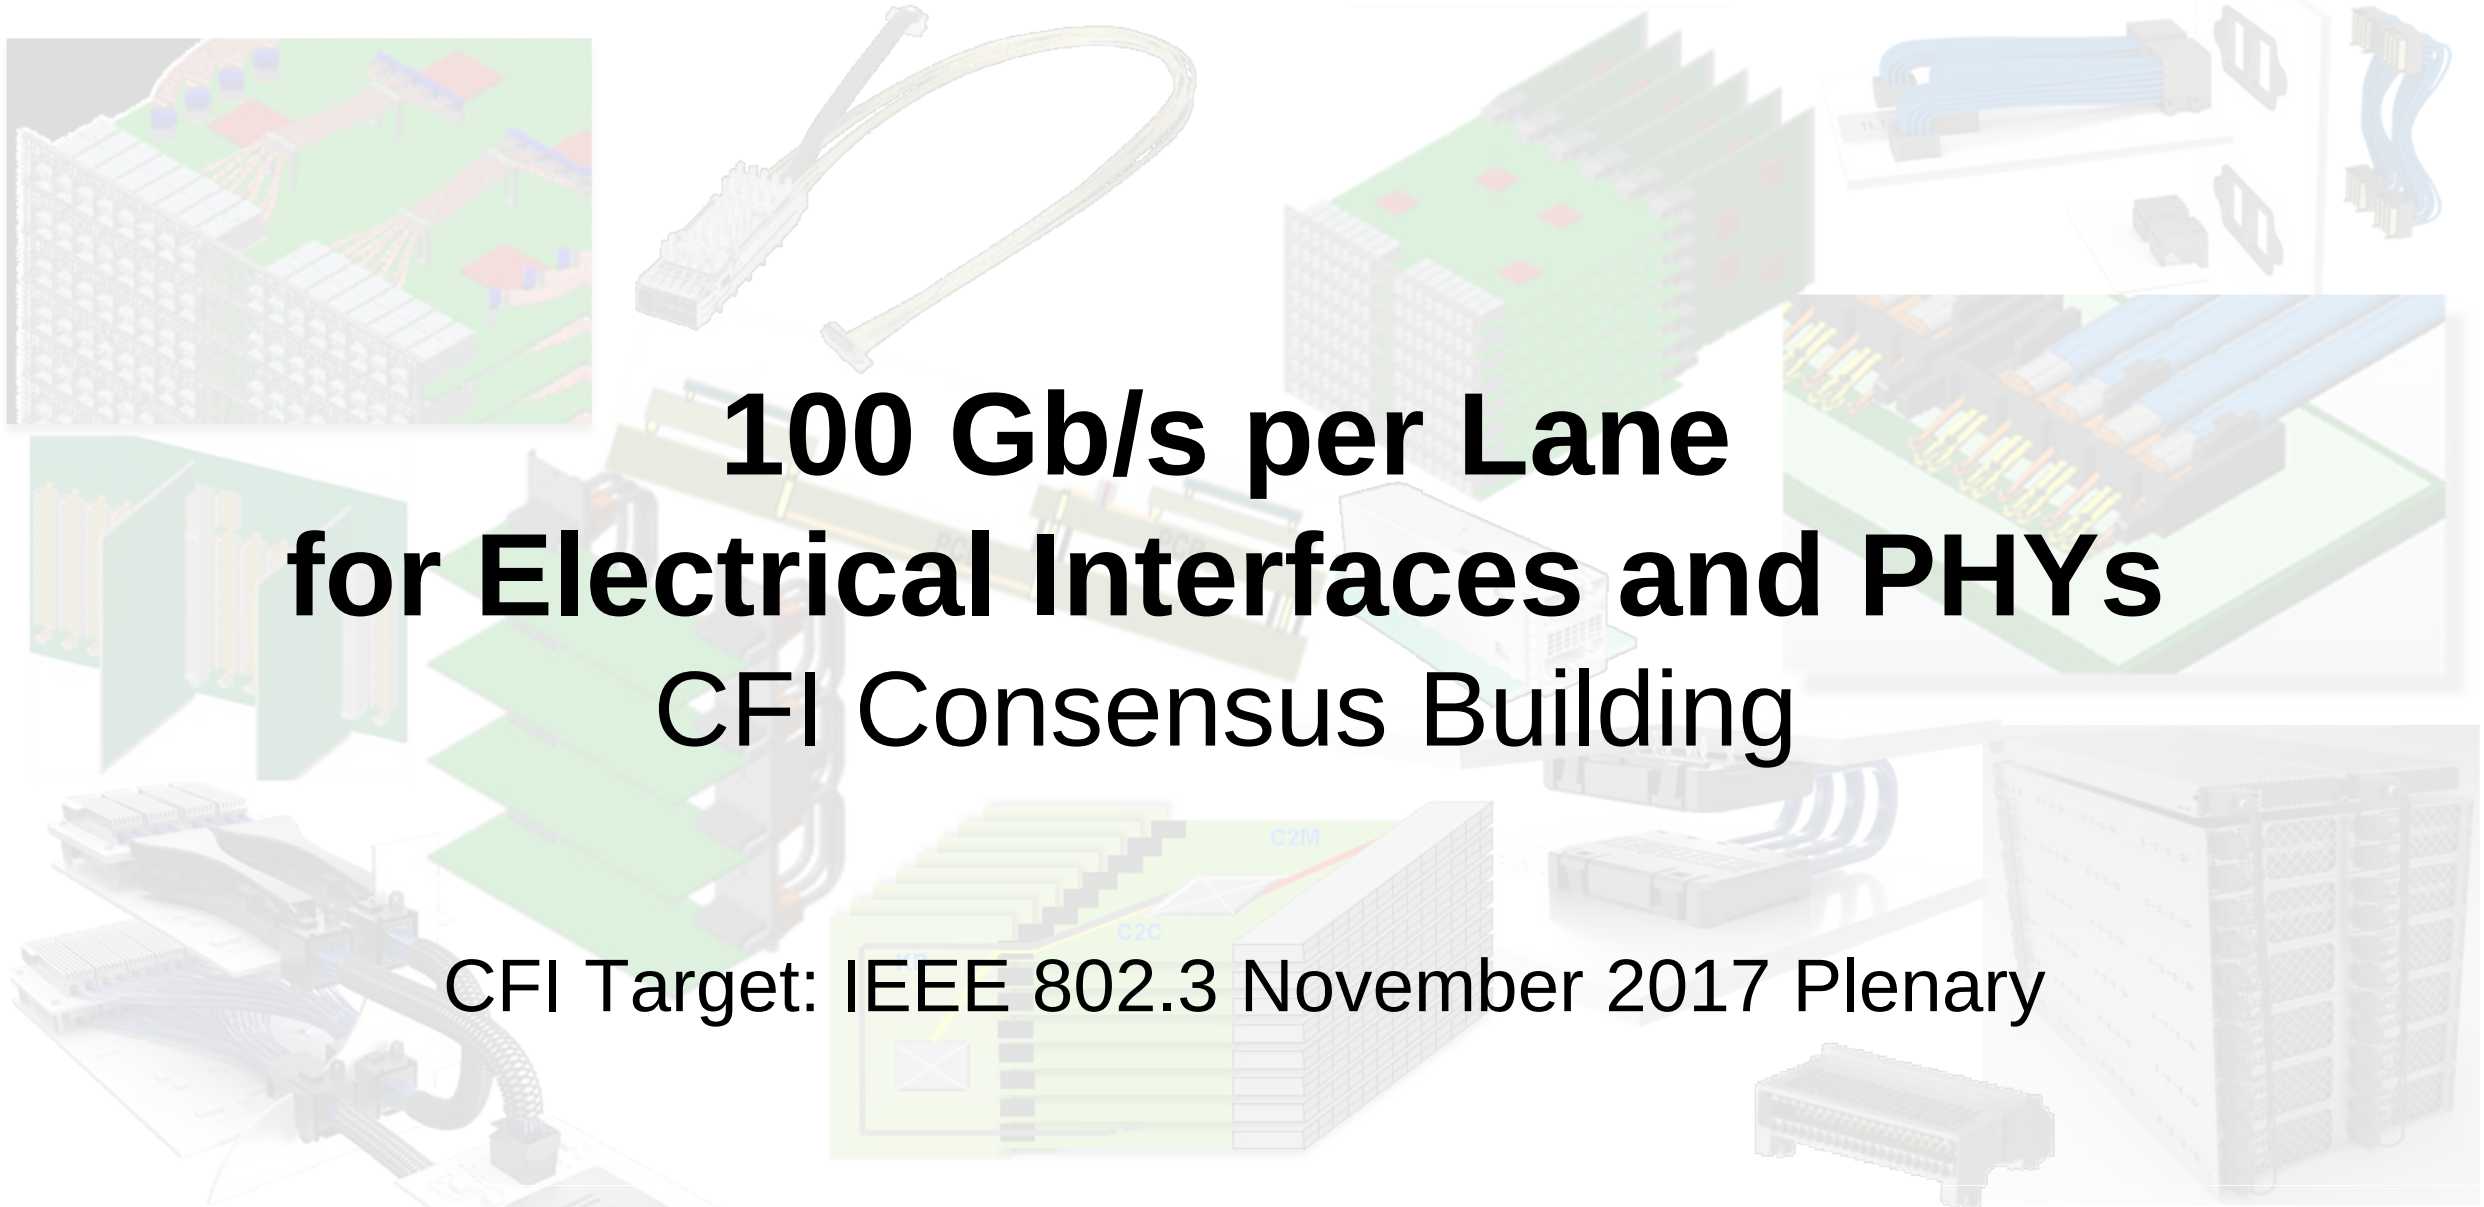

100 Gb/s per Lane for Electrical Interfaces and PHYs

CFI Consensus Building

CFI Target: IEEE 802.3 November 2017 Plenary

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

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!

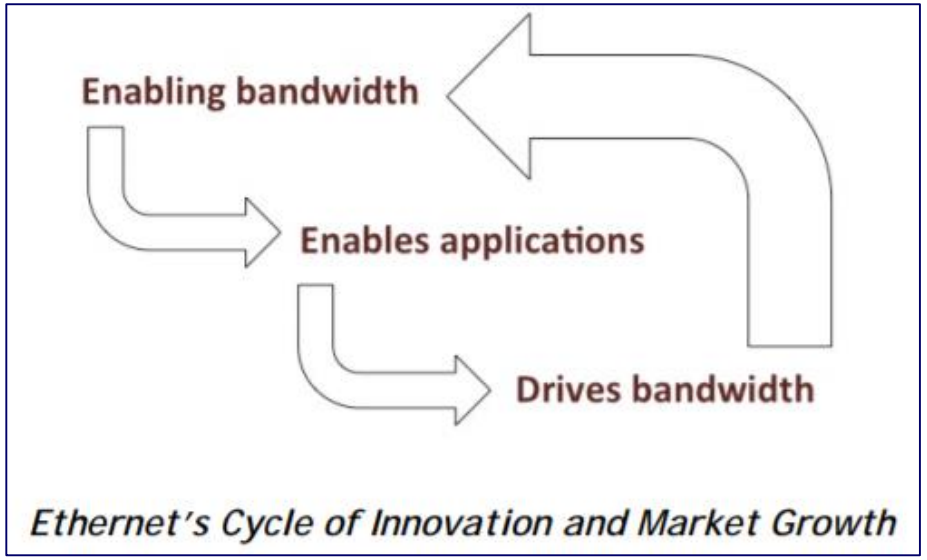

Ethernet's Cycle of Innovation and Market Growth

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

Electrical interfaces come in many shapes and sizes.

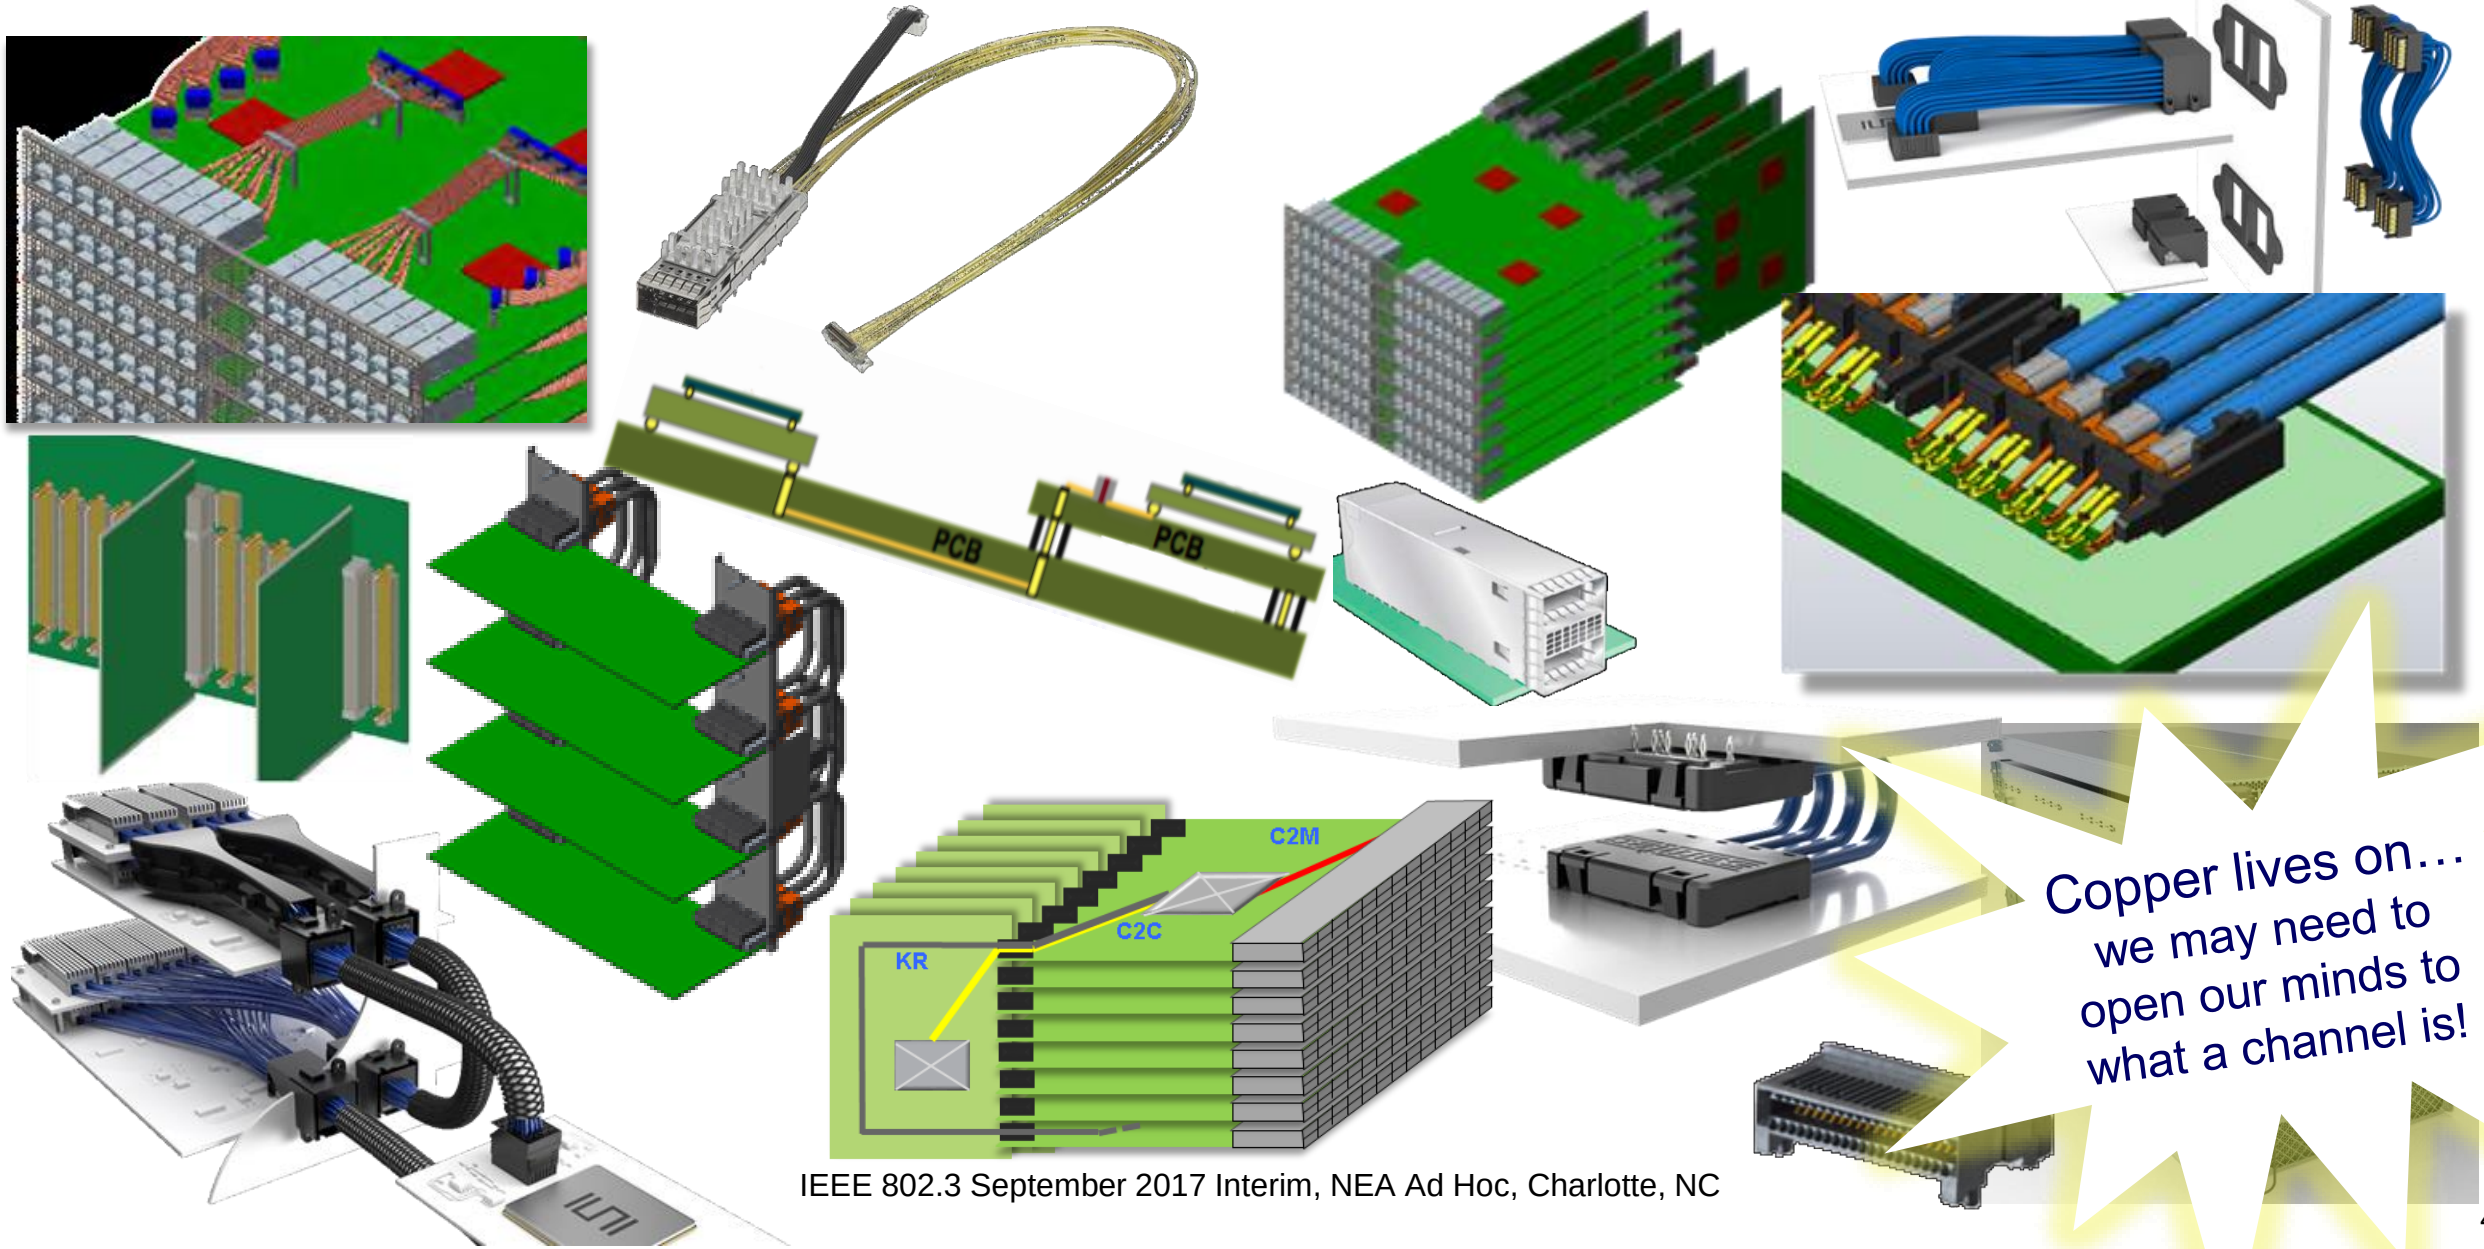

Tonight's Meeting

- To present the

  market ***NEED***,

  technical **Feasibility**,

  and *Why Now??*

  of 100Gb/s per lane of electrical signaling.

- To gain consensus towards Thursday's Call-for-Interest.
- 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

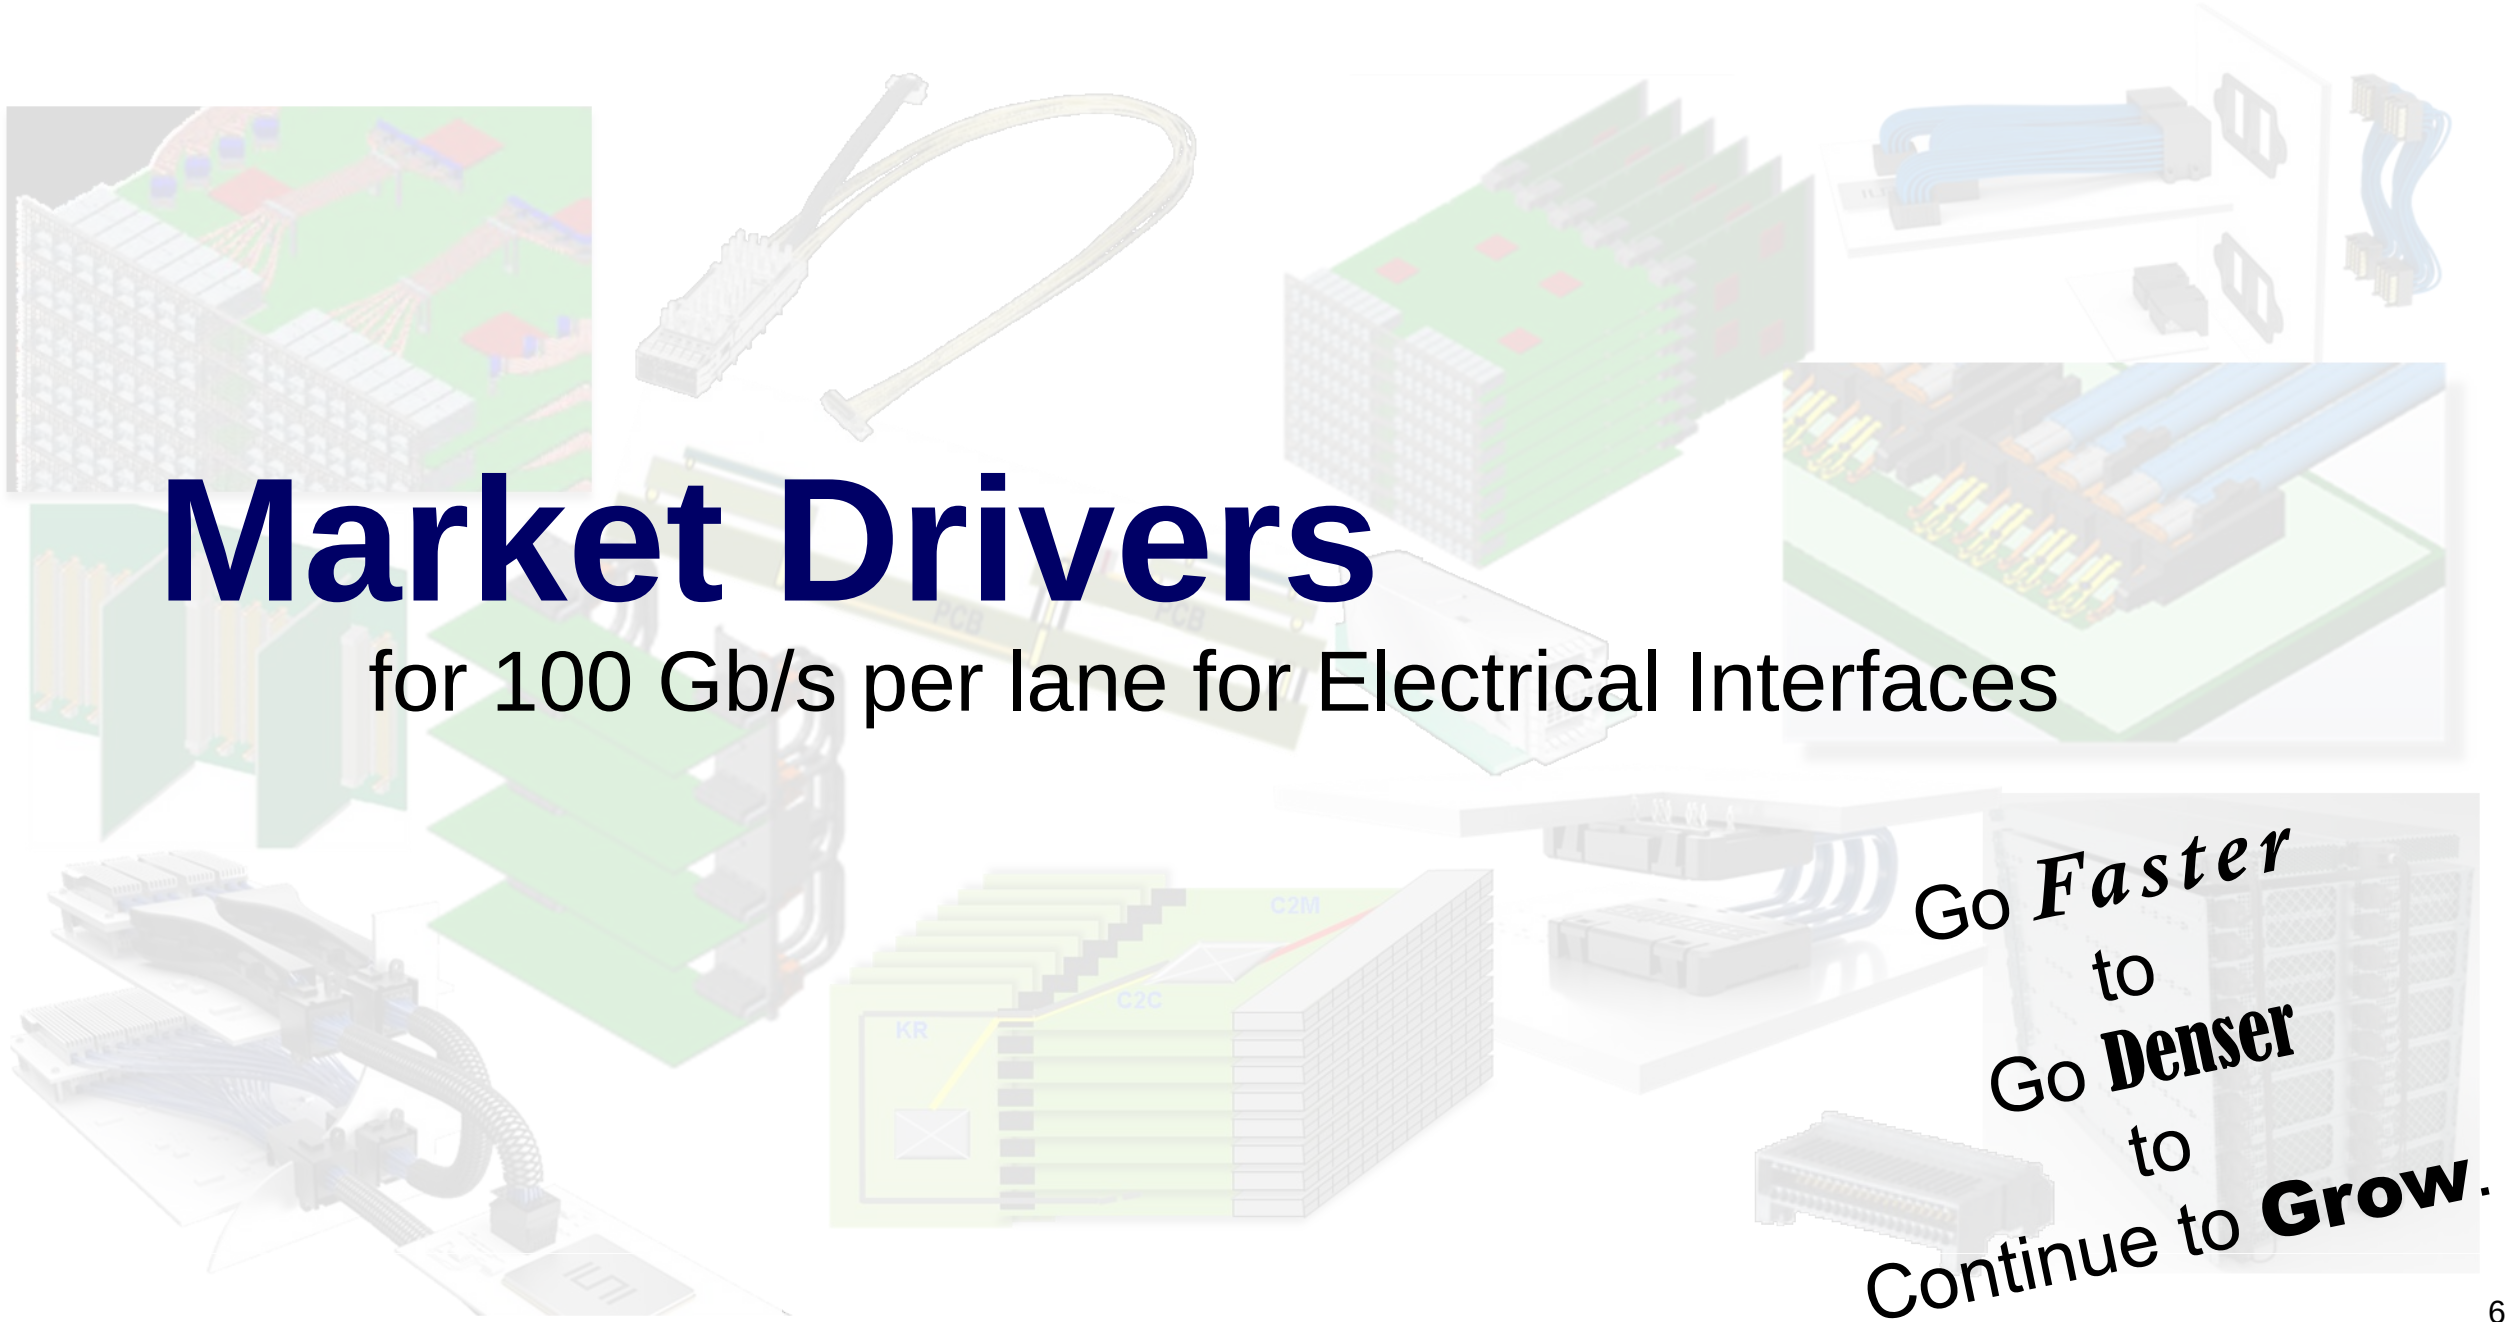

Go *Faster*
to
Go **Denser**
to
Continue to **Grow**.

What Are We Talking About?

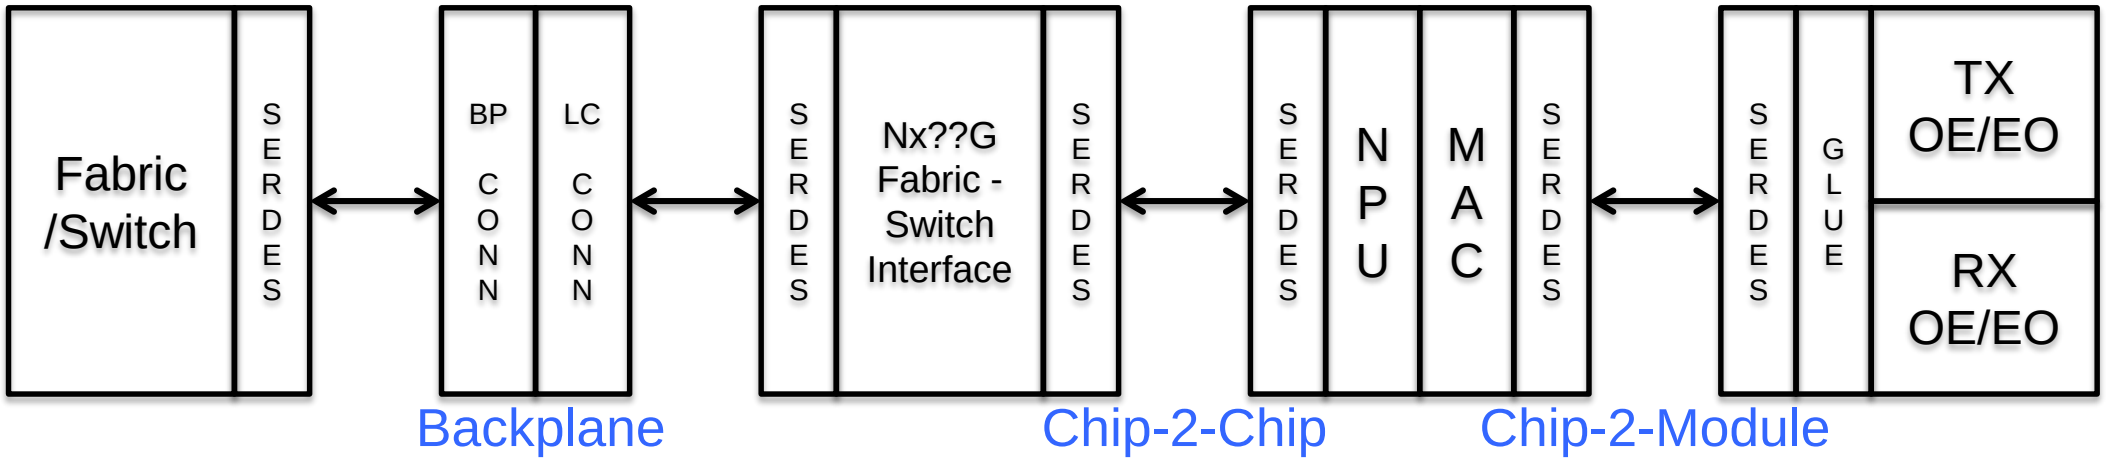

Consider how many instances of the interface can exist in the system…

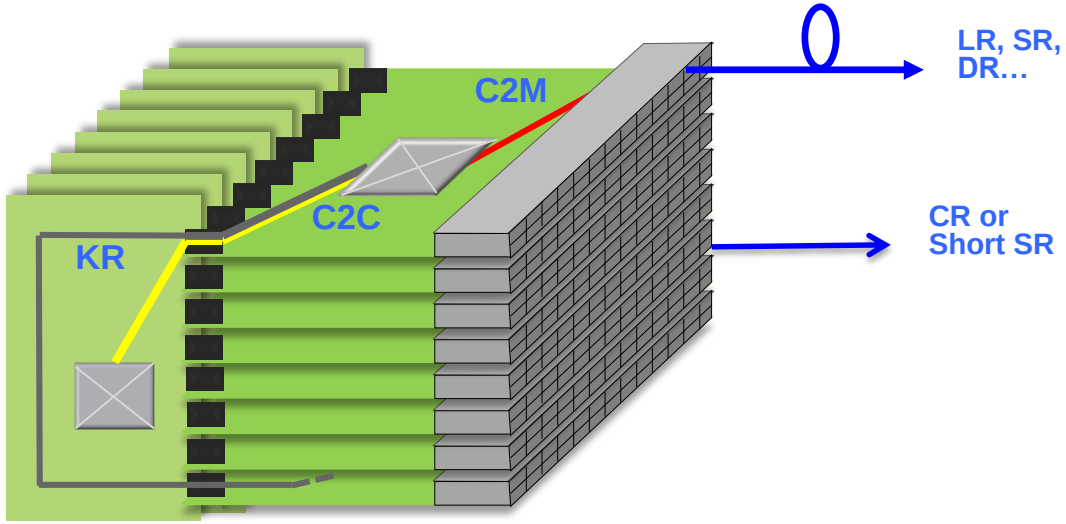

Historical Perspective Shows What's Coming

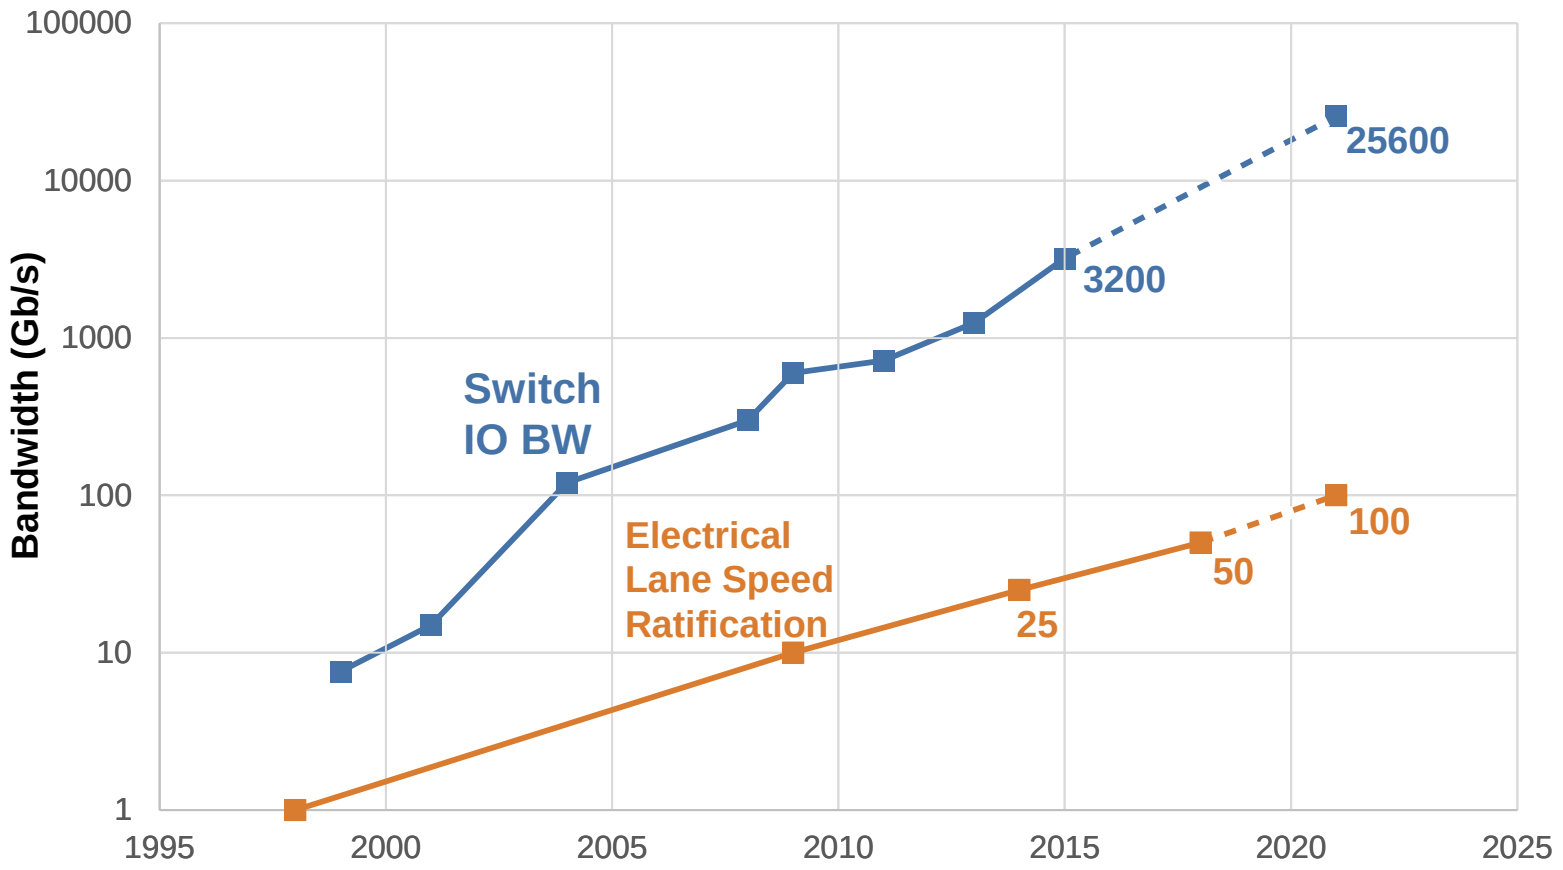

- Historical curve fit to highest rate switch products introduced to market (blue squares)
- Single ASIC IO capacity doubling every ~ 2 years

IO Escape forcing transition to higher lane speeds

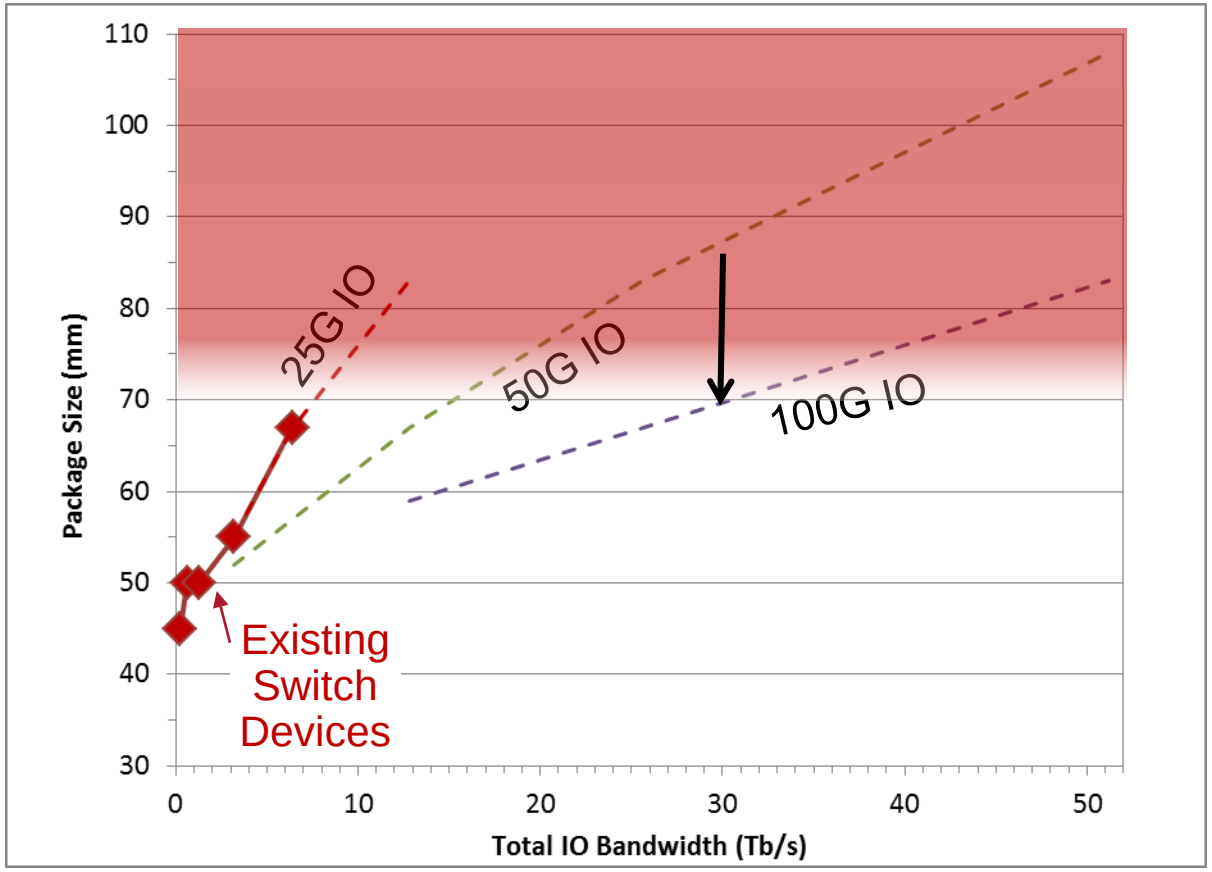

- ~ 70mm package is a current BGA practical maximum (due to coplanarity / warpage)
- This will force BGA devices with > 14Tb/s of aggregate bandwidth to transition to lane rates of higher greater than 50G (possibly 100G?)

Backplane is easily system bottleneck

Existing Systems

Defined backplane and pin count

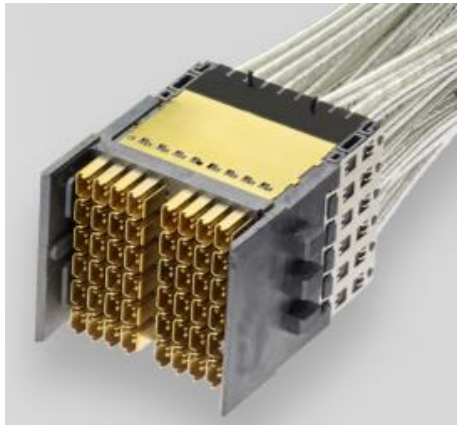

No choice but to put more signal across the pin

New Systems

Can tune pitch and pin count

However, there is limited gain left in this mechanical density

backplane speed needs to scale for bandwidth to grow

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		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
Blue – indicates the areas of interest for this CFI

Technical Feasibility

for 100 Gb/s per lane for Electrical Interfaces

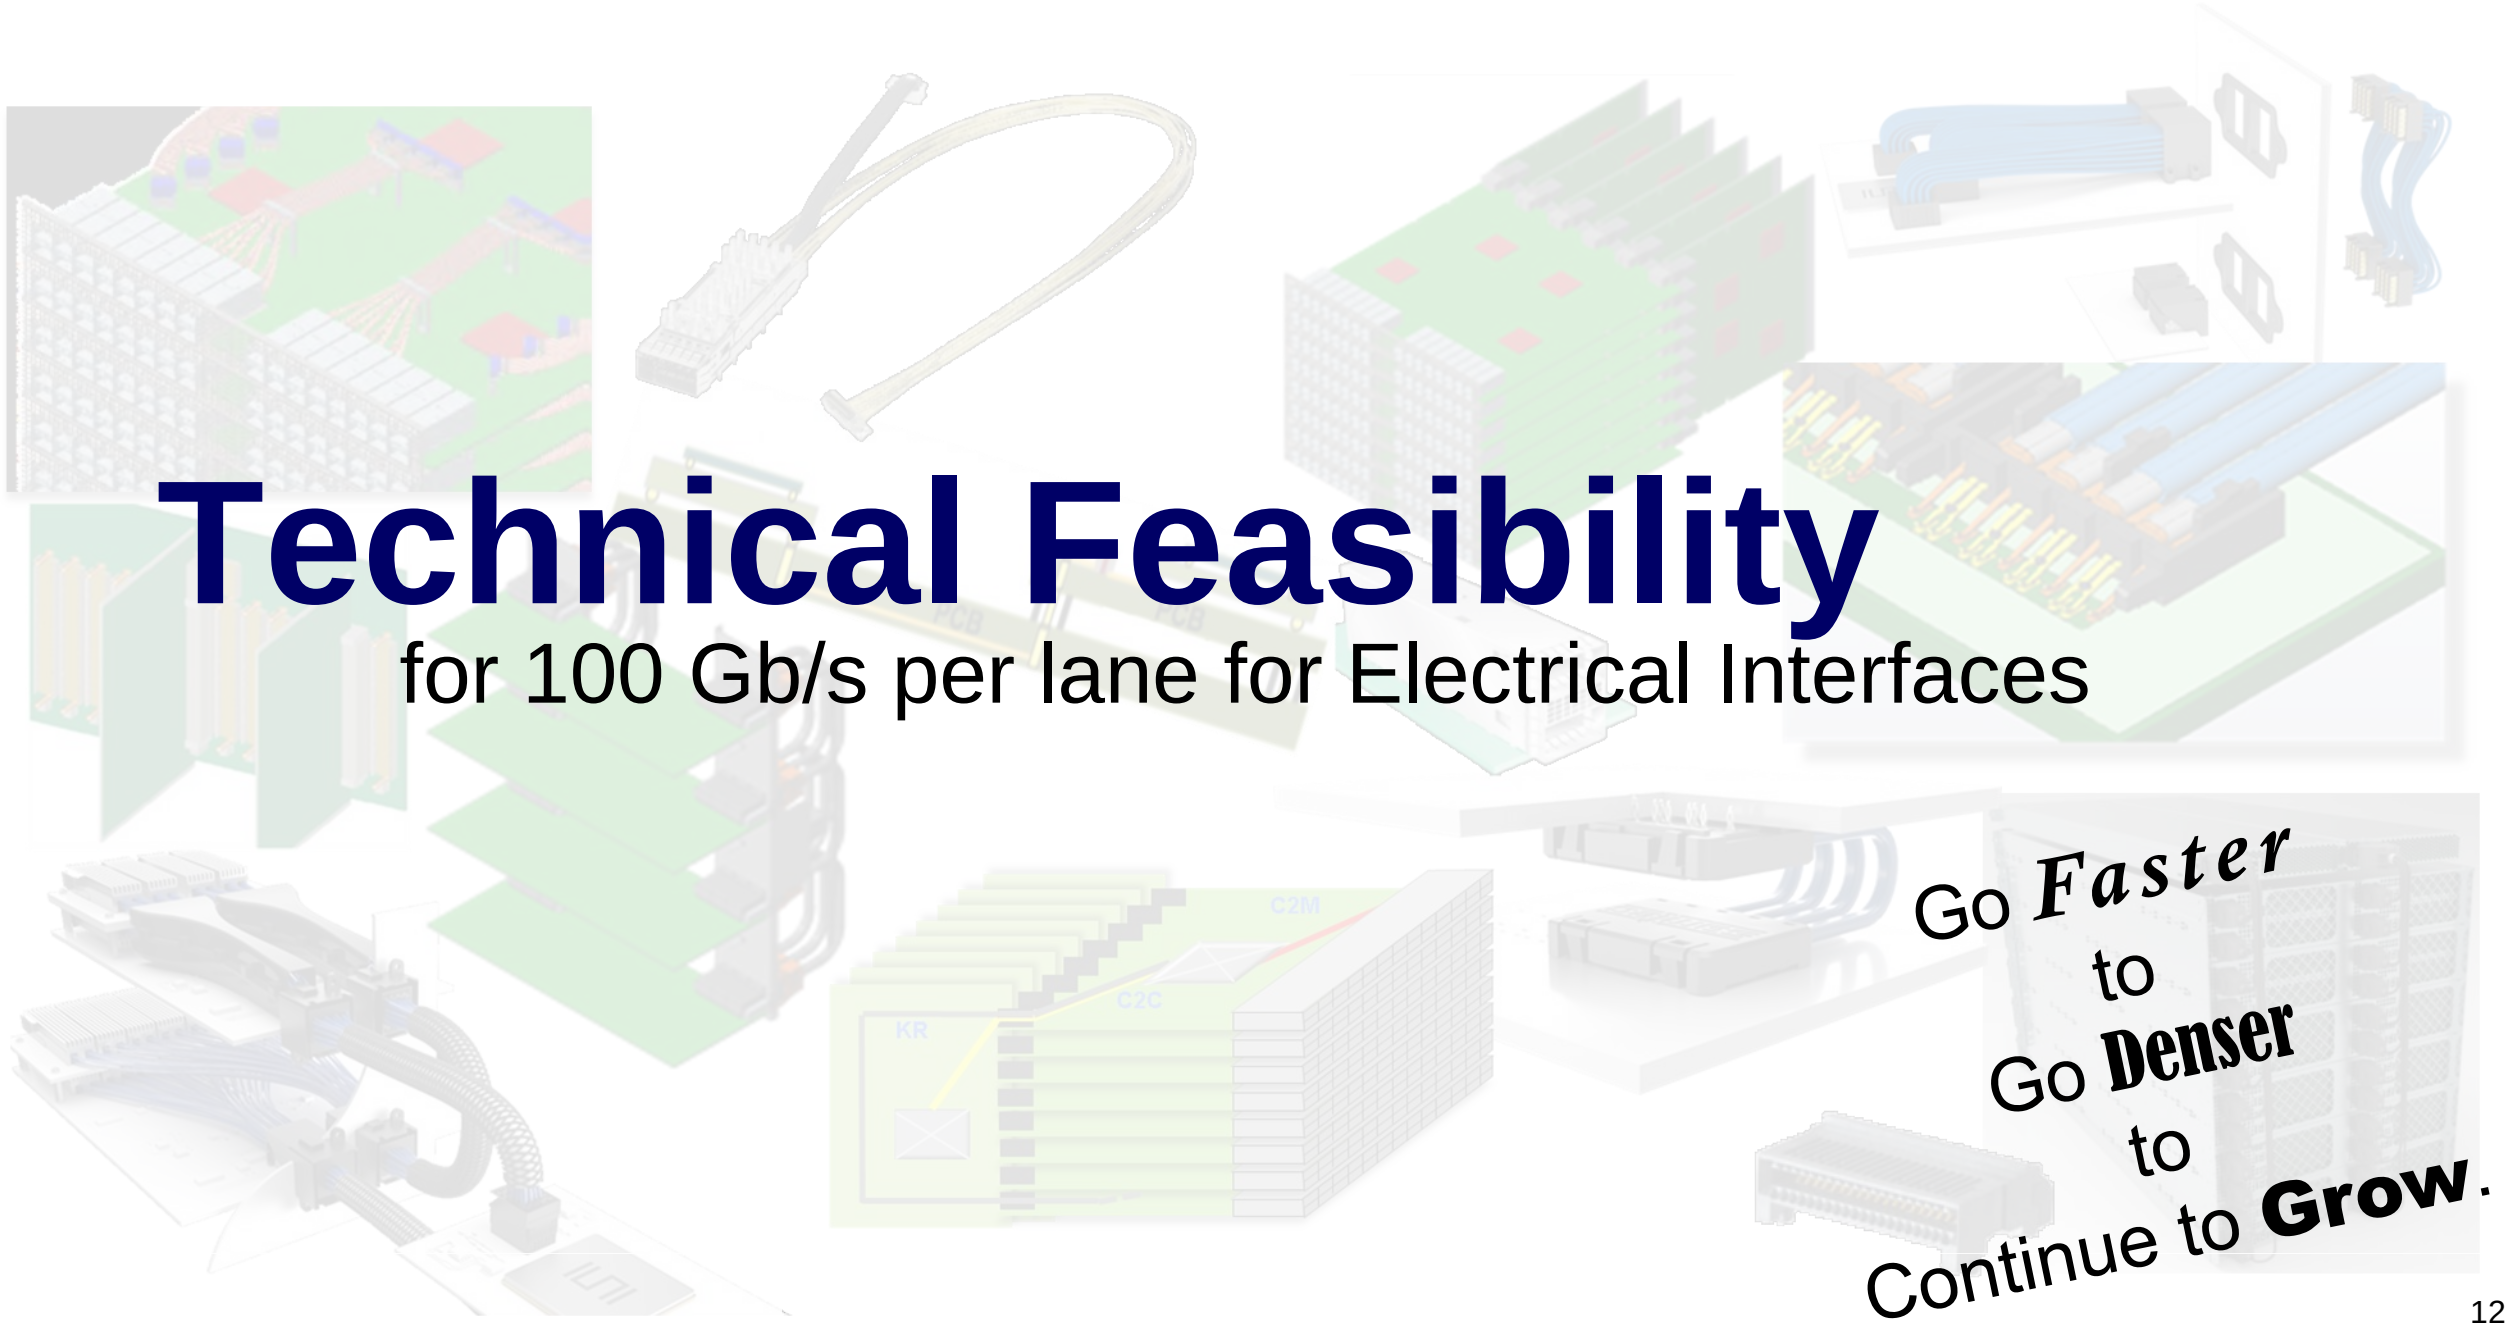

Go *Faster*
to
Go **Denser**
to
Continue to **Grow**.

It’s time to open the toolbox again…

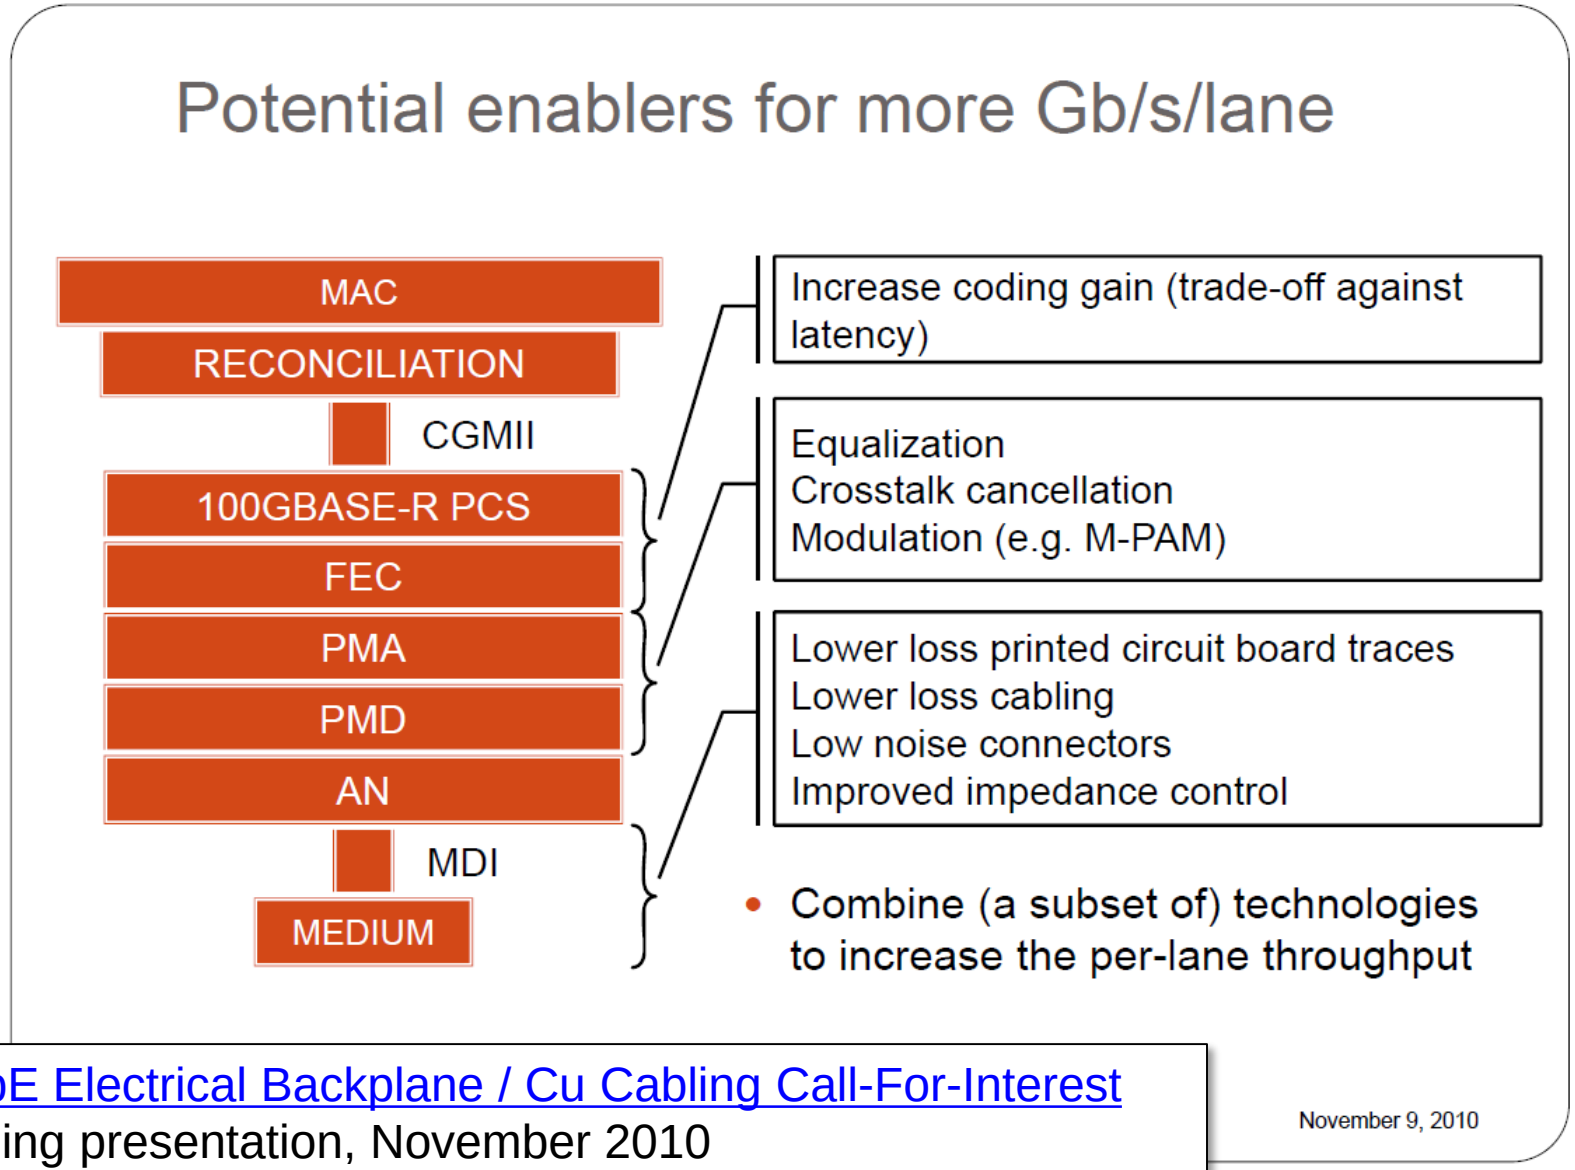

From the [100GbE Electrical Backplane / Cu Cabling Call-For-Interest] consensus building presentation, November 2010

Solutions for each generation

SerDes | Backplane | Cable

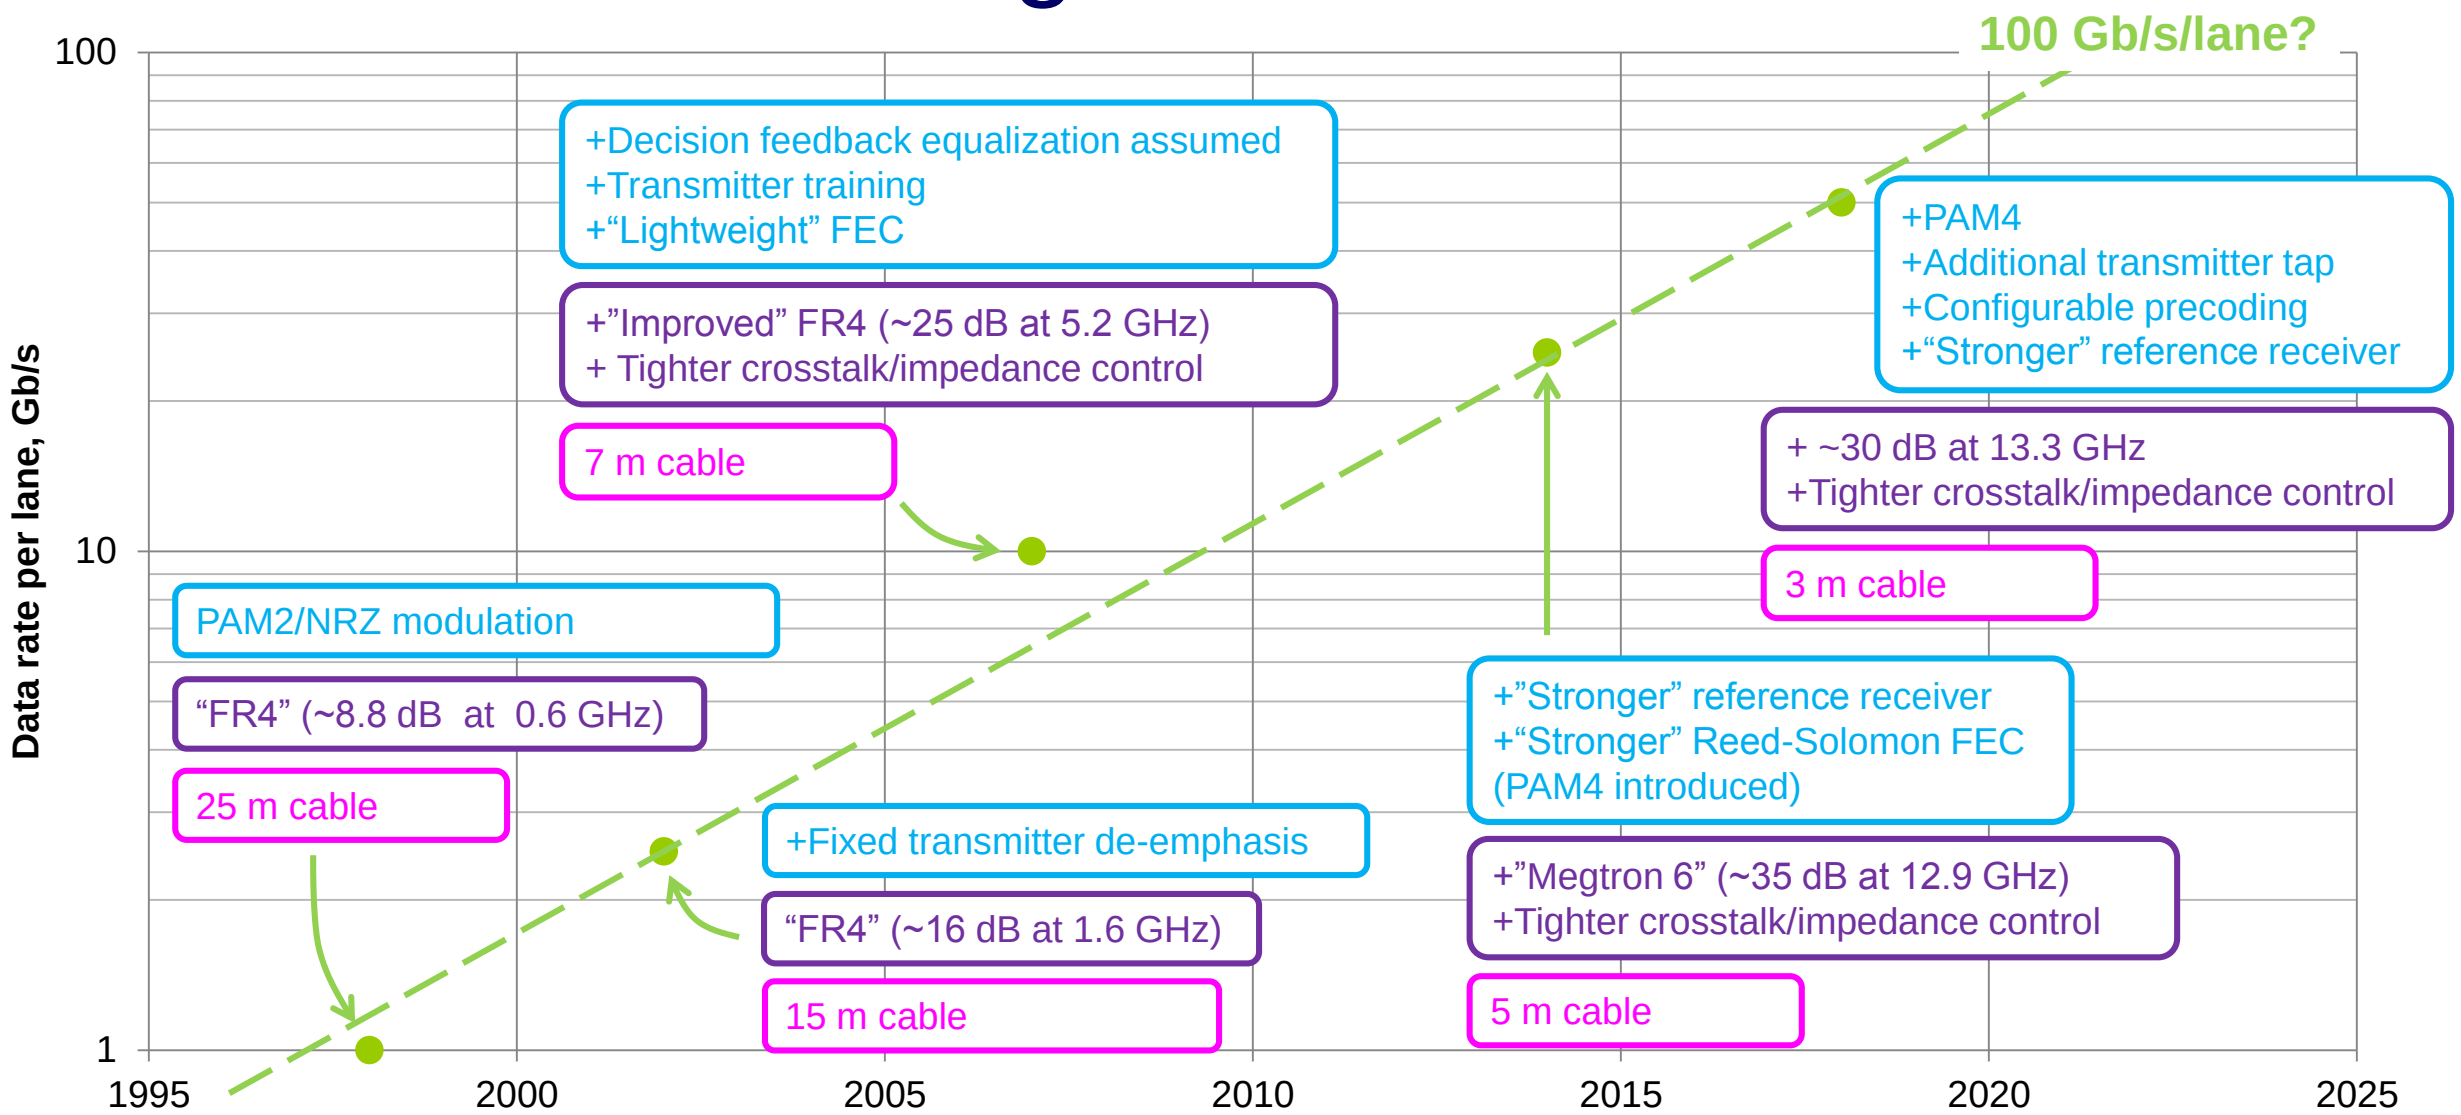

* Dates are approximate

Different constraints for different applications

Chip-to-module

- "Coexistence" with defined PHYs, FEC, PCS?
- What is the minimum insertion loss that supports useful topologies?
- Consider improved PCB materials, PCB vs. cable, improvements in impedance/noise control?

Chip-to-chip and "backplane"

- What is the minimum insertion loss that support useful topologies?
- Consider improved PCB materials, PCB vs. cable, improvements in impedance/noise control?

Cable

- What is the minimum useful reach?
- Consider "middle-of-rack" topologies?

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

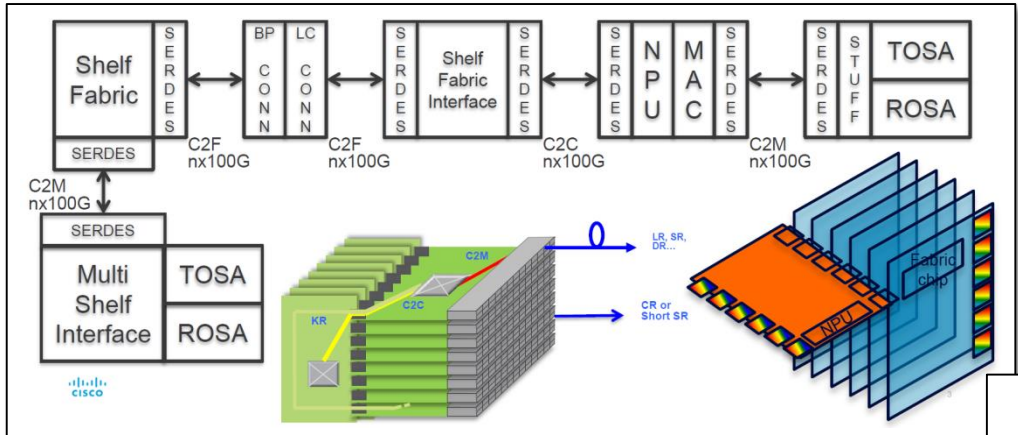

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/goergen_nea_01_0517.pdf

System considerations
Channel options
Higher-speed SerDes

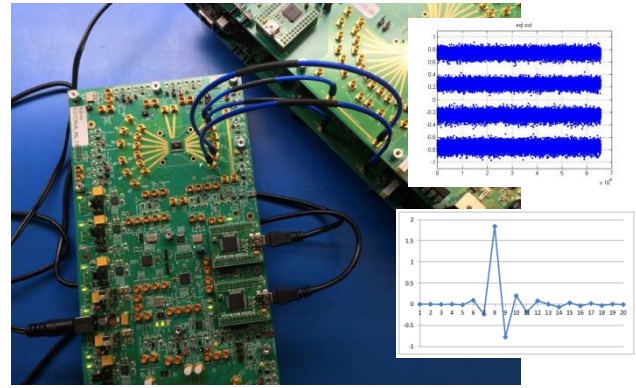

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/palkert_nea_02_0517.pdf

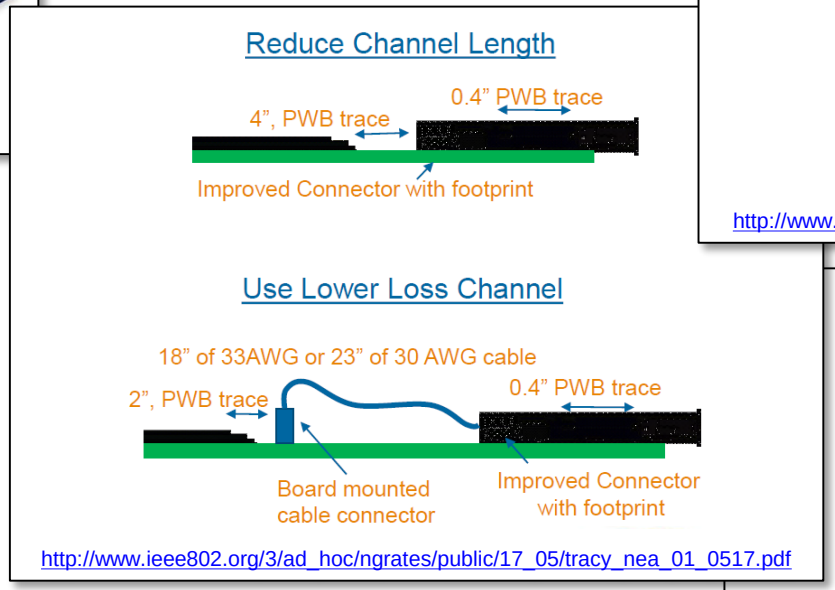

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/tracy_nea_01_0517.pdf

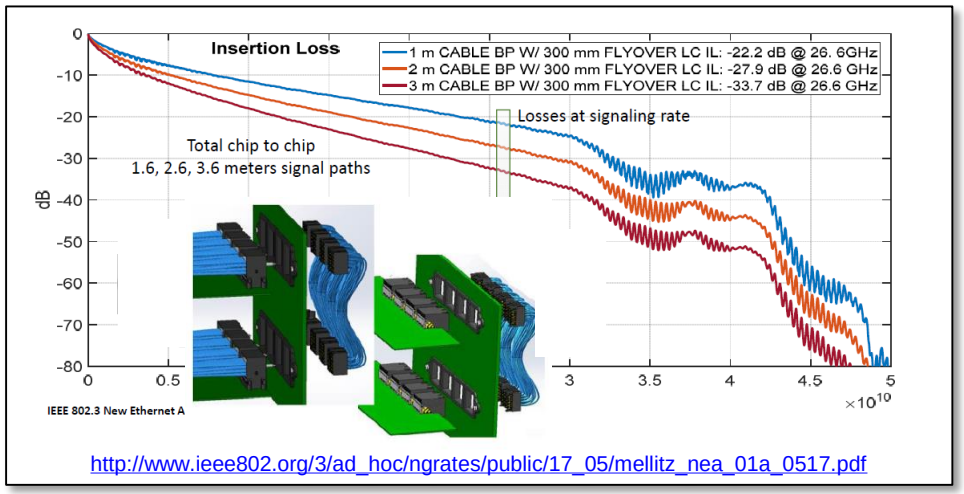

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/mellitz_nea_01a_0517.pdf

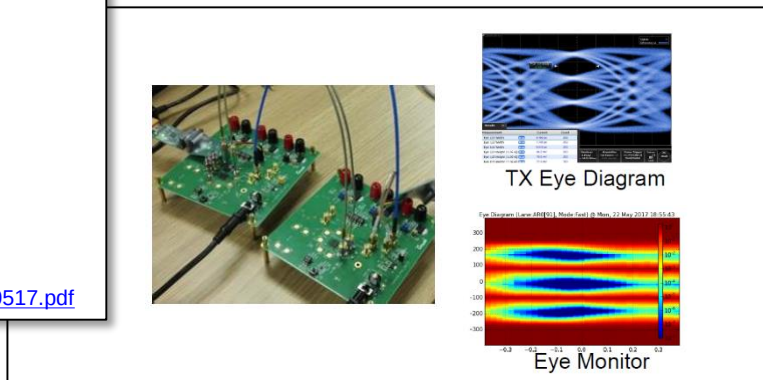

http://www.ieee802.org/3/ad_hoc/ngrates/public/17_05/sun_nea_01a_0517.pdf

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

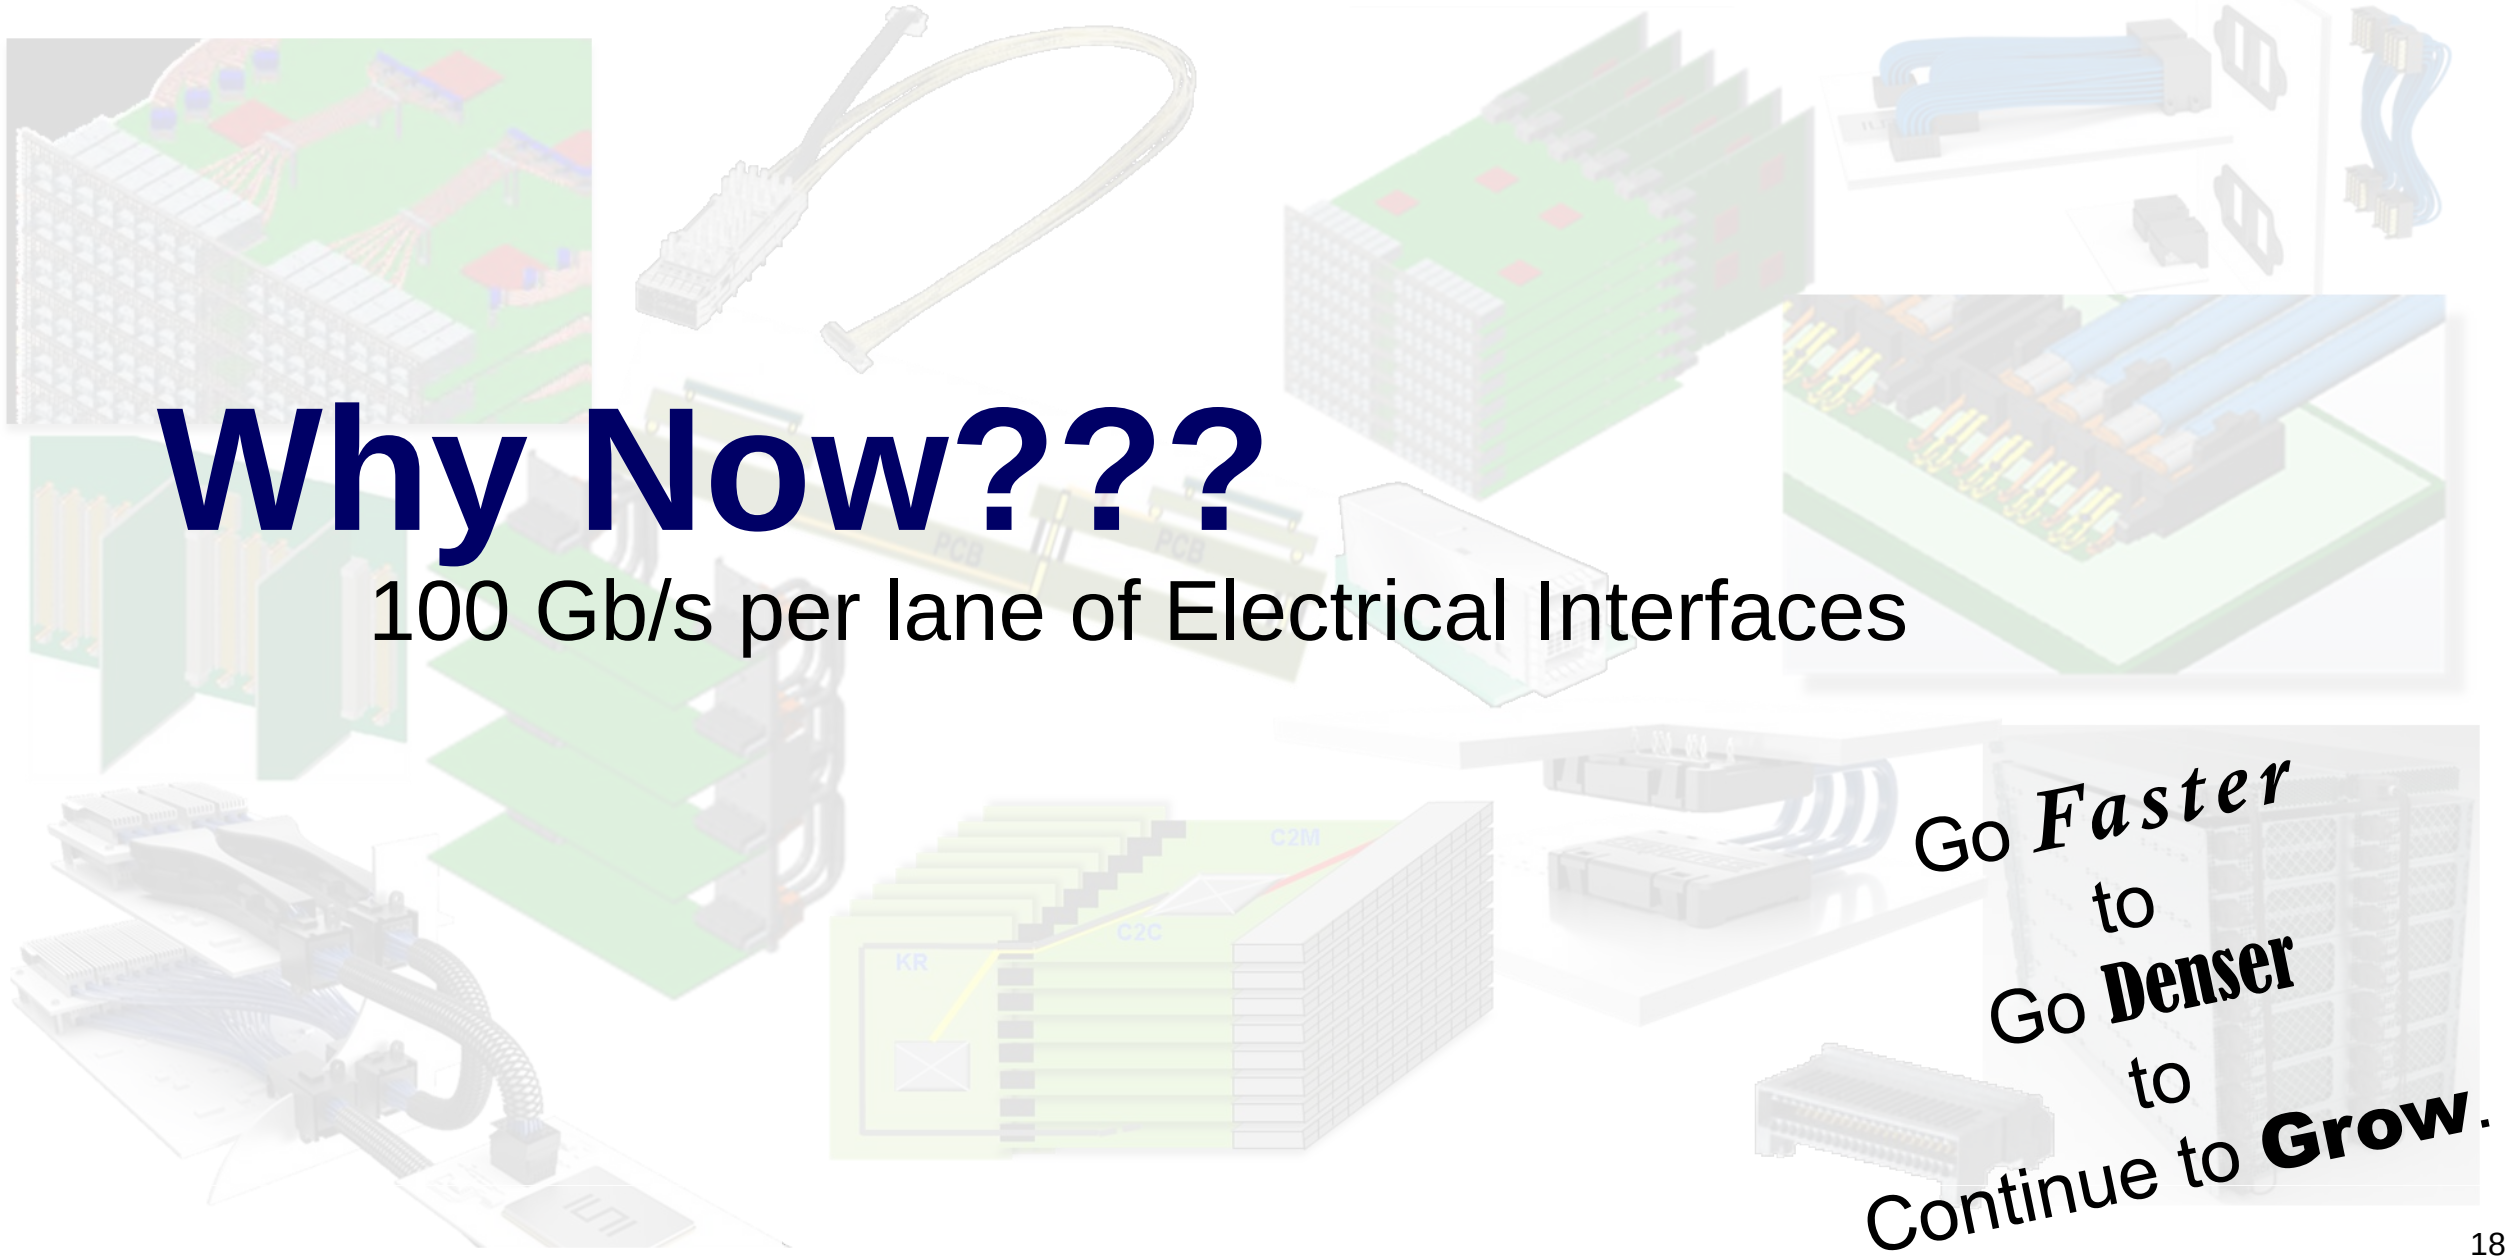

Go *Faster*
to
Go **Denser**
to
Continue to **Grow**.

The Road Map of Port Rates - next logical step

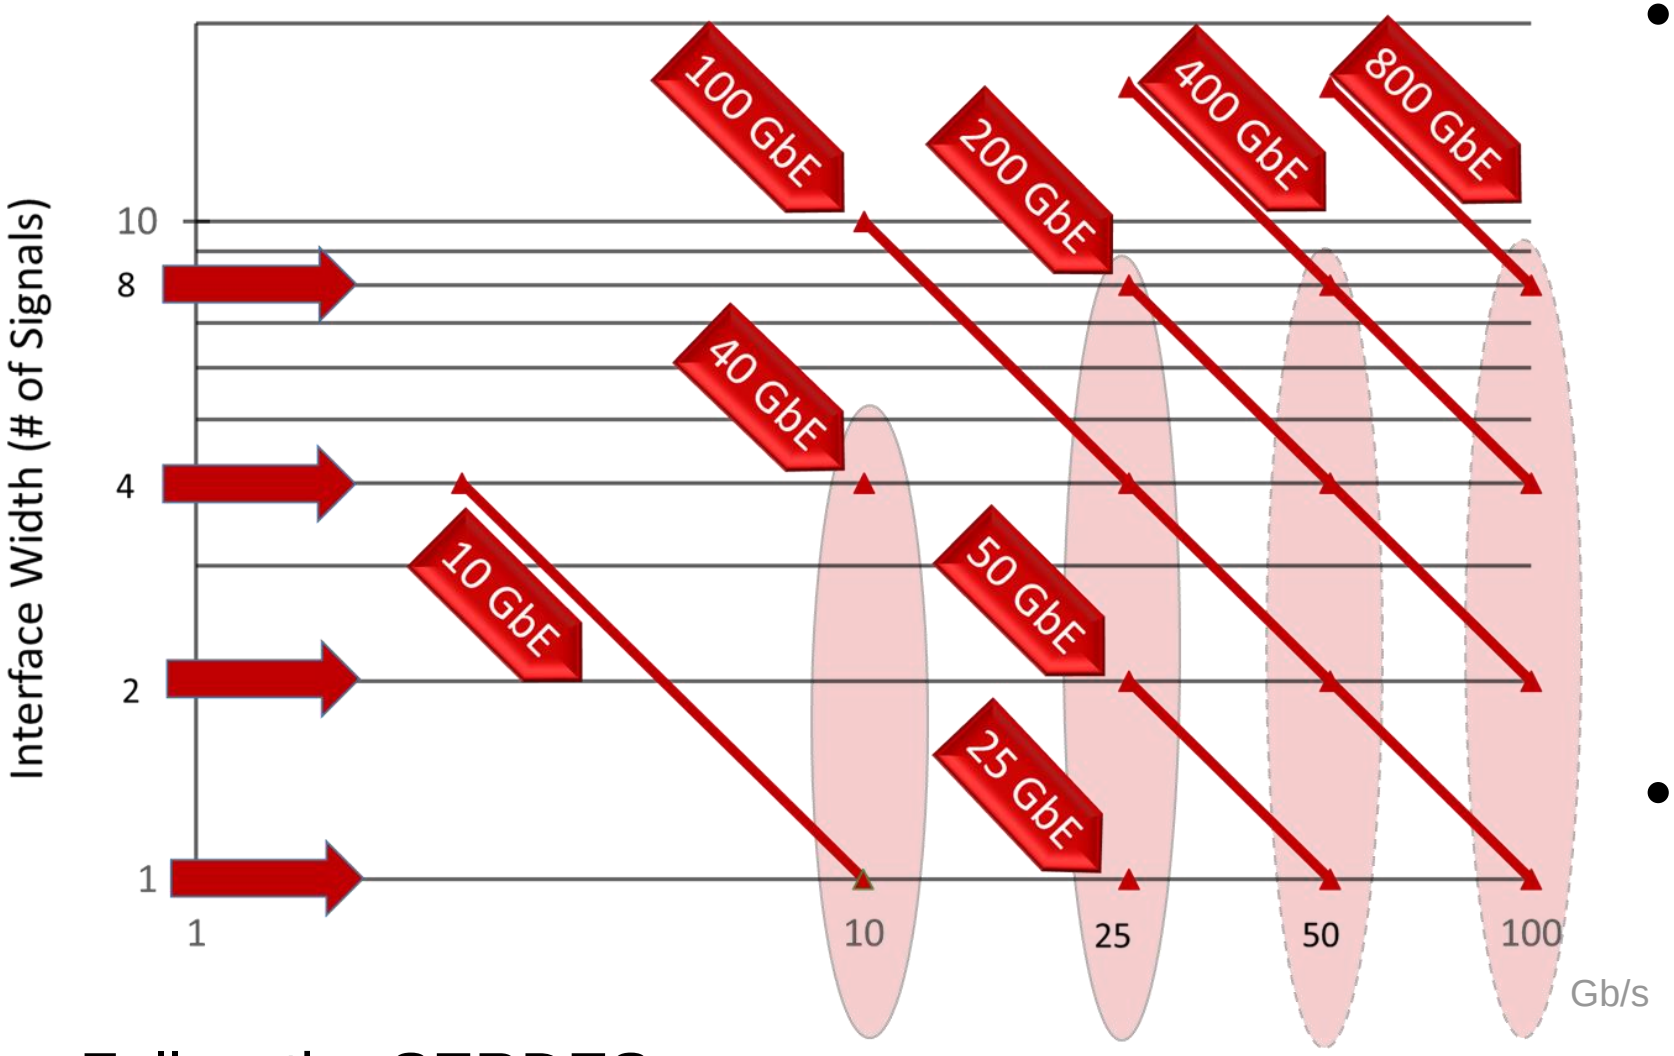

- 100G/lane is coming...
  - 100G/lane optics are here
  - OIF/Inifinband are working on this
- We need study and frame it NOW so the industry can plan

Follow the SERDES

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

Summary

- 100 Gb/s is the next step on “Follow the SerDes” and continues existing market trends
- We’ve moved to the “unknown” before and the industry survived.
- Technical details need to be rebalanced for the next speed.
- Impact of 100 Gb/s Electrical Signaling is wide across the Ethernet Family
- Let’s form a Study Group!!

Thank You!

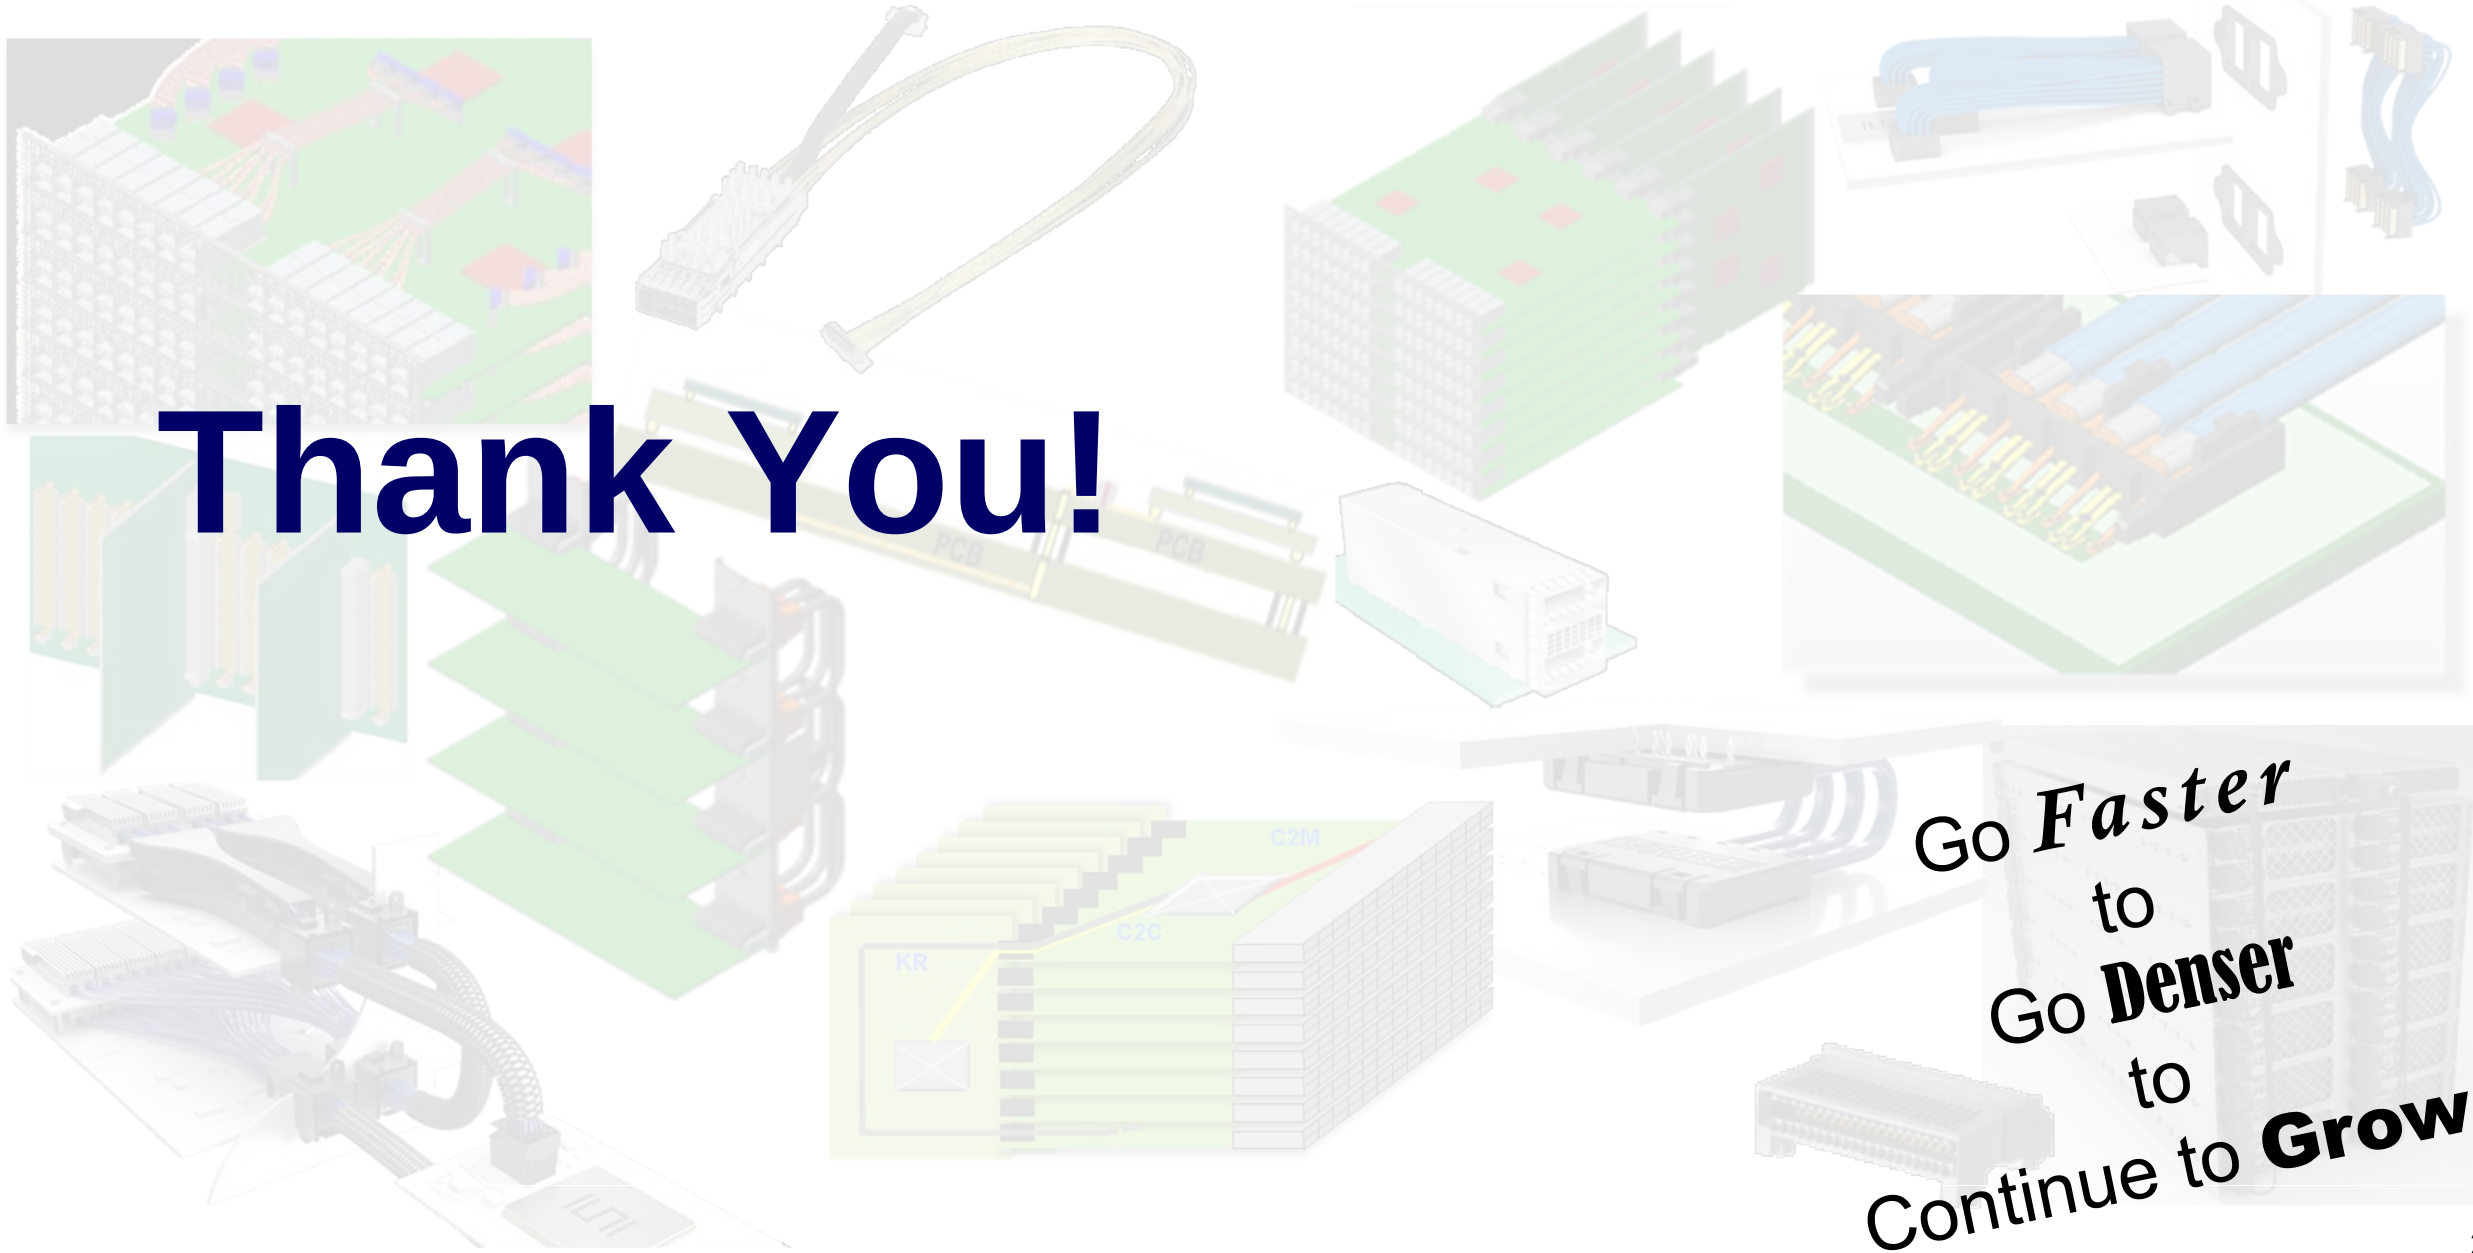

Go *Faster*
to
Go **Denser**
to
Continue to **Grow**.