OIF 448G AI WORKSHOP

Navigating the Path to 448G: Architectural and Ecosystem Considerations Over Electrical Interfaces

Mike Klempa, Behzad Dehlaghi, Tony Carusone, Todd Bermensolo

April 16, 2025

AI Interconnect Scaling

Dominant Driver for 448G

AI Scaling

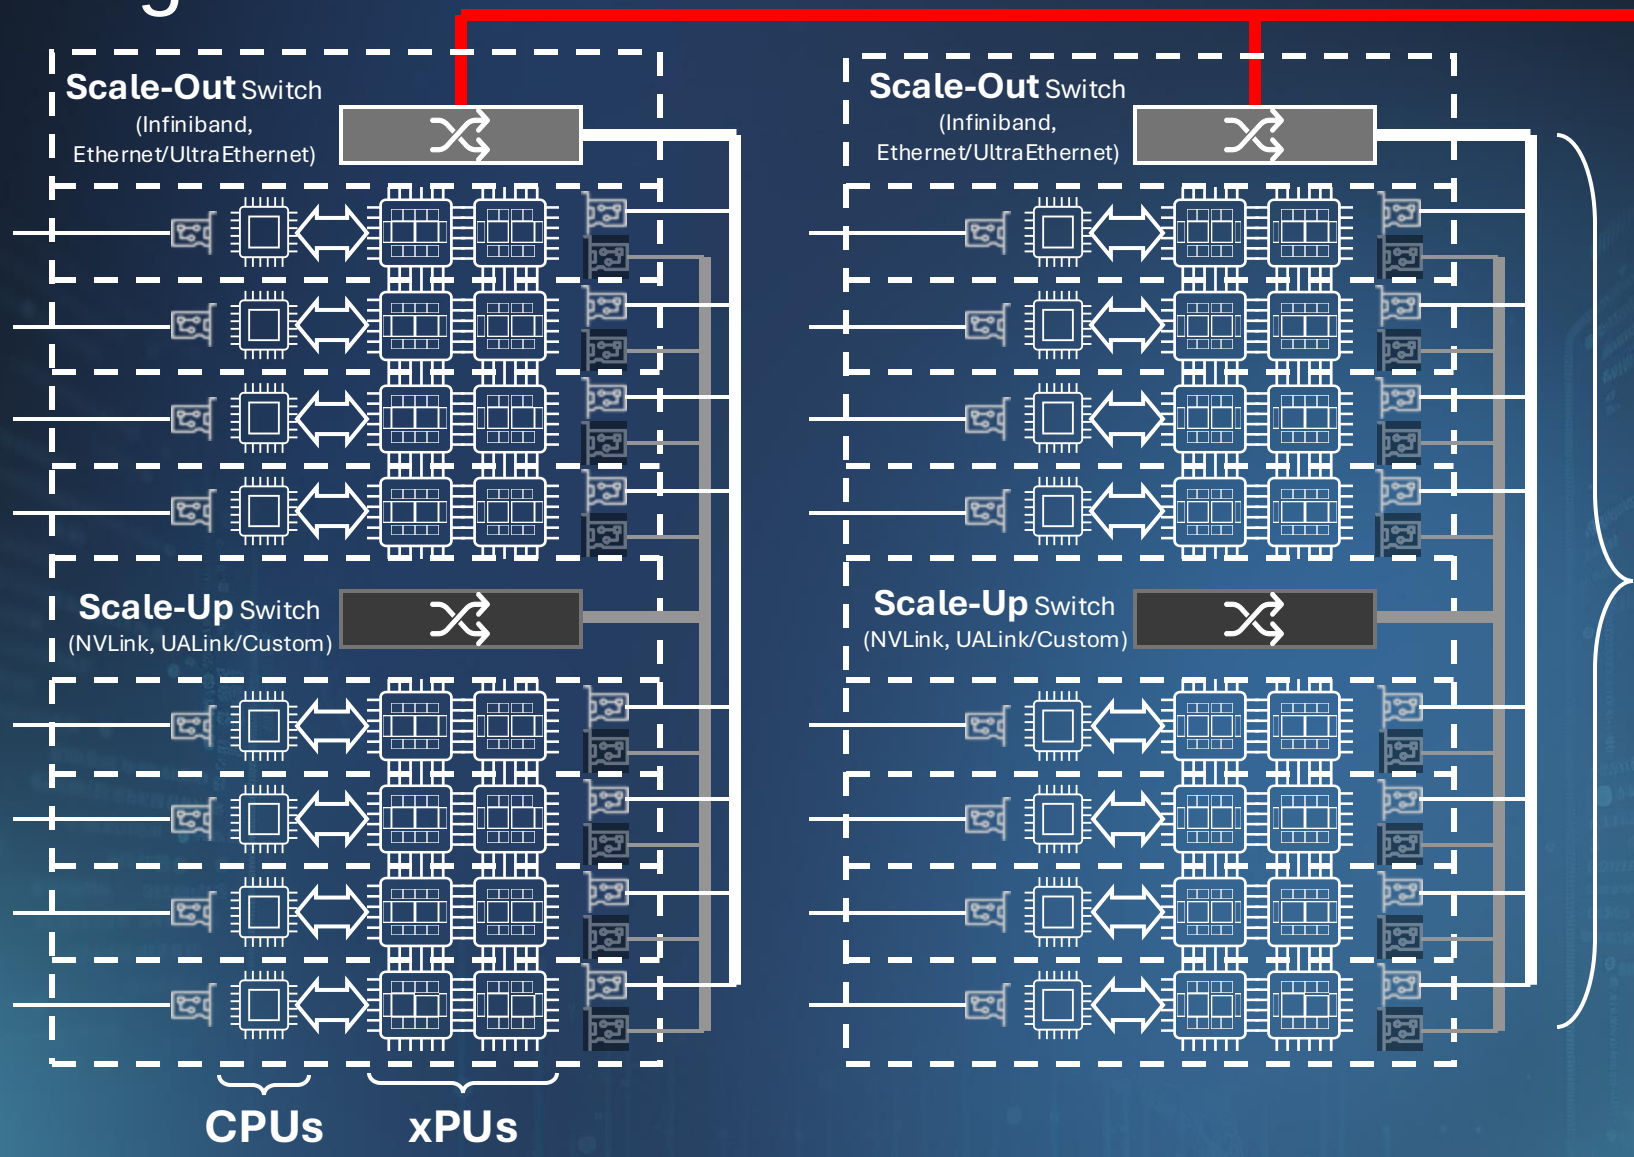

AI Scaling

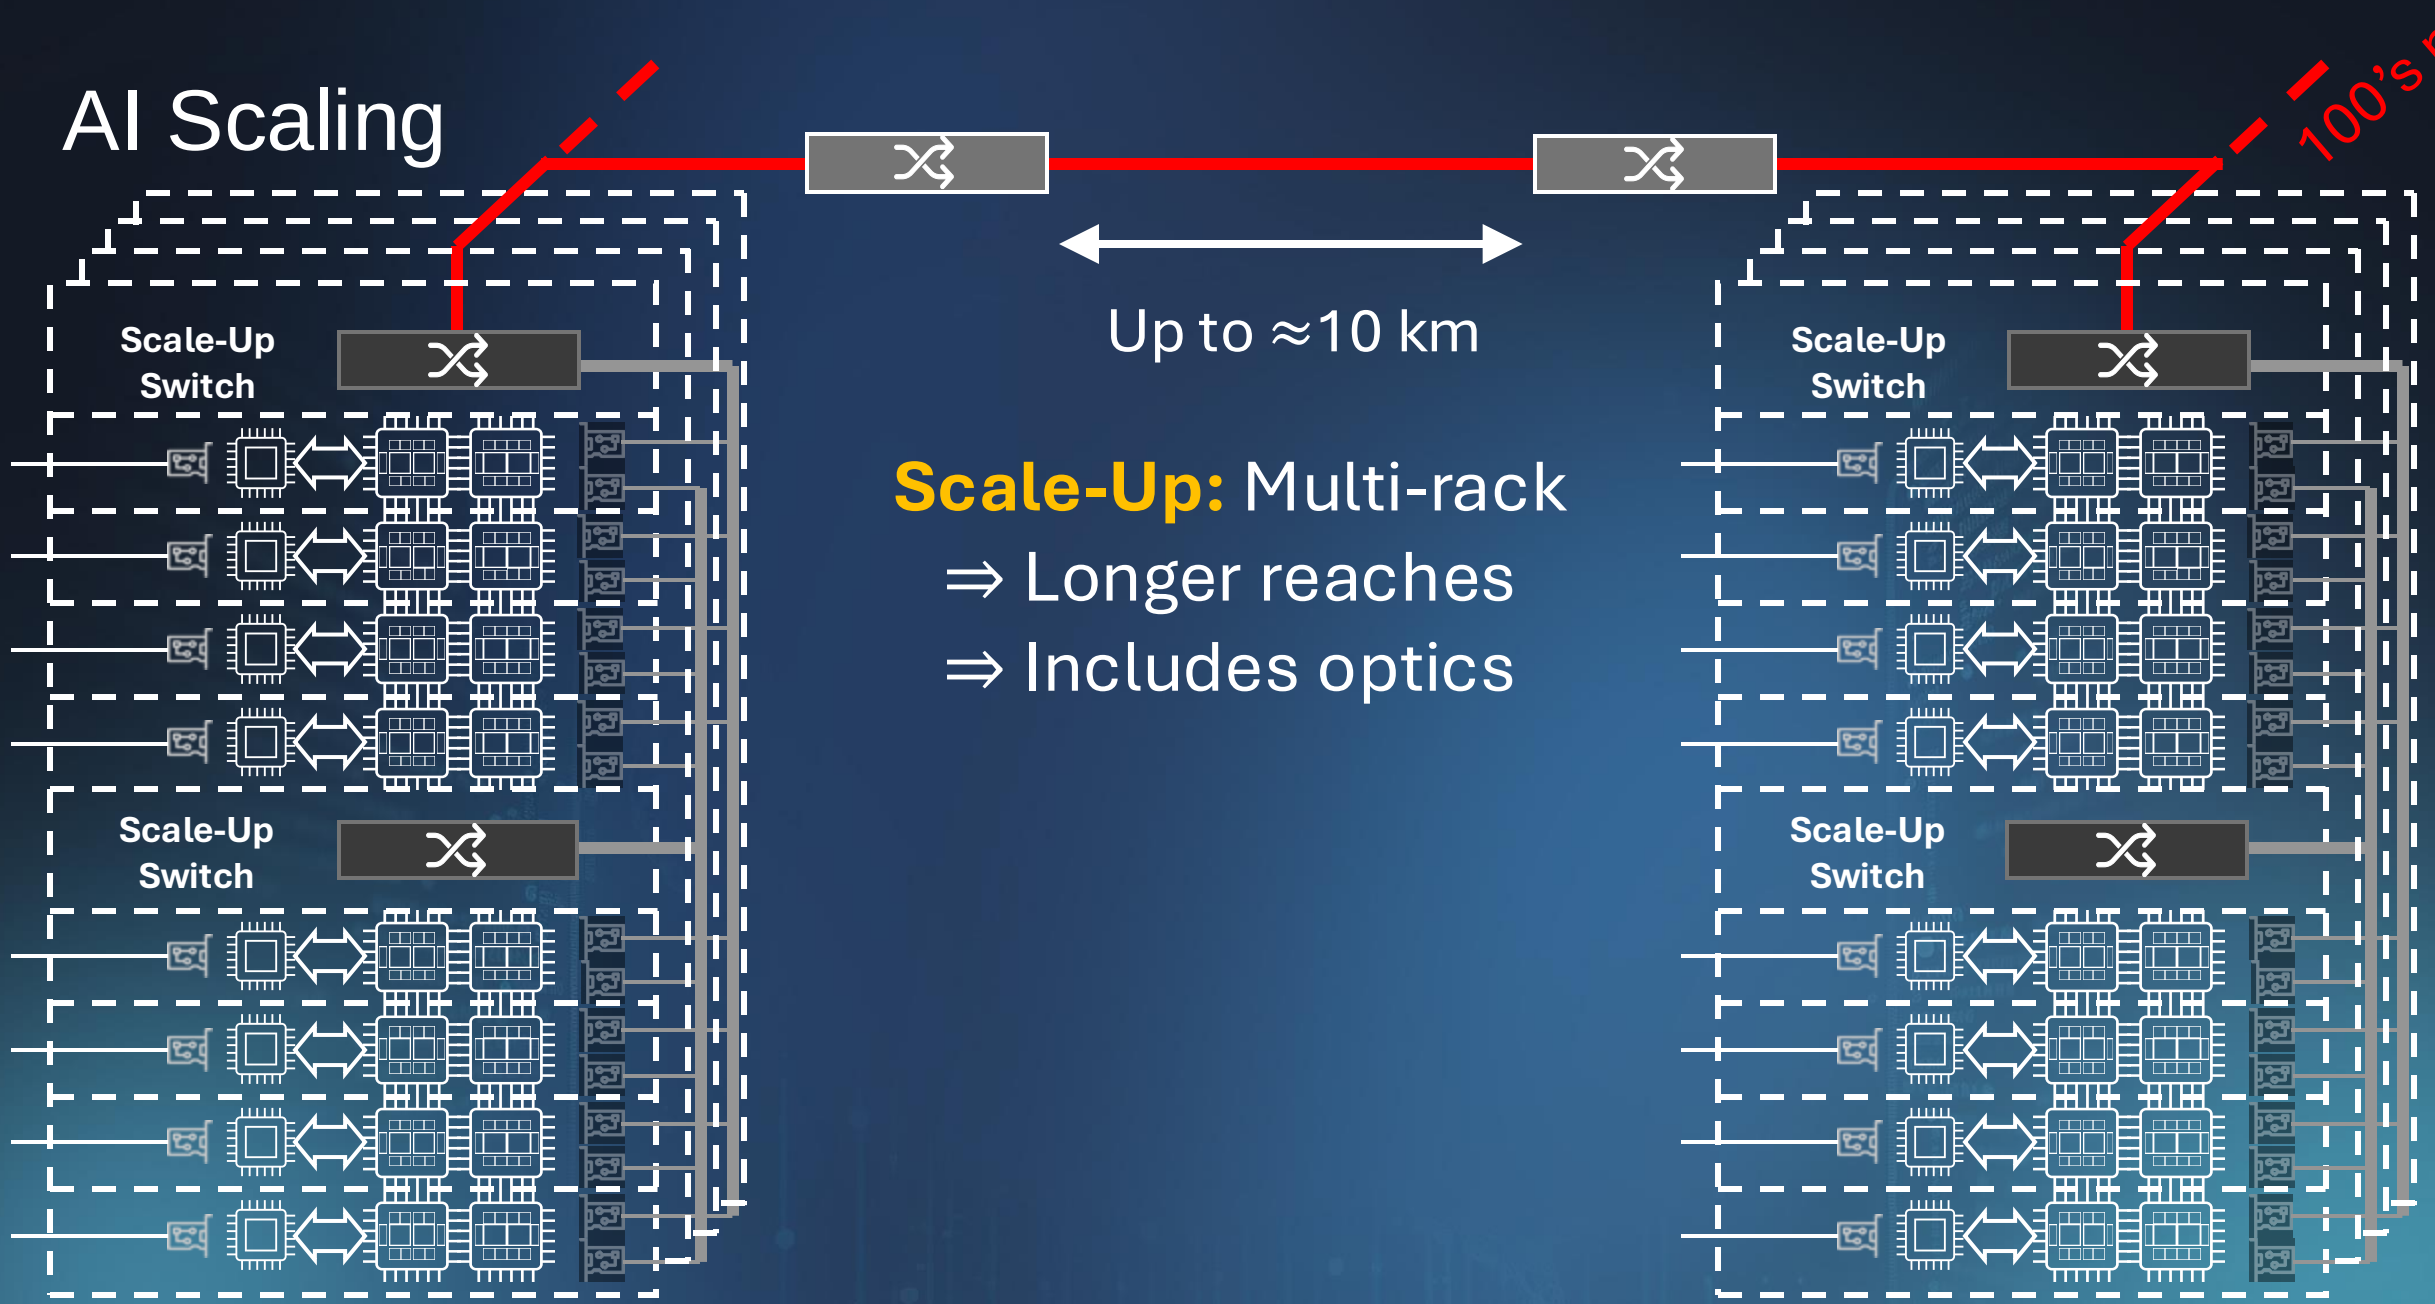

Scale-Up: Multi-rack

⇒ Longer reaches

⇒ Includes optics

AI Scaling

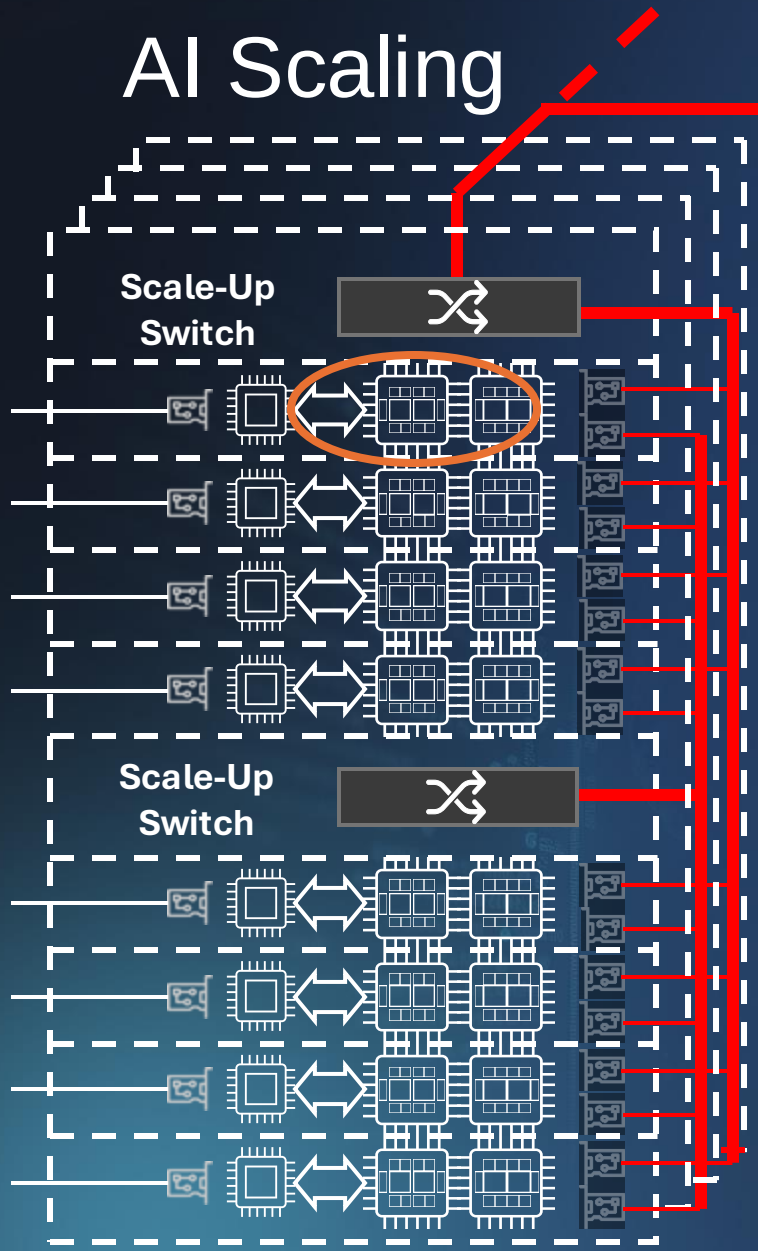

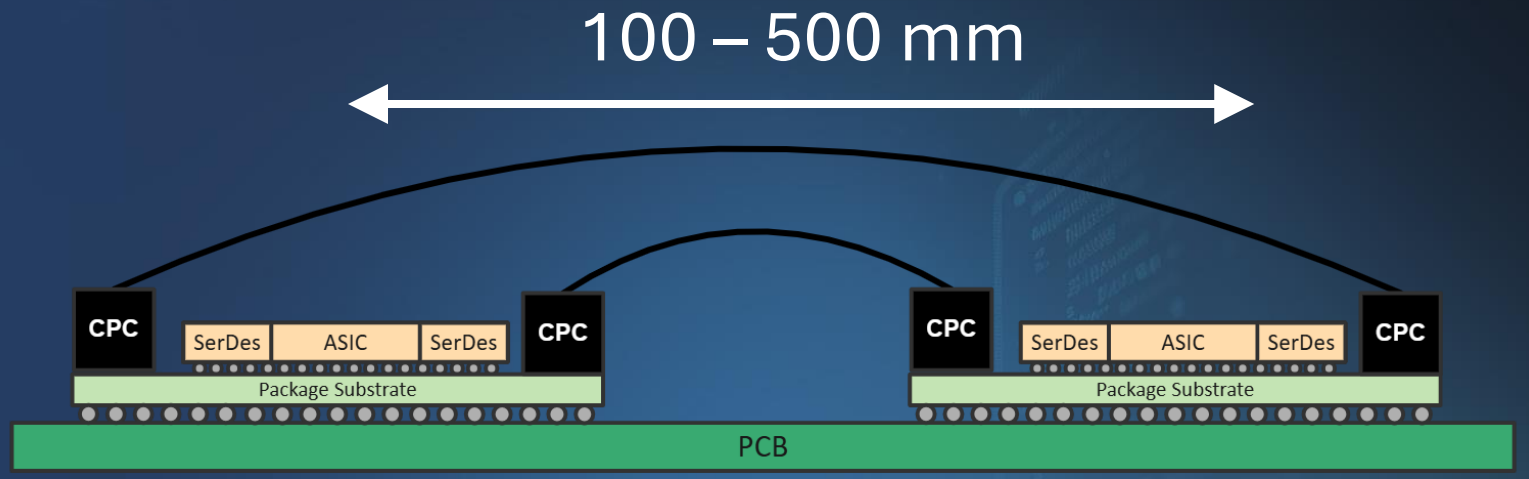

- C2C connection

AI Scaling

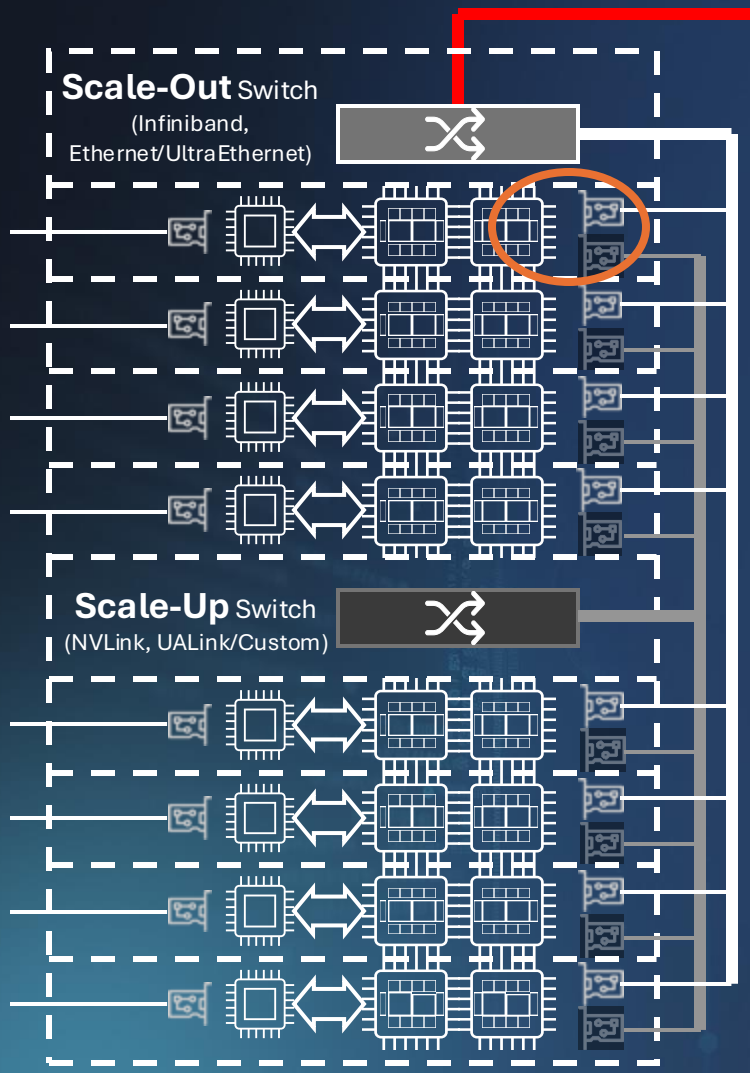

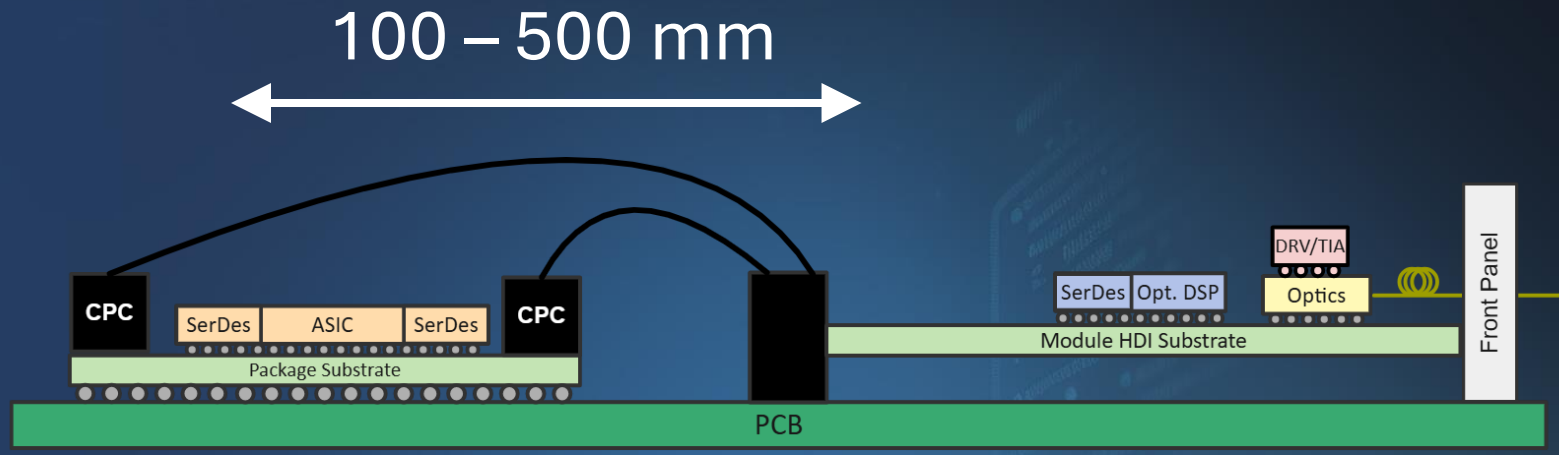

- C2M connection

448 Gb/s Electrical Feasibility using COM

Special thanks to Behzad Dehlaghi for the analysis

Simulation Setup – Channel Models

- Two sets of channel models are used in this study
  1. C2C - 9 mm Package + 500 mm CPC + 11 mm Package
     a) Current generation SiFlyHD
     b) Next generation SiFlyHDplus
  2. C2M – 25 mm Package 400 mm CPC
     a) “Limit of today’s pluggable”
     b) “New pluggable, familiar feel”
     c) “New pluggable, New paradigm”
- Crosstalk is included in the analysis
  - 2x FEXT1 channels are included

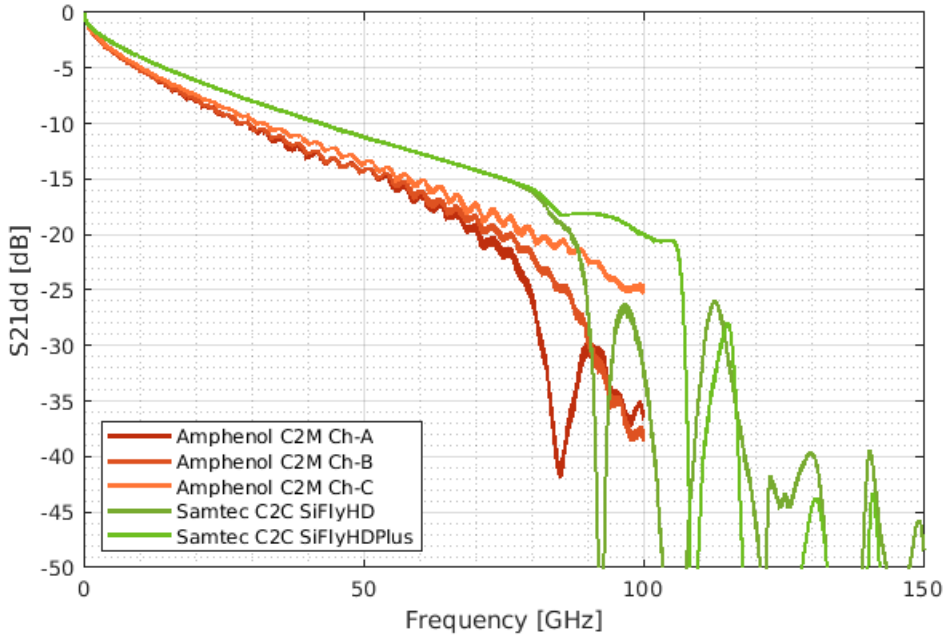

C2M Channel models from **Amphenol**

Kocsis_e4ai_01_250327

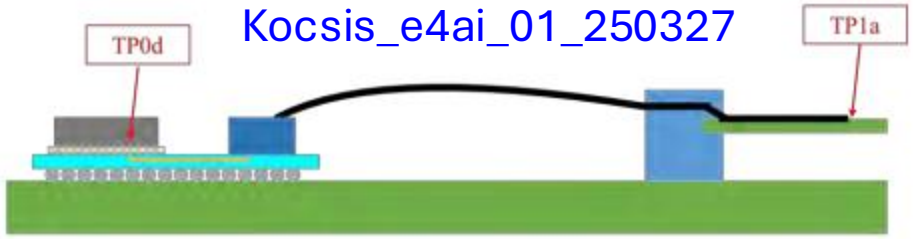

C2C Channel models from **Samtec**

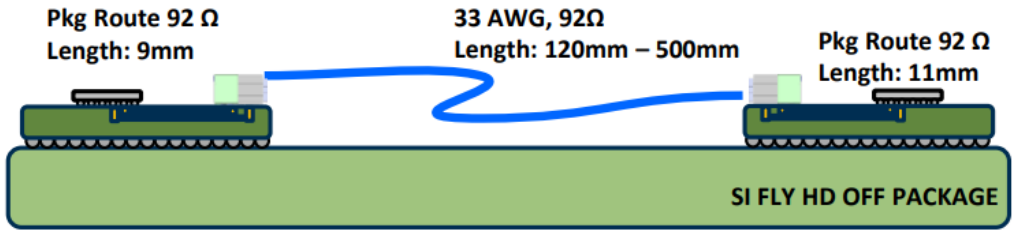

Simulation Setup – COM Sheet

Table 93A-1 parameters			
Parameter	**Setting**	**Units**	**Information**
f_b	212.5	GBd	
f_min	0.05	GHz	
Delta_f	0.01	GHz	
C_d	[0.4e-4 0.9e-4 1.1e-4 ;0.4e-4 0.9e-4 1.1e-4]	nF	[TX RX]
L_s	[0.13 0.15 0.14; 0.13 0.15 0.14]	nH	[TX RX]
C_b	[0.3e-4 0.3e-4]	nF	[TX RX]
R_0	50	Ohm	
R_d	[50 50]	Ohm	[TX RX]
PKG_NAME	PKG_MODEL PKG_MODEL		TX RX
z_p select	1		[test cases to run]
L	4		
M	32		
filter and Eq			
f_r	0.58	*fb	
c(0)	0.55		min
c(-1)	[-0.4:0.05:0]		[min:step:max]
c(-2)	[0:0.05:0.1]		[min:step:max]
c(-3)	0		[min:step:max]
c(1)	0		[min:step:max]
N_b	1	UI	
b_max(1)	0.75		As/dffe1
b_max(2..N_b)	0.3		As/dfe2..N_b
b_min(1)	0		As/dffe1
b_min(2..N_b)	-0.15		As/dfe2..N_b
g_DC	[-10:2:0]	dB	[min:step:max]
f_z	42.5	GHz	
f_p1	42.5	GHz	
f_p2	106.25	GHz	
g_DC_HP	0		[min:step:max]
f_HP_PZ	1.32815	GHz	
Bessel_Thomson	0	logical	Bessel filter
Raised_Cosine	0	logical	RaisedCosine filter
Butterworth	1	logical	Butterworth filter
RC_Start	6.70E+10	Hz	start freq for RCos
RC_end	1.23E+11	Hz	end freq for RCos

.START	PKG_MODEL		

Table 93A-3 parameters			
Parameter	**Setting**	**Units**	Information
package_tl_gamma0_a1_a2	[5e-4 0.00065 0.000293]		
package_tl_tau	0	ns/mm	
package_Z_c	[87.5 87.5; 95 95; 100 100; 100 100]	Ohm	
R_d	[50 50]	Ohm	
z_p (TX)	[0; 0; 0; 0]	mm	[test cases]
z_p (NEXT)	[0; 0; 0; 0]	mm	[test cases]
z_p (FEXT)	[0; 0; 0; 0]	mm	[test cases]
z_p (RX)	[0; 0; 0; 0]	mm	[test cases]
C_p	[0 0]	nF	[TX RX]
A_v	0.413	V	vp/vf=
A_fe	0.413	V	vp/vf=
A_ne	0.45	V	

.END			

I/O control		
DIAGNOSTICS	0	logical
DISPLAY_WINDOW	0	logical
CSV_REPORT	0	logical
RESULT_DIR	.\results\400g\400G_{date}\	
SAVE_FIGURES	0	logical
Port Order	[1 3 2 4]	
RUNTAG		
COM_CONTRIBUTION	0	logical
TDR and ERL options		logical
TDR	0	logical
ERL	0	logical
ERL_ONLY	0	ns
TR_TDR	0.01	
N	4000	logical
TDR_Butterworth	1	
beta_x	0	
rho_x	0.618	
TDR_W_TXPKG	0	UI
N_bx	20	
fixture delay time	[0 0]	
Tukey_Window	1	
Noise, jitter		UI
sigma_RJ	0.01	UI
A_DD	0.02	V^2/GHz
eta_0	4.00E-09	dB
SNR_TX	33	
R_LM	0.95	
11-2022 BenArtsi pkg	oif2022.065.02	
highlighted are under re-consideration		

Table 93A-3 parameters			
Parameter	**Setting**	**Units**	Information
package_tl_gamma0_a1_a2	[5e-4 0.00065 0.0003]		
package_tl_tau	0.006141	ns/mm	
package_Z_c	[92 92 ; 70 70; 80 80; 100 100]	Ohm	
z_p (TX)	[12 30 45; 1 1 1 ; 1 1 1 ; 0.5 0.5 0.5]	mm	[test cases]
z_p (NEXT)	[12 30 45; 1 1 1 ; 1 1 1 ; 0.5 0.5 0.5]	mm	[test cases]
z_p (FEXT)	[12 30 45; 1 1 1 ; 1 1 1 ; 0.5 0.5 0.5]	mm	[test cases]
z_p (RX)	[12 30 45; 1 1 1 ; 1 1 1 ; 0.5 0.5 0.5]	mm	[test cases]
C_p	[0.4e-4 0.4e-4]	nF	[TX RX]
A_v	0.413	V	vp/vf=
A_fe	0.413	V	vp/vf=
A_ne	0.45	V	
Operational			
ERL Pass threshold	10	dB	
COM Pass threshold	3	db	
DER_0	2.40E-04		
T_r	2.35E-03	ns	
FORCE_TR	1	logical	
PMD_type	C2C		
EW	1		
MLSE	1		
ts_anchor	1		
sample_adjustment	[-8 8]		
Local Search	2		
DER_CDR	1.00E-02		Maximimum DER_DFE that MLSE will be evaluated
Q	0.00E+00	ns	MLSE Implementation penalty
Filter: RX FFE			
ffe_pre_tap_len	10	UI	
ffe_post_tap_len	80	UI	
ffe_pre_tap1_max	1		
ffe_post_tap1_max	1		
ffe_tapn_max	1		
FFE_OPT_METHOD	MMSE		FV-LMS or MMSE
num_ui_RXFF_noise	1024		
T_O	0	mUI	Needed for C2M VEC calculations
Floating Tap Control			0 1 2 or 3 groups
N_bg	0		taps per group
N_bf	4		UI span for floating taps
N_f	80		max DFE value for floating taps
bmaxg	0.2		rss tail tap limit
B_float_RSS_MAX	0.1		start of tail taps limit
N_tail_start	25	UI	

SAVE_CONFIG2MAT	0	
Receiver testing		
RX_CALIBRATION	0	logical
Sigma BBN step	5.00E-03	V
ICN parameters		
f_v	0.139	Fb
f_f	0.139	Fb
f_n	0.139	Fb
f_2	123.250	GHz
A_ft	0.450	V
A_nt	0.450	V

Parameter	**Setting**	
board_tl_gamma0_a1_a2	[0 6.44084e-4 3.6036e-05]	1.5 db/in @ 56G
board_tl_tau	5.790E-03	ns/mm
board_Z_c	100	Ohm
z_bp (TX)	32	mm
z_bp (NEXT)	32	mm
z_bp (FEXT)	32	mm
z_bp (RX)	32	mm
C_0	[0.2e-4 0]	nF
C_1	[0.2e-4 0]	nF
Include PCB	**0**	**logical**
Seletions (rectangle, gaussian,dual_rayleigh,triangle		
Histogram_Window_Weigh	gaussian	selection
Qr	0.02	UI

Simulation Setup – COM Parameters

- COM 4.8 is used in our analysis
  - N_qb is used to model ADC quantization noise
- Nominal values, parameters in our analysis:
  a) Percentage of improvement compared to 200G LC network
     - C_d/C_b/Ls parameters are simply scaled down together
     - More complicated LC networks might be needed
  b) Assuming a 4th order Butterworth filter
  c) DC gain values are [-10:2:0] and up to 10 dB of boost is achieved by moving zero to a location below this frequency
- COM implementation penalty of 3 dB is not included in the BER, therefore we need to build margin into our BER target
  - 1e-7 is roughly equivalent to 3 dB COM margin for BER=2.4e-4

Parameter	Nominal Value
Front-end Improvement[a]	40%
TX SNR	33 dB
TX RLM	0.95
RX Bandwidth[b]	100 GHz
CTLE P1/Z[c]	75 GHz
CTLE P2	140 GHz
No of FFE Pre-cursors	20
No of FFE Post-cursors	50
ADC ENOB	7 bit
RX Noise Density	4e-9 V^2/GHz
Random Jitter	70 fs
Dual-Dirac Jitter	150 fs

Simulation Setup – COM Extrapolation

- S-parameters only go up to 100 GHz for C2M channels
  - Missing channel data >100 GHz leaves uncertainty in the results and potentially wrong conclusions
  - **The industry needs data with reasonable certainty out to 150 GHz from the ecosystem**
- Two different extrapolation methods are used in COM to show the impact of the method chosen
  - ~1.5 dB difference in COM resulting in 1-2 orders of magnitude difference in BER for PAM4
- Extrapolation methods used for the analysis shown in this presentation:
  - Magnitude extrapolation = “trend_to_DC”
  - Phase extrapolation = “trend_and_shift_to_DC”

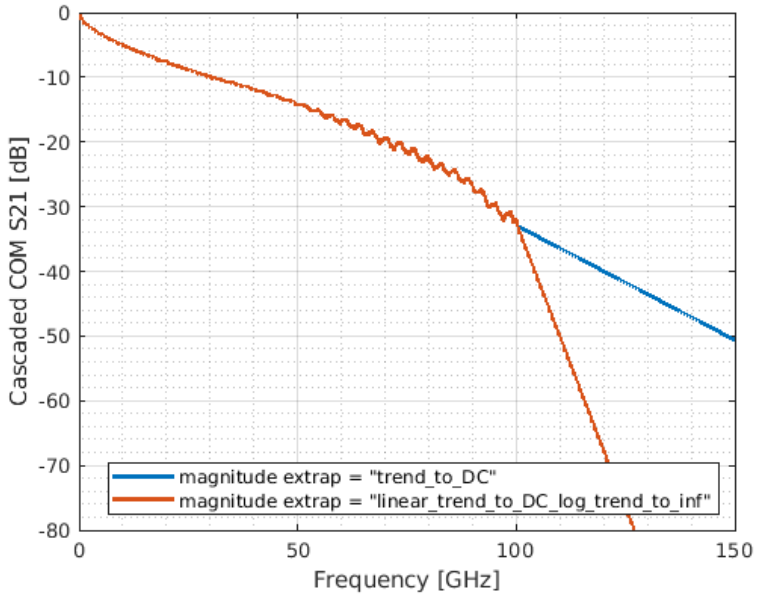

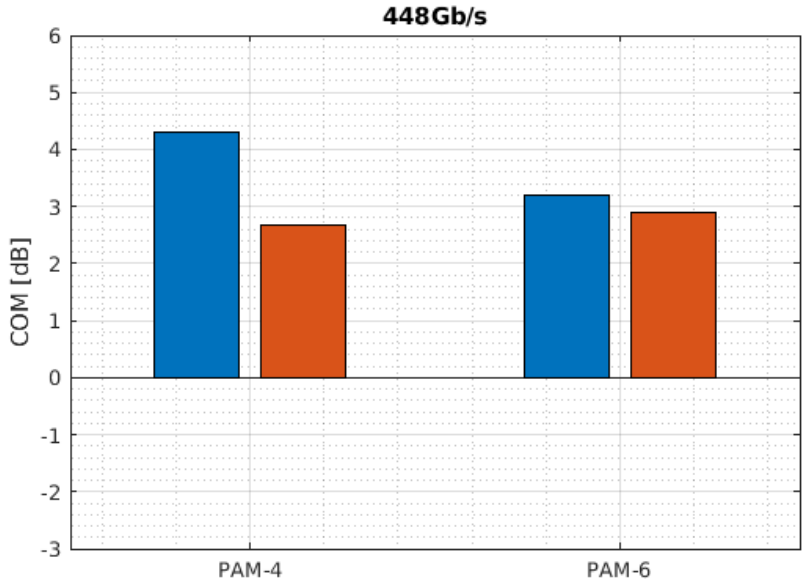

COM Results – 448 Gb/s with and without MLSD

- PAM6 outperforms PAM4 across channels with limited bandwidth
- Further improvements in C2C channels needed to enable PAM4
- C2M Ch-C shows the best recipe for success
  - This could change once channel data above 100 GHz becomes available
- MLSD is necessary to give us the margin we need
  - Samtec C2C SiFlyHDPlus results for PAM4 with MLSD are worse than without
    - Needs further investigation but it appears to have something to do with MLSD and quantization noise

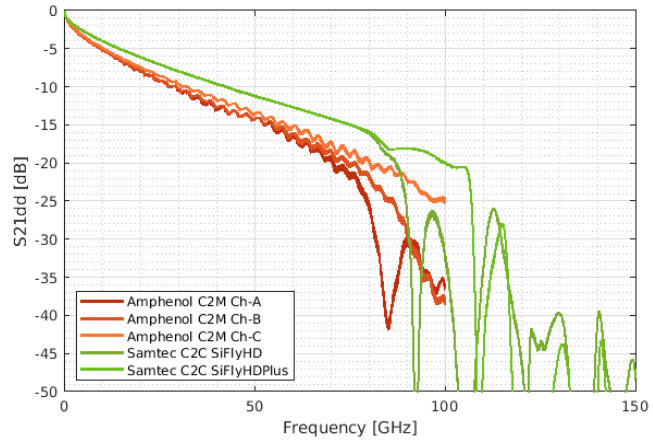

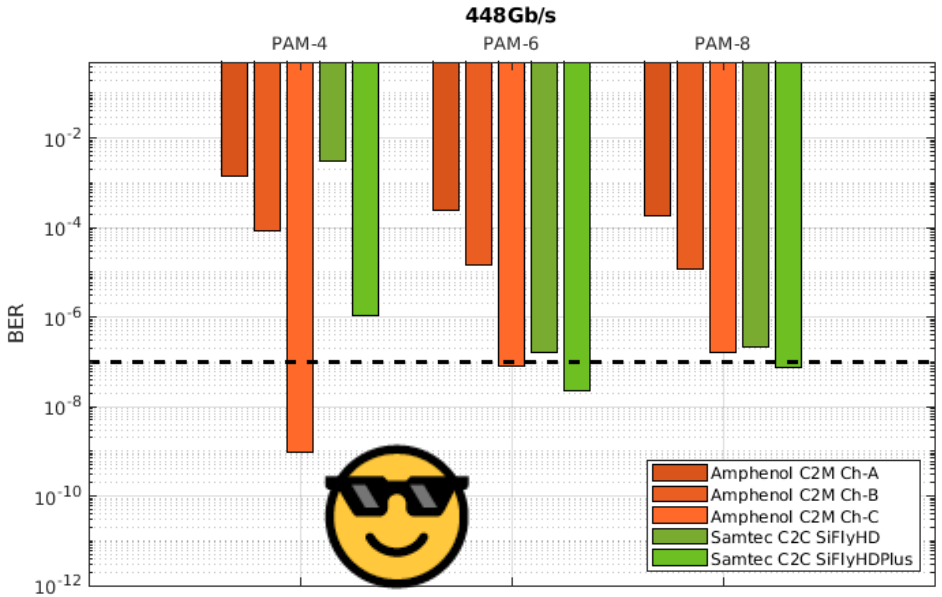

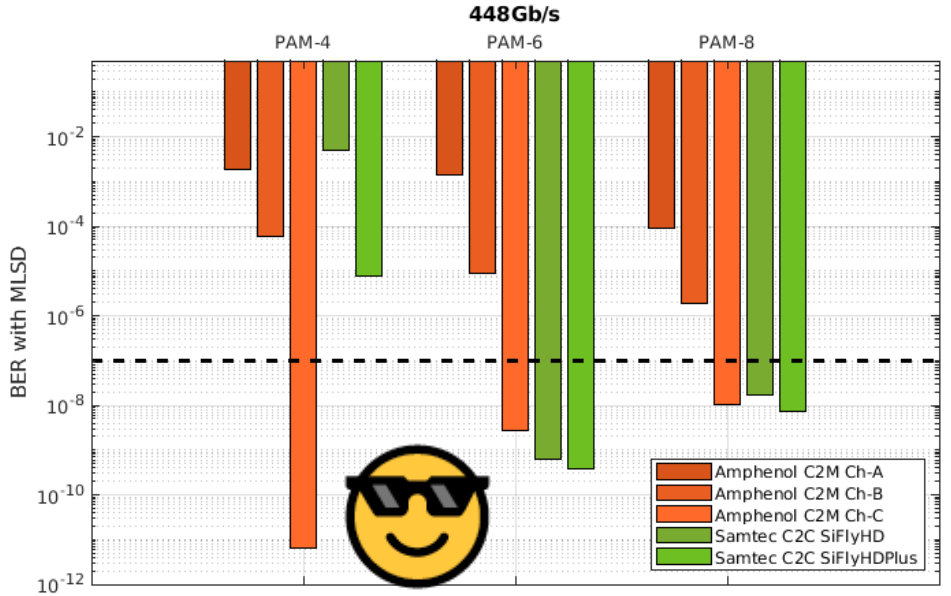

COM Results – 425 Gb/s vs 448 Gb/s

- Trade-off between the FEC overhead and gain is more pronounced than EVER at 400G!

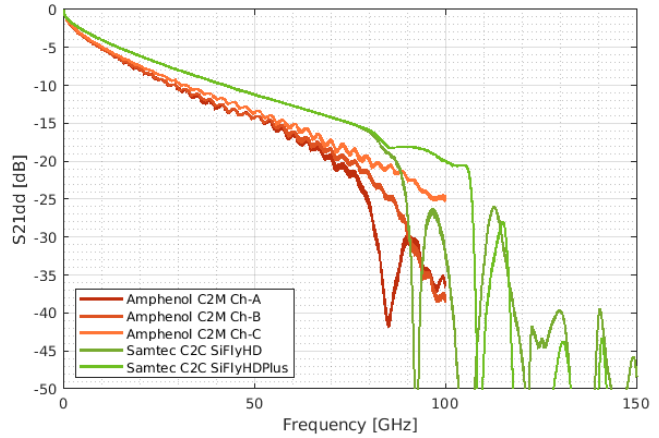

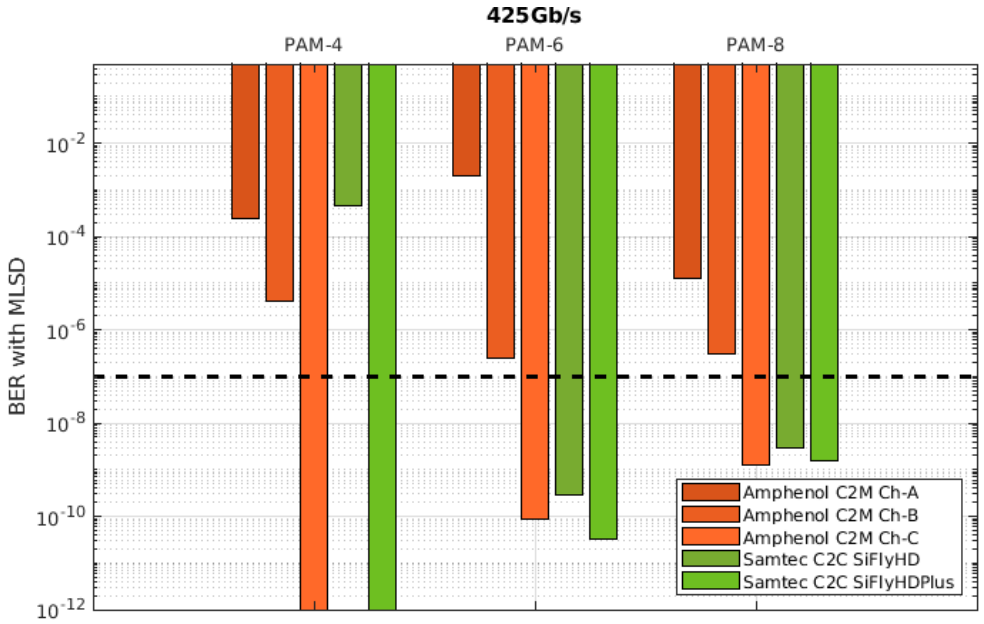

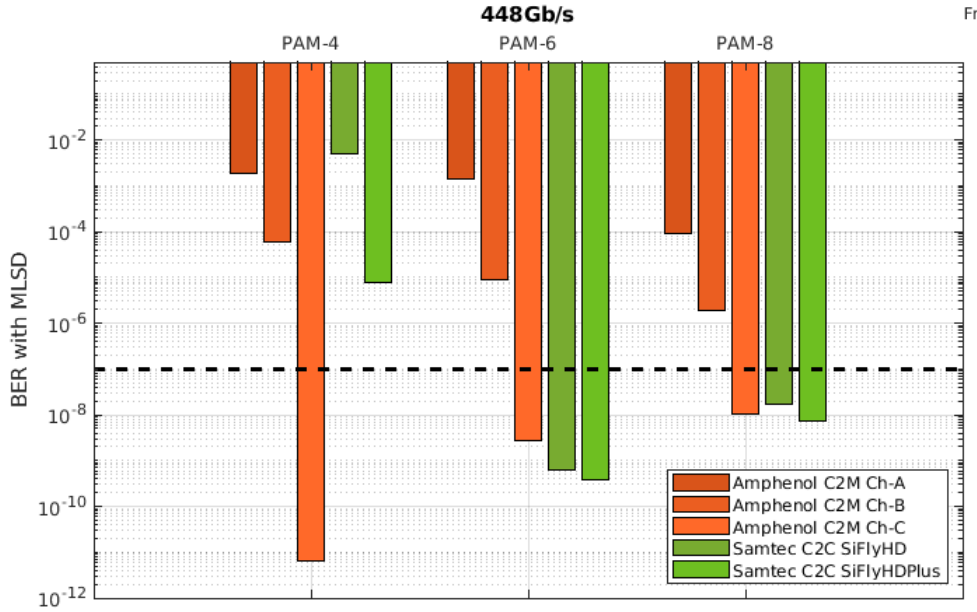

COM Results – Sensitivity to Different Parameters C2M Amphenol Channel C [448 Gb/s PAM4]

- Taking that previous performance across various channels (nominal) and sweeping SerDes parameters finding additional margin at 448 Gb/s PAM4
  - Assuming the extrapolation assumption stands, and the real channel data 100-150 GHz follows it
- The challenges:
  - **Analog bandwidth of front-end networks and CTLE/VGA**
    - With the same ESD capacitance, can we get the bandwidth we need for 448 Gb/s?
  - **Jitter**
    - AFE doesn't benefit as much as the digital as we go to more advanced nodes – how can we keep jitter low at these data rates? Especially in 2nm and below?
  - **ADC ENOB**
    - Need to maintain ENOB as sampling frequency is increased

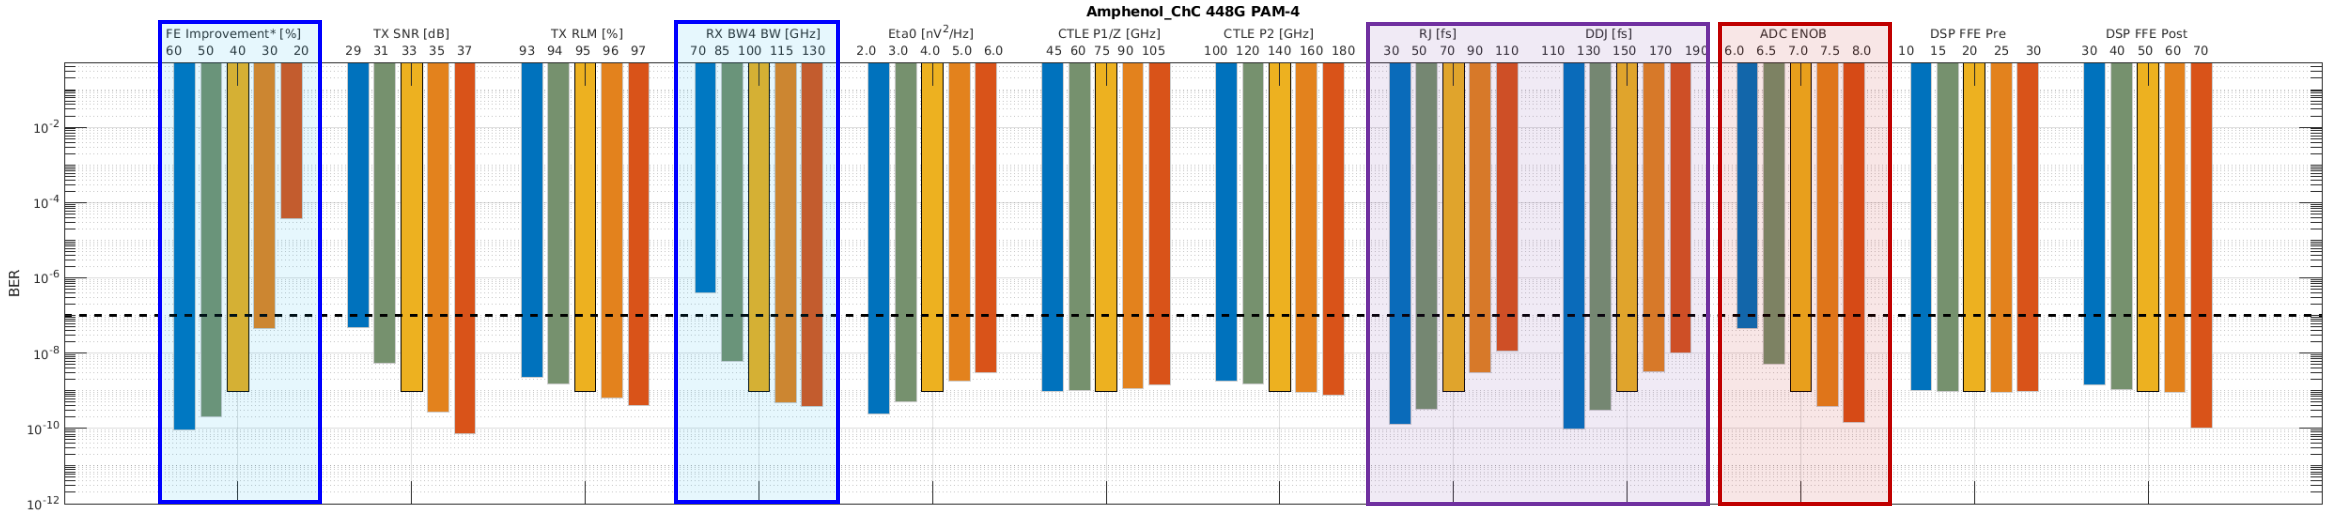

COM Results – Sensitivity to Different Parameters
C2C Samtec SiFlyHDPlus [448 Gb/s PAM4]

- There is not much we can do on the SerDes side to meet the performance requirement at 448 Gb/s
  - Using MLSD would help, but may not have enough margin at 448 Gb/s
  - As shown in the previous slides, at 425 Gb/s there is still a chance for PAM4
- Without further improving the interconnect, we'd need to go to PAM6 at 448 Gb/s here

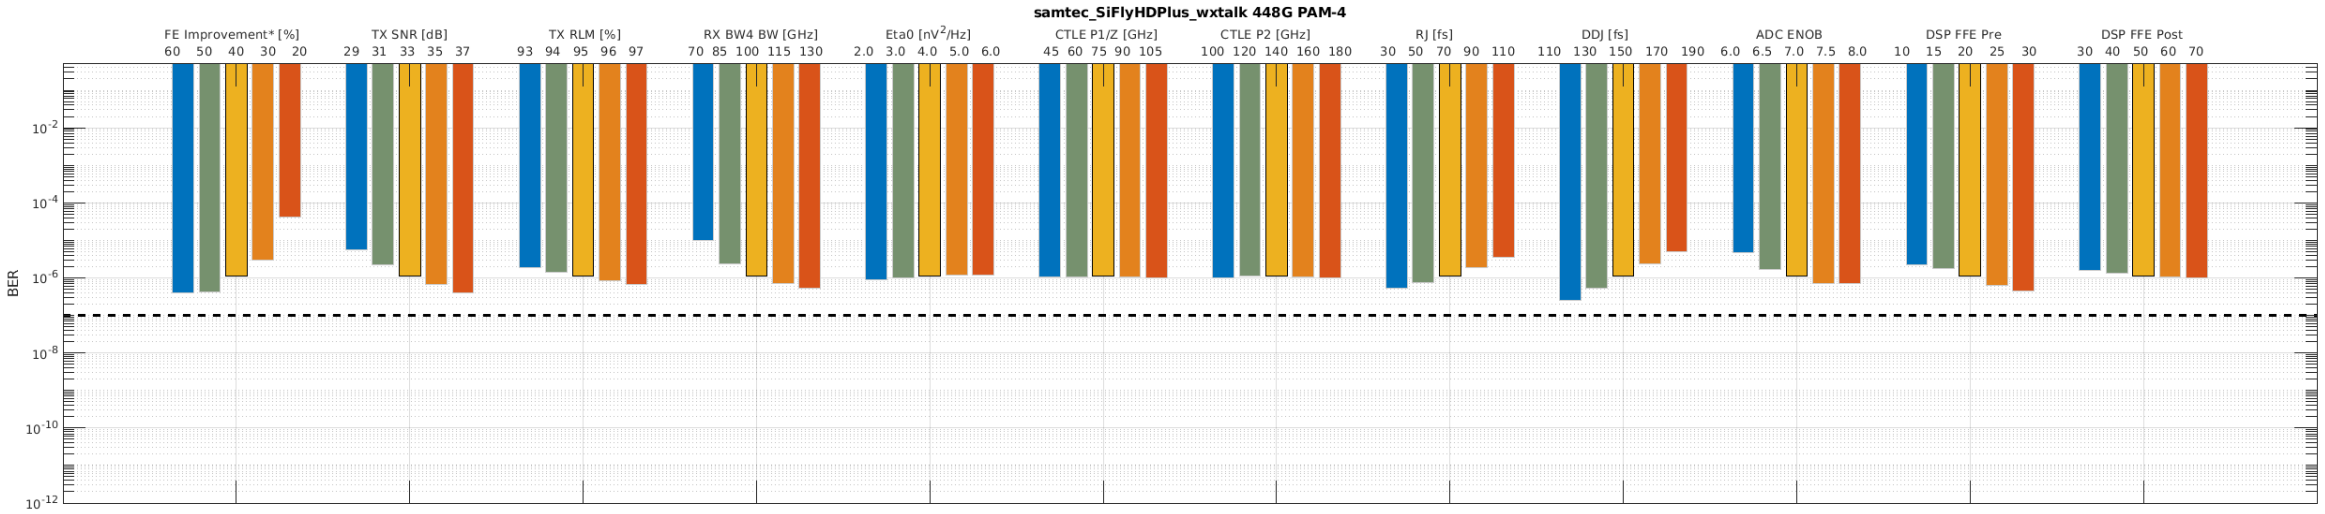

COM Results – Sensitivity to Different Parameters
C2C Samtec SiFlyHDPlus [448 Gb/s PAM6]

- PAM6 relaxes the sensitivity to front-end network and VGA/CTLE significantlly
- The main challenges for PAM-6 are **jitter** and **TX SNR**
  - The more we can push the SNR, the more benefit we get from higher order modulations such as PAM6

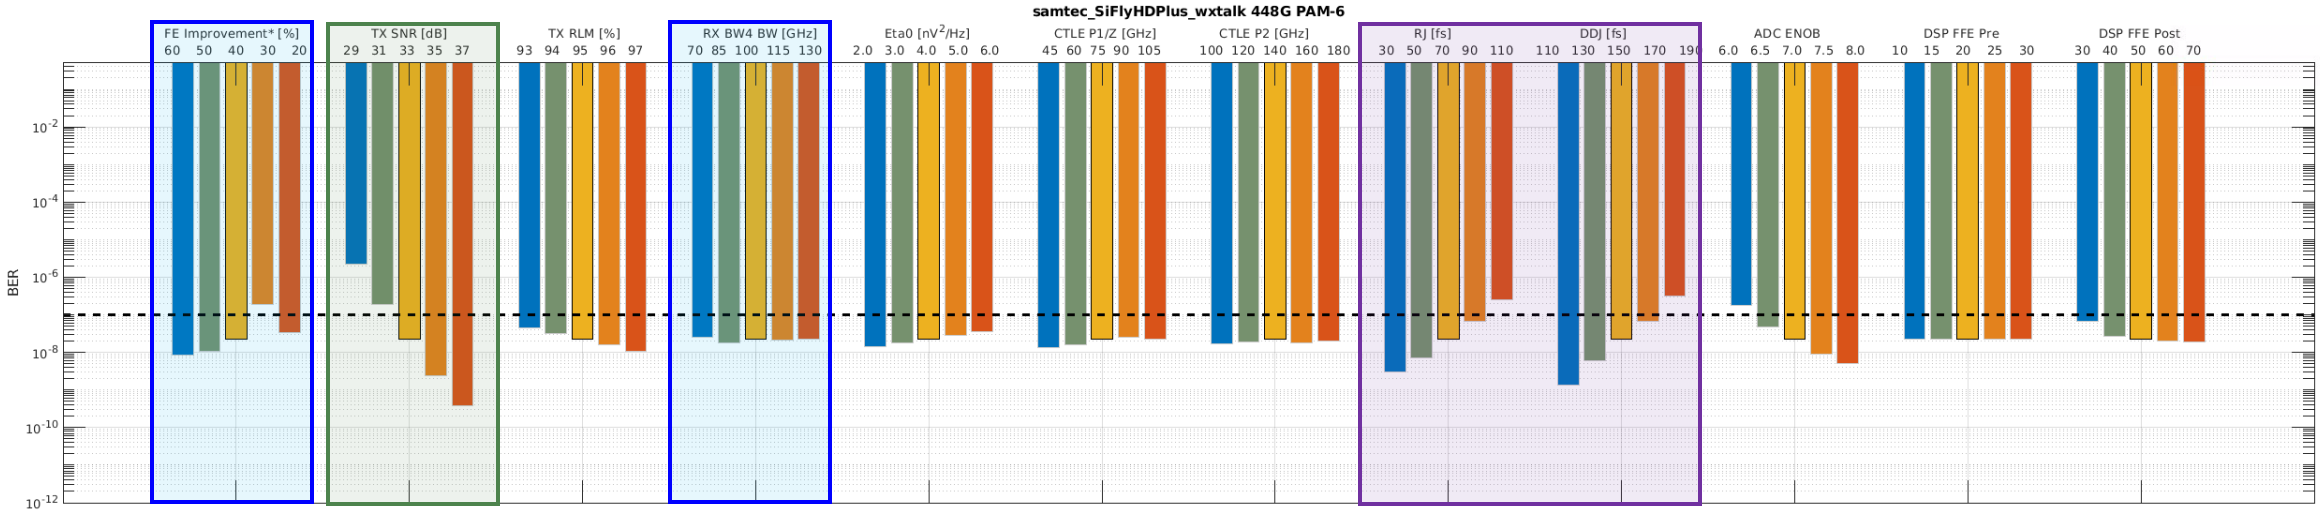

The Role of Standards in the Rapid Evolution of AI Networks

A message from the OIF PLL Interop Chair

Timeline of IEEE Ethernet Standards Development

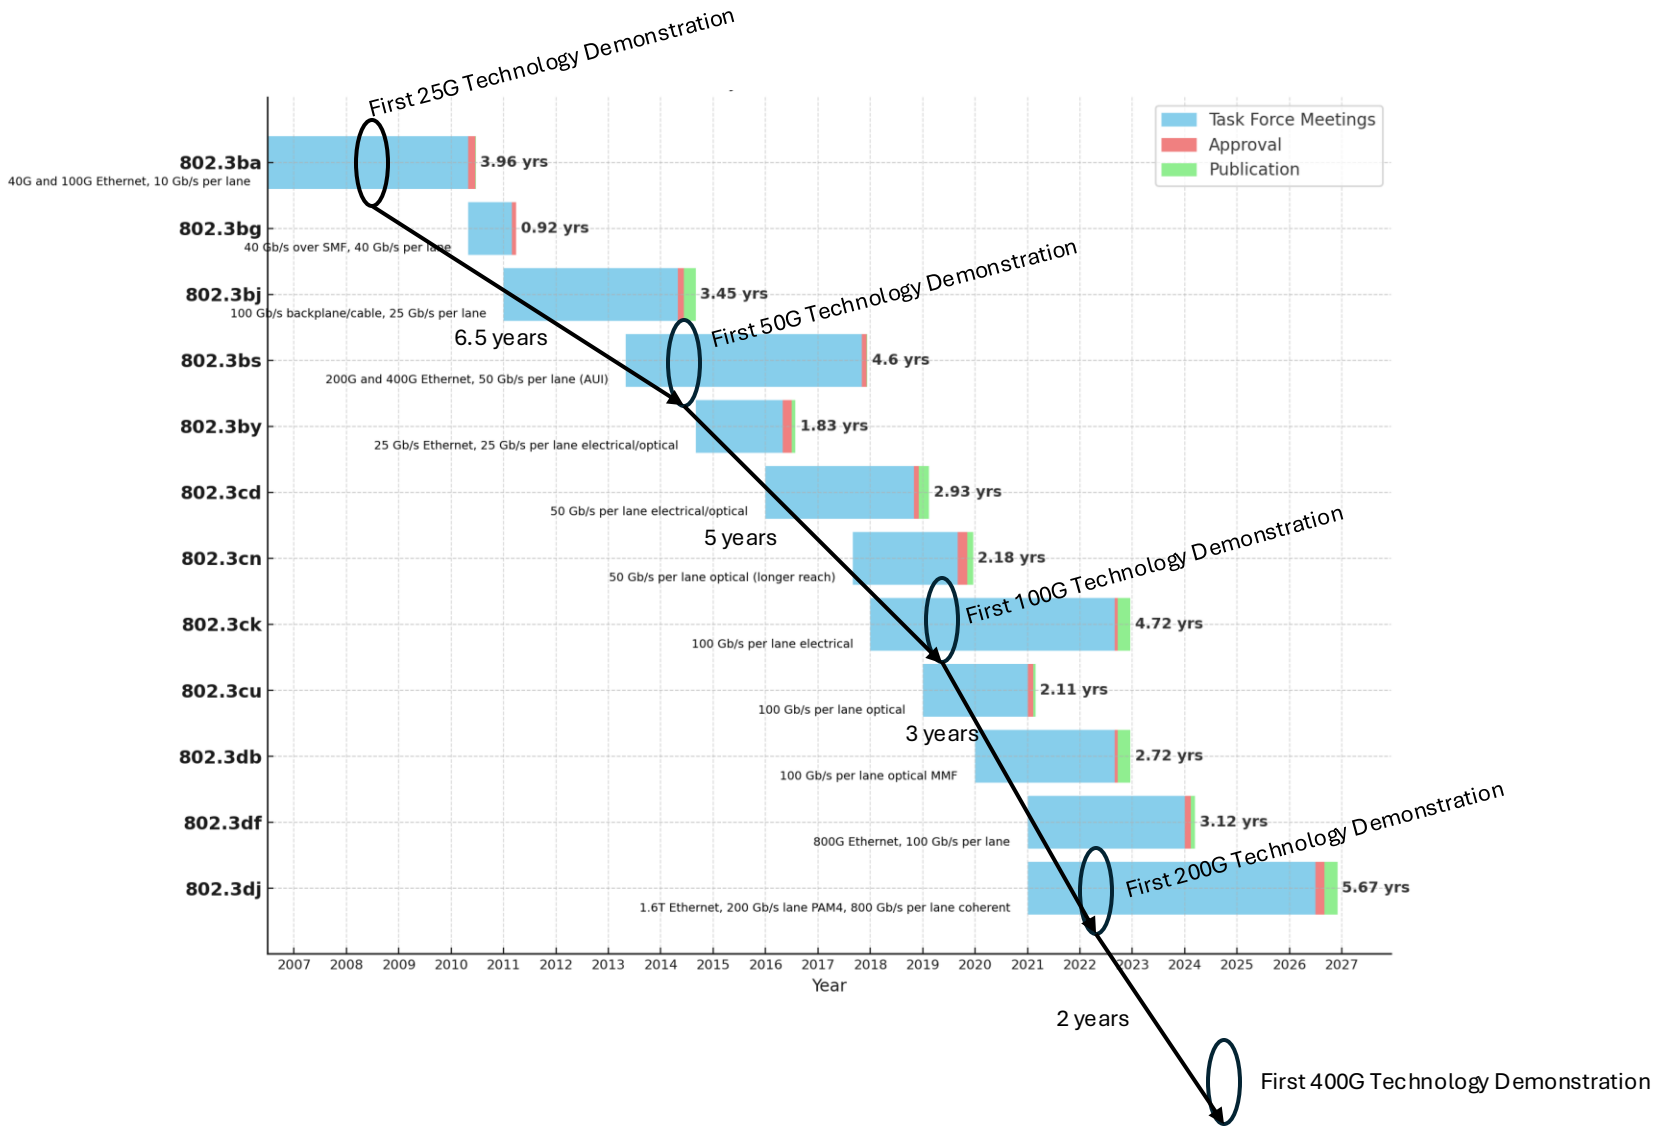

Accelerating Timelines vs. Massive Scaling

If you want to go fast, go alone.
If you want to go far, go together.

AI Hardware upgrading on 12-month cadence

Typical standard development timeline: 3 years for each generation

OIF
448G AI WORKSHOP

Thank you!

Simulation Setup - Front-End Network Improvements

- What improvements to the LC network mean in terms of bandwidth
  - 60% improvements mean getting close to 200 GHz bandwidth from the front-end networks!
- Improvements are quantified with respect to the numbers from the 200G COM sheet
  - The details can be found on the COM sheet