
ETHERNET ALLIANCE ROADMAP UPDATE

Moderator: David J. Rodgers, Teledyne LeCroy

Presenters:

Scott Kipp – Brocade/Broadcom

Mark Gustlin – Xilinx

Jeffery Maki – Juniper Networks

Matt Traverso – Cisco

13 March 2018



Presentation Disclaimer

The views expressed in this presentation are the views of the presenter and not necessarily those of the Ethernet Alliance.

The Ethernet Alliance

**A Global Community of End Users, System Vendors,
Component Suppliers and Academia**

■ Our Mission

- To promote industry awareness, acceptance and advancement of technology and products based on, or dependent upon, both existing and emerging IEEE 802 Ethernet standards and their management.
- To accelerate industry adoption and remove barriers to market entry by providing a cohesive, market-responsive, industry voice.
- Provide resources to establish and demonstrate multi-vendor interoperability.



The Ethernet Alliance

**A Global Community of End Users, System Vendors,
Component Suppliers and Academia**

▪ **Activities**

- Promotion and marketing
- Education awareness
- Interoperability testing and demonstration
- Industry consensus building
- Technology and standards incubation
- Power over Ethernet (PoE) Certification program

- **For more information – see**
www.ethernetalliance.org



Ethernet Alliance Strategy

- **Expanding the Ethernet Ecosystem and Supporting Ethernet Development**
 - **Facilitate interoperability testing**
 - Industry Plug Fests supporting member and technology initiatives
 - **Collaborative Interaction with other Industry Organizations**
 - Multiple SIGs, Applications and MSAs
 - **Global Outreach**
 - Worldwide Membership
 - **Host Technology Exploration Forums (TEFs)**
 - **Thought Leadership**
 - PoE Certification Program
 - **Promotion of Ethernet**
 - Industry Analysts
 - Education
 - Marketing
 - *Trade show*
 - *White Papers*

Learn More About Ethernet Alliance

- Ethernet Alliance:

visit www.ethernetalliance.org;



Follow @EthernetAllianc on Twitter



Join the Ethernet Alliance [LinkedIn group](#)



Visit the Ethernet Alliance [Facebook](#) page

- PoE Certification Program:

visit www.ethernetalliance.org/poecert/

- On The OFC 2018 Show Floor at Booth 2648

2018 ETHERNET OVERVIEW

Scott Kipp – Brocade/Broadcom

Scott Kipp has developed Ethernet and Fibre Channel standards and worked in industry associations for 17 years at Brocade Storage Networking within Broadcom. He was the President of the Ethernet Alliance for 6 years and is the Roadmap Chair who developed the Ethernet Roadmaps since 2014.

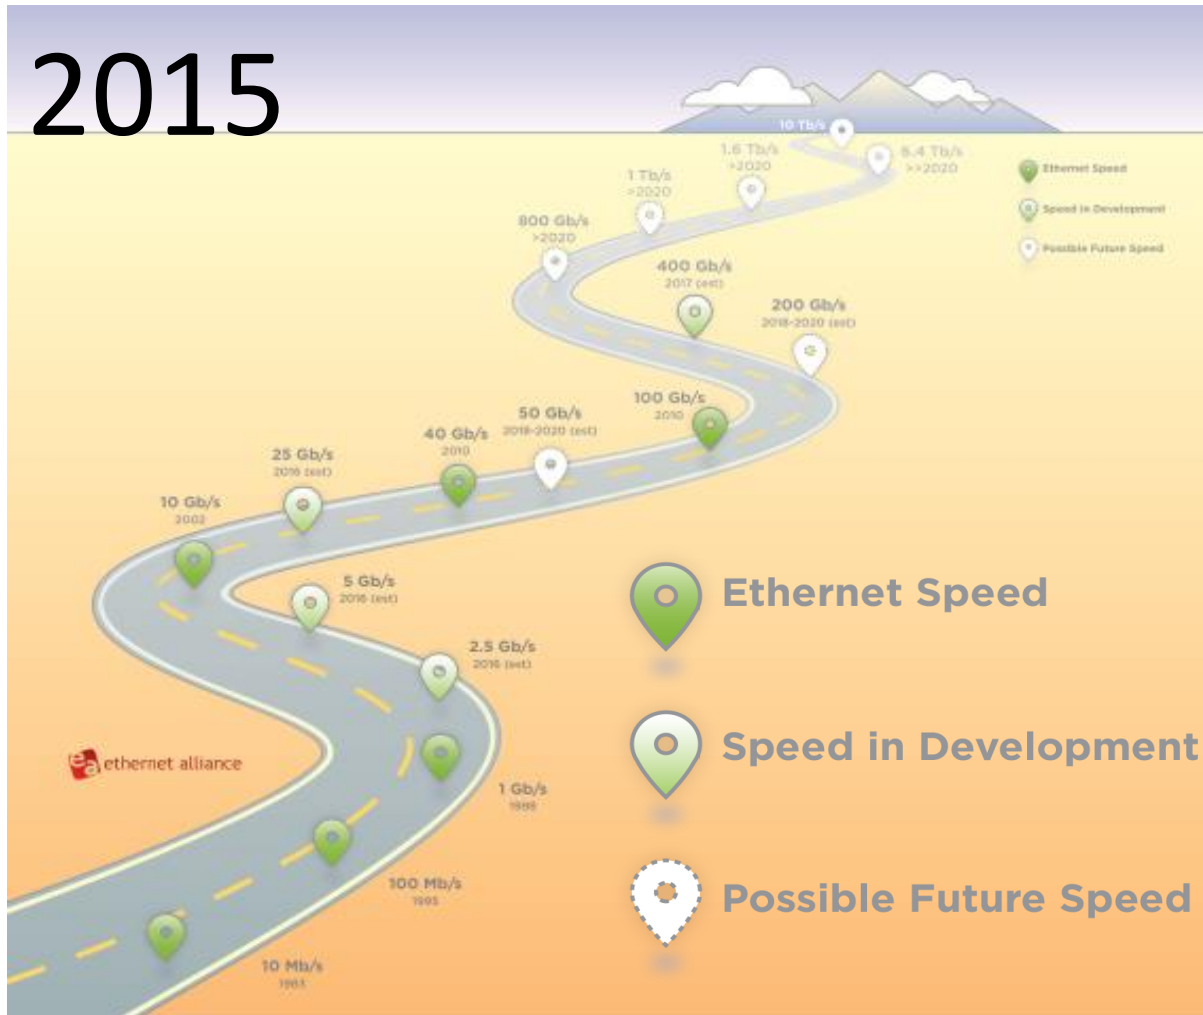


2018 Ethernet Roadmap Graphics


- Download all of the graphics from the 2018 Ethernet Roadmap at: www.ethernetalliance.org/roadmap
- All the graphics in this presentation are available to you!
- All graphics in this presentation are from the new roadmap – except the next page...

Past Roadmaps

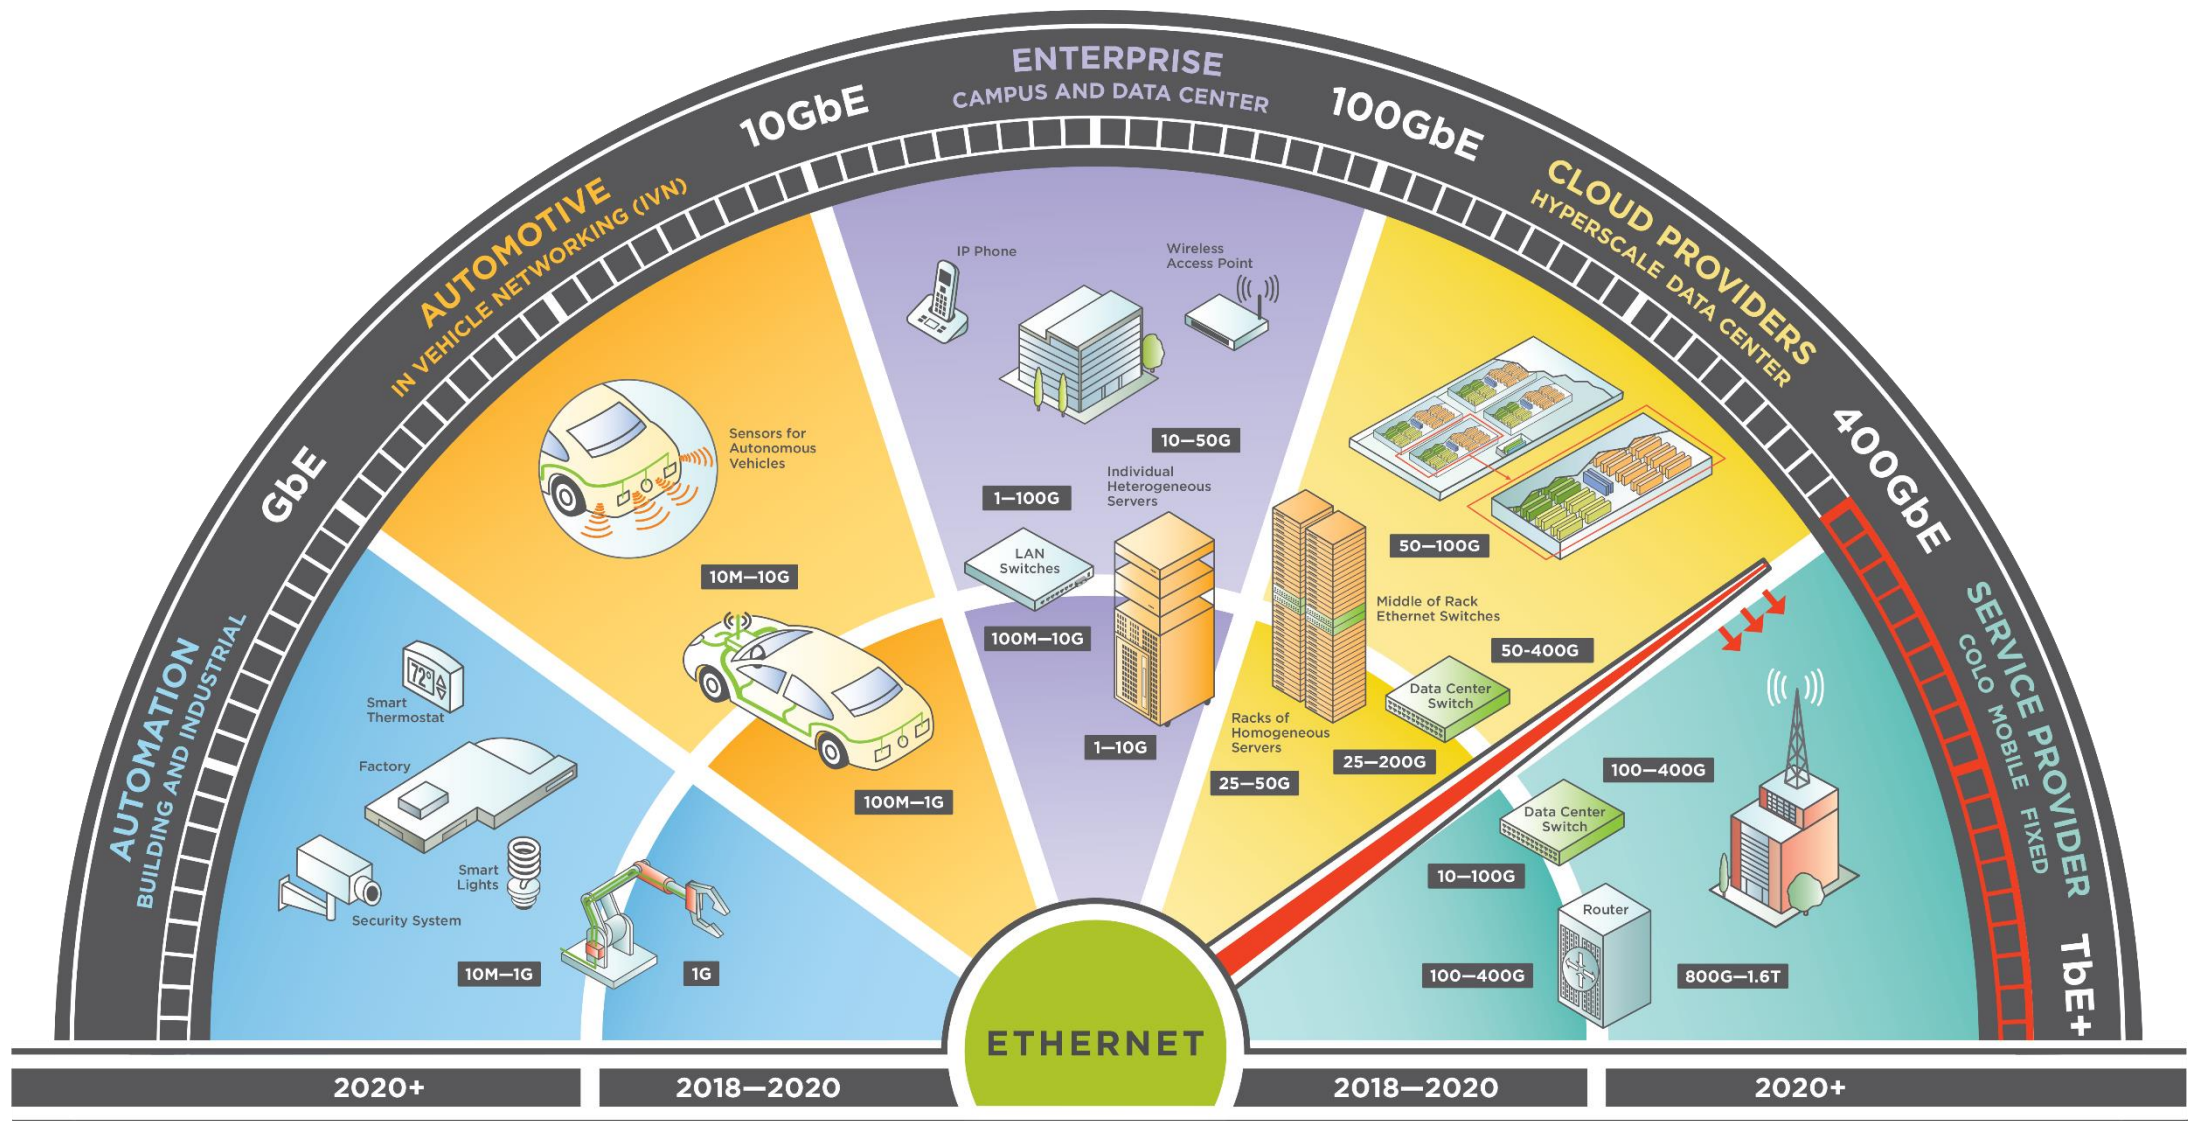
2015



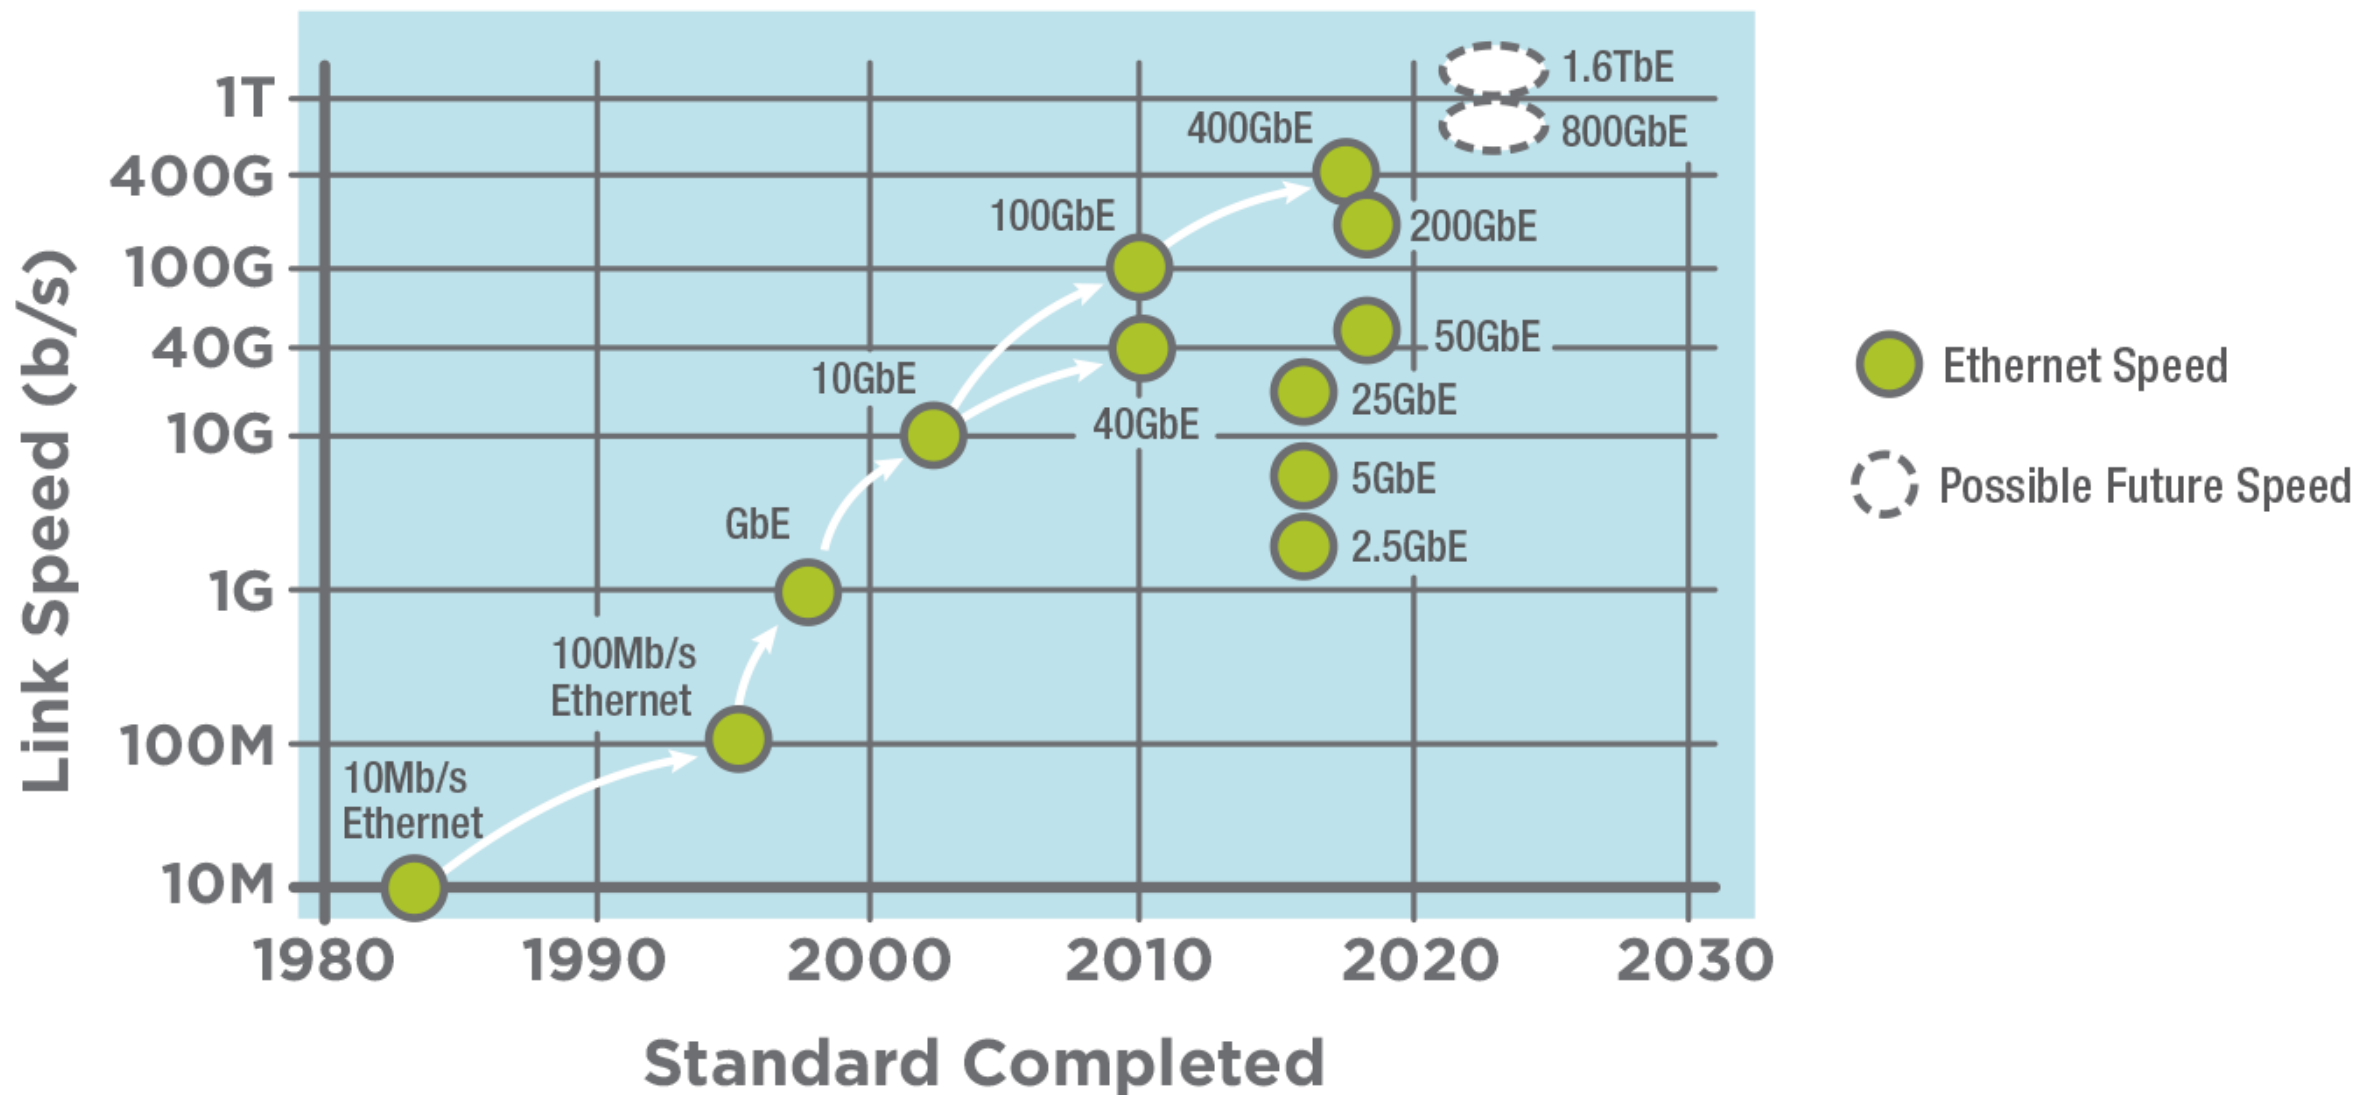
2016



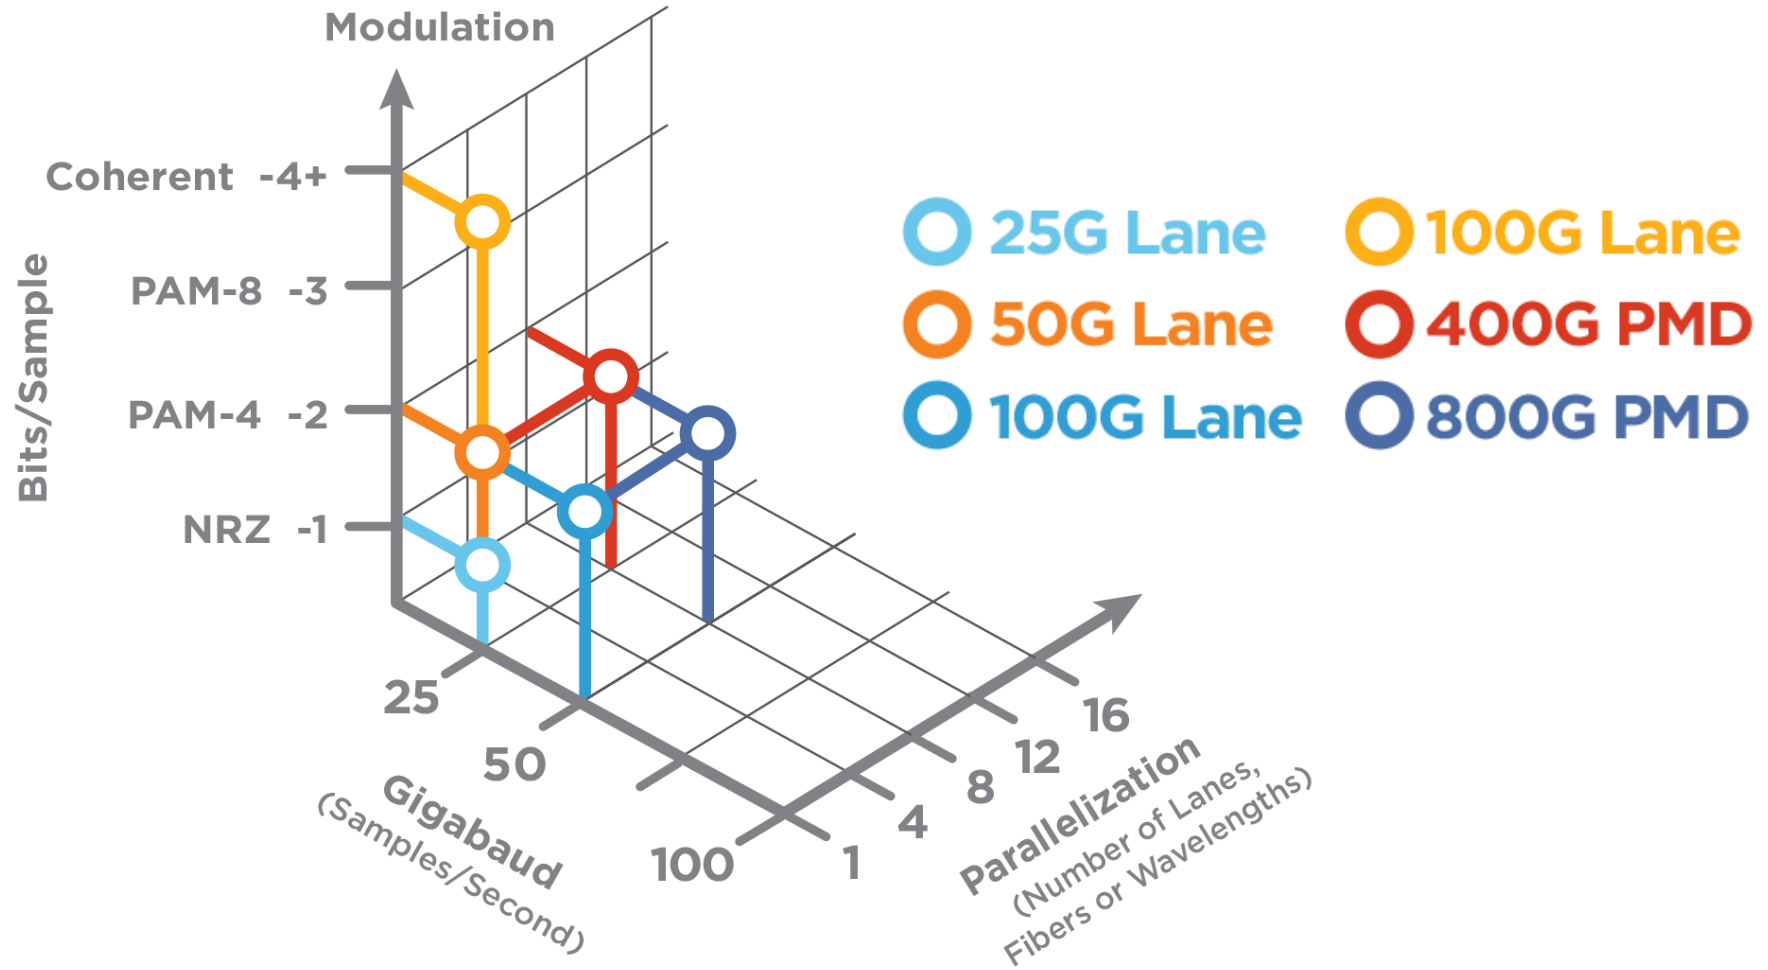
2018 Ethernet Roadmap



ETHERNET SPEEDS




HOW TO GO FASTER




How Many Lanes?

- 400GbE has used 4, 8 and 16 lanes for various links
- 100GbE electrical lanes being developed now to enable 400GbE in QSFP




FORM FACTORS

1-4 Lane Interfaces

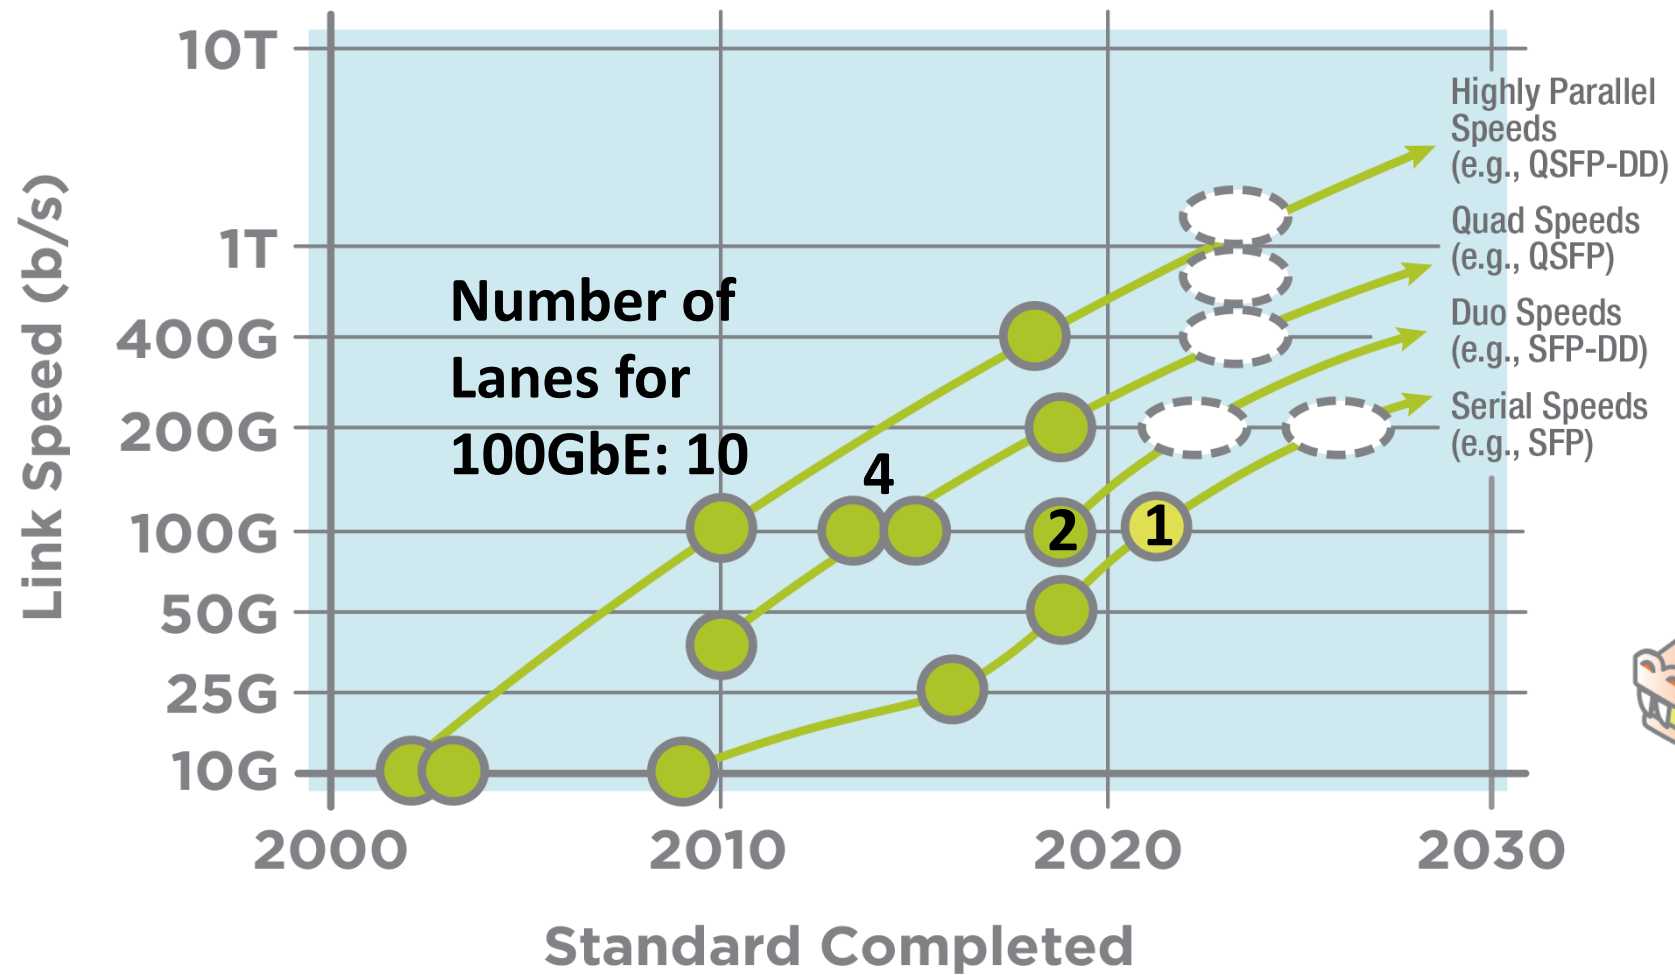


*Square inches of top surface of the module

4-16 Lane Interfaces



TO TERABIT SPEEDS



Ethernet Speed
 Speed in Development
 Possible Future Speed

CHALLENGES ON THE ROAD TO FASTER ETHERNET

Mark Gustlin – Xilinx

Mark is a director of product planning at Xilinx working on next generation silicon architecture. Mark previously spent 15 years at Cisco Systems architecting networking hardware and has over 30 years in communication hardware design. Mark is an active participant in the IEEE 802.3, and has been a past editor of 100, 200, and 400GbE standards.




Agenda

- **Insatiable appetite for bandwidth**
- **The need for Forward Error Correction**
- **100Gb/s per lane thoughts**
- **What happens beyond 100Gb/s/lane**

Insatiable Appetite for Bandwidth

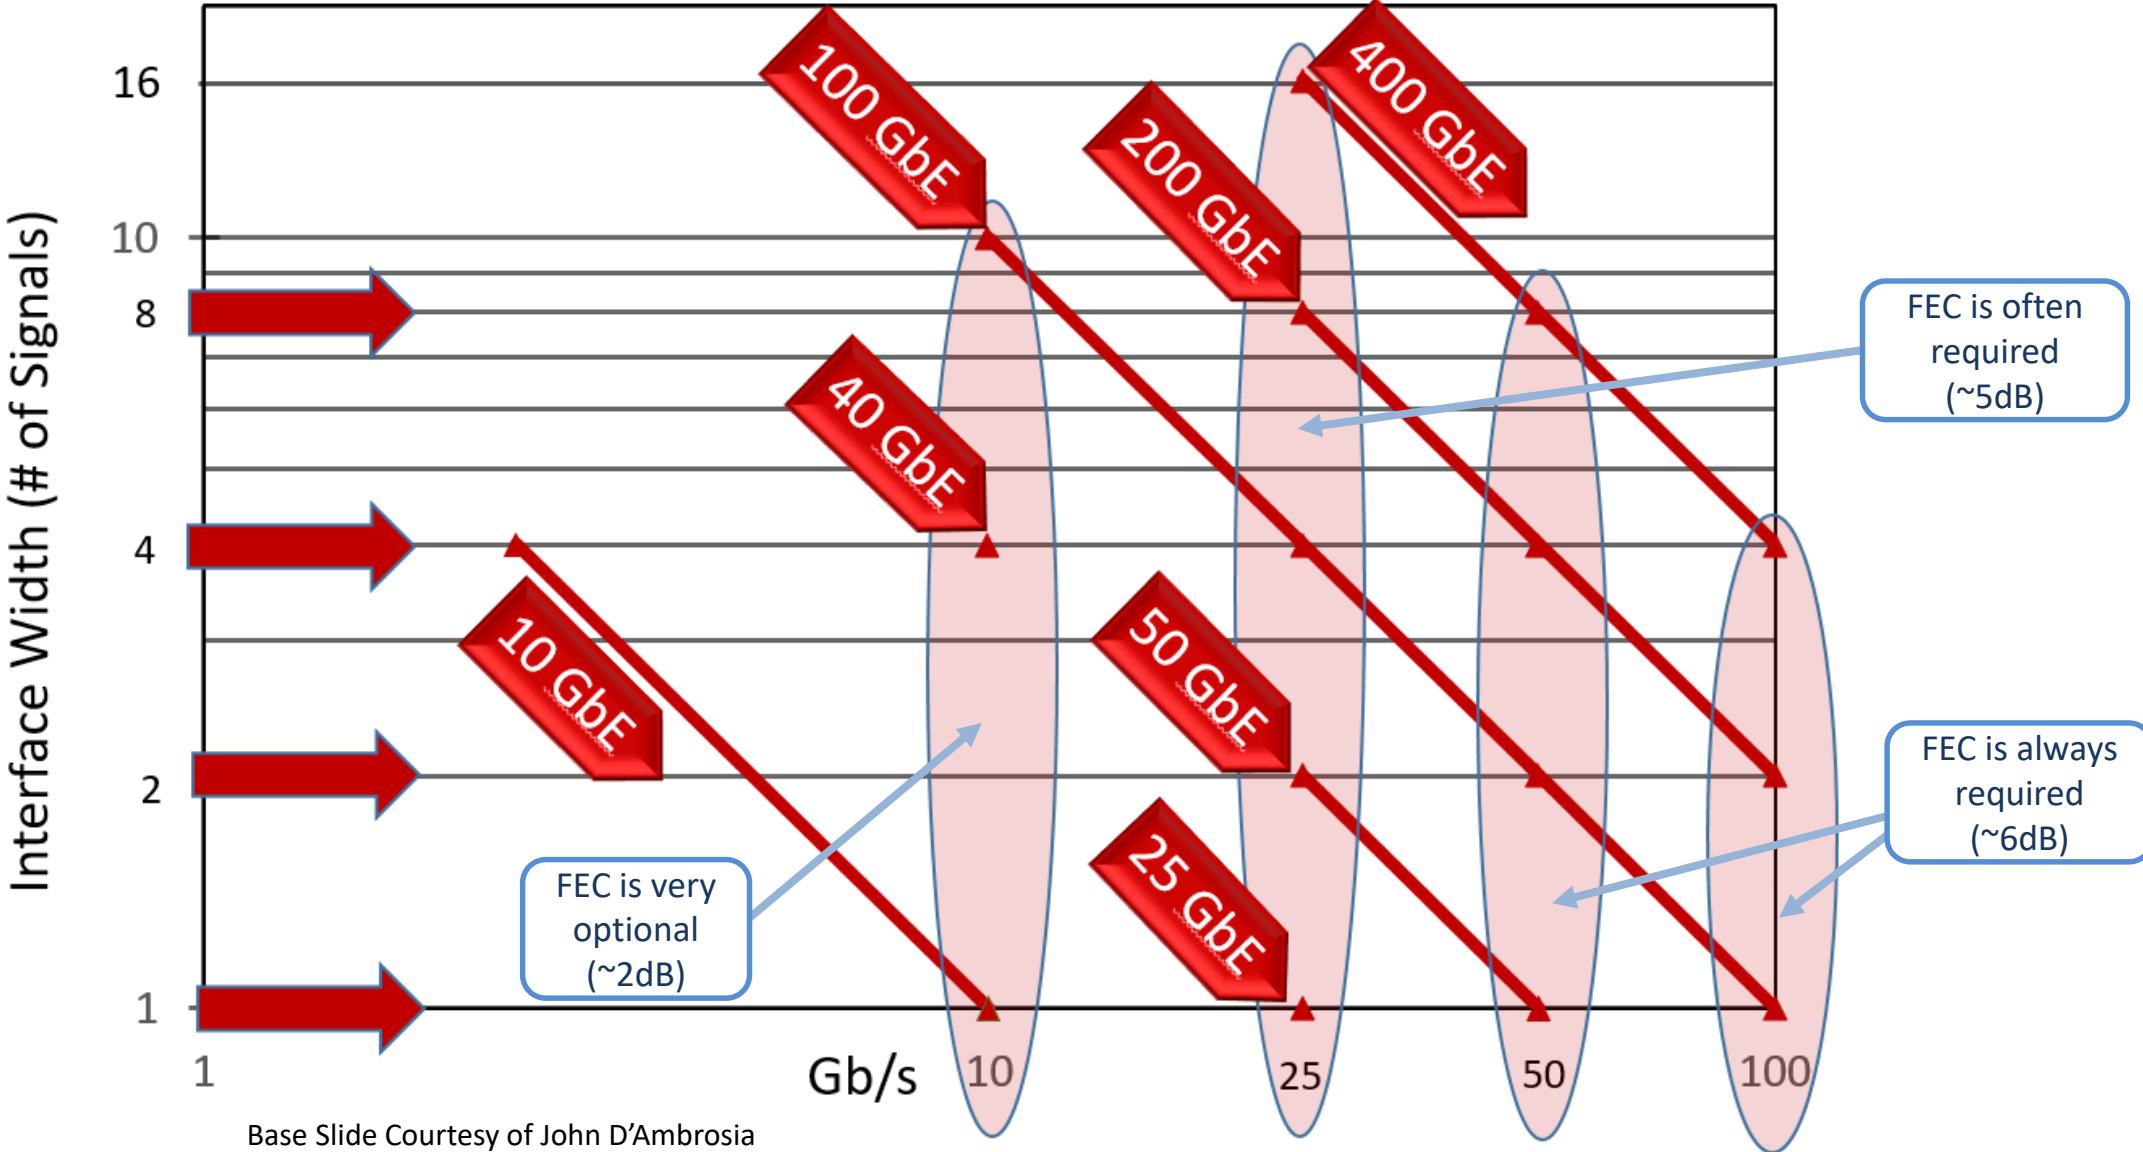
- **Constant growth in bandwidth requirements**
- **Requires constant innovation in:**
 - Switching and related silicon
 - Optical technology
 - SerDes technology
 - Underlying technologies such as FEC
- **Pushing physics forward**
 - Becoming painfully difficult!
- **Higher Speed ToR to Server**
 - 25GbE -> 50GbE -> 100GbE
- **Higher Speed Leaf/spine**
 - 100GbE -> 400GbE -> 800/1.6Tb

Figure 6. Cloud data center traffic growth



Source: Cisco Global Cloud Index, 2016-2021.

Families of Ethernet Lane Rates




Base Slide Courtesy of John D'Ambrosia


100G Per Lane Thoughts

- **100Gb/s per Electrical Lane study group just started**

- Looking to do chip to module, chip to chip, backplane and copper cable interfaces, all likely based on PAM4 signaling
- Things are much more difficult at 100Gb/s per lane compared to previous speeds
 - *Channels will be shorter than we are used to*
 - *From 35dB -> 30dB -> 28dB*
- Might need to utilize flyover cabling, orthogonal backplanes, cabled backplanes or other lower loss mediums to build the systems we are used to



Pictures courtesy of:





100G Per Lane FEC

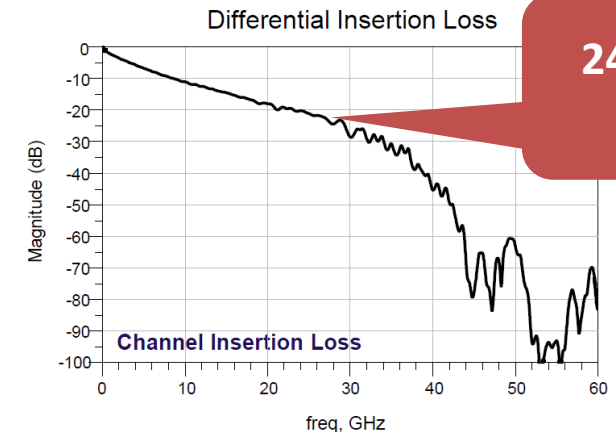

- **Strong desire to re-use the same FEC as we used for 50Gb/s per lane interfaces**
 - This is the RS(544,514,10) “KP4” FEC
 - Already used for some 100Gb/s Optical interfaces (the DR/DR4 PMDs)
 - We will re-use this FEC for chip to module interfaces per the adopted objective
 - We should be able to use it for backplane, copper cable and chip to chip interfaces
 - *Depends on if we want to compromise on the reach/loss budgets*
 - *Verses adding more complexity to switch and other silicon devices*
 - *Adding a stronger FEC cost > power, longer latency etc.*
 - Also strong desire to stay with PAM4 encoding (vs. a higher level encoding)
 - *We will decide once in task force*
 - *At least one proposal for bidirectional signaling*

Early 112G Simulations

- Early 112G per lane sims are being done
- This work is from zhang_100GEL_adhoc_01_022618
- These eyes are actually pretty open
- Challenging problem
 - Need to choose our channels wisely
 - Limit noise/crosstalk etc.
 - 2m copper cable and 28dB backplane channels were just adopted



Orthogonal Backplane Channel Results



Beyond 100G Per Lane

- **To 200G+ and beyond, yikes....**
- **Physics gets really difficult**
 - Channels will get shorter again?
 - *What are the minimum useful channels?*
 - *Passive copper cable disappears (could Active be used)?*
 - 7m (10G), 5m (25G), 3m (50G), 2m (100G)
 - We could look at higher order modulation (beyond PAM4 to PAM8/16)?
 - *Can cost more power, stronger FEC is required*
 - *More equalization*
 - Power per lane is critical
 - *Some devices have up to 256 SerDes today*
 - Stronger FEC has higher latency
 - *At odds with low latency DC requirements*
 - *Has higher power also (reduces chip densities)*
 - Do optical backplanes become a reality finally?
 - *Very short channels to nearby optics will still work*

FORM FACTORS AND OPTICAL MSA'S

Jeffery Maki – Juniper Networks

Dr. Jeffery Maki is a Distinguished Engineer II at Juniper Networks working on cloud optics. He has a Ph.D. in Optics from The Institute of Optics, University of Rochester, Rochester, New York. He is a member of the board of directors of the Ethernet Alliance and an IEEE 802.3 Ethernet voter. In the OIF, he is PLUG Working Group Chairman and technical voter for Juniper Networks. For the Consortium for On-Board Optics (COBO), he is Vice President and Data Center Networking Working Group chairman. He is co-chairman of the 100G Lambda MSA Group and a participant in the QSFP-DD and SFP-DD MSA groups. He is an OSA and IEEE member.



Agenda

- 100G and 400G Optical Specifications (PMDs)
- Single-Mode 100G Ethernet Evolution
- 400G by Form Factor
- Potential Bandwidth Density Progression

4WDM MSA (4 x 25G-λ NRZ)

- 100G-4WDM-10 Specification
 - 10 km reach with 802.3bj KR4 FEC mandatory
 - CWDM4 grid and PIN receiver
- 100G-4WDM-20 Specification
 - 20 km reach with 802.3bj KR4 FEC mandatory
 - LAN-WDM grid and PIN receiver
- 100G-4WDM-40 Specification
 - 40 km reach with 802.3bj KR4 FEC mandatory
 - LAN-WDM grid and APD receiver
- See http://www.4wdm-msa.org/wp-content/uploads/2017/01/4WDM10_MSA_Spec_R1.0.pdf
- See http://4wdm-msa.org/wp-content/uploads/2017/07/4WDM20_40_MSA_R1.0_Jul28_2017.pdf

100G Lambda MSA

- 100G-FR Specification
 - 100G- λ PAM-4 serial, 2 km reach with 802.3cd “KP4” FEC mandatory
- 100G-LR Specification
 - 100G- λ PAM-4 serial, 10 km reach with 802.3cd “KP4” FEC mandatory
- See <http://100glambda.com/specifications/download/2-specifications/2-100g-fr-and-100g-lr-technical-specs-d1p0>
- 400G-FR4 Specification
 - 4 x 100G- λ PAM-4, CWDM4 grid, 2 km reach with 802.3bs “KP4” FEC mandatory
- See <http://100glambda.com/specifications/download/2-specifications/1-400g-fr4-technical-spec-d1p0>

CWDM8 MSA

- 400G-CWDM8-2 Specification
 - 8 x 50G- λ NRZ, CWDM grid, 2 km reach with 802.3bs “KP4” FEC mandatory
 - See https://www.cwdm8-msa.org/media/400G_CWDM8_MSA_2km_Optical_Interface_Specifications_r1_1.pdf
- 400G-CWDM8-10 Specification
 - 8 x 50G- λ NRZ, CWDM grid, 10 km reach with 802.3bs “KP4” FEC mandatory
 - See https://www.cwdm8-msa.org/media/400G_CWDM8_MSA_10km_Optical_Interface_Specifications_r1_1.pdf

100GE – PMDs

Reach	Media	10 x 10G-λ NRZ (no FEC)	4 x 25G-λ NRZ (no FEC)	4 x 25G-λ NRZ (802.3bj KR4 FEC)	2 x 50G-λ PAM-4 (802.3cd FEC)	100G-λ PAM-4 (802.3cd FEC)
SR (100 m)	Parallel MMF	100GBASE-SR10		100GBASE-SR4	100GBASE-SR2	
	Duplex MMF			100G-SWDM4		
	Duplex MMF				100G-BiDi [†]	
DR (500 m)	Parallel SMF			100G-PSM4		
	Duplex SMF					100GBASE-DR
FR (2 km)	Duplex SMF	10x10-2km		100G-CWDM4		100G-FR
LR (10 km)	Duplex SMF	10x10-10km	100GBASE-LR4	100G-4WDM-10		100G-LR
LR+ (20 km)	Duplex SMF			100G-4WDM-20		
ERlite (30 km)	Duplex SMF		100G-ER4lite			
ER (40 km)	Duplex SMF	10x10-40km	100GBASE-ER4	100G-4WDM-40		

- Bold font indicates IEEE 802.3 standard
- No **BASE** means defined outside of the IEEE 802.3 in a MSA or other SSO
- [†]Note: 100G-BiDi uses proprietary FEC in the module

100G Ethernet Evolution – Back Fill to QSFP28

... ← 100GBASE-DR, 100G-FR, and 100G-LR

(mux inside)

Single Port

Dual Ports

50G SERDES/Lanes

100G SERDES/Lanes


Quad Ports

Single Port

Eight Ports

Dual Port

Single Ports



Module width:

82 41.5 21.7 18.4

18.4

18.4

13.8

18.4

13.8

13.8

Electrical lanes:

10x10G

10x10G
4x25G

4x25G

8x25G

8x50G

2x50G

8x100G

2x100G

100G



ethernet alliance

400GE – PMDs

Reach	Media	16 x 25G-λ NRZ (802.3bs FEC)	8 x 50G-λ PAM-4 (802.3bs FEC)	4 x 100G-λ PAM-4 (802.3bs FEC)	400G-λ Coh. (400ZR FEC)
SR (100 m)	Parallel MMF	400GBASE-SR16	<i>400GBASE-SR8/SR4.2</i>		
DR (500 m)	Parallel SMF			400GBASE-DR4	
FR (2 km)	Duplex SMF		400GBASE-FR8 (PAM-4) 400G-CWDM8-2 (NRZ)	400G-FR4 (PAM-4)	
LR (10 km)	Duplex SMF		400GBASE-LR8 (PAM-4) 400G-CWDM8-10 (NRZ)	400G-LR4 (PAM-4) (under study)	
ER (40 km)	Duplex SMF				
ZR (80 km)	Duplex SMF				400G-ZR (OIF)

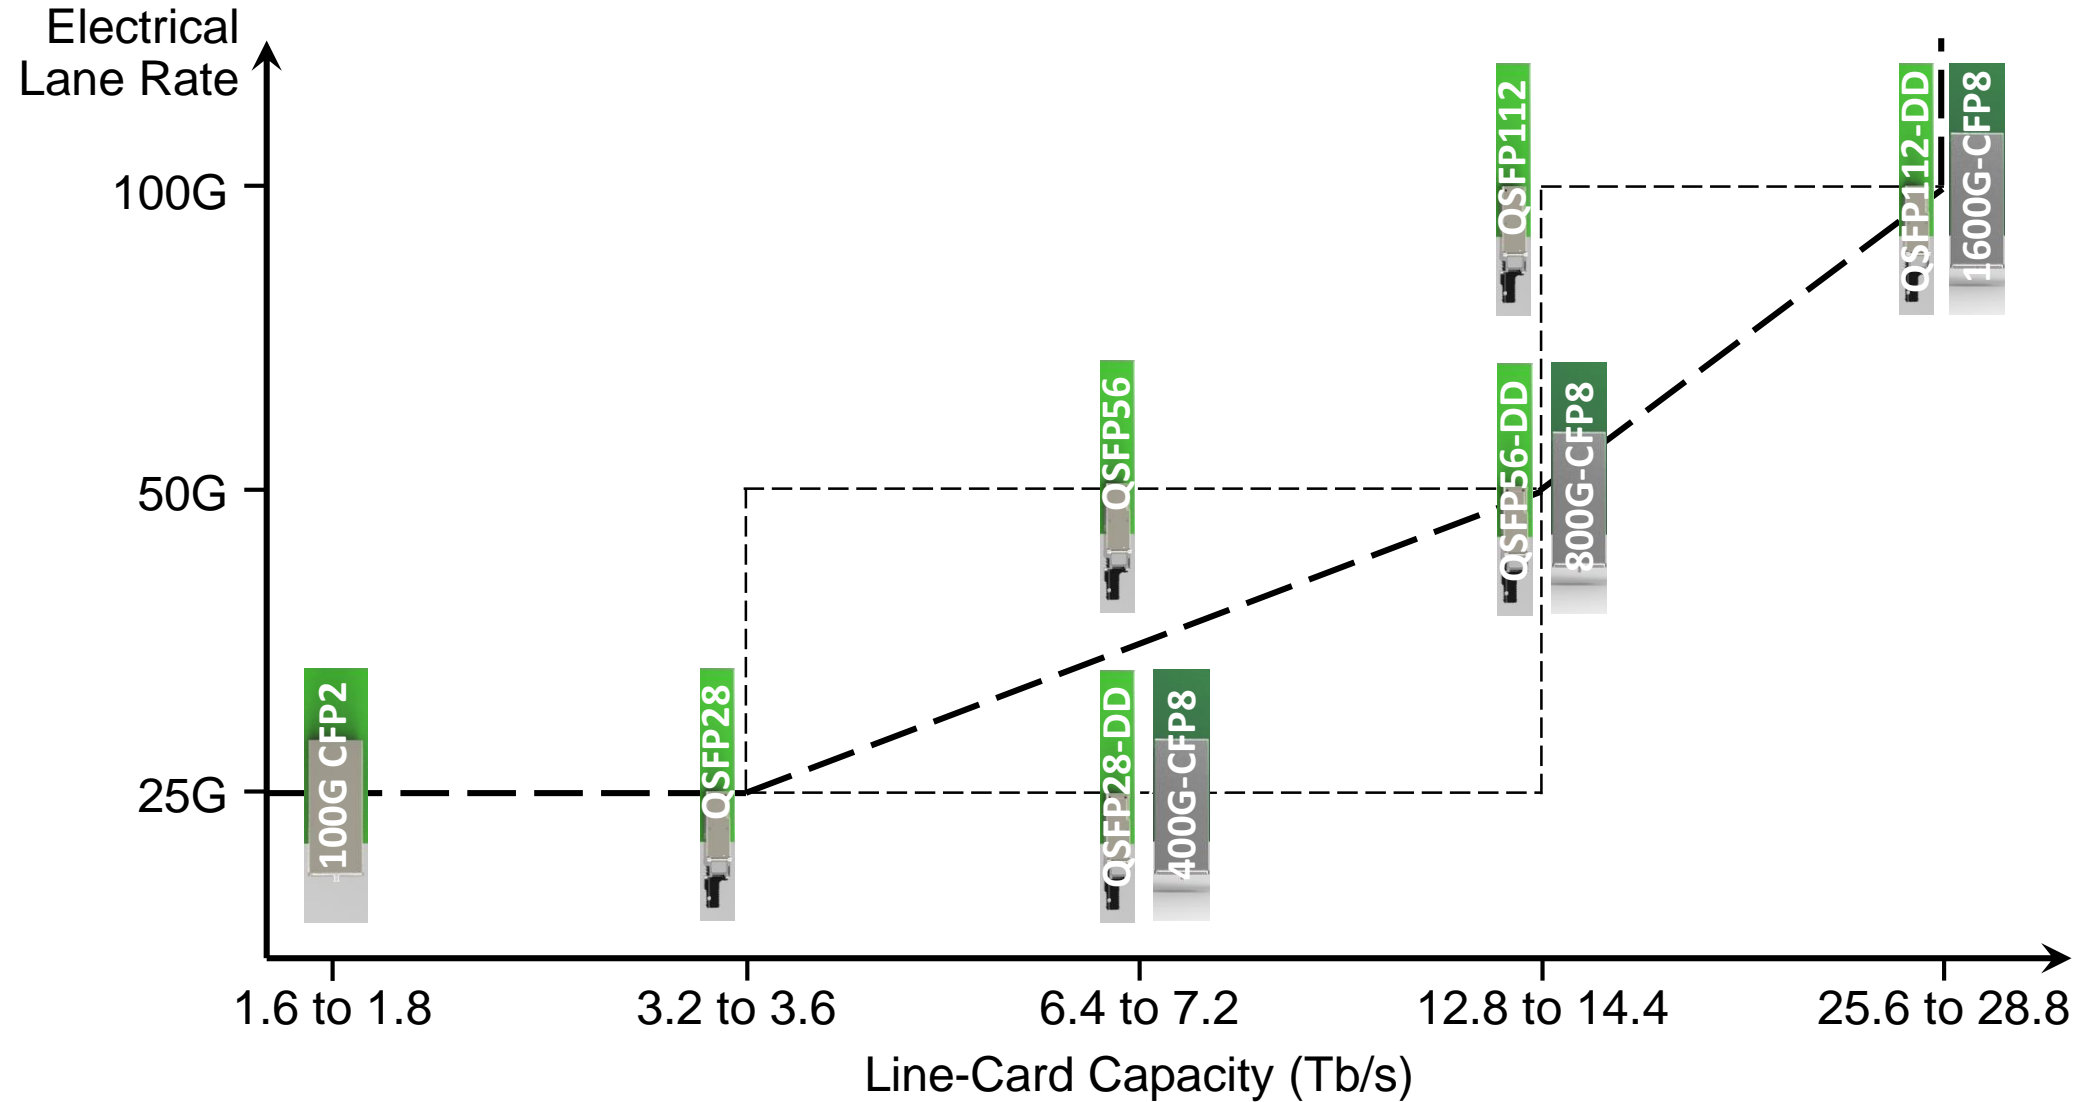
- Bold font indicates IEEE 802.3bs standard
- Italic-font indicates potential IEEE 802.3 standard via NGMMF Study Group
- No **BASE** means defined outside of the IEEE 802.3 in a MSA or other SSO

400GE by Form Factor

Reach	Media	PMD	CFP8	QSFP-DD	OSFP	COBO
SR (100 m)	Parallel MMF	400GBASE-SR16	Yes*			
		<i>400GBASE-SR8</i>		Yes	Yes	
DR (500 m)	Parallel SMF	400GBASE-DR4		Yes	Yes	Yes
FR (2 km)	Duplex SMF	400GBASE-FR8	Yes	Yes		
		400G-CWDM8-2		Yes	Yes	
		400G-FR4		Yes	Yes	
LR (10 km)	Duplex SMF	400GBASE-LR8	Yes	Yes		
		400G-CWDM8-10		Yes	Yes	
ZR (80 km)	Duplex SMF	400G-ZR		Yes	Yes	Yes

*Prototypes for demonstration. No plans from vendor base to provide GA.

Potential Bandwidth Density Progression



HIGHER ORDER MODULATION IN ETHERNET

Matt Traverso – Cisco


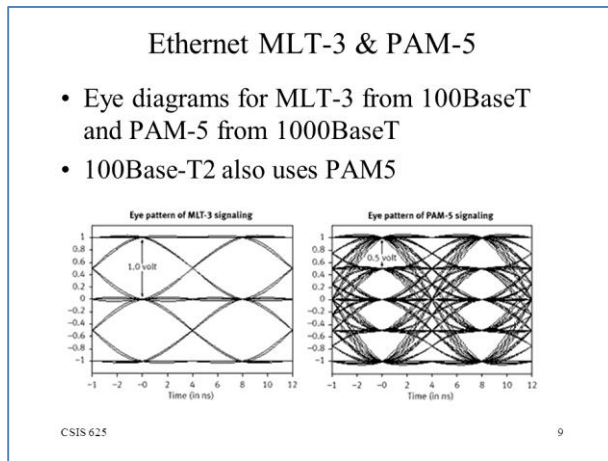


Matt Traverso is a Principal Engineer for Cisco Systems with a focus on next generation optical architectures based on silicon photonics. Matt has been active in the development and definition of optical networking standards and module form factors since 2000 including as the original editor of the CFP MSA (Multi-Source Agreement). He has been a frequent contributor to the IEEE Standards Association, Optical Internetworking Forum and other optical communications forums. Matt graduated from Stanford University in Materials Science & Engineering.



IEEE embracing higher order modulation







- IEEE 802.3bs & IEEE 802.3cd
- ≥ 50 Gbps elec & optical using PAM4 encoding
- Note 100BaseT & 1000BaseT used PAM...



Higher Order Modulation

3 ways to increase data-rate:

- Speed: Increase transmission frequency over each fiber,
- Parallelize: Increase number of fibers or wavelengths, or
- Complexity: Increase number of bits per symbol

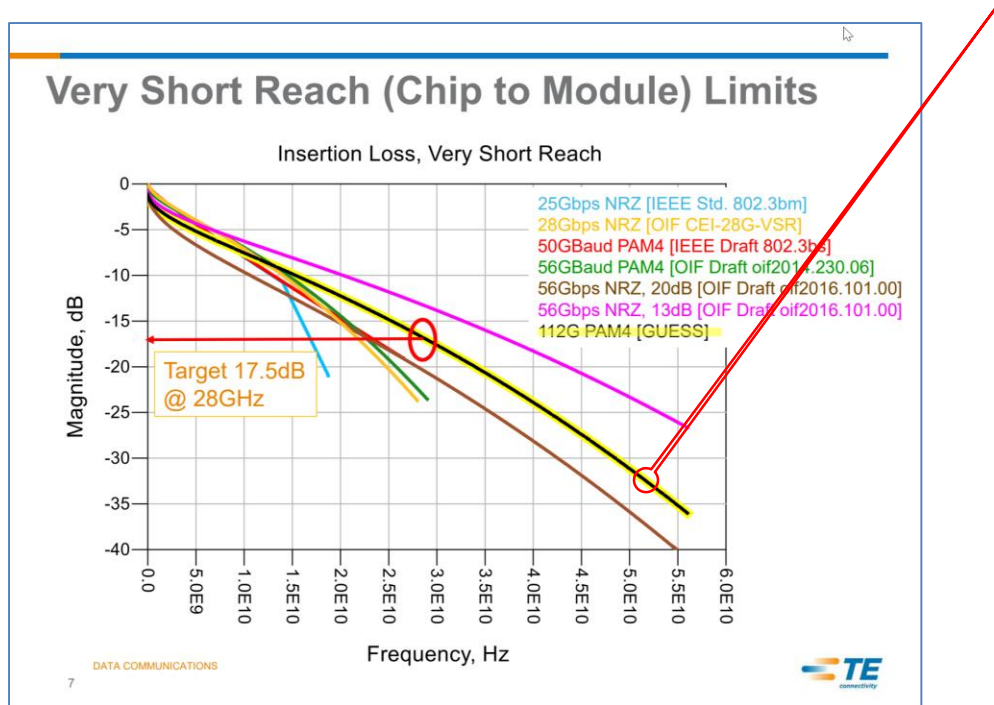


Same Data and Data Rate; however, PAM-4 can achieve it at half the frequency.

Why Higher Order Modulation

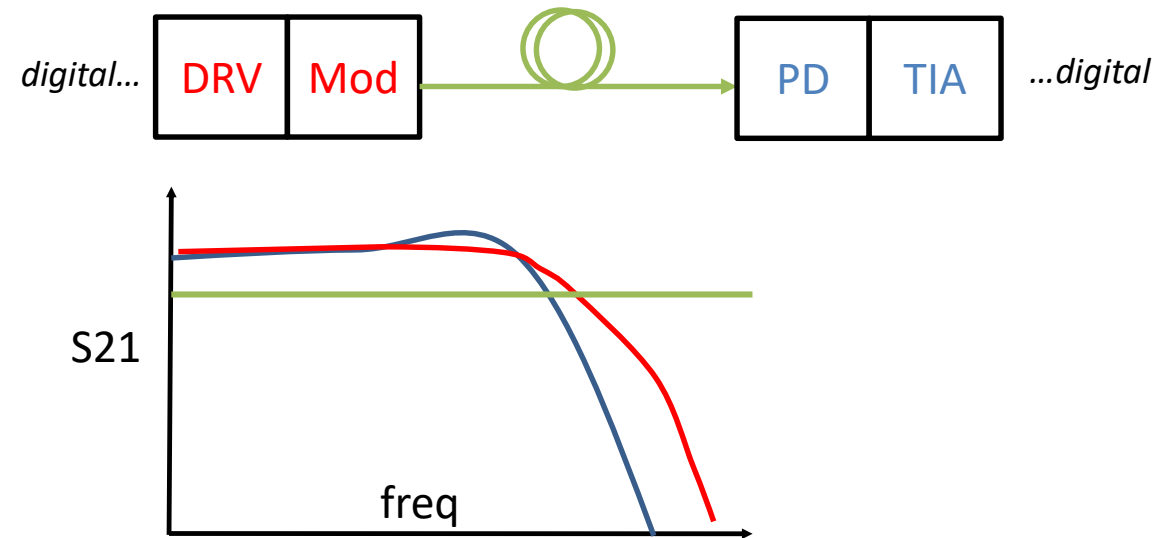
Electrical Channel

- Using PAM4 a ~17.5dB channel loss is approximated for 112G VSR length
- Using PAM2 the loss would be >30dB




Optical Channel

- The fiber insertion loss does not change with data rate
- Components (Driver/Modulator/TIA) introduce bandwidth dependent loss
 - Electrical channel within optical packaging also has BW dependent loss



Digression on understanding PAM2*





- The decision of a '1' or a '0' is based on the histogram (see red)
- The wider the spacings in the histogram, the easier for the receiver to correctly decide
- The decision of a '1' or a '0' is timed based on the phase transitions (see pink)
- These edge transitions indicate to the receiver when to time the amplitude decision

*In industry parlance, PAM2 is commonly referred to as NRZ

TDECQ and PAM4

- In PAM4, there are now twice the number of levels and four times the number of transitions
- IEEE invented a parameter known as TDECQ (Transmitter and dispersion eye closure – quaternary)
- This created a single transmitter penalty to allow designer to tradeoff different technologies and design choices

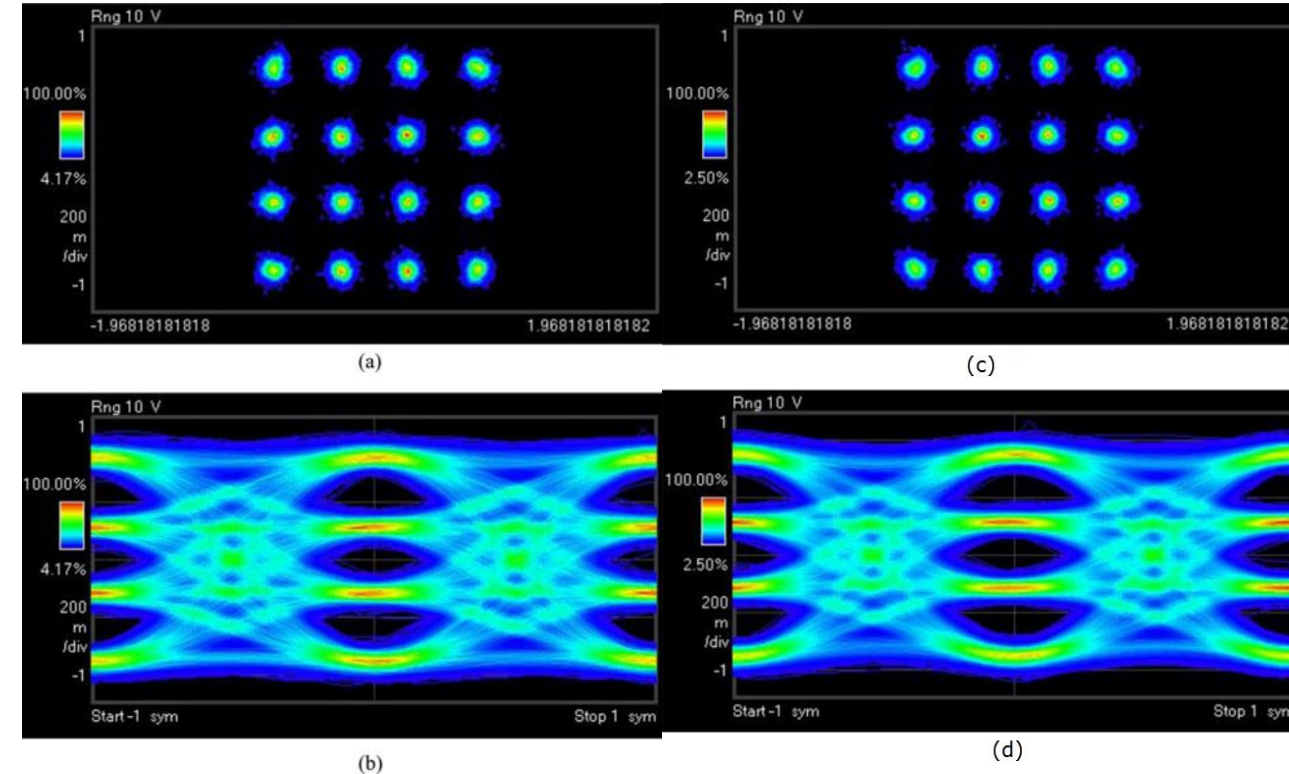


What's next? Ethernet poised to exploit DWDM

- Last week at the IEEE 802.3 Standards meeting in Chicago, the B10K study group passed a motion adding a DWDM objective
- “Provide a physical layer specification supporting 100 Gb/s operation on a single wavelength capable of at least 80 km over a DWDM system.” – motion language
- The modulation is not yet specified but is expected to be at least as complex as PAM4...

25 Gsymbols/s

28 Gsymbols/s



QAM-16 at 25/28 Gsymbols/s: (a/c) constellation, and (b/d) eye diagram of the Q component of the electric field – see figures 7 & 8 of reference

ETHERNET ALLIANCE ROADMAP UPDATE

Matt Traverso – Cisco



Thank You and Questions!!

If you have any questions or comments, please email
admin@ethernetalliance.org

Ethernet Alliance: visit www.ethernetalliance.org

 Join the Ethernet Alliance [LinkedIn group](#)

 Follow @EthernetAllianc on Twitter

 Visit the Ethernet Alliance on [Facebook](#)