BROCADE

"Here is more rough stuff on the ALTO ALOHA network." Memo sent by Bob Metcalfe on May 22, 1973.

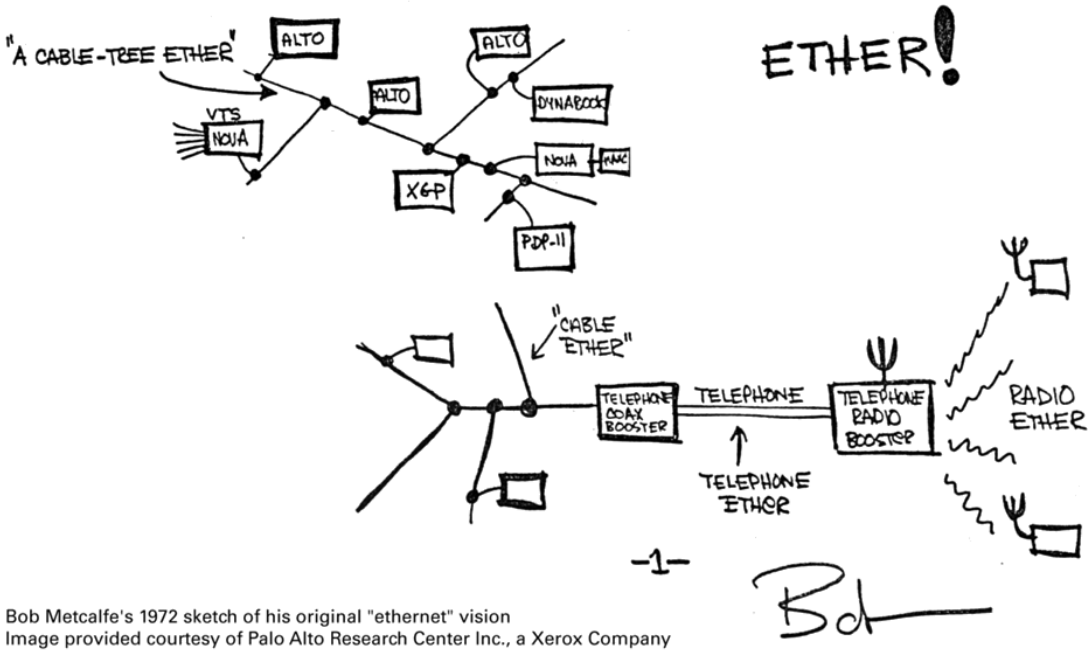

Bob Metcalfe's 1972 sketch of his original "ethernet" vision
Image provided courtesy of Palo Alto Research Center Inc., a Xerox Company

100 GBE - WHAT'S NEW AND WHAT'S NEXT

Greg Hankins <ghankins@brocade.com>

UKNOF25

40 Gigabit and 100 Gigabit Ethernet Developments

- IEEE 802.3ba standard approved June 17, 2010
  - 340 pages added to IEEE Std 802.3-2012
- Shipping 1st generation media, test equipment, router interfaces, and optical transport gear in 2010/2011
- Mature, interoperable technology with broad vendor support
- 2nd generation technology projects for both 40 and 100 GbE have started
  - Expected on the market in 2013 – 2015+

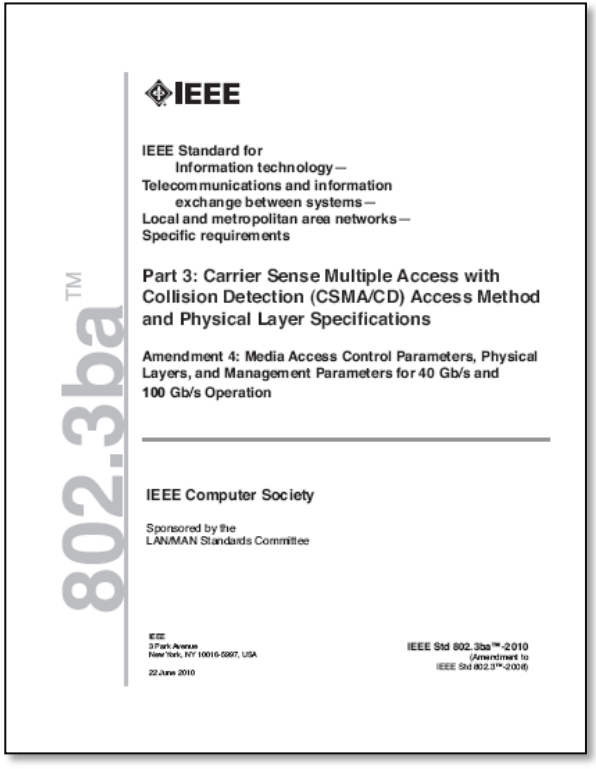

100 GbE Technology Adoption Timelines

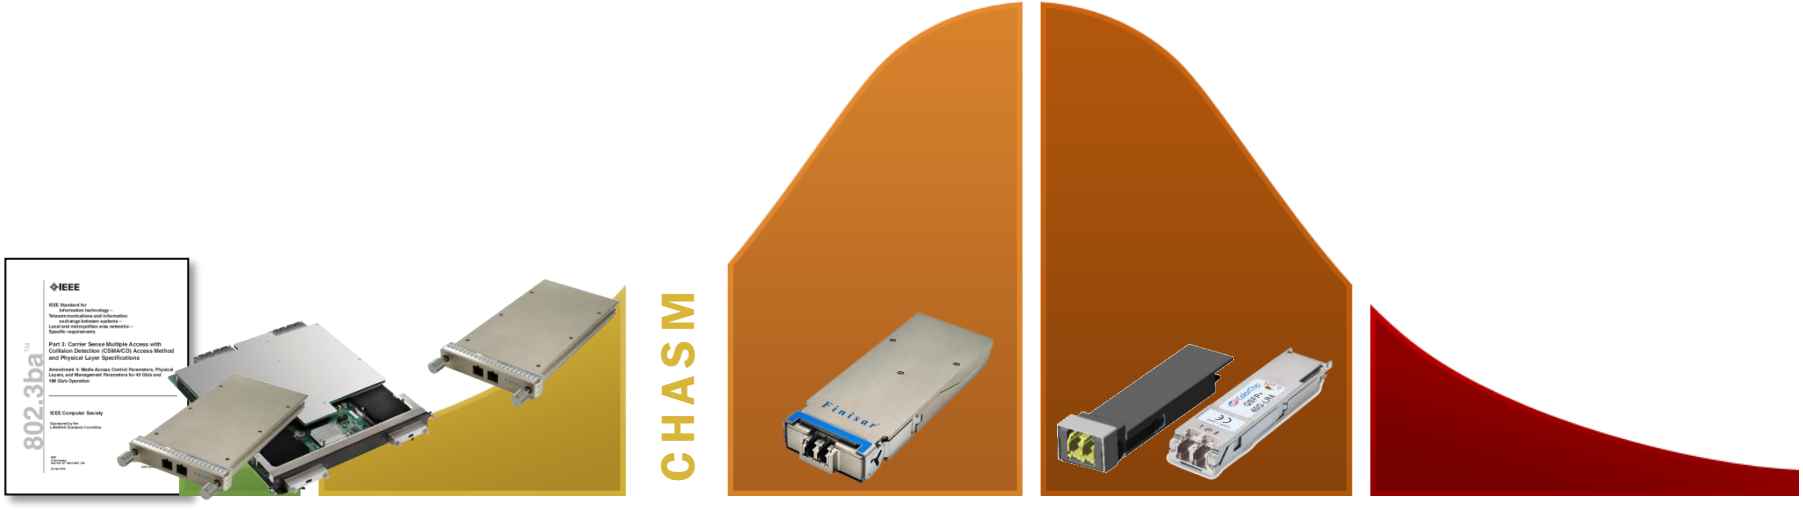

Innovators | **Early Adopters** | **Early Majority** | **Late Majority** | **Laggards**

2010/2011
1 - 2 Premium Ports/Slot
1st Generation (CFP)

2012
2 - 4 Premium Ports/Slot
Lower-Cost 1st Generation (CFP)

2013/2014
4 – 8 Lower-Cost Ports/Slot
2nd Generation (CFP2)

2015+
8 – 16 Commodity Ports/Slot
2nd Generation (CFP4, QSFP28)

2022+
24 – 48 Commodity Ports/Slot
3rd Generation (Serial 100 Gbps)

CFP images courtesy of Finisar.

100 GbE CFP Modules

C (100) Form-factor Pluggable

- New module optimized for 100 GbE long reach applications
- Used for 40GBASE-SR4, 40GBASE-LR4, 40GBASE-FR, 100GBASE-SR10, 100GBASE-LR4, 100GBASE-ER4, and 10x10 MSA
- Complex electrical and optical components need a large module
- Large module form factor and power consumption limits front panel density (larger than an iPhone 4)

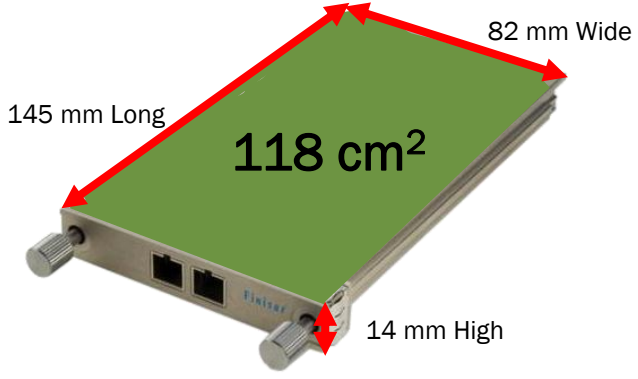

CFP image courtesy of Finisar.

- IEEE 802.3ba standards have a gap
  - 100GBASE-SR10 supports up to 150 m on OM4 MMF
  - 100GBASE-LR4 supports up to 10 km on SMF
  - Missing a shorter SMF reach
- 100GBASE-LR4 100 GbE optics are complex and expensive
- 10x10 MSA bridges the gap
  - Support for 2 km, 10 km and 40 km on SMF
  - Considerably more economical
  - Eliminates expensive components
  - Consumes lesser power
- Network operator members!

Members

BROCADE, EXFO, KAIAM, Google, HUAWEI, INNOLIGHT, JDSU, KOTURA, Verdant Networks, Santur, facebook, VITESSE, Enablence, EFFDON, MAXIM, amsix, Viscore, AFOP, bti systems, MRV, CORTINA, OPLINK, HitachiCable, source photonics

www.10x10msa.org

100 GbE Optics Overview

Physical Layer Reach	100 m OM3/ 150 m OM4 MMF	2 km SMF	10 km SMF		40 km SMF	
CFP Module	 100GBASE-SR10	 10x10-2km	 10x10-10km	 100GBASE-LR4	 10x10-40km	 100GBASE-ER4
Media Type	 Parallel MMF (MPO24)	 Duplex SMF (LC)	 Duplex SMF (LC)	 Duplex SMF (SC)	 Duplex SMF (LC)	 Duplex SMF (LC)
Standard	June 2010 IEEE 802.3ba	March 2011 10x10 MSA	August 2011 10x10 MSA	June 2010 IEEE 802.3ba	August 2011 10x10 MSA	June 2010 IEEE 802.3ba
Electrical Signaling (Gbps)	10 x 10	10 x 10	10 x 10	10 x 10	10 x 10	10 x 10
Optical Signaling (Gbps)	10 x 10	10 x 10	10 x 10	4 x 25	10 x 10	4 x 25
Power (W)	12	19	19	24	TBD	30
Availability	January 2013	June 2011	June 2011	June 2011	TBD	January 2013
Relative List Price	£	~~£££~~	~~££££~~	~~£££££~~	TBD	~~£££££££££££~~

Current State of the Industry

- Fundamental 1st generation technology constraints limit higher 100 GbE density and lower cost
- Electrical signaling to the CFP
  - 100 Gbps Attachment Unit Interface (CAUI) uses 10 x 10 Gbps lanes
- Optical signaling on the media
  - 10x10 MSA: 10 x 10 Gbps wavelengths
  - 100GBASE-LR4 and 100GBASE-ER4: 4 x 25 Gbps wavelengths
- CFP module size and power consumption

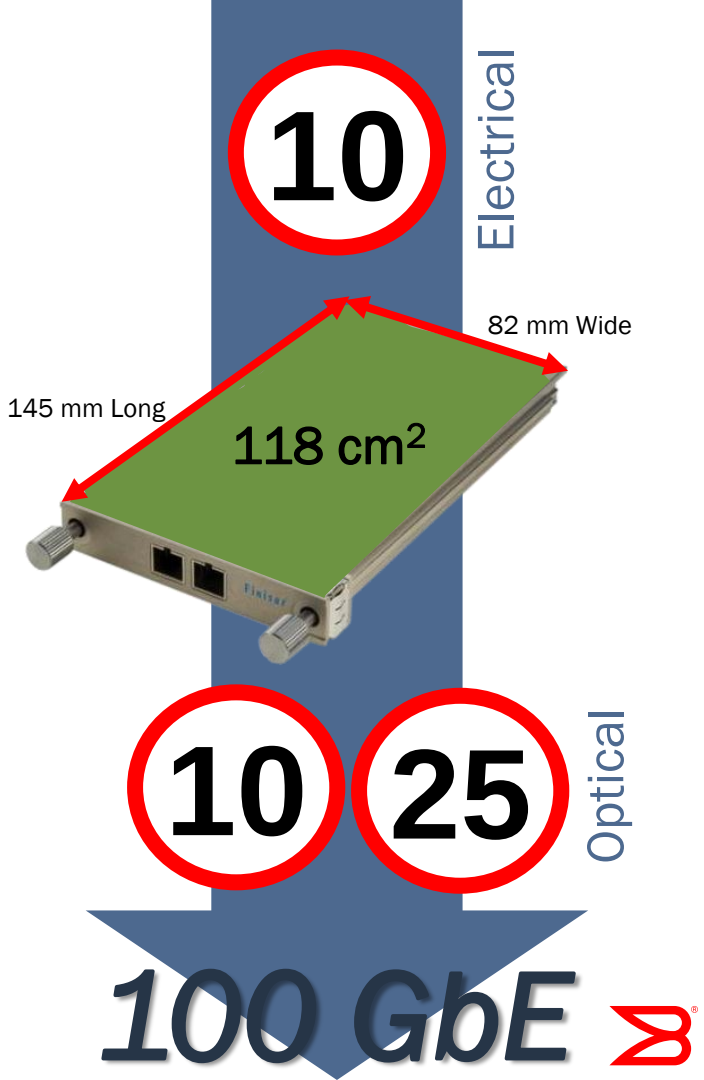

CFP image courtesy of Finisar.

1st Generation vs 2nd Generation 100 GbE

2nd Generation 100 GbE Needs Faster Electrical Signaling

1st Generation: 10 x 10 Gbps Electrical and 4 x 25 Gbps Optical

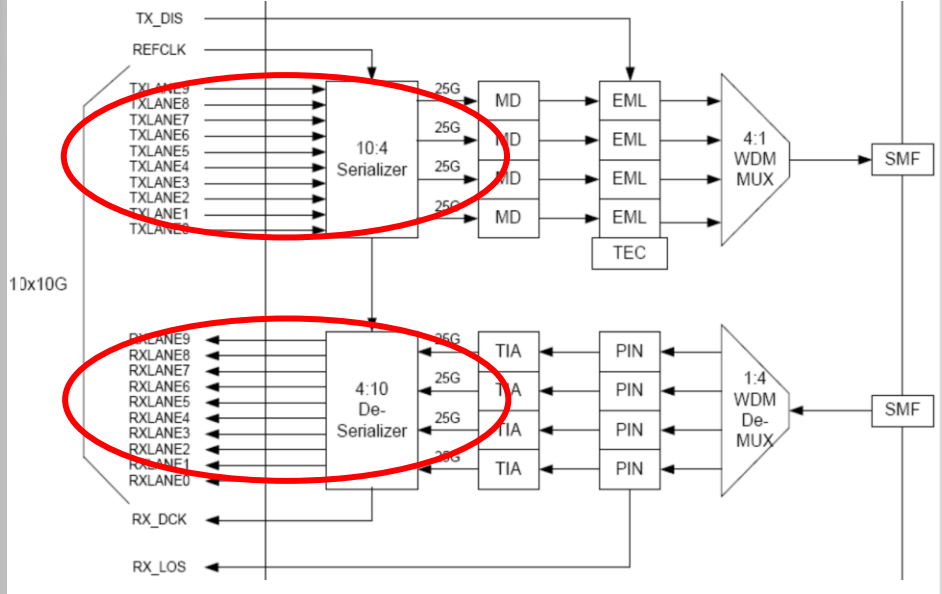

10 Gbps Electrical Signaling and 10:4 Gearbox Adds Complexity, Cost, Space, Power…

2nd Generation: 4 x 25 Gbps

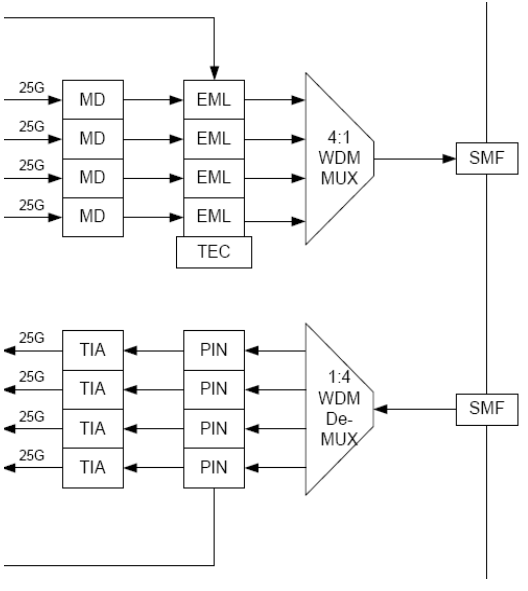

25 Gbps Electrical and Optical Signaling

Source: http://grouper.ieee.org/groups/802/3/ba/public/jul08/cole_03_0708.pdf

100 GbE Module Evolution

Each Module Increases Density, While Reducing Cost and Power

Generation	1st Generation		2nd Generation		
Year Available	2010	2010	2013	2015	2014
Electrical Signaling	CAUI 10 x 10 Gbps	CPPI 10 x 10 Gbps	CAUI 10 x 10 Gbps CAUI-4 4 x 25 Gbps	CAUI-4 4 x 25 Gbps	CAUI-4 4 x 25 Gbps
Front Panel Capacity	4 Ports	12 Ports	8 –10 Ports	16 Ports	22 Ports
Industry Standard Path	 CFP (82 mm Wide)	 CXP MMF (27 mm Wide)	 CFP2 (41.5 mm Wide)	 CFP4 (21.5 mm Wide)	 QSFP28 MMF (18.35 mm)
Cisco Proprietary Path			 CPAK (35 mm Wide)	?	?

CFP Module Evolution

Graphical View of Module Form Factors

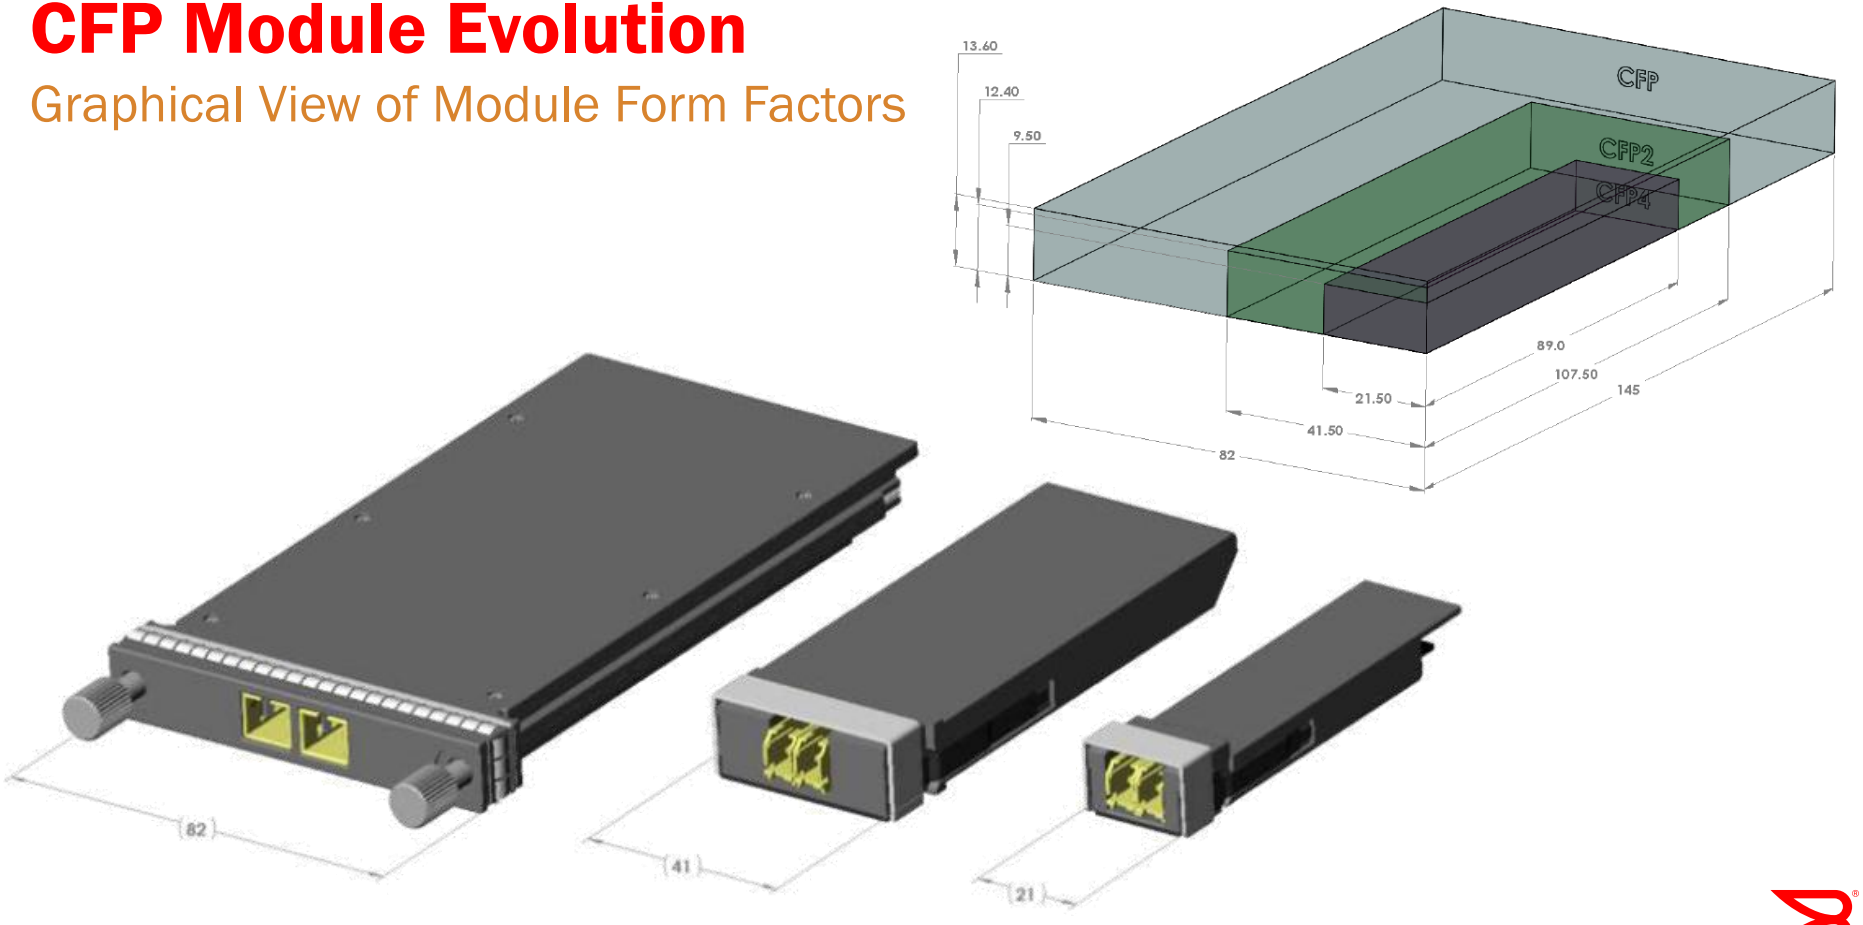

CFP2 Media Module

- External universal gearbox converts electrical signaling so that all current IEEE and MSA standards can be supported in the CFP2 module
- Specification is technically stable and manufacturing started in August 2012
- CFP2 supports 10 electrical lanes that can run at multiple speeds
  - 10 x 10 Gbps lanes
  - 8 x 25 Gbps lanes
- Smaller size and lower cost, complexity and power consumption than the CFP

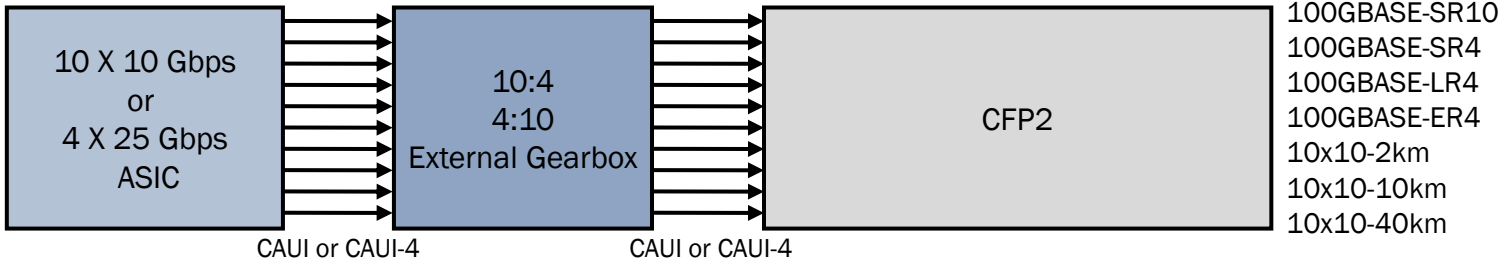

100GBASE-SR10
100GBASE-SR4
100GBASE-LR4
100GBASE-ER4
10x10-2km
10x10-10km
10x10-40km

Recent 100 GbE Developments

- 2nd generation projects based on 4 x 25 Gbps electrical signaling have started and are expected to be available in 2014+
- IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force started in September, 2011
  - 100GBASE-KR4: 4 x 25 Gbps NRZ 35 dB at 12.9 GHz backplane
  - 100GBASE-KP4: 4 x 25 Gbps PAM-4 33 dB at 7 GHz backplane
  - 100GBASE-CR4: 4 x 25 Gbps over 5 m copper twinax cable
  - Task Force web page: http://www.ieee802.org/3/bj/
  - Standard expected in March, 2014

Recent 100 GbE Developments

- IEEE P802.3bm 40 Gb/s and 100 Gb/s Operation Over Fiber Optic Cables Task Force started in August, 2012
  - 40GBASE-ER4: 4 x 10 Gbps over 40 km SMF
  - 100GBASE-nR4: 4 x 25 Gbps over 500 m parallel SMF
  - 100GBASE-UR4: 4 x 25 Gbps over 20 m OM3?/OM4 parallel MMF
  - 100GBASE-SR4: 4 x 25 Gbps over 100 m OM3?/OM4 parallel MMF
  - CAUI-4: electrical signaling to the CFP2, CFP4 and QSFP28
  - Task Force web page http://www.ieee802.org/3/bm/
  - Standard expected in March, 2015

100 GbE Technology Generations

1st Generation: 10 Gbps and 25 Gbps Signaling

- Developed in 2010
- More expensive, higher complexity and power consumption
- Lower port density
- supports 400 Gbps/slot

Interfaces:

- 100GBASE-SR10
- 10x10-2km
- 100GBASE-LR4
- 10x10-10km
- 100GBASE-ER4
- 10x10-40km

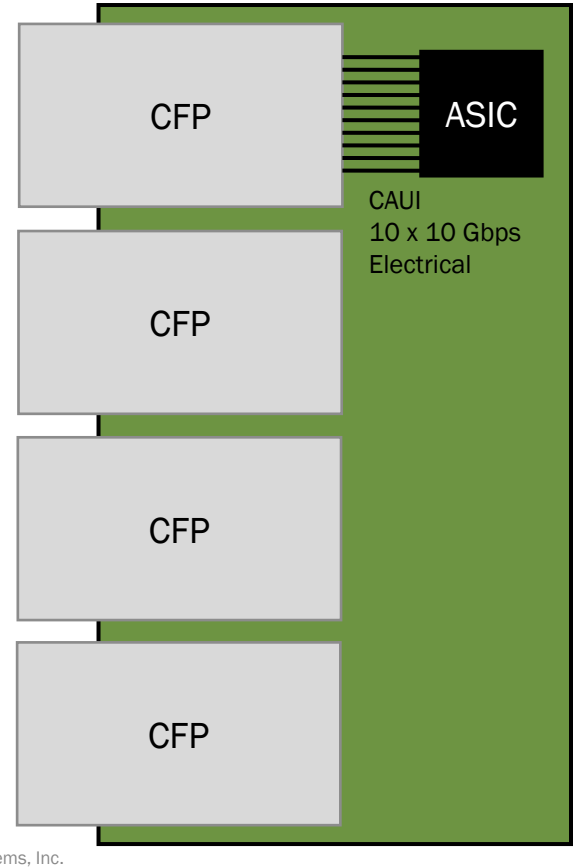

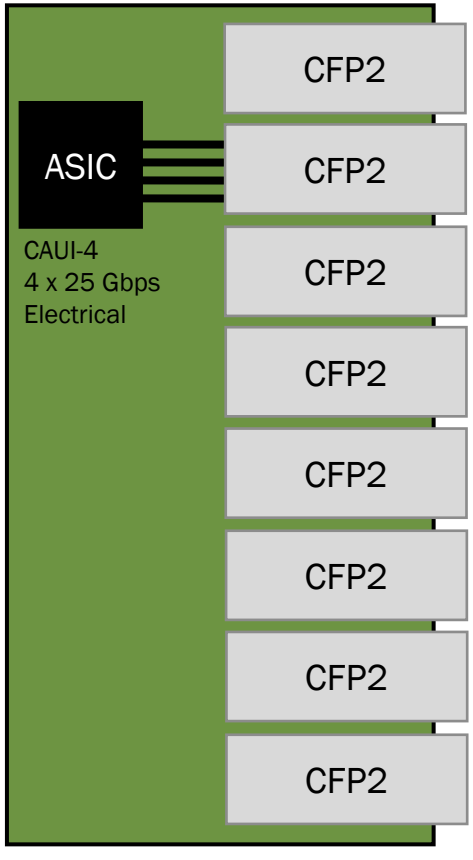

2nd Generation: 25 Gbps Signaling

- Expected in 2013 – 2015
- Less expensive, lower complexity and power consumption
- Higher port density supports 800 Gbps/slot

Supports all of the previous interfaces on the left plus:

- 100GBASE-nR4
- 100GBASE-UR4
- 100GBASE-SR4

100 GbE Technology Reference

1st Generation IEEE
1st Generation 10x10 MSA
2nd Generation IEEE

Physical Layer Reach	Back-plane	5 m Copper Cable	7 m Copper Cable	20 m OM3?/ OM4 MMF	100 m OM3?/ OM4 MMF	100 m OM3, 150 m OM4 MMF	500 m SMF	2 km SMF	10 km SMF		40 km SMF	
Name	100GBASE -KP4 and 100GBASE -KR4	100GBASE -CR4	100GBASE -CR10	100GBASE -UR4	100GBASE -SR4	100GBASE -SR10	100GBASE -nR4	10x10-2km	10x10-10km	100GBASE -LR4	10x10-40km	100GBASE -ER4
Standard Status	2014 IEEE 802.3bj	2014 IEEE 802.3bj	June 2010 IEEE 802.3ba	2015 IEEE 802.3bm	2015 IEEE 802.3bm	June 2010 IEEE 802.3ba	2015 IEEE 802.3bm	March 2011 10x10 MSA	August 2011 10x10 MSA	June 2010 IEEE 802.3ba	August 2011 10x10 MSA	June 2010 IEEE 802.3ba
Electrical Signaling (Gbps)	4 x 25 PAM-4 and NRZ	4 x 25	10 x 10	4 x 25	4 x 25	10 x 10	4 x 25	10 x 10	10 x 10	10 x 10	10 x 10	10 x 10
Media Signaling (Gbps)	4 x 25	4 x 25	10 x 10	4 x 25	4 x 25	10 x 10	4 x 25	10 x 10	10 x 10	4 x 25	10 x 10	4 x 25
Media Type	Backplane	Twinax	Twinax	Parallel MMF	Parallel MMF	Parallel MMF	Parallel SMF	Duplex SMF	Duplex SMF	Duplex SMF	Duplex SMF	Duplex SMF
Media Module	Backplane	CFP2, CFP4, CPAK, QSFP28	CFP, CXP	CFP2, CFP4, CPAK, QSFP28	CFP2, CFP4, CPAK, QSFP28	CFP, CFP2, CPAK, CXP	CFP2, CFP4, CPAK, QSFP28	CFP, CFP2	CFP, CFP2	CFP, CFP2, CFP4, CPAK	CFP, CFP2?	CFP, CFP2?
Availability	2014+	2014+	2010	2015+	2015+	2012	2015+	2011	2011	2010	2012	2012

Industry Challenges for 100 GbE and Beyond

All Solutions are Good, Fast, or Cheap – Pick Any Two

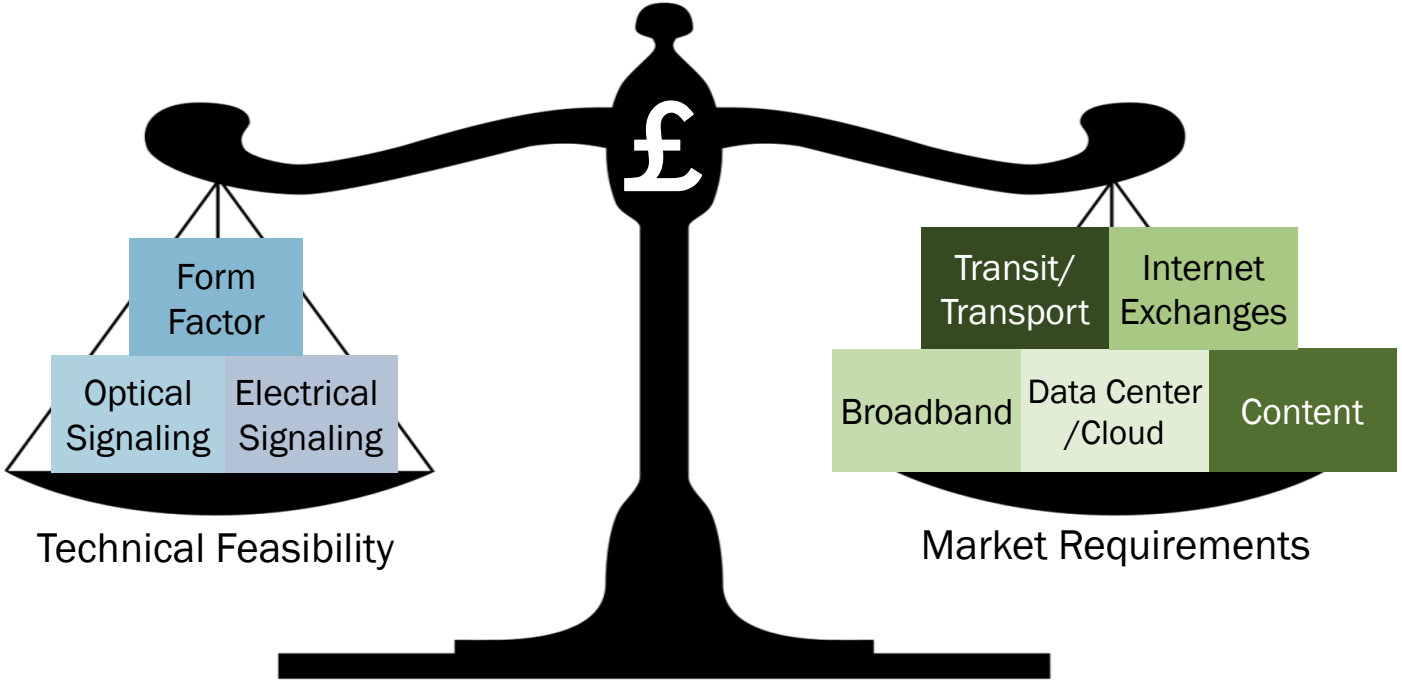

Economics Dictate the Solution

IEEE Provides an Open Industry Forum to Make Decisions

Bandwidth Requirements Projection

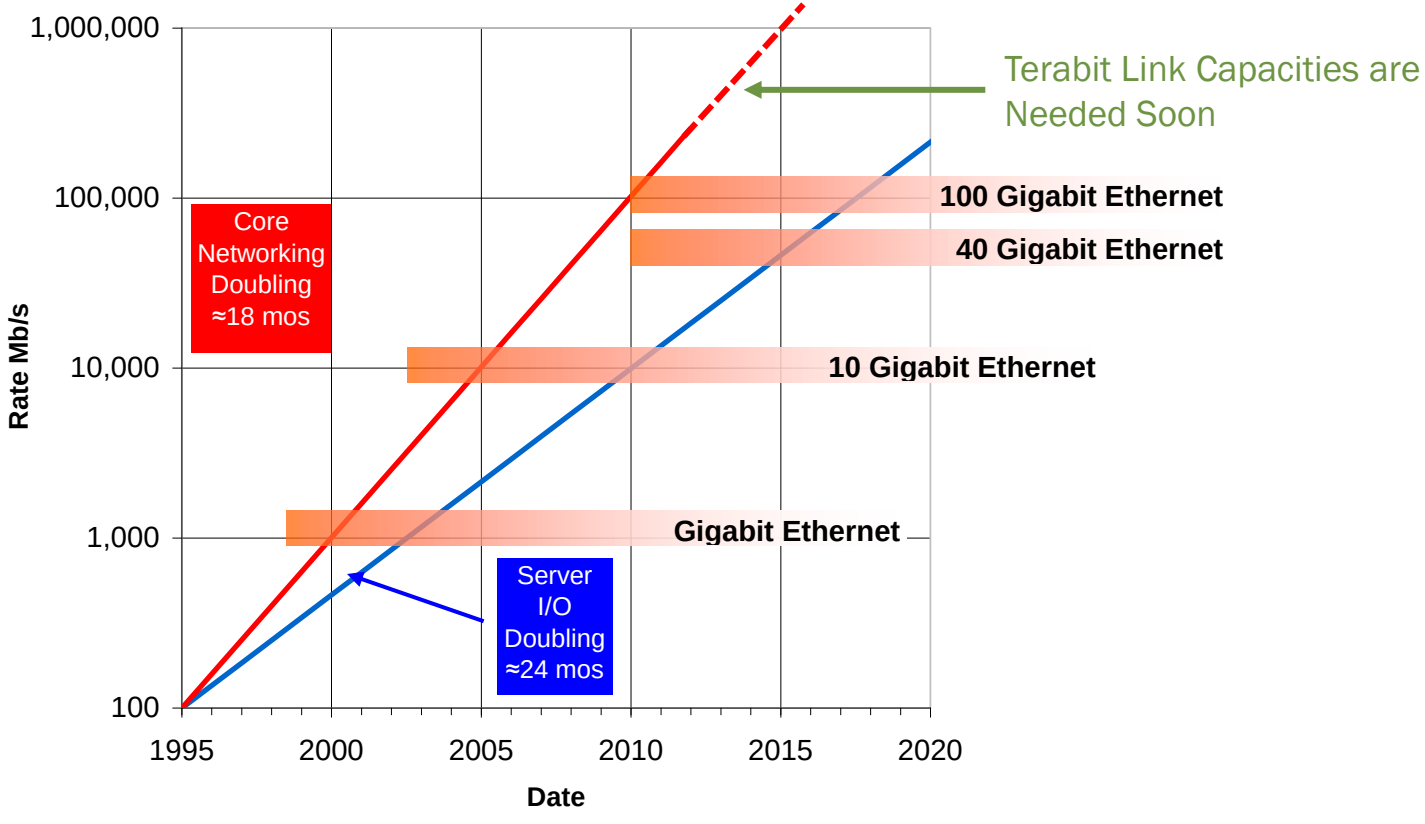

Diagram source: http://www.ieee802.org/3/hssg/public/nov07/HSSG_Tutorial_1107.zip

Why 400 GbE? Why not TbE?

- Given that TbE is technically and economically impractical to develop until 2020+ we had to make a choice
  - Wait… for >10 years between Ethernet speed increases (100 GbE June 2010)
  - Start a feasible higher speed Ethernet standard now that can be ready by 2016 when the market needs something faster
- IEEE 802.3 Ethernet Bandwidth Assessment (BWA) Ad Hoc and the IEEE 802.3 Higher Speed Ethernet (HSE) Consensus Ad Hoc spent a lot of time analyzing market demand and technical options
- High degree of consensus in the IEEE that 400 GbE should be the next Ethernet speed
  - CFI straw poll: yes 132, no 0, abstain 1
  - Motion for Study Group: yes 87, no 0, abstain 4
- This will make the standardization process faster compared to 802.3ba
  - IEEE 802.3 Higher Speed Study Group first had to analyze market demand and technical options
  - Eventually decided on 40 GbE and 100 GbE

Theoretical M (1000) Form-factor Pluggable

Using Today's Technology is Impractical

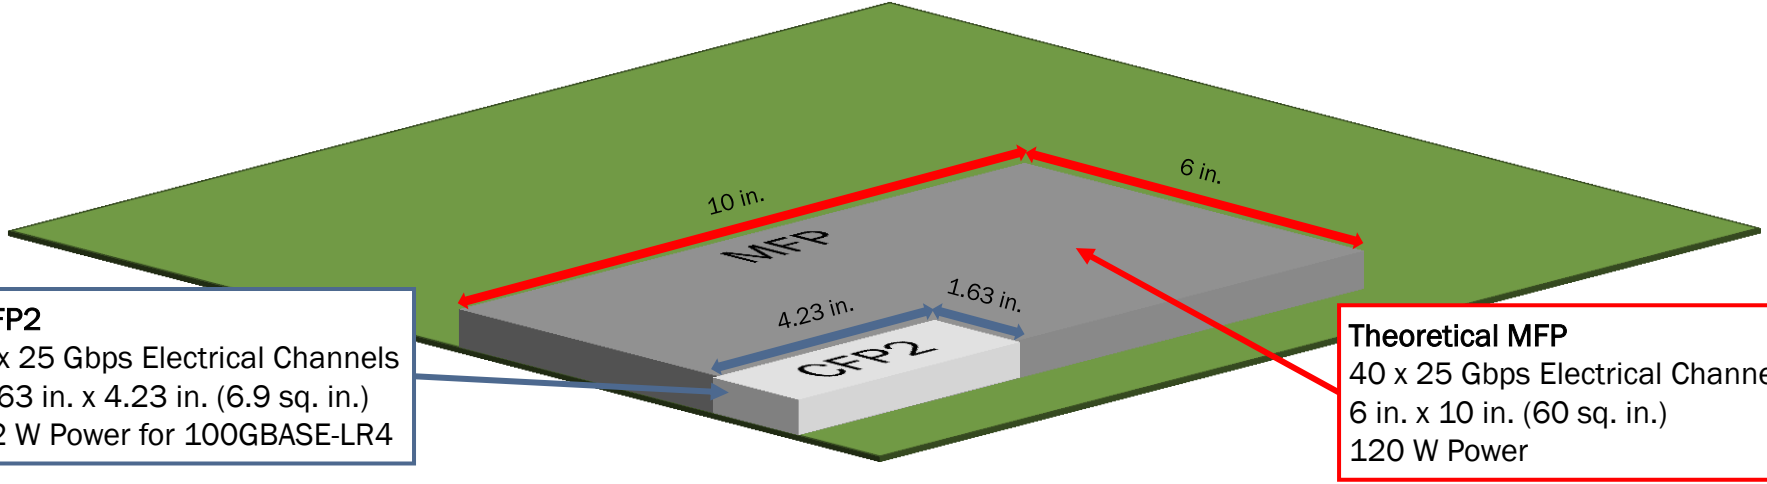

CFP2
4 x 25 Gbps Electrical Channels
1.63 in. x 4.23 in. (6.9 sq. in.)
12 W Power for 100GBASE-LR4

Theoretical MFP
40 x 25 Gbps Electrical Channels
6 in. x 10 in. (60 sq. in.)
120 W Power

- Using current technology 25 Gbps electrical signaling, a TbE module (MFP) would use 40 channels
- The size of the module can be estimated at ½ sq. in. per W of power consumed
- If a CFP2 module consumes 12 W, then the MFP could consume 120 W
- This would require 60 sq. in. to cool and could make the module 6 in. by 10 in., taking up over 25% of the board space

IEEE 802.3 400 Gb/s Ethernet Study Group

- 400 GbE Call-For-Interest (CFI) was presented at the March, 2013 IEEE Plenary
- Approved to be an official IEEE Study Group on March 22, 2013
  - 400 GbE Study Group will define the objectives for reach and media
  - 400 GbE standard expected in 2016
- Study Group web page http://www.ieee802.org/3/400GSG/
- CDFP (CD (400) Form-factor Pluggable) Multi-Sourcing Agreement (MSA) announced to develop specifications for high-density 400 GbE modules
- First interfaces expected to be available in 2016+

400 GbE Module Evolution Estimates

Each Module Increases Density, While Reducing Cost and Power

	1st Generation		2nd Generation	3rd Generation
Year	2016	2016	2019+	2022+
Electrical Signaling	CDAUI-16 16 x 25 Gbps	CDAUI-16 16 x 25 Gbps	CDAUI-8 8 x 50 Gbps	CDAUI-4 4 x 100 Gbps
Media Module	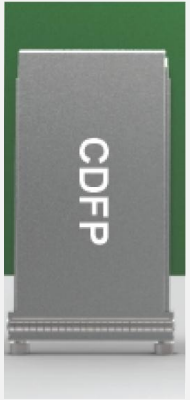	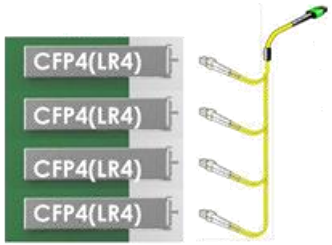		
	CDFP	4 x CFP4	CDFP2	CDFP4

Images courtesy of TE Connectivity.

400 GbE CDFP Module Overview

- The CDFP MSA will define a specification for a module and the supporting hardware for a high-density 400 Gigabit Ethernet (400 GbE) based on 16 x 25 Gbps signaling
  - CD = 400 in Roman numerals, C = 100 and D = 500
  - FP = Form-factor Pluggable, follows the naming convention started in 1999 with the SFP – Small Form-factor Pluggable
  - MSA = Multi-Sourcing Agreement, a legal document that defines a group that will define a specification for how to make and license the technology

Electrical Interface
16 x 25 Gbps in Each Direction

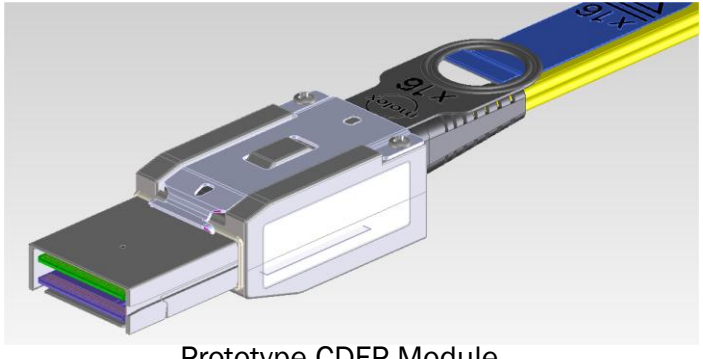

Optical Interface
With Two 16 Fiber Ribbons
(16 Transmit Fibers + 16 Receiver Fibers)

Prototype CDFP Module

Image courtesy of the CDFP MSA.

CDFP Switch Overview

- The CDFP module will be defined to support several modules in a 1 RU (19”) switch or modular switch blade to yield 4 Tbps of throughput

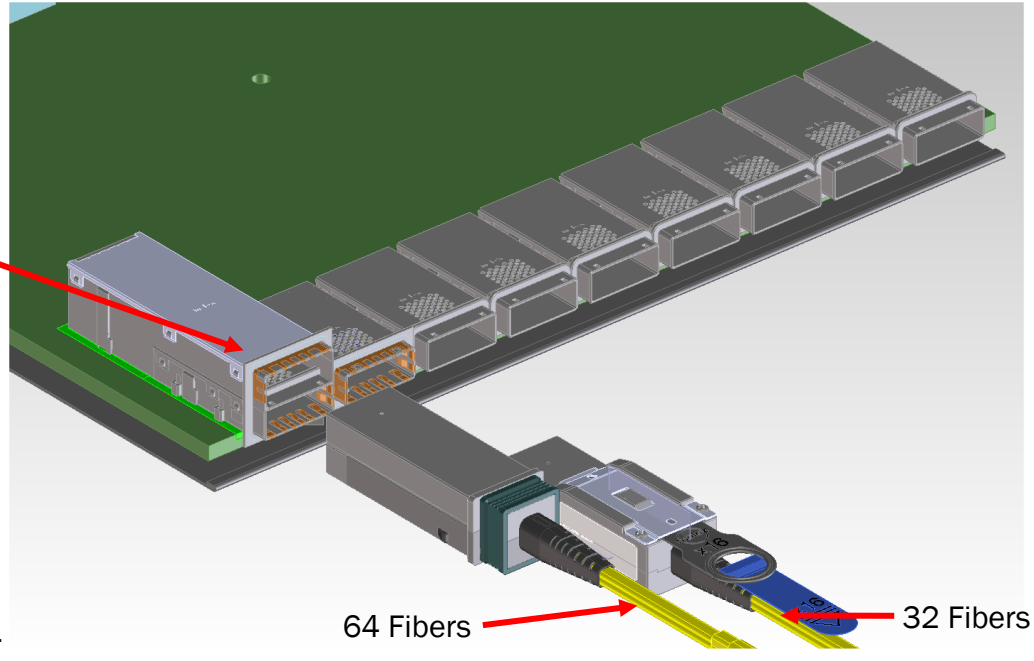

4 Tbps Slot Capacity vs.
400 Gbps Today
(4 x 100 GbE or 48 x 10 GbE)

Image courtesy of the CDFP MSA.

More Information

- IEEE 802.3 Industry Connections Ethernet Bandwidth Assessment Ad Hoc
  http://www.ieee802.org/3/ad_hoc/bwa/index.html
- IEEE 802.3 Industry Connections Higher Speed Ethernet Consensus Ad Hoc
  http://www.ieee802.org/3/ad_hoc/hse/public/index.html
- IEEE 802.3 Higher Speed Study Group
  http://www.ieee802.org/3/hssg/index.html
- IEEE 400 GbE Call-For-Interest presentation
  http://www.ieee802.org/3/cfi/0313_1/CFI_01_0313.pdf

Ethernet Industry Summary

- 10 GbE is being widely deployed to servers today and driving the need for 40 GbE and 100 GbE in data centers
- 40 GbE is mature QSFP+ technology and will be available in a smaller form-factor in the future
- 100 GbE is still a couple of generations away from serial signaling
  - 25 Gbps signaling is challenging
  - We'll get a better idea of what is possible as 25 Gbps technology matures
- 400 GbE is just starting and will leverage 100 GbE
- TbE is still science fiction from a low cost and low power perspective, but give it some time (2020+)

BROCADE

QUESTIONS?