100Gb/s Electrical Signaling

IEEE 802.3 New Ethernet Applications Adhoc
IEEE 802 March 2017 Plenary
March 14, 2017
Vancouver, BC, Canada

Contributors

- Joel Goergen, Cisco
- John D’Ambrosia, Futurewei (Subsidiary of Huawei)
- Rob Stone, Broadcom
- Nathan Tracy, TE Connectivity
- Mark Gustlin, Xilinx
- David Ofelt, Juniper
- Gary Nicholl, Cisco

Supporters

- Andreas Bechtolsheim, Arista

Tonight's Meeting

- To measure the interest in 100 Gb/s electrical signaling as related to the following topic areas:
  - Impact of 100Gb/s electrical signaling on existing relevant PHYs for existing Ethernet data rates
  - Potential of 100Gb/s electrical signaling for new relevant PHYs for existing Ethernet data rates
  - Consider impact of 100Gb/s electrical signaling on new PHYs for new Ethernet rates
- To discuss future of potential Call-for-Interest on areas that have significant consensus

The Never Ending Saga – Face Plate Capacity & Optical Modules

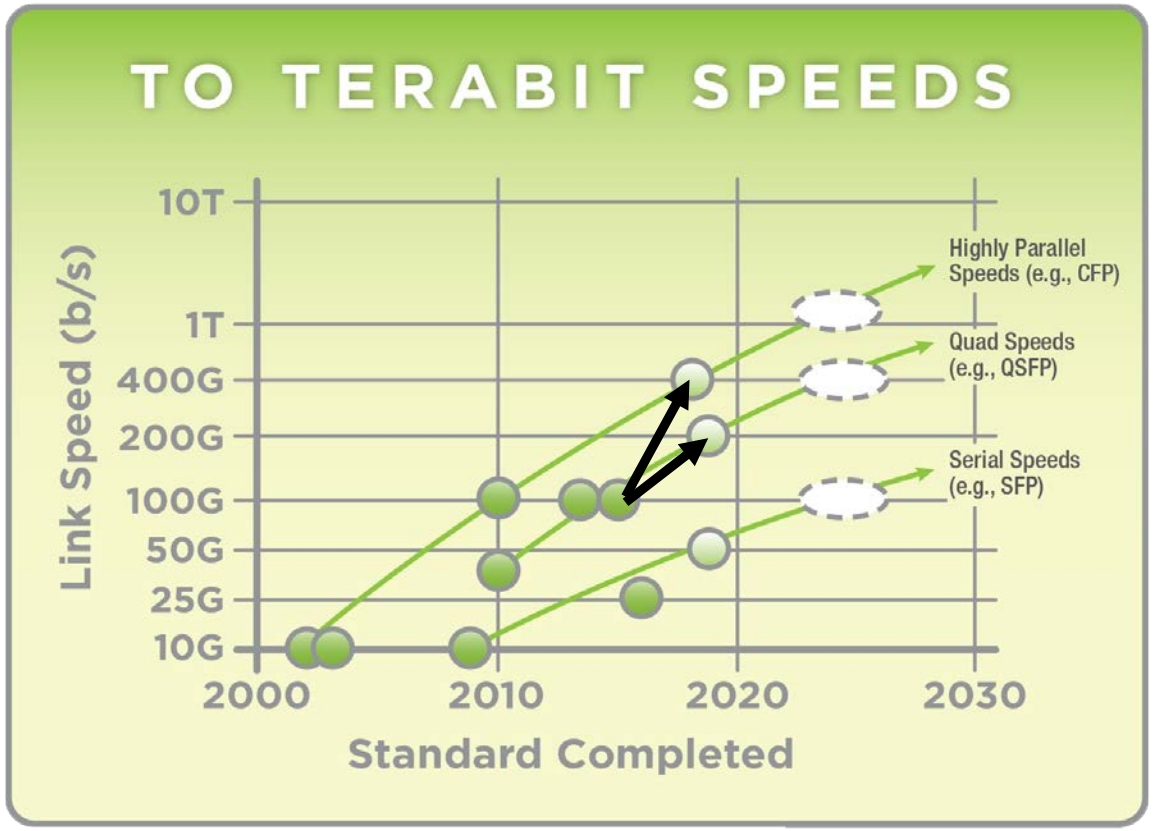

- 32 Ports – Typical Target
- QSFP – 2x BW Increase going from 100GbE to 200GbE
- QSFP-DD / OSFP - 4x BW increase going from 100GbE to 400GbE

ethernet alliance

Multiple Instances of an “Interface”

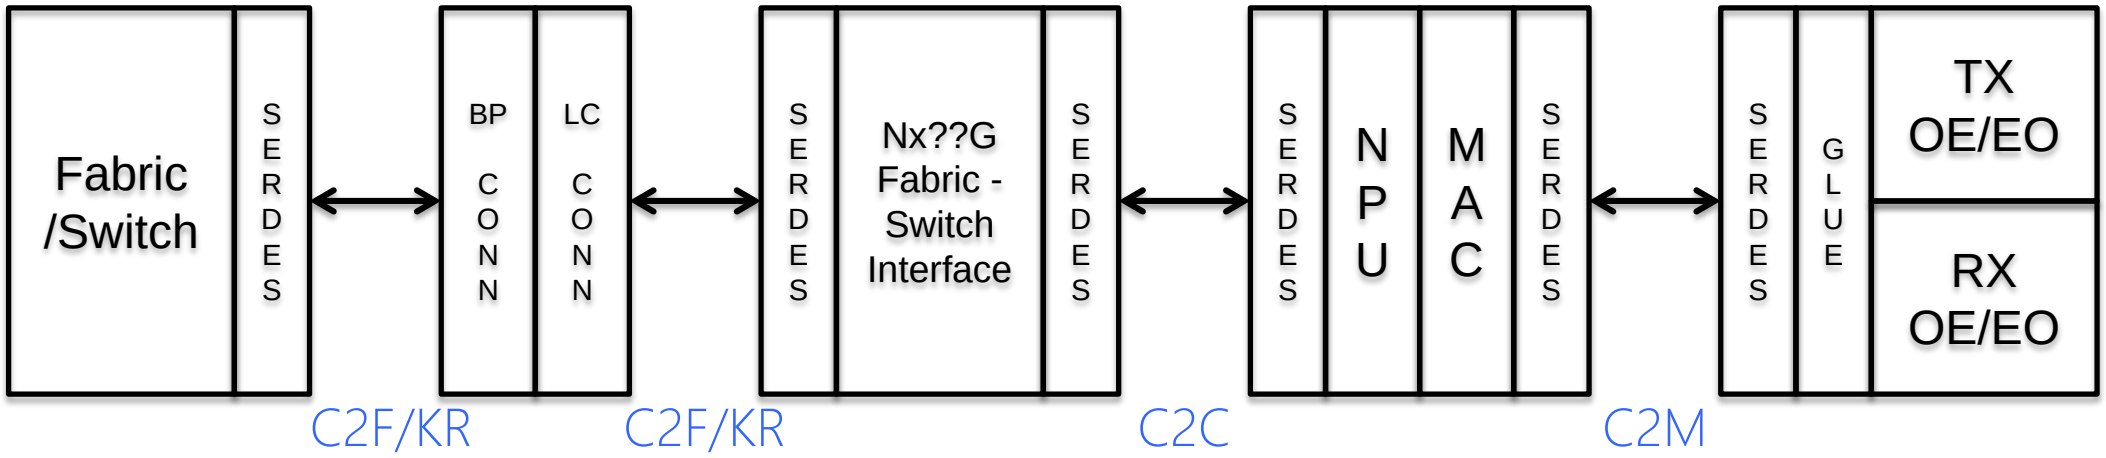

Consider how many instances of the interface can exist in the system…

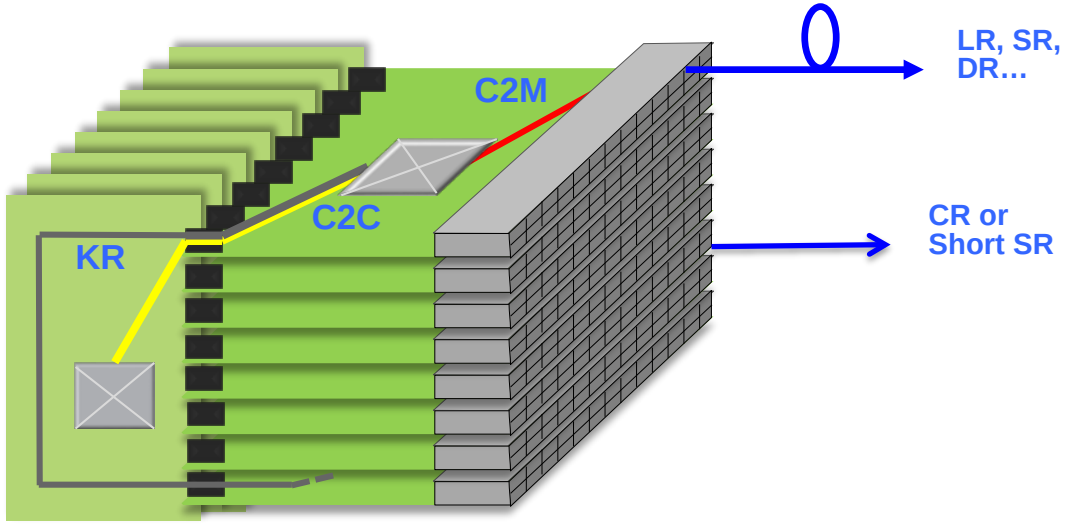

Impact of Interface Width on System Implementations

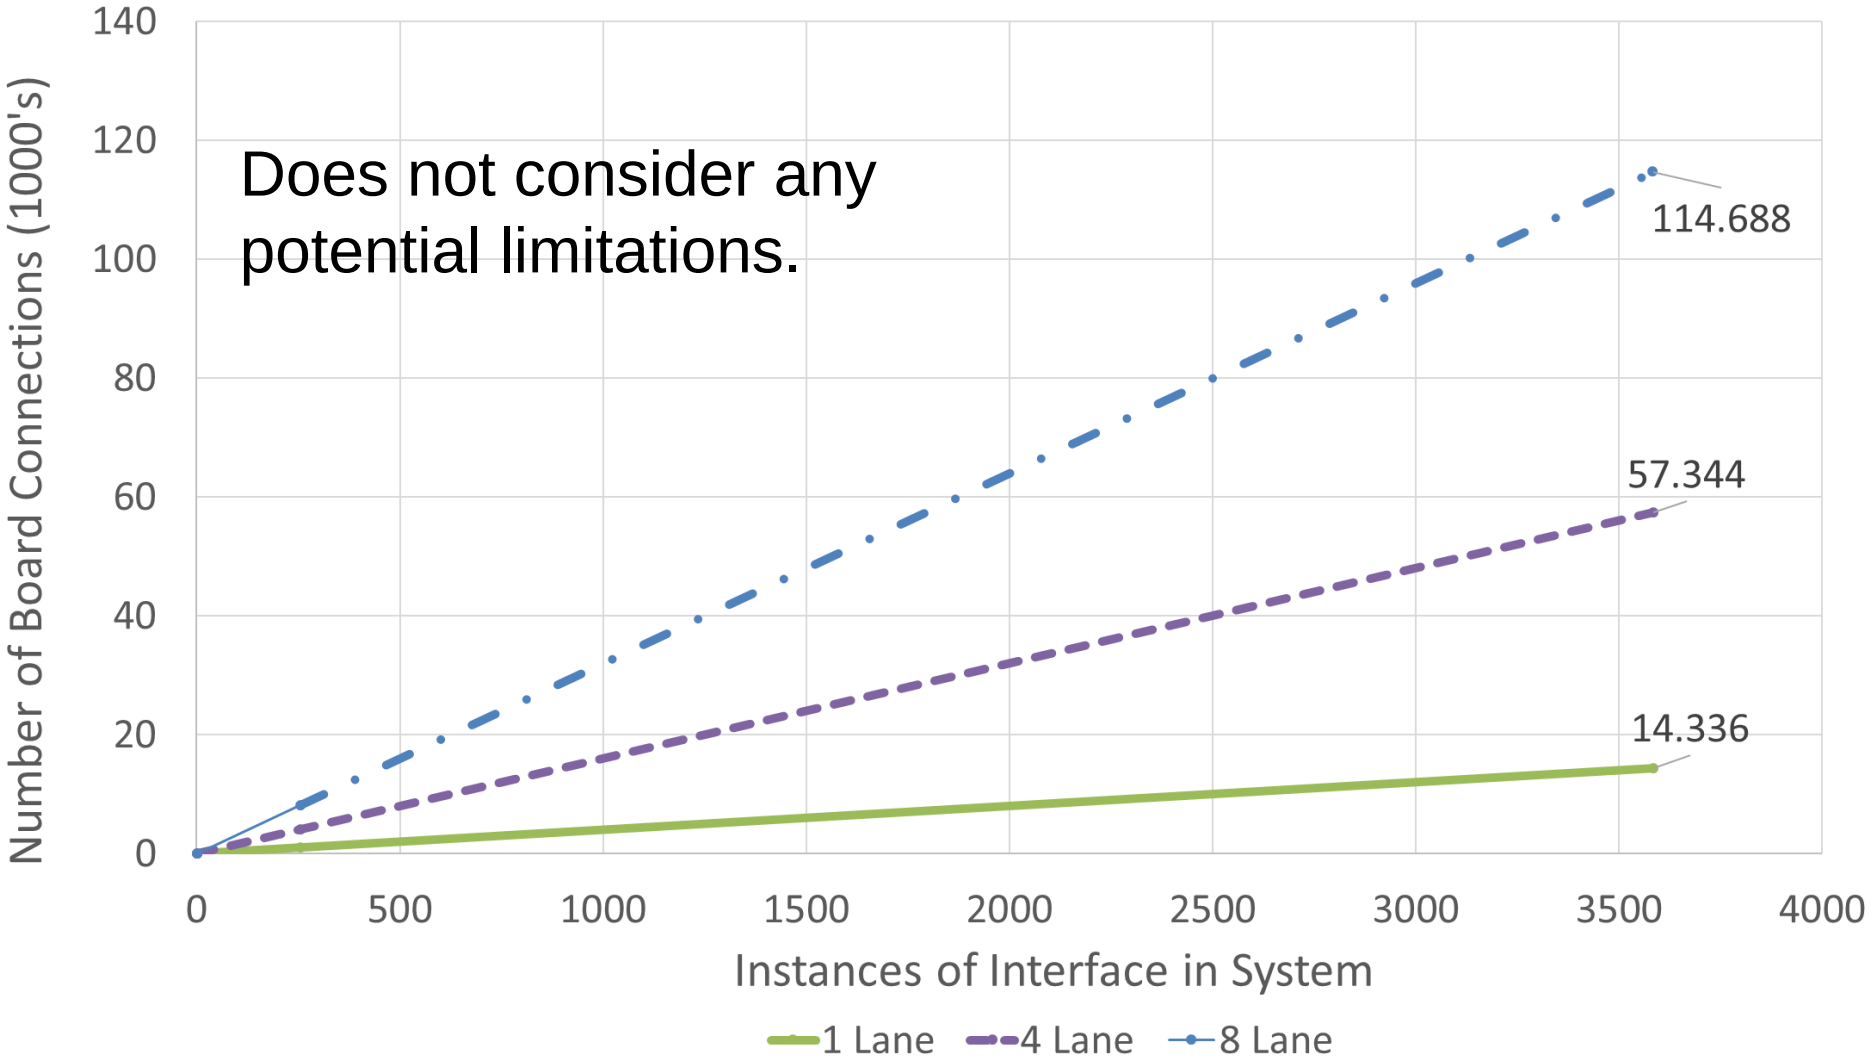

Historical Perspective – Why 100Gb/s Now?

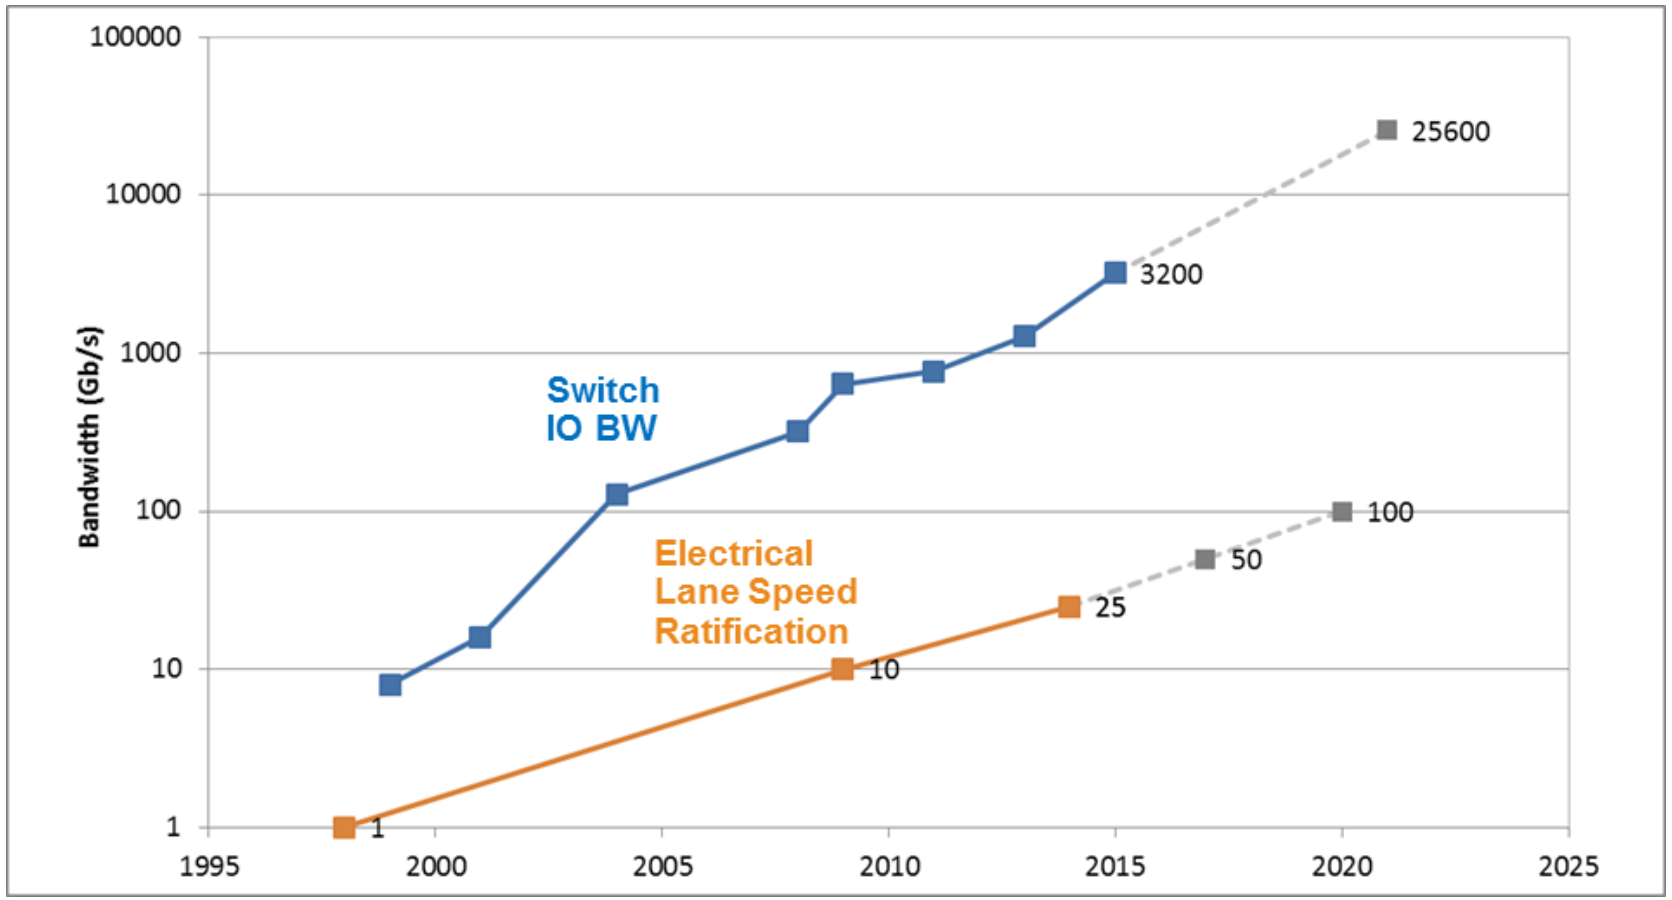

- Historical curve fit to highest rate switch products introduced to market (blue squares)
- Single ASIC IO capacity doubling every ~ 2 years

IO Escape forcing transition to higher lane speeds

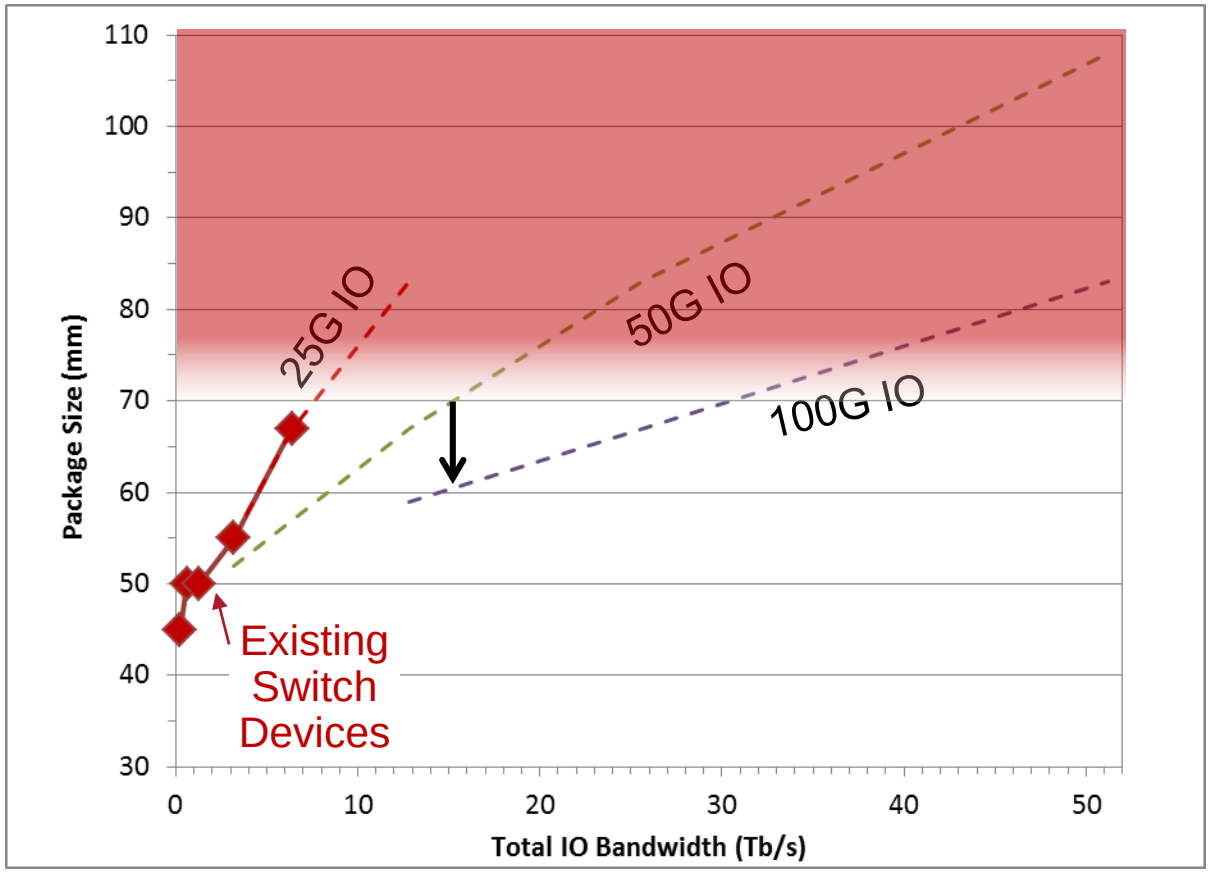

- ~ 70mm package is a current BGA practical maximum (due to coplanarity / warpage)
- This will force BGA devices with > 14Tb/s of aggregate bandwidth to transition to lane rates of higher greater than 50G (possibly 100G?)

Backplane Considerations

Traditional Backplane

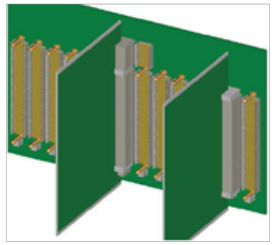

- 1.0m (40”) of Meg6
- 2 BP connectors
- 5.1mm (0.200”) thick BP
- 2.8mm (0.110”) thick DCs

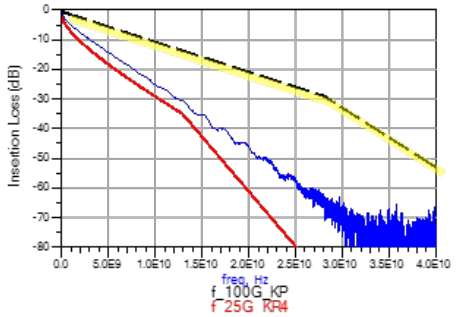

25G limit
100G limit
Existing channel

Epic Fail

Orthogonal

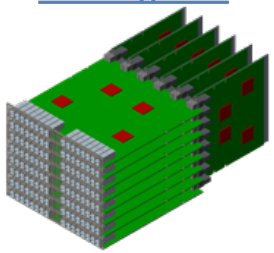

- 0.5m (20”) of Meg6
- 1 DPO connector
- 2.8mm (0.110”) thick DCs

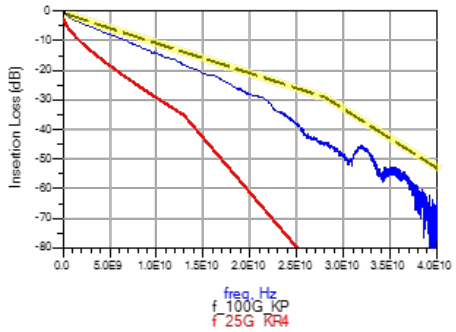

25G limit
100G limit
Existing channel

Fail

Cabled Backplane

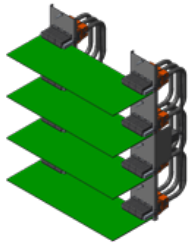

- 0.3m (12”) of Meg6
- 1.0m of 30AWG HS cable
- 2 cable connectors
- 2.8mm (0.110”) thick DCs

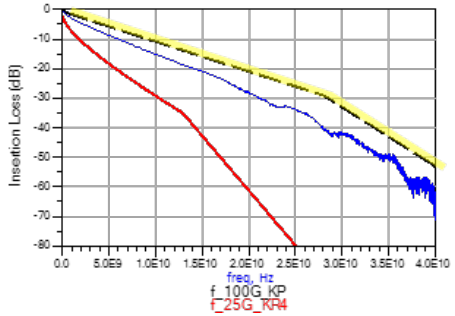

25G limit
100G limit
Existing channel

Fail

Questions to ask

- What is a “backplane”?
- What is the channel?
- What will work?

Source: Nathan Tracy, TE Connectivity, DesignCon 2017 - CEI-112G: Considering Electrical Channels.”

From a System Perspective
One Type of Implementation

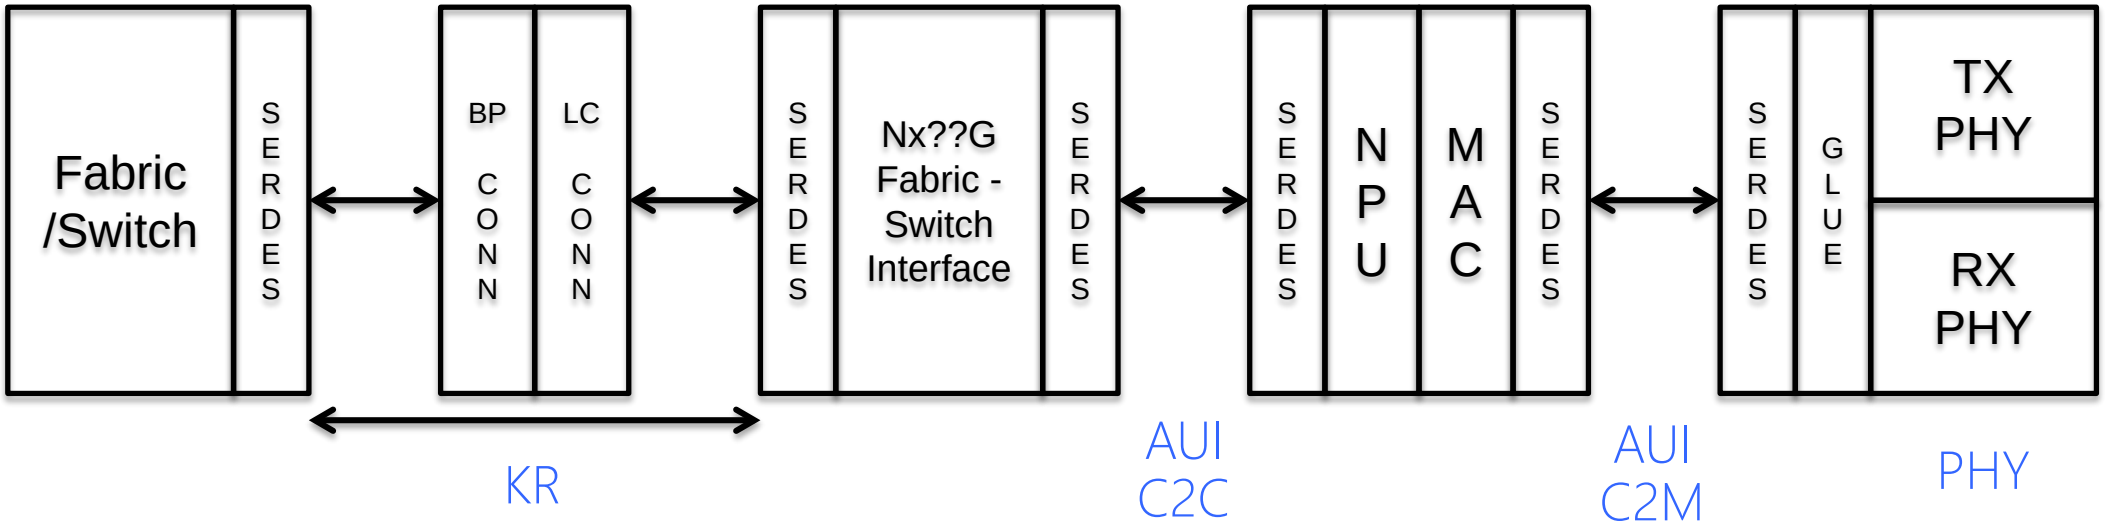

Breaking it Down

1. PHYs
2. AUIs (each with its own issues) –
    1. Chip-to-Chip
    2. Chip-to-Module
3. Backplane

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

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	?	?	

PHY Impacted by new AUI?

New PHY?

Includes Ethernet standards in development

Underlined – indicates industry MSA or proprietary solutions

Consideration of 100 Gb/s Electrical signaling and PHYs

- Existing PHYs
  - FEC defined to cover end-to-end link, based on 25Gb/s or 50 Gb/s electrical signaling. Would same FEC support end-to-end link with 100Gb/s electrical signaling?
  - Only 100GBASE-DR and 400GBASE-DR4 use 100Gb/s signaling per fiber. Other solutions based on 25 Gb/s or 50 Gb/s – inverse mux to support other existing PHYs?
  - Backwards compatibility
- New Copper PHYs (Backplane / Cu Cable)
  - Can they be done?
  - New PCS / FEC requirements?
- New Optical PHYs?
  - Unknown at this time – PCS / FEC requirements unknown.

From an Architectural Perspective

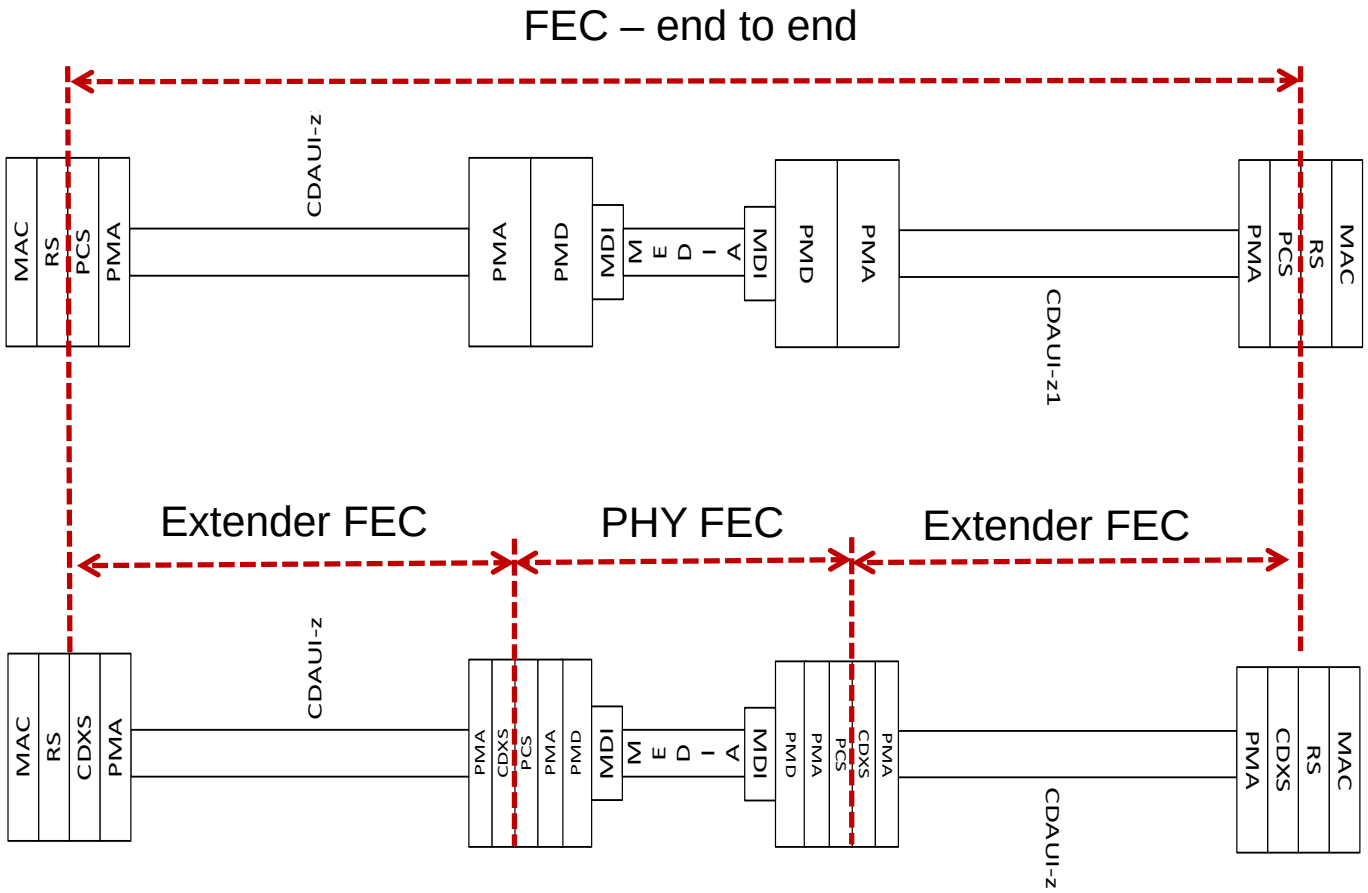

- Understanding FEC requirements will be key aspect to architecture
- Use of Extender allows future proofing and flexibility
- Impact on optical modules?

High Speed Front End Interconnects
The Road Map of Port Rates

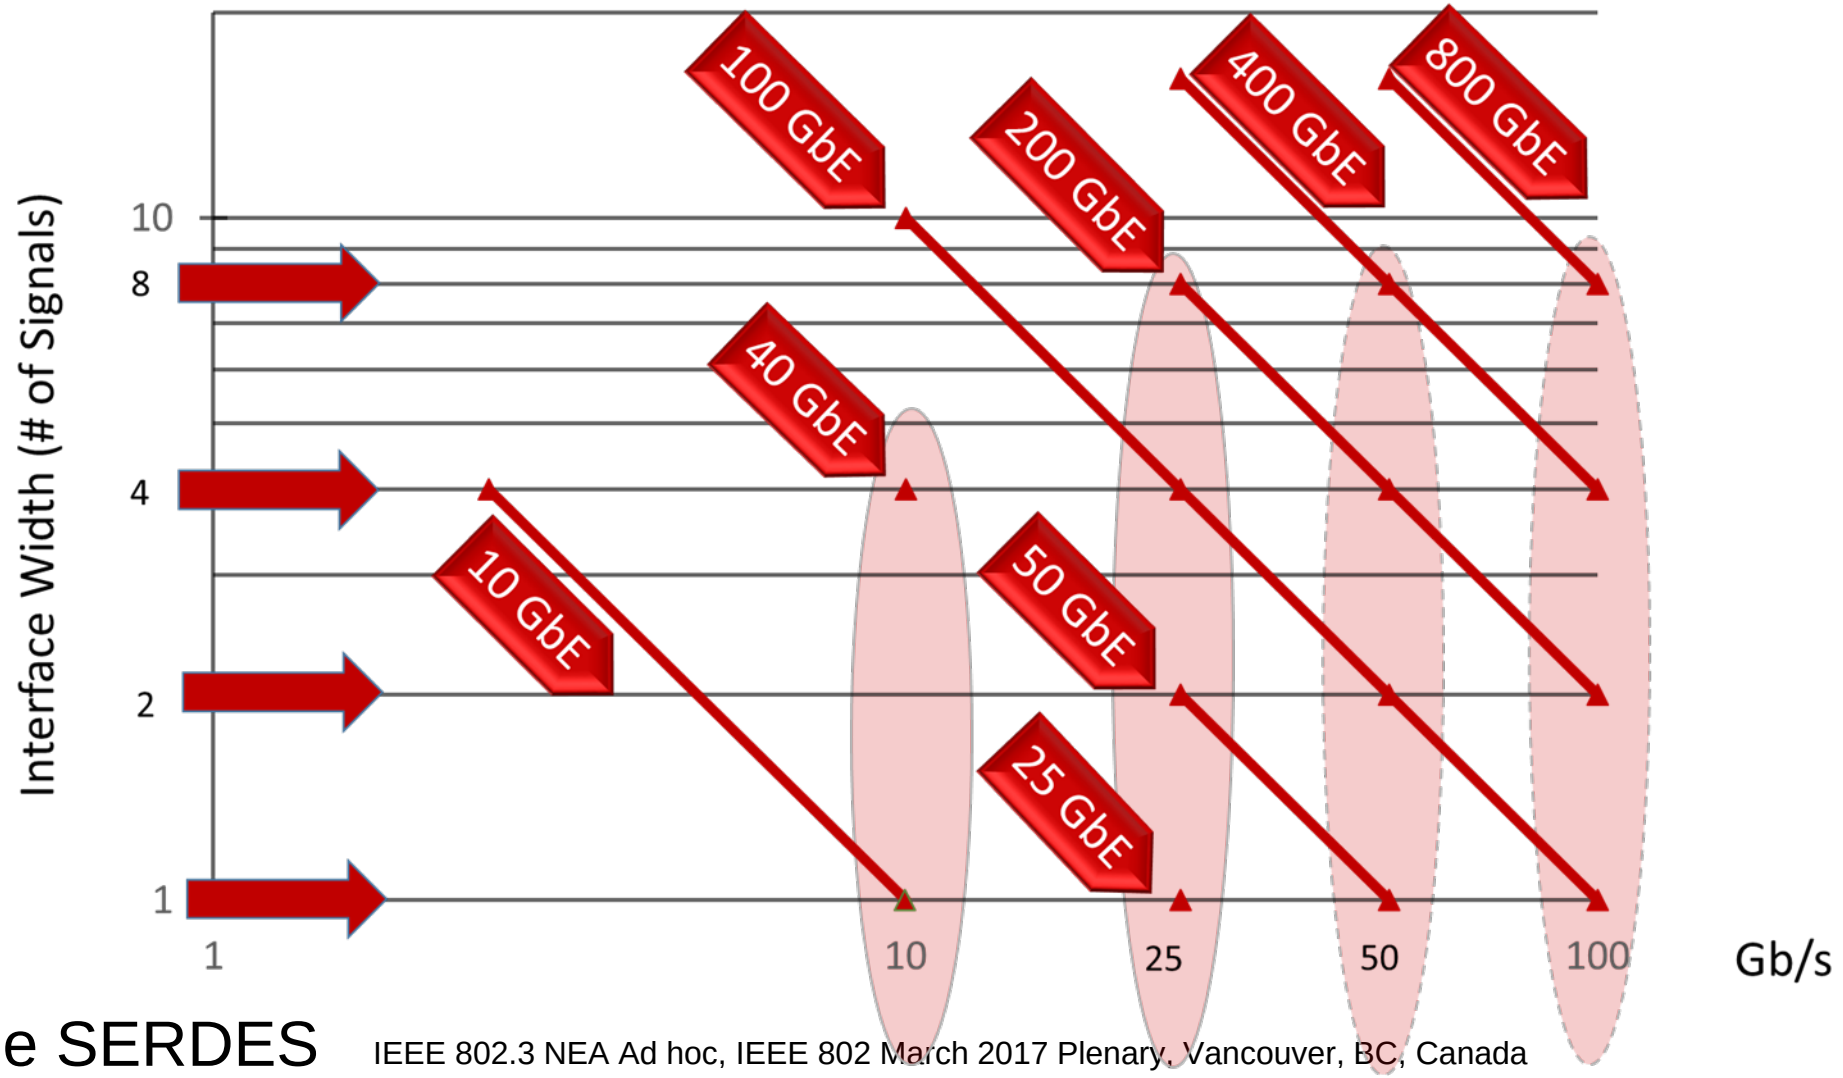

Follow the SERDES

Tomorrow's Possible Interfaces

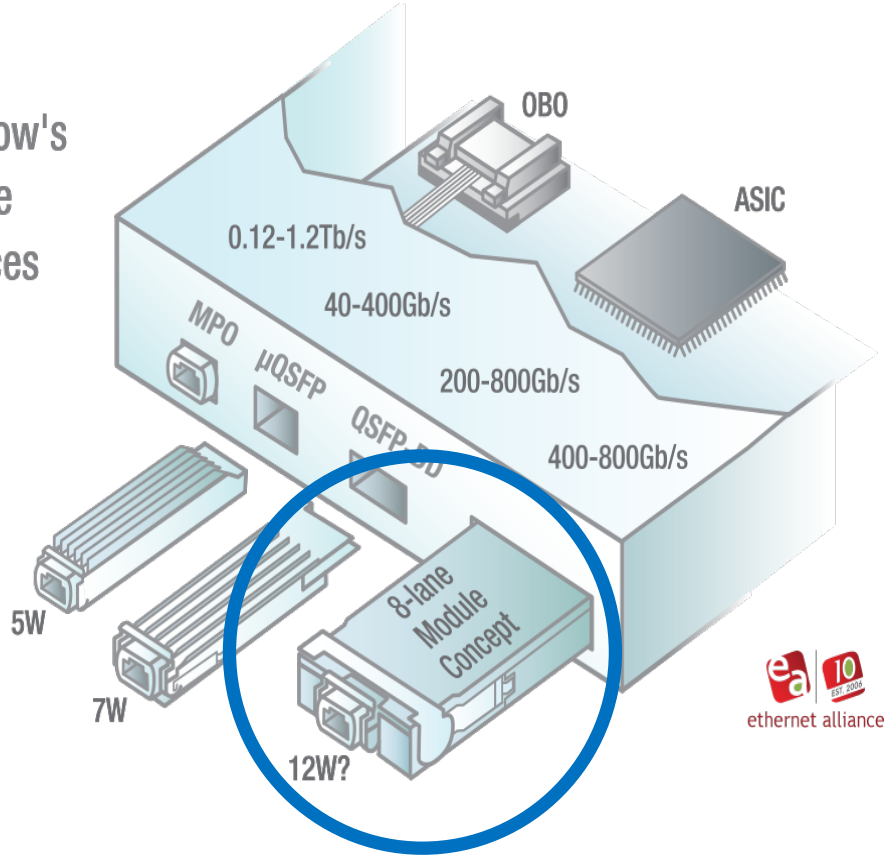

800 GbE?

100 Gb/s Signaling

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	?	?	
800GBASE-	100	?	?	?	?	?	?	?	

Includes Ethernet standards in development

Underlined – indicates industry MSA or proprietary solutions

PHY Impacted by new AUI?

New PHY?

Possible 800G Architecture

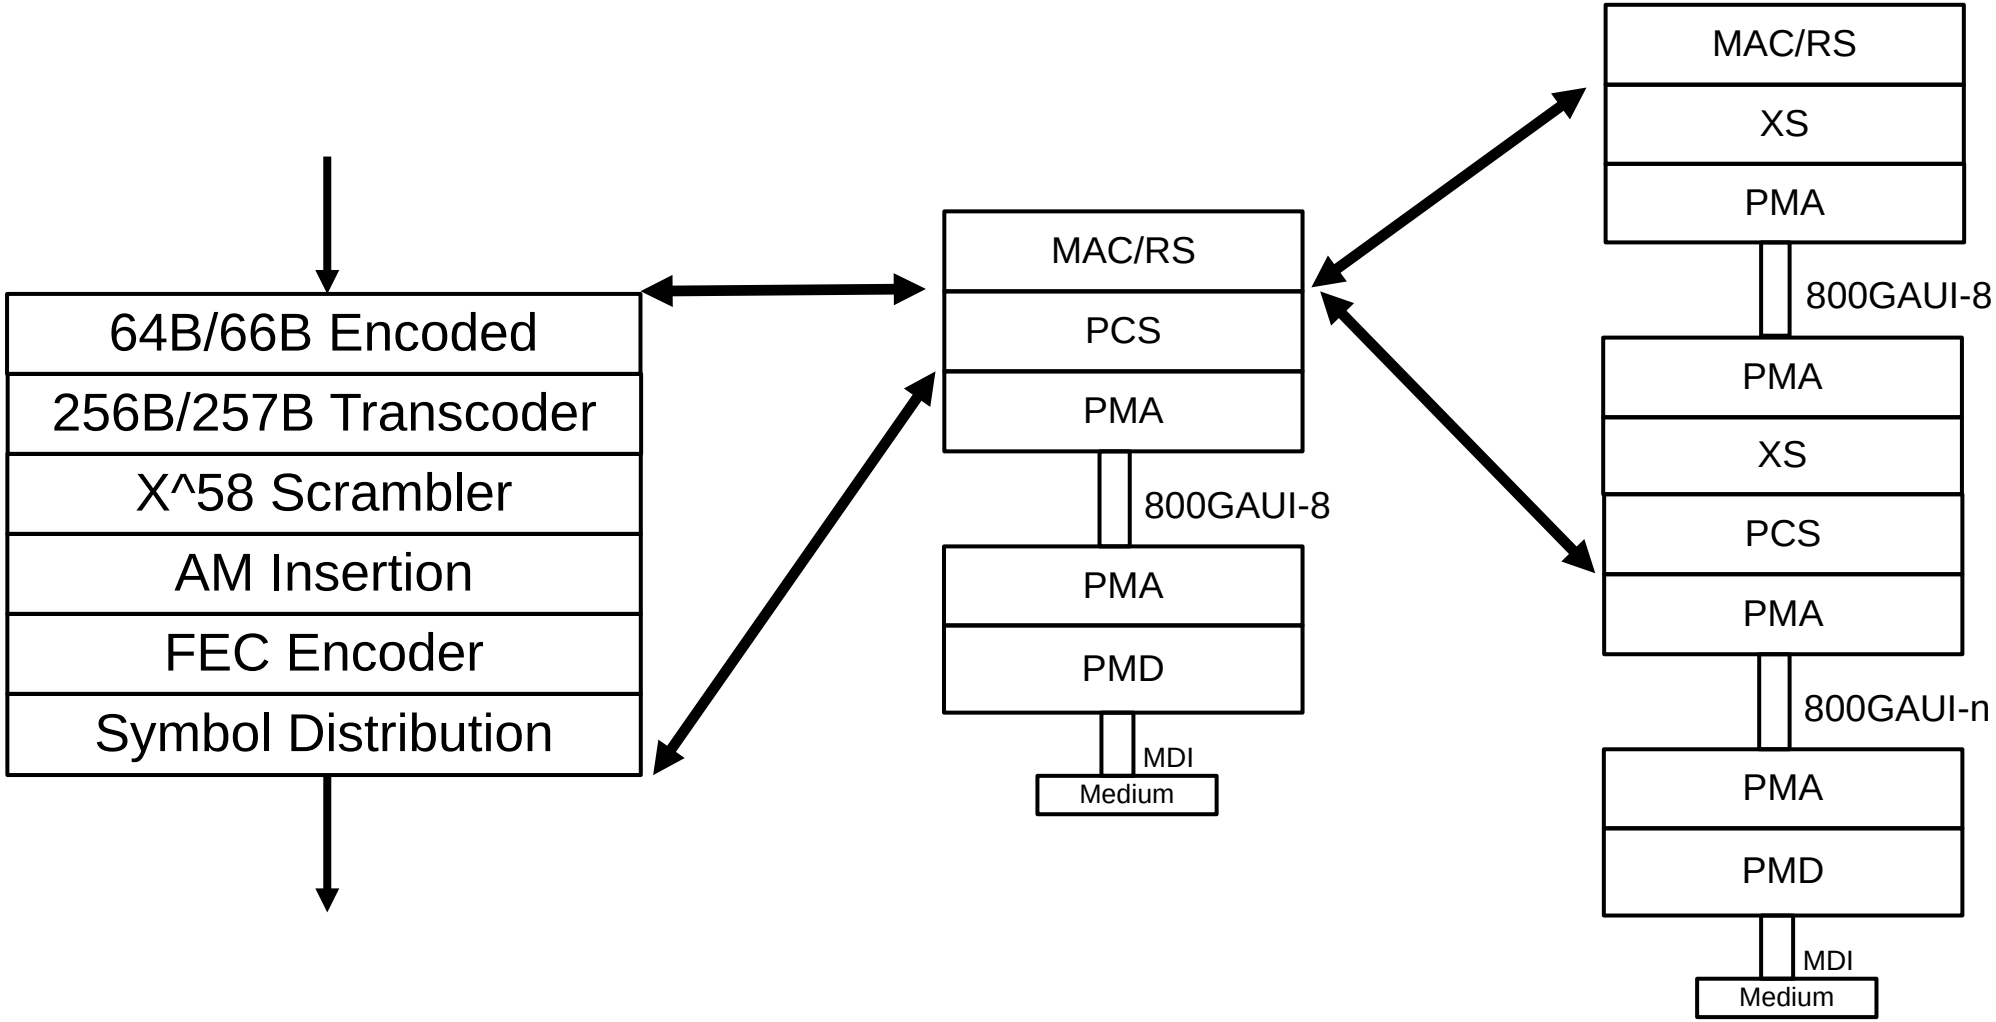

nx100Gb/s MAC/PCS Technical Feasibility

- 100G/200G/400G MACs have been implemented in the industry already, 600G and above FlexE MACs are being implemented today
- An 800G MAC is feasible in existing technology (16nm) in either ASIC or FPGA technology
- A PCS for each possible speed (100G, 200G, 400G and 800G) is feasible and can leverage existing technology, a possible PCS choices is:
  - Leverage 802.3bs/cd logic architectures
  - Would include either KP4 FEC, or something stronger is that is required
  - Number of PCS lanes = 16 or 8?
- Compact IP is possible, taking a small fraction of ASICs or FPGAs

800 Gb/s Universal Ports?

- Support 800 GbE?
  - What PHYs?
- Enable 800 Gb/s Port Capacity?
  - Increased density for previous Ethernet rates
    - 8 x 100GbE
    - 4 x 200GbE
    - 2 x 400GbE

Summary

- 100 Gb/s is the next step on “Follow the Serdes” and continues existing market trends
  - Switching capacity progression
  - Reduction of interface width
- 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?
- Is there interest in any one or combination of the above?
- What is the impact of 100 Gb/s electrical signaling?

Inventory of Issues (add as needed!)

- Develop 100 Gb/s electrical signaling
  - AUI C2C
  - AUI C2M
  - Backplane
  - Copper Cable
- FEC
  - How much coding gain is needed?
  - Reuse of existing FEC?
  - FEC architecture? - End-to-end versus Segmented
- Backplane – definition & channel requirements
- Support of legacy PHYs
  - Inverse Muxing?
- Can we get 100 Gb/s electrical signals to traditional pluggable optical modules?
- Is 800 GbE or 800 Gb/s capacity needed?

Going Forward

- Where is the interest?
  - Impact of 100Gb/s electrical signaling on existing relevant PHYs for existing Ethernet data rates
  - Potential of 100Gb/s electrical signaling for new relevant PHYs for existing Ethernet data rates
  - Consider impact of 100Gb/s electrical signaling on new PHYs for new Ethernet rates
- Should we start something?