THE ETHERNET ROADMAP

APPLICATIONS, SERVERS, OPTICS, AND SWITCHES

Scott Kipp

April 15, 2015

ethernet alliance

Agenda

- 4:40-:52 – The 2015 Ethernet Roadmap – Scott Kipp, Brocade
- 4:52-5:04 – Optical Ethernet Roadmap – Chris Cole, Finisar
- 5:04-5:16 – Copper Connectivity in the 2015 Ethernet Roadmap - David Chalupsky, Intel
- 5:16-5:28 – Implications of 50G SERDES on Ethernet Speeds - Kapil Shrikhande, Dell
- 5:28-5:40 – Q&A

Disclaimer

- Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance.

THE 2015 ETHERNET ROADMAP

Scott Kipp

April 15, 2015

ethernet alliance

2015 ETHERNET ROADMAP

As shown on the long and winding road, Ethernet could have 12 speeds before 2020 with 6 new speeds introduced in the next 5 years. The progression of speeds is not in chronological order because 40GbE and 100GbE were primarily based on multiple lanes of 10Gb/s technology that was available before 25Gb/s serial technology enabled 25GbE. Lanes running at 25Gb/s are becoming practical in 2015 and will be used in 25GbE SFP+ and 4x25Gb/s 100GbE QSFP28. The next serial lane speed is expected to be 50Gb/s and enable 50GbE SFP28, 200GbE QSFP28 (4X50G) and 400GbE CFP2 (8X50G).

Beyond 400GbE, the map shows the unknown distant future that will become clearer as we approach 2020. Terabit links are expected when single lanes can be modulated at 100Gb/s and grouped into 10 or 16 lanes to form TbE or 1.6TbE. Significant investments in technology are needed before 100Gb/s lanes are economically feasible.

Low cost 100Gb/s lane technology that can fit in an SFP+ is not expected to be available until after 2020. The Ethernet Alliance will award the first company that produces a 100GbE SFP+ with the Holy Grail of the 100GbE SFP+.

The twisted pair or BASE-T roadmap in the lower right corner of the map shows how 10GBASE-T technology is being used in 4 new speeds — 2.5, 5, 25 and 40Gb/s. All four of these speeds are expected to be standardized in 2016 but they are targeting different cabling infrastructure. 2.5 and 5GBASE-T are being designed for Cat 5e cabling up to 100 meters while 25 and 40GBASE-T are being designed for 30 meters of Cat 8 cabling.

10 Tb/s
1.6 Tb/s >2020
6.4 Tb/s >>2020
1 Tb/s >2020
800 Gb/s >2020
400 Gb/s 2017 (est)
200 Gb/s 2018-2020 (est)
100 Gb/s 2010
50 Gb/s 2018-2020 (est)
40 Gb/s 2010
25 Gb/s 2016 (est)
10 Gb/s 2002
5 Gb/s 2016 (est)
2.5 Gb/s 2016 (est)
1 Gb/s 1998
100 Mb/s 1995
10 Mb/s 1983

Ethernet Speed
Speed in Development
Possible Future Speed

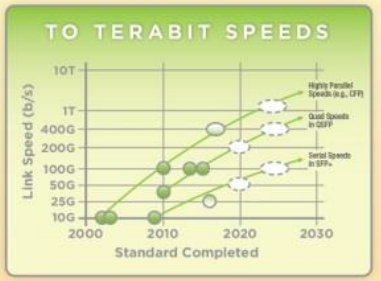

FIBER OPTIC LINKS

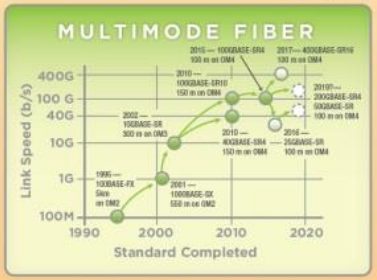

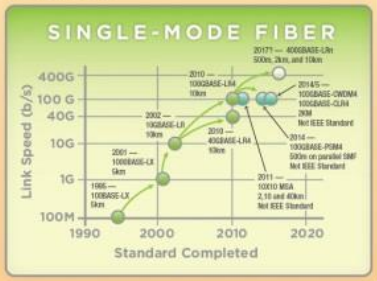

COPPER LINKS

To get a pdf version of the roadmap and to find out more about the roadmap, please go to: **www.ethernetalliance.org/roadmap/**

MEDIA AND MODULES

Ethernet is wired technology and supports a variety of media including backplanes, twisted pair, twinax, multimode fiber and single-mode fiber. Most people know Ethernet by the twisted pair or Cat "x" cabling with RJ45 connectors because close to a billion ports a year are sold. Cat 8 is the latest generation of twisted pair cabling that will be used in 25GBASE-T and 40GBASE-T.

Another popular copper interface is Twinax copper cables that are also known as direct attach cables (DAC)s. DACs may be passive or active and provide very low cost connectivity to servers. Passive DACs are limited to 25 meters or less while active optical cables can go hundreds of meters.

For links longer than 100 meters, fiber optics are required and the graphic below shows three of many module types. The SFP family is the most popular module and supports a single channel or lane in each direction and duplex fibers. The QSFP family supports 4 channels while the CFP2 supports up to 10 channels and duplex or parallel fibers. For 40GbE and beyond, the electrical interface to the module is being defined in IEEE and supports a variety of optical interfaces from IEEE and other sources.

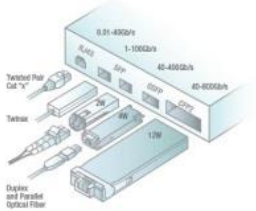

2015

ETHERNET ROADMAP

THE PAST, PRESENT AND FUTURE OF ETHERNET

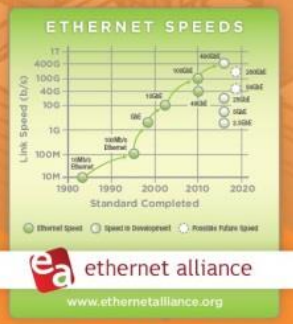

ethernet alliance
www.ethernetalliance.org

Designed by Scott Kipp

$9.95

The Near Term Speeds

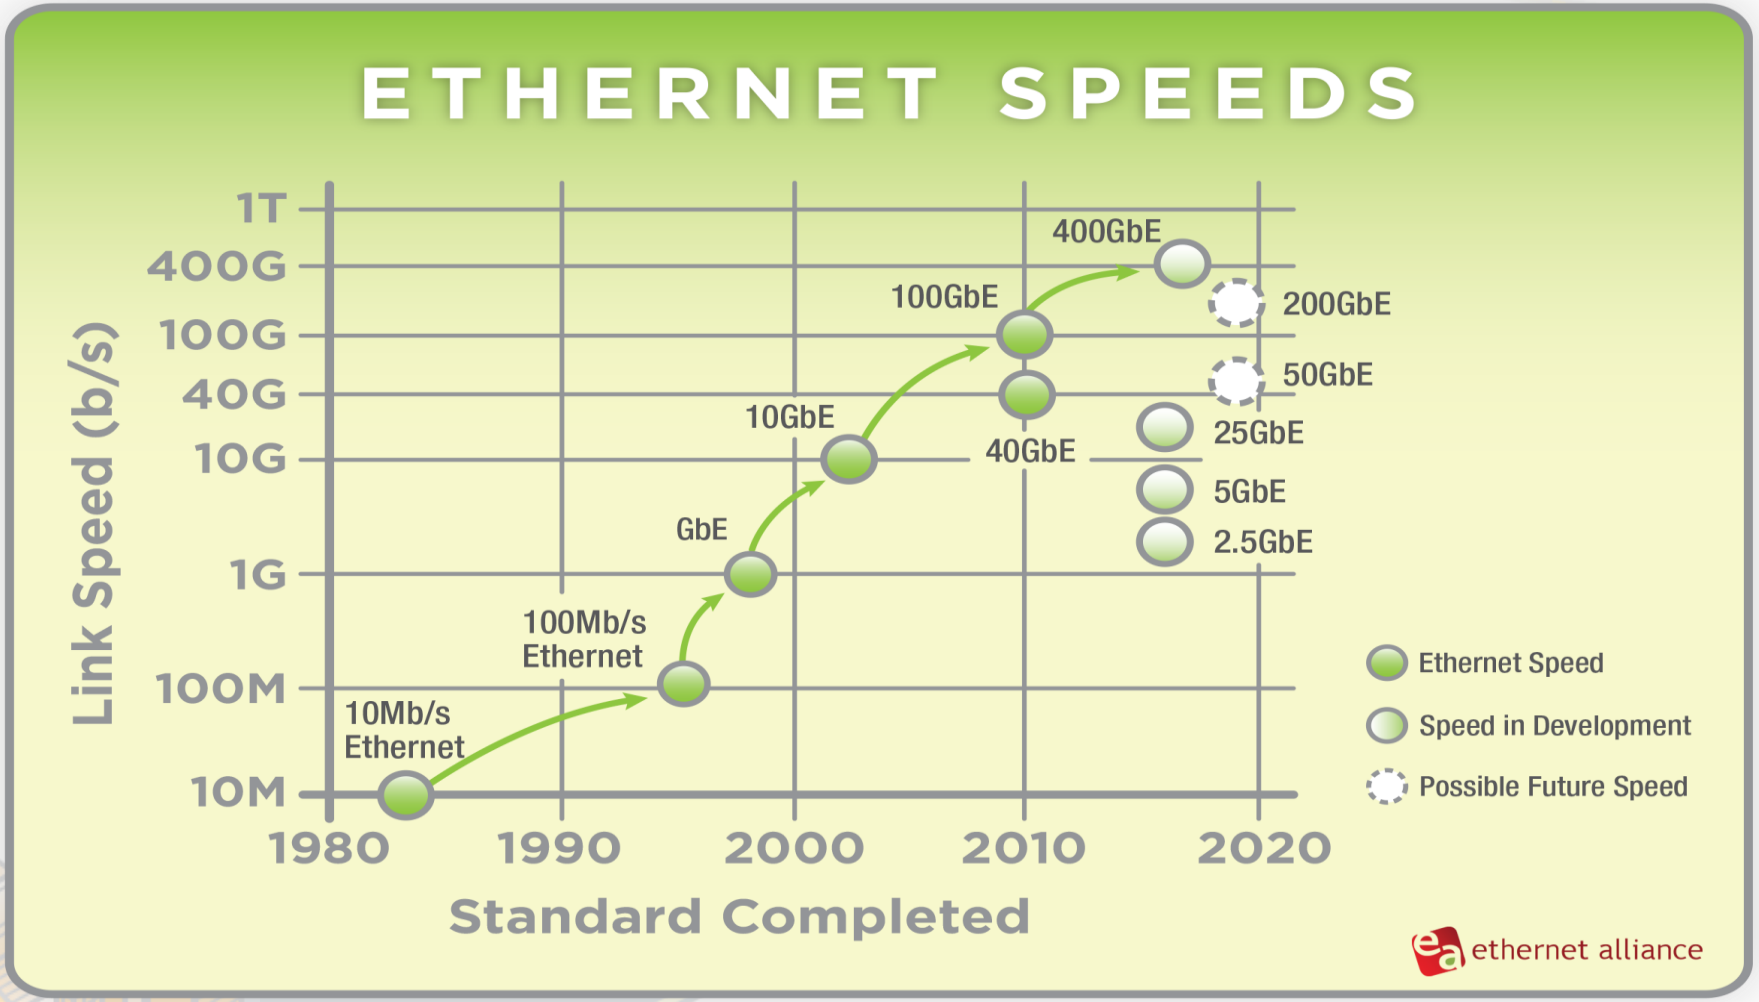

Media and Modules

- Most common port types through 2020

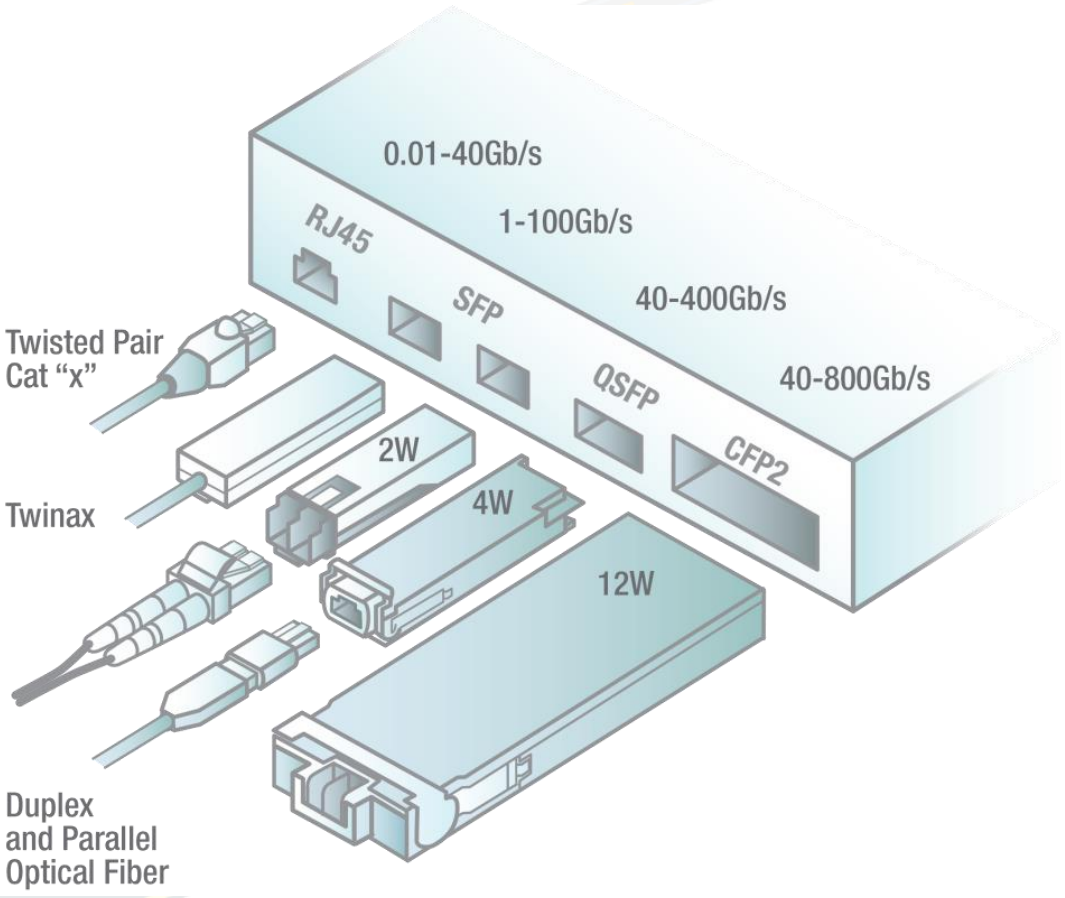

Switch Port Densities

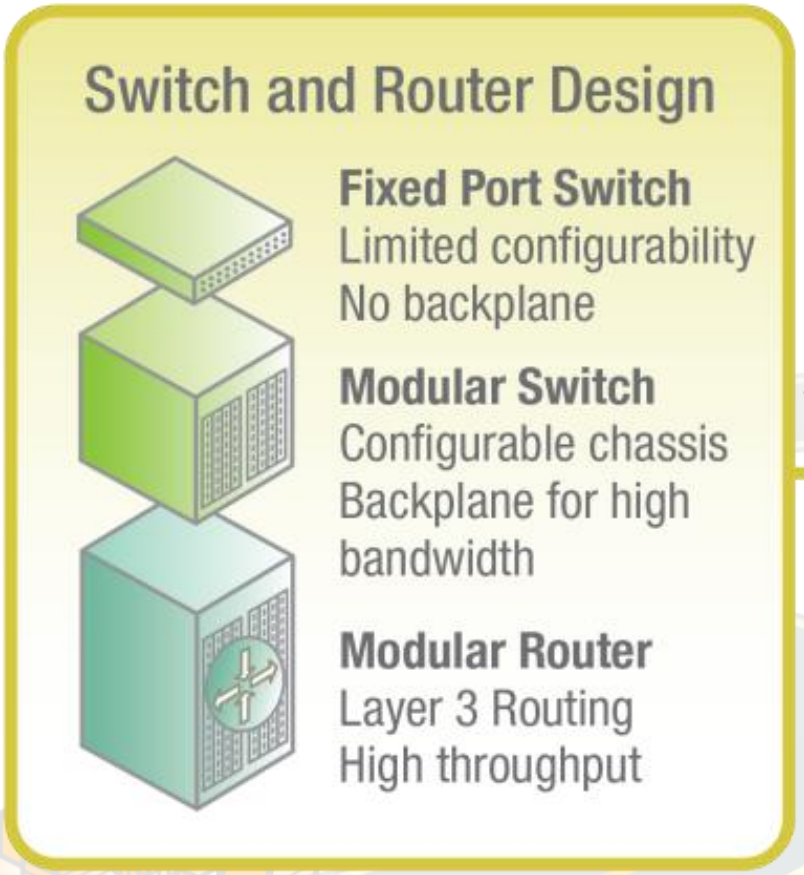

94% of Ports
435M

<5% of Ports
21M

~1% of Ports
5M

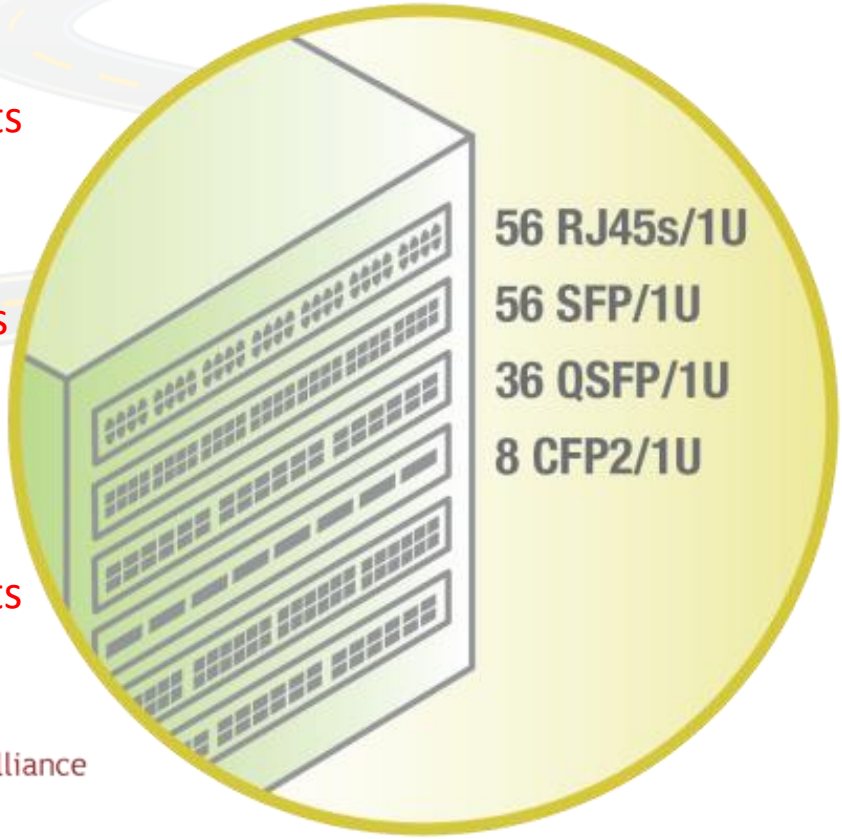

ethernet alliance

Source: Dell'Oro

Router Port Shipments

- According to Dell'Oro, what Router port speed shipped the most in 2014?

Port Speed	(000s)	Revenue Share
– T1/E1/T3/E3	1,570	4
– 100M	287	5
– 1GbE	4,460	1
– 10GbE	725	2
– 40GbE	1.1	6
– 100GbE	17	3

#1 through 2019

ETHERNET ECOSYSTEM

Represented as a city, the Ethernet Ecosystem is divided into four quadrants that are interconnected by multiple MANs that are typically not Ethernet. While each quadrant has overlapping technologies and requirements, this map organizes the environments with a broad brush. Specific implementations may vary considerably.

The top half of the map represents applications where cost and connectivity are driving concerns. In the home, small office and car, link distances are less than 100 meters and speeds are typically under 10Gb/s, so copper cabling and wireless are ideal. As enterprises scale in size and requirements, they shift towards fiber and 10Gb/s speeds and beyond.

The lower half of the map captures applications that consistently push the bounds of Ethernet and require higher speeds and massive scalability. For example, service providers and hyperscale data centers will be the early adopters of 400GbE. These users may deploy hundreds of thousands of servers in data centers that span multiple football fields and consume hundreds of megawatts of power.

ENTERPRISE AND CAMPUS

Enterprises consume more Ethernet ports than the other environments by connecting desktop computers, devices and Voice over IP (VoIP) phones. The wired Ethernet networks are supplemented with wireless access points (WAPs) that are connected to Ethernet cables. 802.11ac WAPs are driving the need for 2.5 and 5GBASE-T and eventually 10GBASE-T. Most enterprise data centers are less than 10,000 sq ft and use Cat "x" cabling to connect to servers.

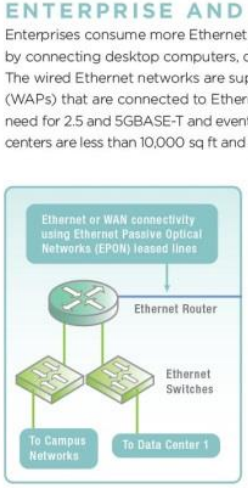

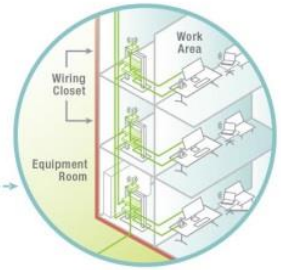

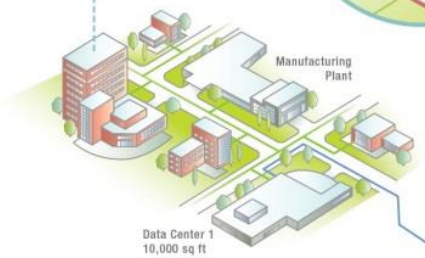

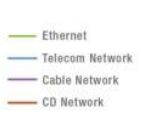

RESIDENTIAL AND CONSUMER

Ethernet Passive Optical Networking (EPON) delivers Internet service to millions of residential customers around the world. Regardless of how the Internet reaches the home, residents may wire their home with Ethernet or use wireless connectivity to connect devices. From cameras to cars, Ethernet provides the network to enable sharing resources and content.

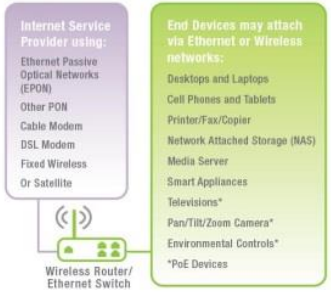

BACKBONE TO OTHER CITIES

BACKBONE TO OTHER CITIES

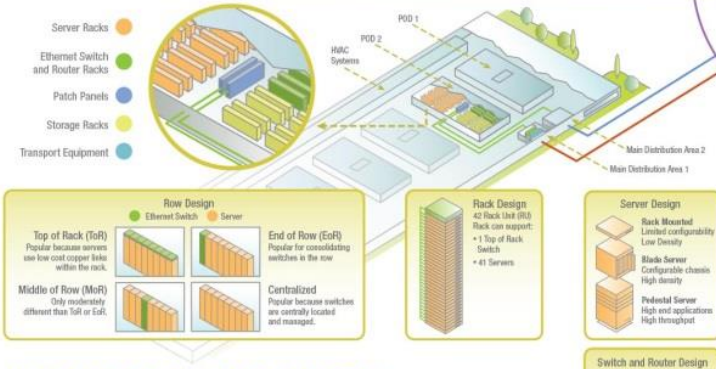

HYPERSCALE DATA CENTER

Hyperscale data centers, also known as warehouse scale computing and mega data centers, are known by their massive size and scalability. Cloud service providers, large enterprises and service providers pack over 100,000 servers that are often divided into several pods. Thousands of 25GbE servers and eventually 50GbE servers in these data centers drive the need for 400GbE to the MAN and WAN.

MANs

Metropolitan Area Networks (MANs) come in many varieties and deliver services to a variety of enterprises, organizations and consumers. Some MANs are based on Ethernet, but the largest MANs are based on Optical Transport Networks (OTN) technologies.

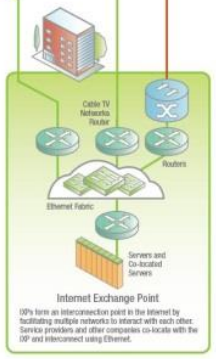

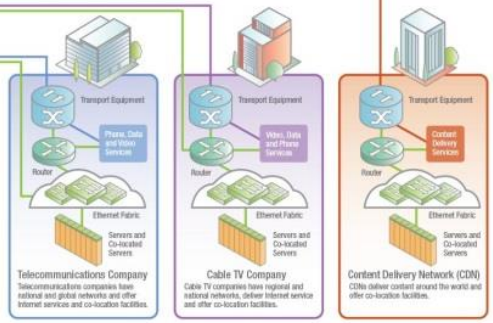

SERVICE PROVIDERS

Service providers deploy MANs and Wide Area Networks (WANs) to deliver a variety of services including Carrier Ethernet. Service providers may use Ethernet Passive Optical Networks (EPON) and cable companies send EPON Protocol over Coax (EPOC). Service providers use routers to interconnect various networks.

Hyperscale Data Centers

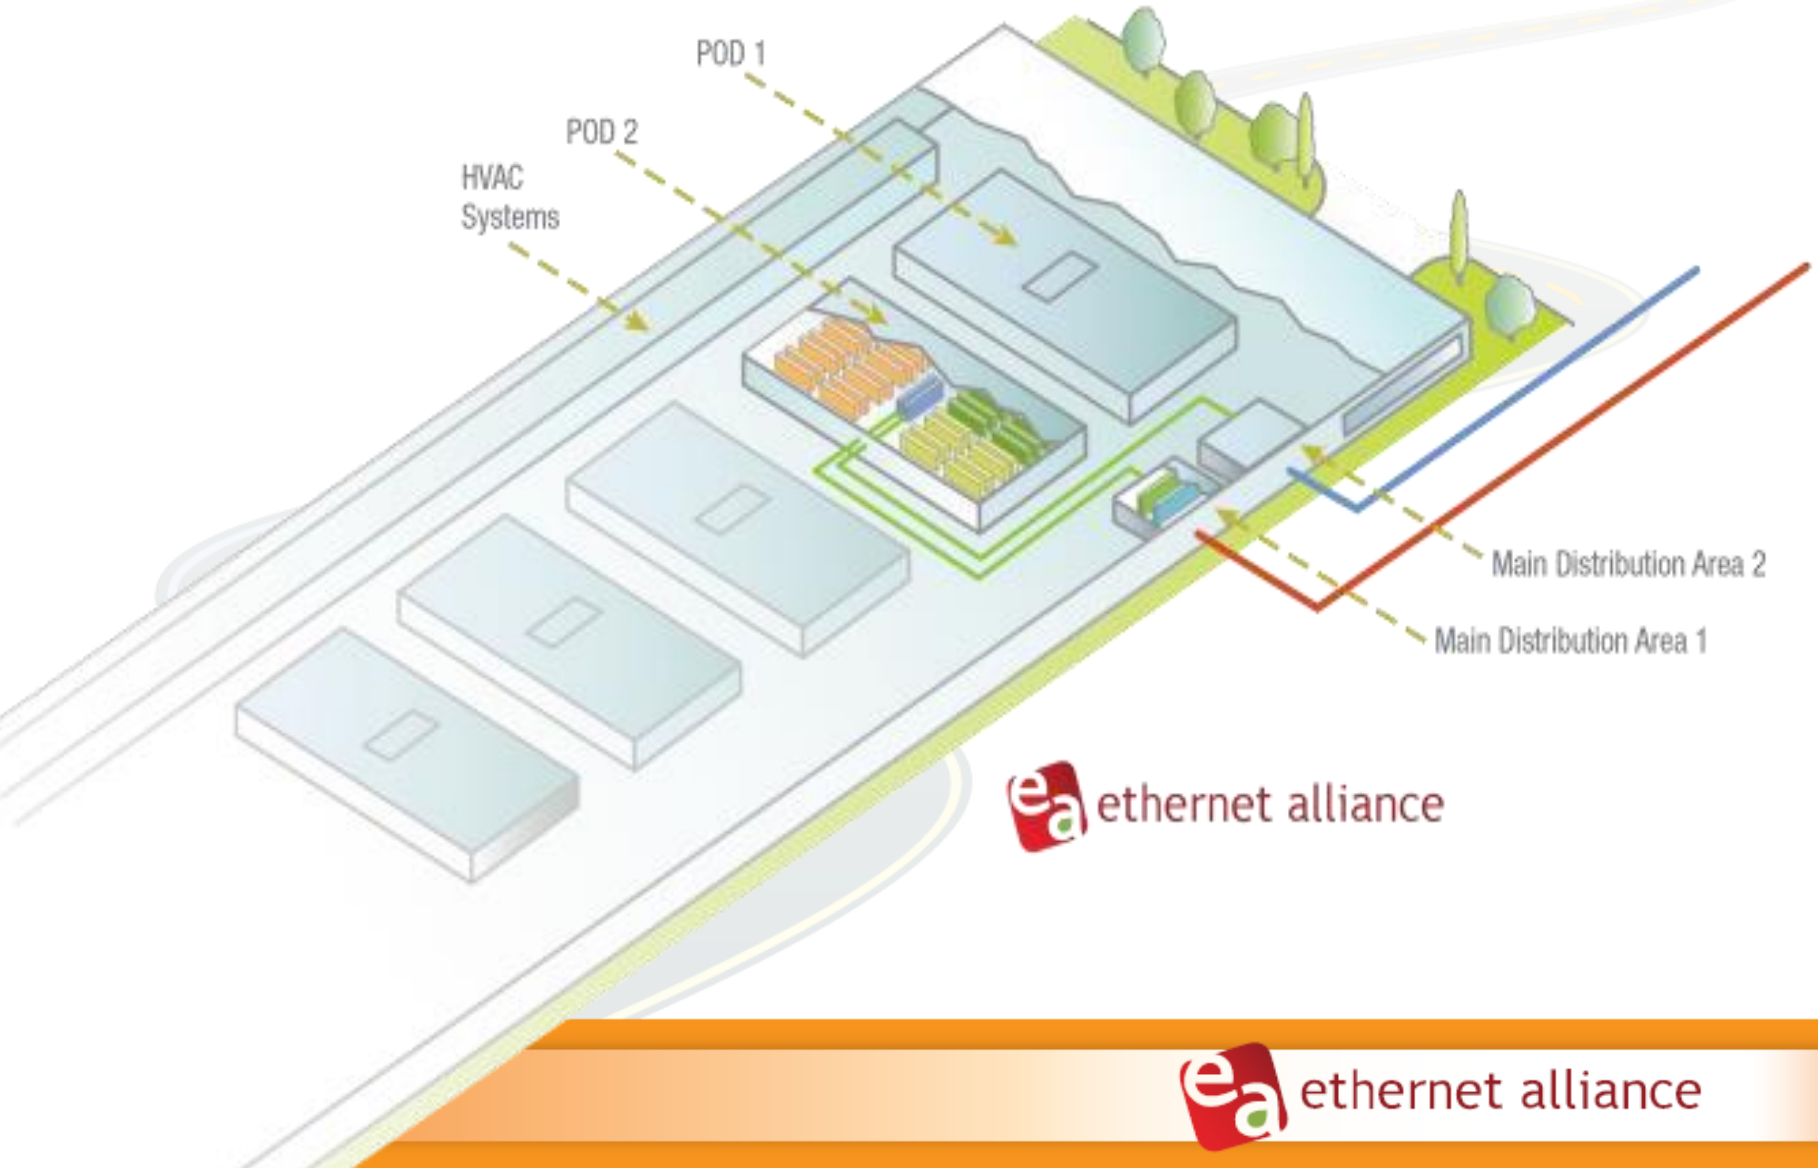

Service Providers

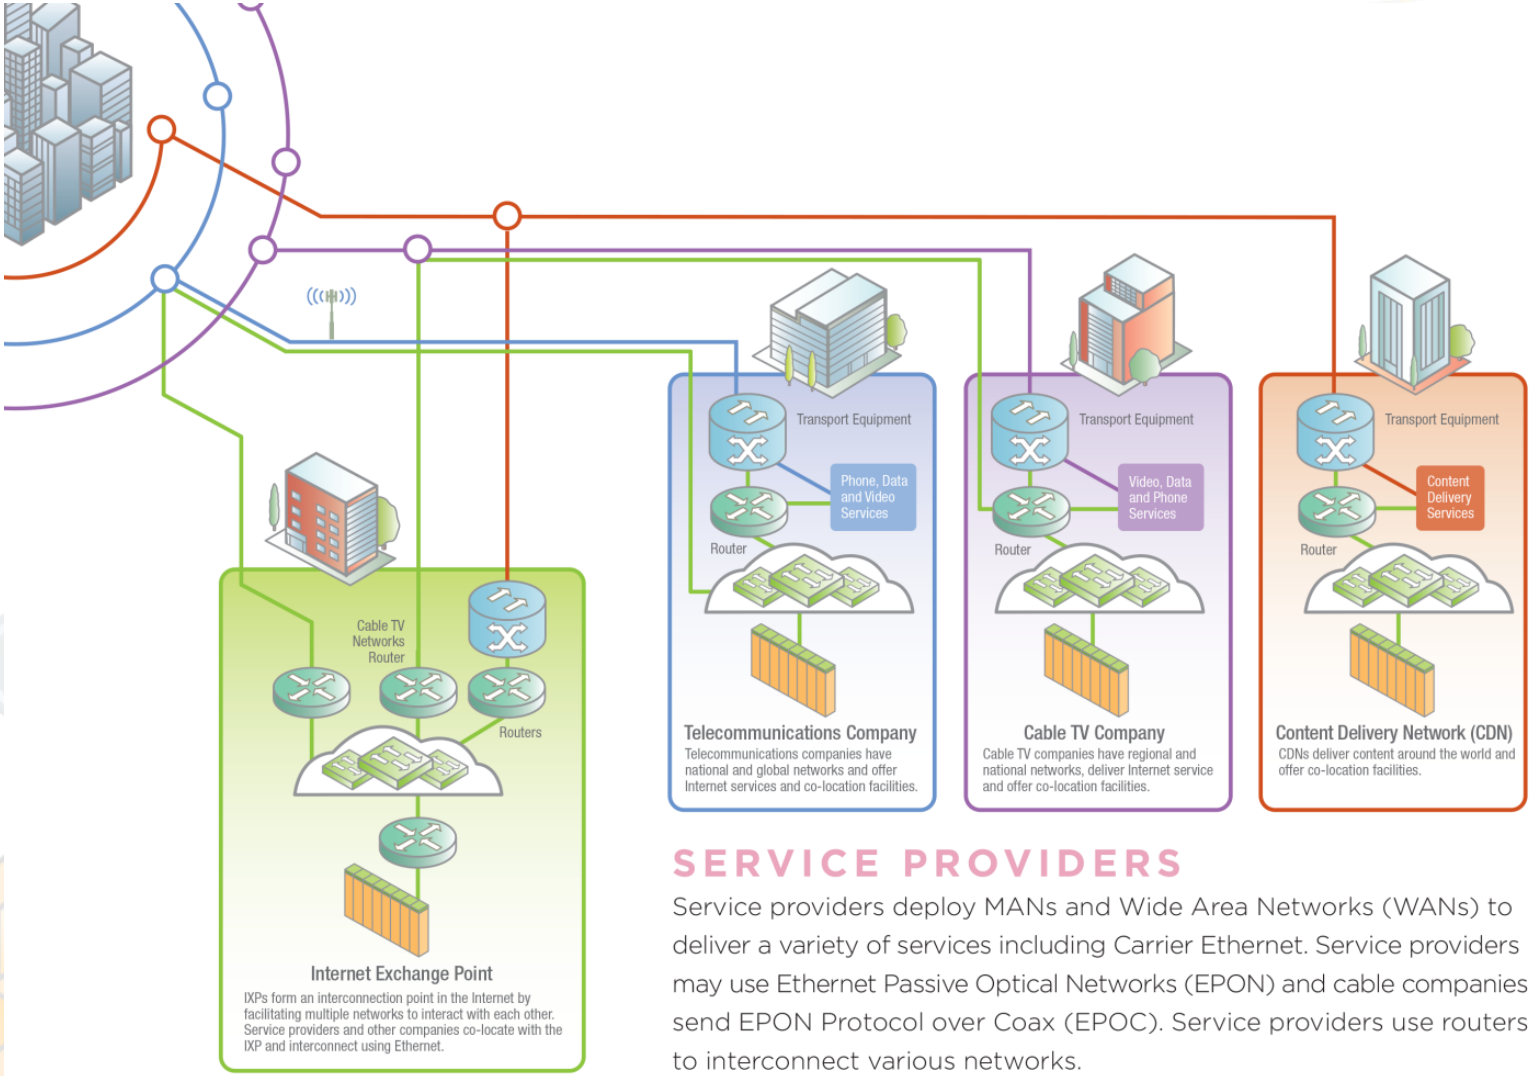

SERVICE PROVIDERS

Service providers deploy MANs and Wide Area Networks (WANs) to deliver a variety of services including Carrier Ethernet. Service providers may use Ethernet Passive Optical Networks (EPON) and cable companies send EPON Protocol over Coax (EPOC). Service providers use routers to interconnect various networks.

New Infographic

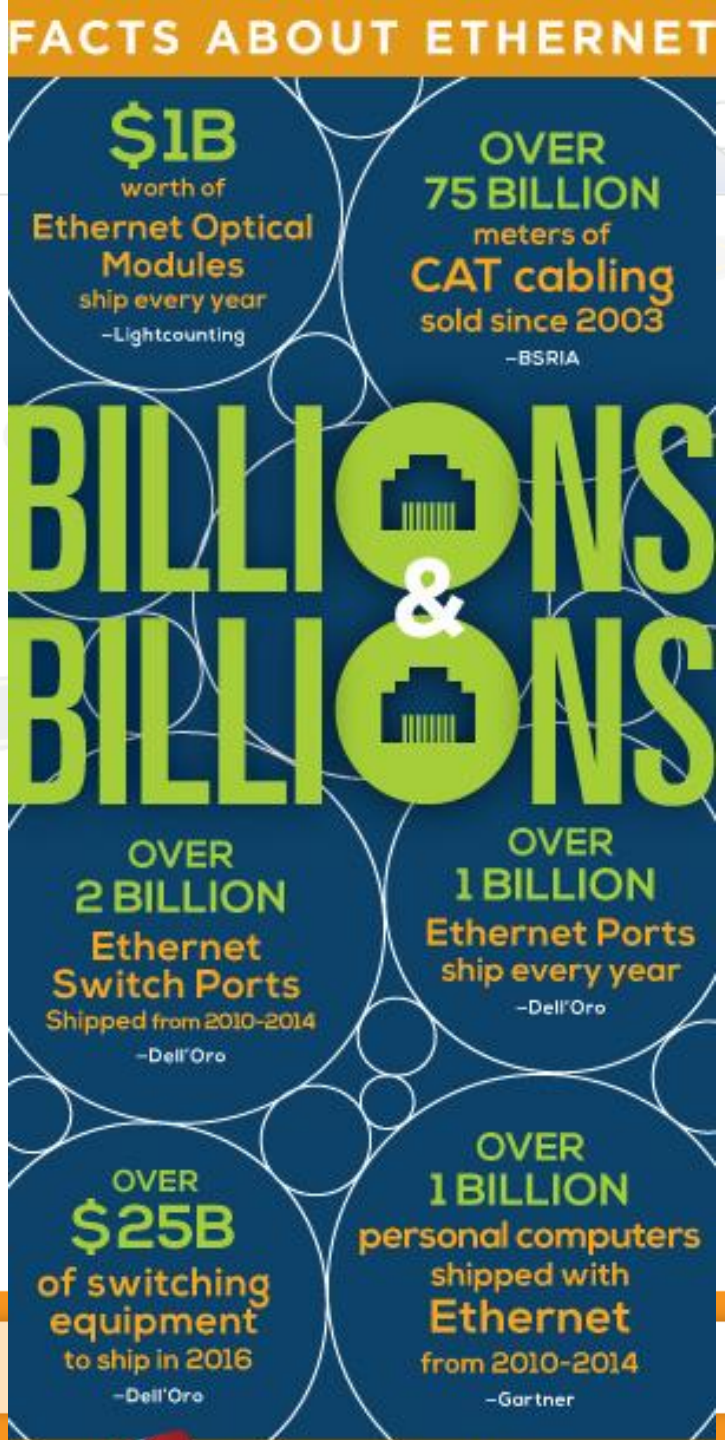

More Roadmap Information

- Your free map is available after the panel
- Free downloads at www.ethernetalliance.org/roadmap/
  - Pdf of map
  - White paper
  - Presentation with graphics for your use

OPTICAL ETHERNET ROADMAP

Session A-104: Ethernet Roadmap: Applications, Servers, Optics, and Switches

Santa Clara, CA

Chris Cole, Finisar Corp.

April 15, 2015

ethernet alliance

Disclaimer

- Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance.

Ex. 40G Data Rate SW (Parallel) Transceiver

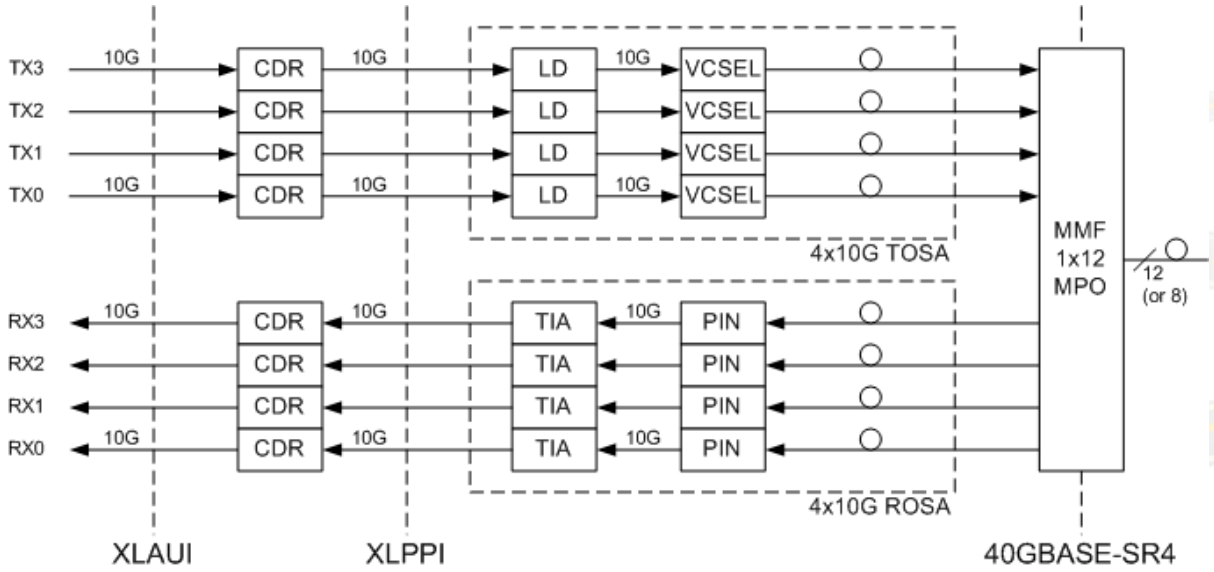

Data Rate	No. of Lanes		Lane Rate
Gb/s	fiber pairs	λ	Gb/s
40	4	1	10

QSFP+

MPO

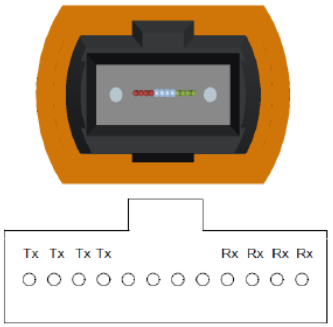

100GBASE-SR4 is 4 lanes of 25 Gb/s

Ex. 40G Data Rate LW (WDM) Transceiver

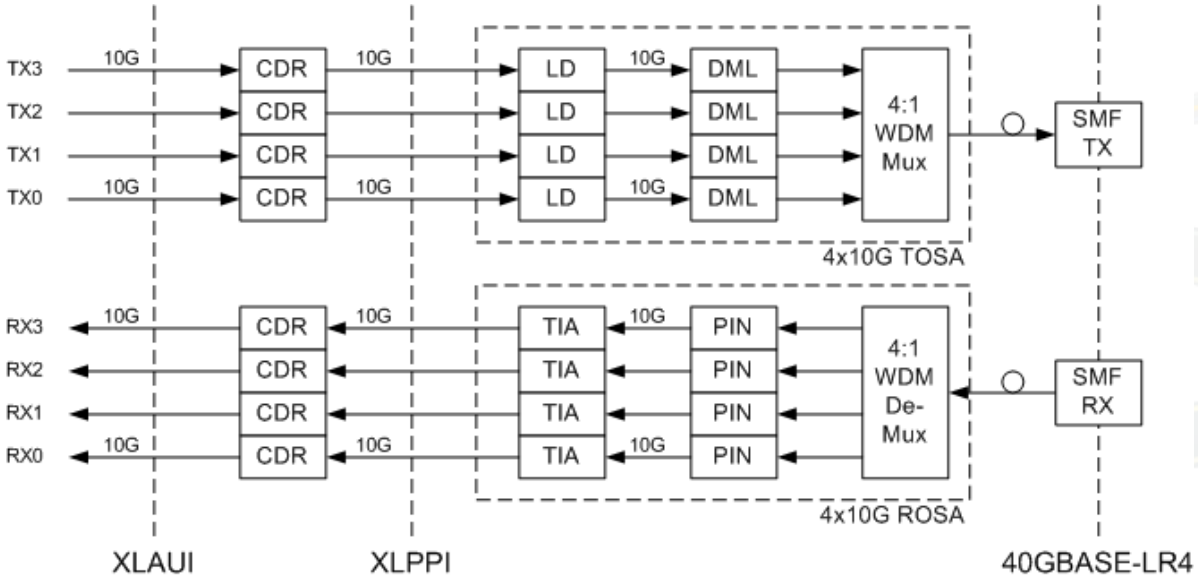

Data Rate	No. of Lanes		Lane Rate
Gb/s	fiber pairs	λ	Gb/s
40	1	4	10

QSFP+

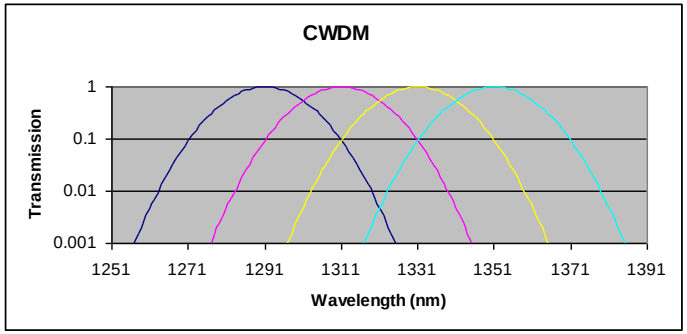

100GBASE-LR4 is 4 lanes of 25 Gb/s

Existing 40G & 100G Optics

Data Rate	No. of Lane Pairs		Lane Rate	SW code	LW code
Gb/s	Fiber	λ	Gb/s	(MMF)	(SMF)
40	4	1	10	**SR4**	*PSM4*
40	1	4	10	*SWDM4*	**LR4**
40	1	1	40		**FR**
100	10	1	10	**SR10**	
100	4	1	25	**SR4**	*PSM4*
100	1	4	25	*SWDM4*	**LR4** *CWDM4*

IEEE standards in BOLD; all others in *ITALICS* are proprietary

Future 25G, 50G & 100G Optics

Data Rate	No. of Lane Pairs		Lane Rate	SW code	LW code
Gb/s	Fiber	λ	Gb/s	(MMF)	(SMF)
25	1	1	25	**SR**	*LR*
50 (& 40)	1	1	50 (& 40)	*SR*	*LR*
100	2	1	50	*SR2*	*PSM2*
100	1	2	50	*SWDM2*	*CWDM2*
100	1	1	100		*FR*

IEEE standards in BOLD; all others in *ITALICS* are proposals

40G vs. 50G Ethernet Deployment

- 40G Ethernet is now deployed in volume supported by 4x10G Parallel & WDM Transceivers
- Limited 40G Serial Transceiver deployment
- Limited 40G deployment using proprietary 2x20G (“BiDi”)
- Future 40G & 50G Serial technology will be common

  40G Transceivers will likely be dual-rate 40/50G Transceivers
- 40G & 50G Serial Transceivers will have same cost:

  i.e. 50G will have 25% more bandwidth at same cost as 40G
- 50G Ethernet volume will quickly surpass 40G Ethernet volume because 25% more bandwidth will be “free”
- Total 40G Ethernet volume will plateau and decline

How did 40G Ethernet Get Stranded?

- 100G Ethernet was first specified for core networking applications using forward looking 25G technology
- 40G Ethernet was then added for cost sensitive Switch and Server applications using 4 lanes of existing 10G technology
- Single lane 25G technology became more cost effective than 4 lanes of 10G technology
- Single lane 50G technology is now in development and well enable low cost 50GbE offering more bandwidth than 40GbE
- Lesson learned:
  - Increase MAC rates in 2x steps; i.e. no 1Tb/s Ethernet
  - Ethernet Roadmap: 100 → 200 → 400 → 800 → 1600Gb/s

Future 200G & 400G Optics

Data Rate	No. of Lane Pairs		Lane Rate	SW code	LW code
Gb/s	Fiber	λ	Gb/s	(MMF)	(SMF)
200	4	1	50	*SR4*	*PSM4*
200	1	4	50	*SWDM4*	*FR4, LR4*
400	16	1	25	**SR16**	
400	4	2	50	*SR4.2*	*PSM4.2*
400	1	8	50	*SWDM8*	*FR8, LR8*
400	4	1	100		*PSM4*
400	1	4	100		*FR4*

IEEE standards in BOLD; all others in *ITALICS* are proposals

OPTICAL ETHERNET ROADMAP

Thank you

COPPER CONNECTIVITY IN THE 2015 ETHERNET ROADMAP

David Chalupsky

April 15, 2015

ethernet alliance

Agenda

- Active copper projects in IEEE 802.3
- Roadmaps
  - Twinax & Backplane
  - Base-t
- Use cases –
  - Server interconnect: TOR, MOR/EOR
  - WAP

Disclaimer

- Opinions expressed during this presentation are the views of the presenters, and should not be considered the views or positions of the Ethernet Alliance.

Current IEEE 802.3 Copper Activity

- High Speed Serial
  - P802.3by 25Gb/s TF: twinax, backplane, chip-to-chip or module. NRZ
  - P802.3bs 400Gb/s TF: 50Gb/s lanes for chip-to-chip or module. PAM4
- Twisted Pair (4-pair)
  - P802.3bq 40GBASE-T TF
  - P802.3bz 2.5G/5GBASE-T
  - 25GBASE-T study group
- Single twisted pair for automotive
  - P802.3bp 1000BASE-T1
  - P802.3bw 100BASE-T1
- PoE
  - P802.3bt – 4-pair PoE
  - P802.3bu – 1-pair PoE

Twinax Copper Roadmap

- 10G SFP+ Direct Attach is highest attach 10G server port today
- 40GBASE-CR4 entering the market
- Notable interest in 25GBASE-CR for cost optimization
- Optimizing single-lane bandwidth (cost/bit) will lead to 50Gb/s

BASE-T Copper Roadmap

- 1000BASE-T still ~75% of the 47M server ports shipped in 2014

Optimizing for the Future

- Data center:
  - 10G->25G->40GBASE-T roadmap on compatible infrastructure
- Enterprise horizontal spaces:
  - 2.5G/5GBASE-T squeeze more usable bandwidth from the 70B meters of Cat5e/6 cabling sold in the last 10 years

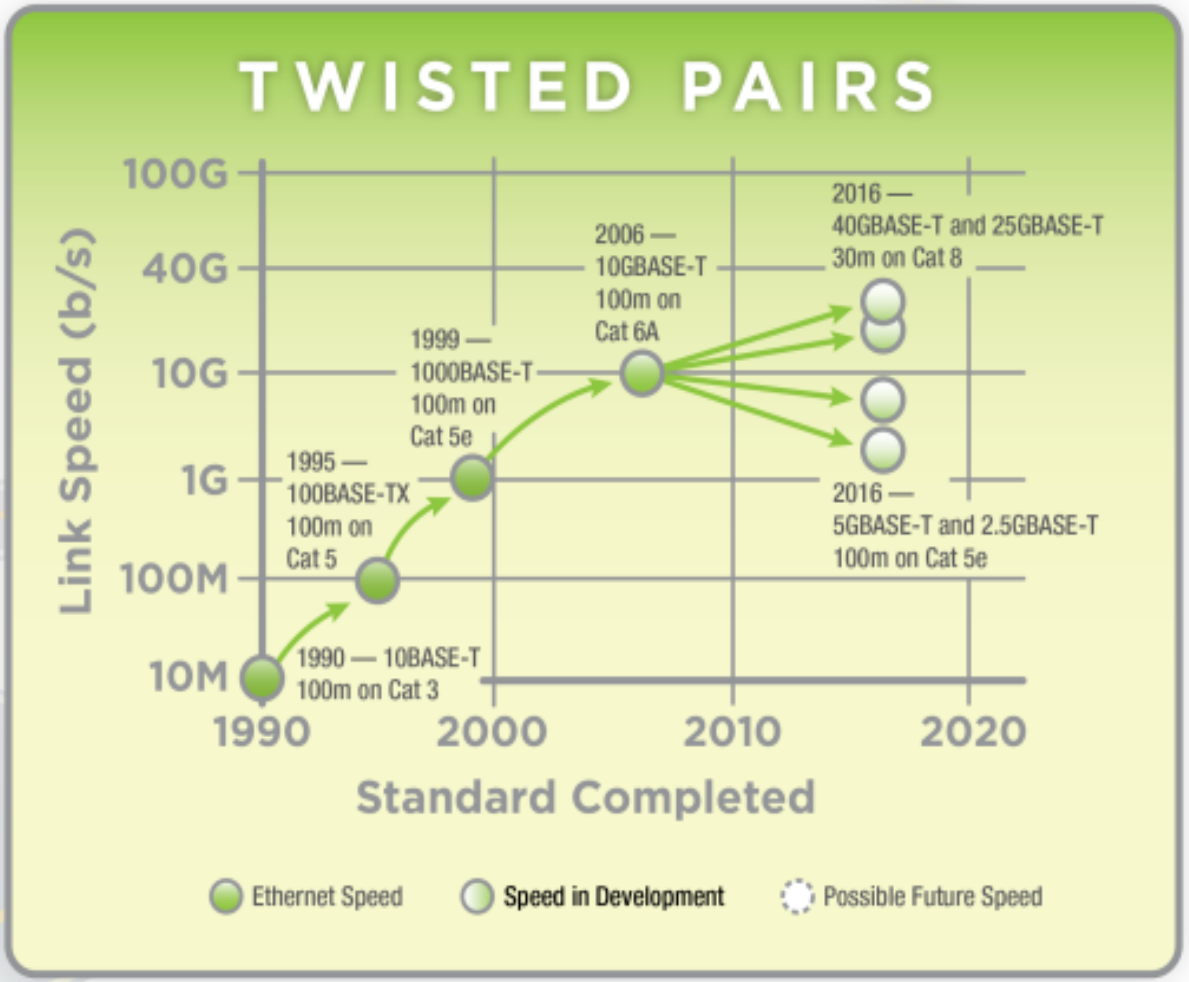

The Applications Spaces of BASE-T

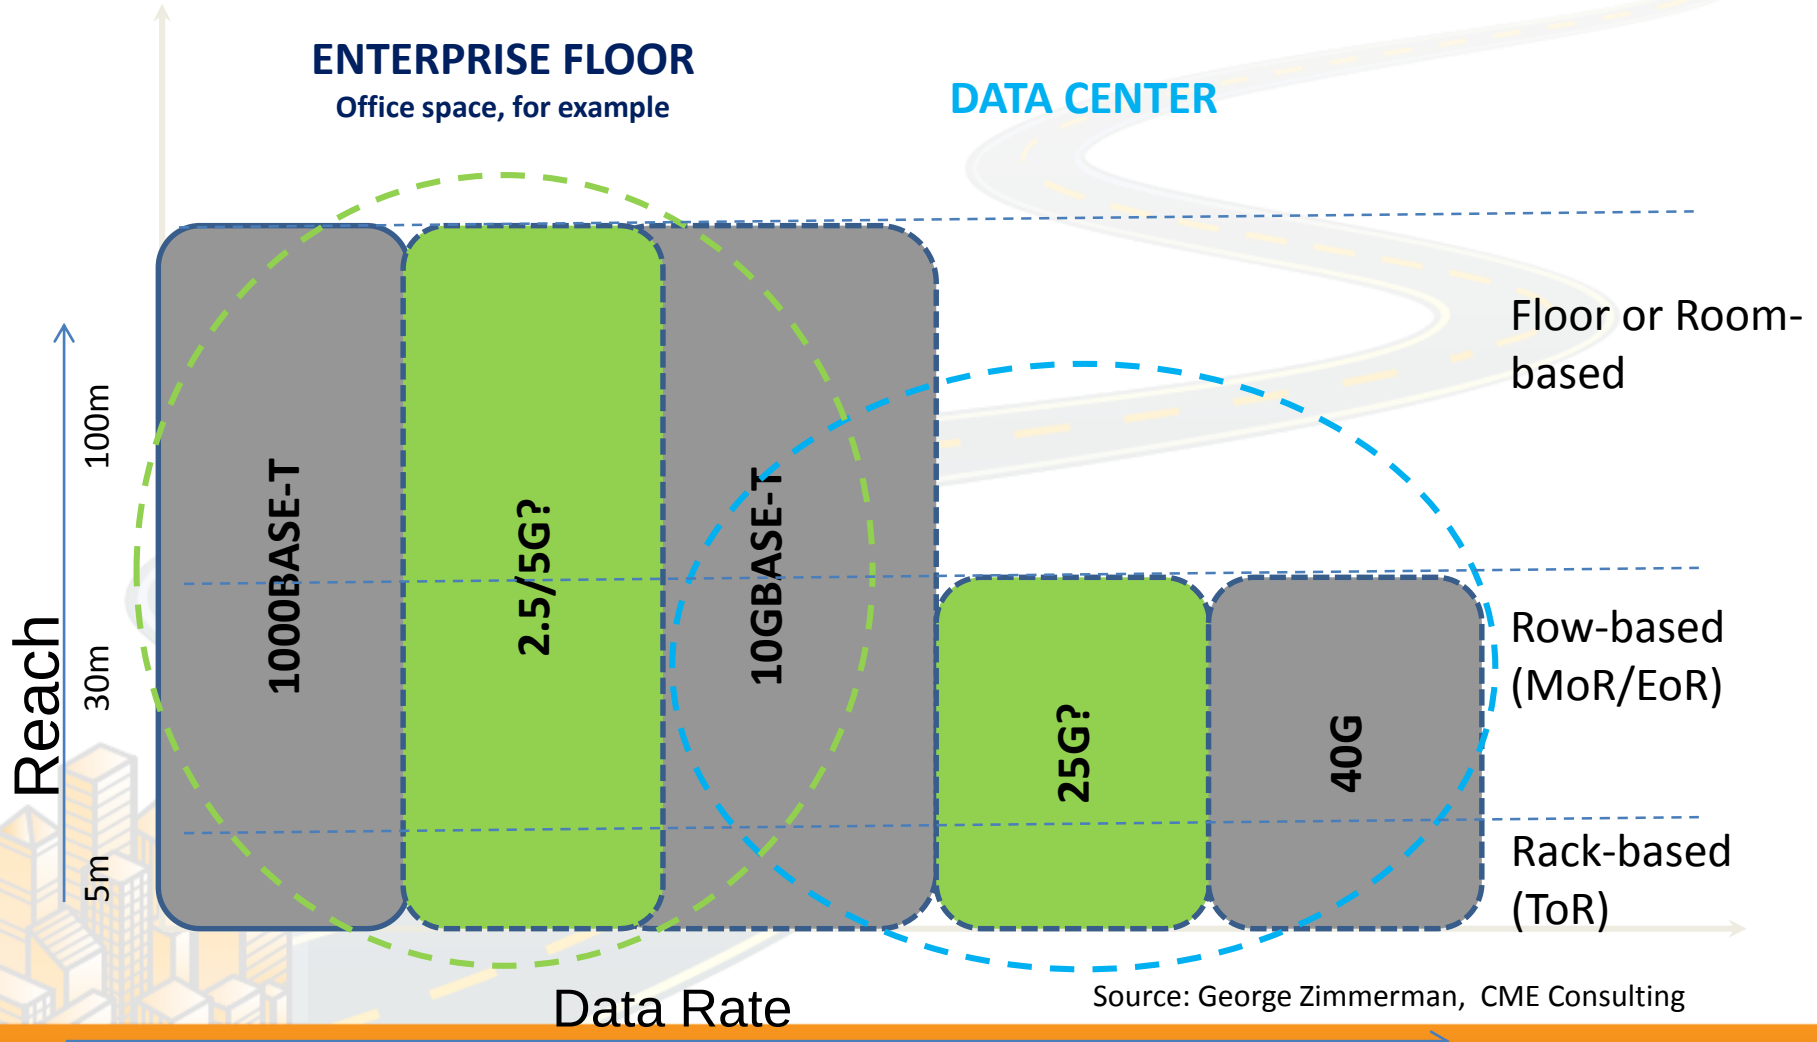

Source: George Zimmerman, CME Consulting

ToR, MoR, EoR Interconnects

ToR

MoR

EoR

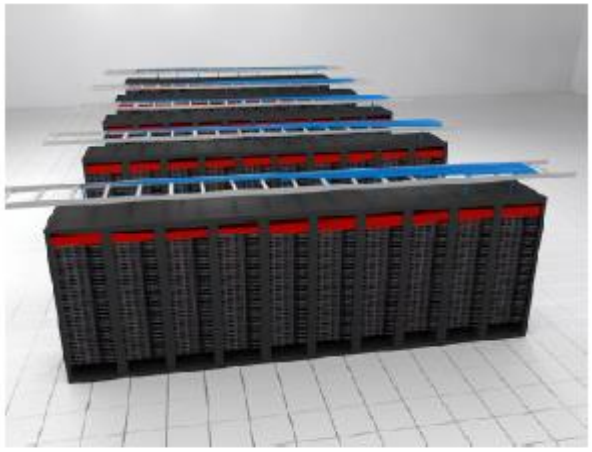

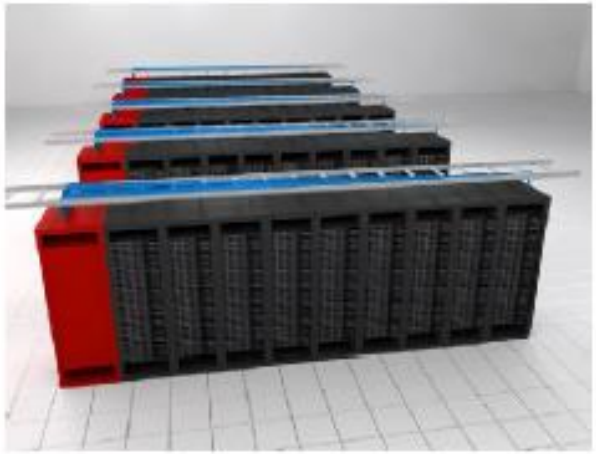

Intra-rack can be addressed by twinax copper direct attach

Reaches addressed by BASE-T and fiber

Pictures from jimenez_3bq_01_0711.pdf, 802.3bq

802.3 Ethernet and 802.11 Wireless LAN

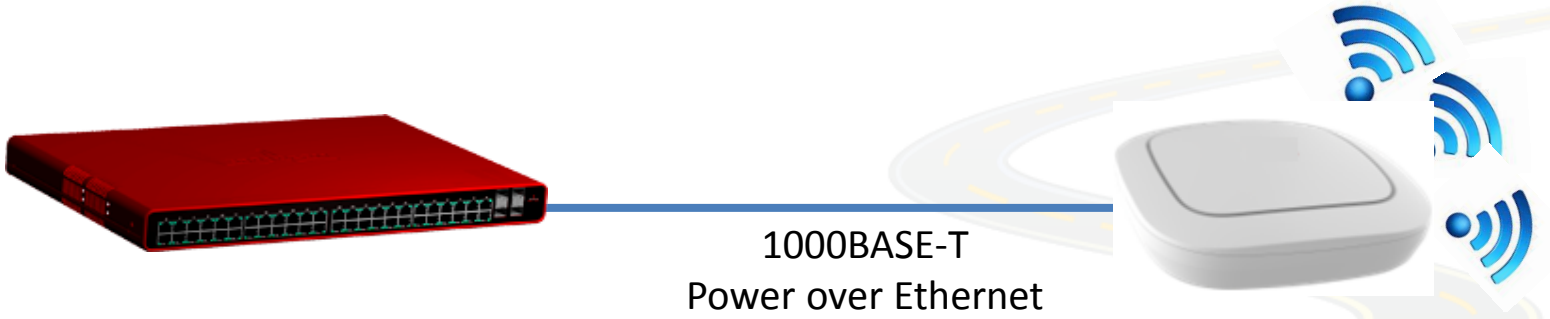

Ethernet Access Switch

- Dominated by 1000BASE-T ports
- Power over Ethernet Power Sourcing Equipment (PoE PSE) supporting 15W, 30W, *4PPoE: 60W-90W*

Cabling

- 100m Cat 5e/6/6A installed base.
- New installs moving to Cat 6A for 10+yr life.

Wireless Access Point

- Mainly connects 802.11 to 802.3
- Normally PoE powered
- Footprint sensitive (e.g. power, cost, heat, etc.)
- Increasing 802.11 radio capability (11ac Wave1 to Wave2) drives Ethernet backhaul traffic beyond 1 Gb/s.
- Link Aggregation (Nx1000BASE-T) or 10GBASE-T only options today

IMPLICATIONS OF 50G SERDES ON ETHERNET SPEEDS

Kapil Shrikhande

ethernet alliance

Ethernet Speeds: Observations

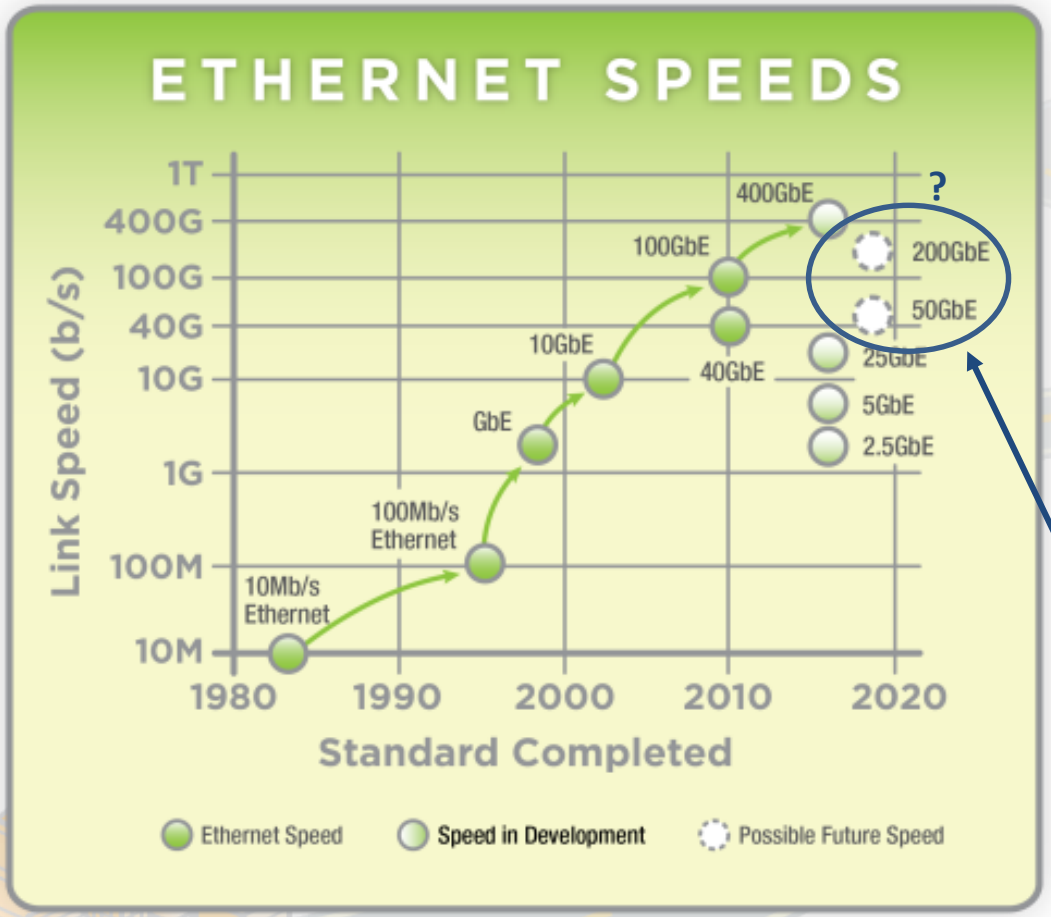

- Data centers driving speeds differently than Core networking
  - 40GE (4x10G) took-off in Data center networks with 10GE servers
  - 25GE being defined for next-gen server IO speed > 10GE
  - 100GE (4x25G) will take off with 25GE servers
    - And 50G (2x25G) servers
  - What's beyond 25/100GE? *Follow the Serdes* ☺

SerDes / Signaling, Lanes and Speeds

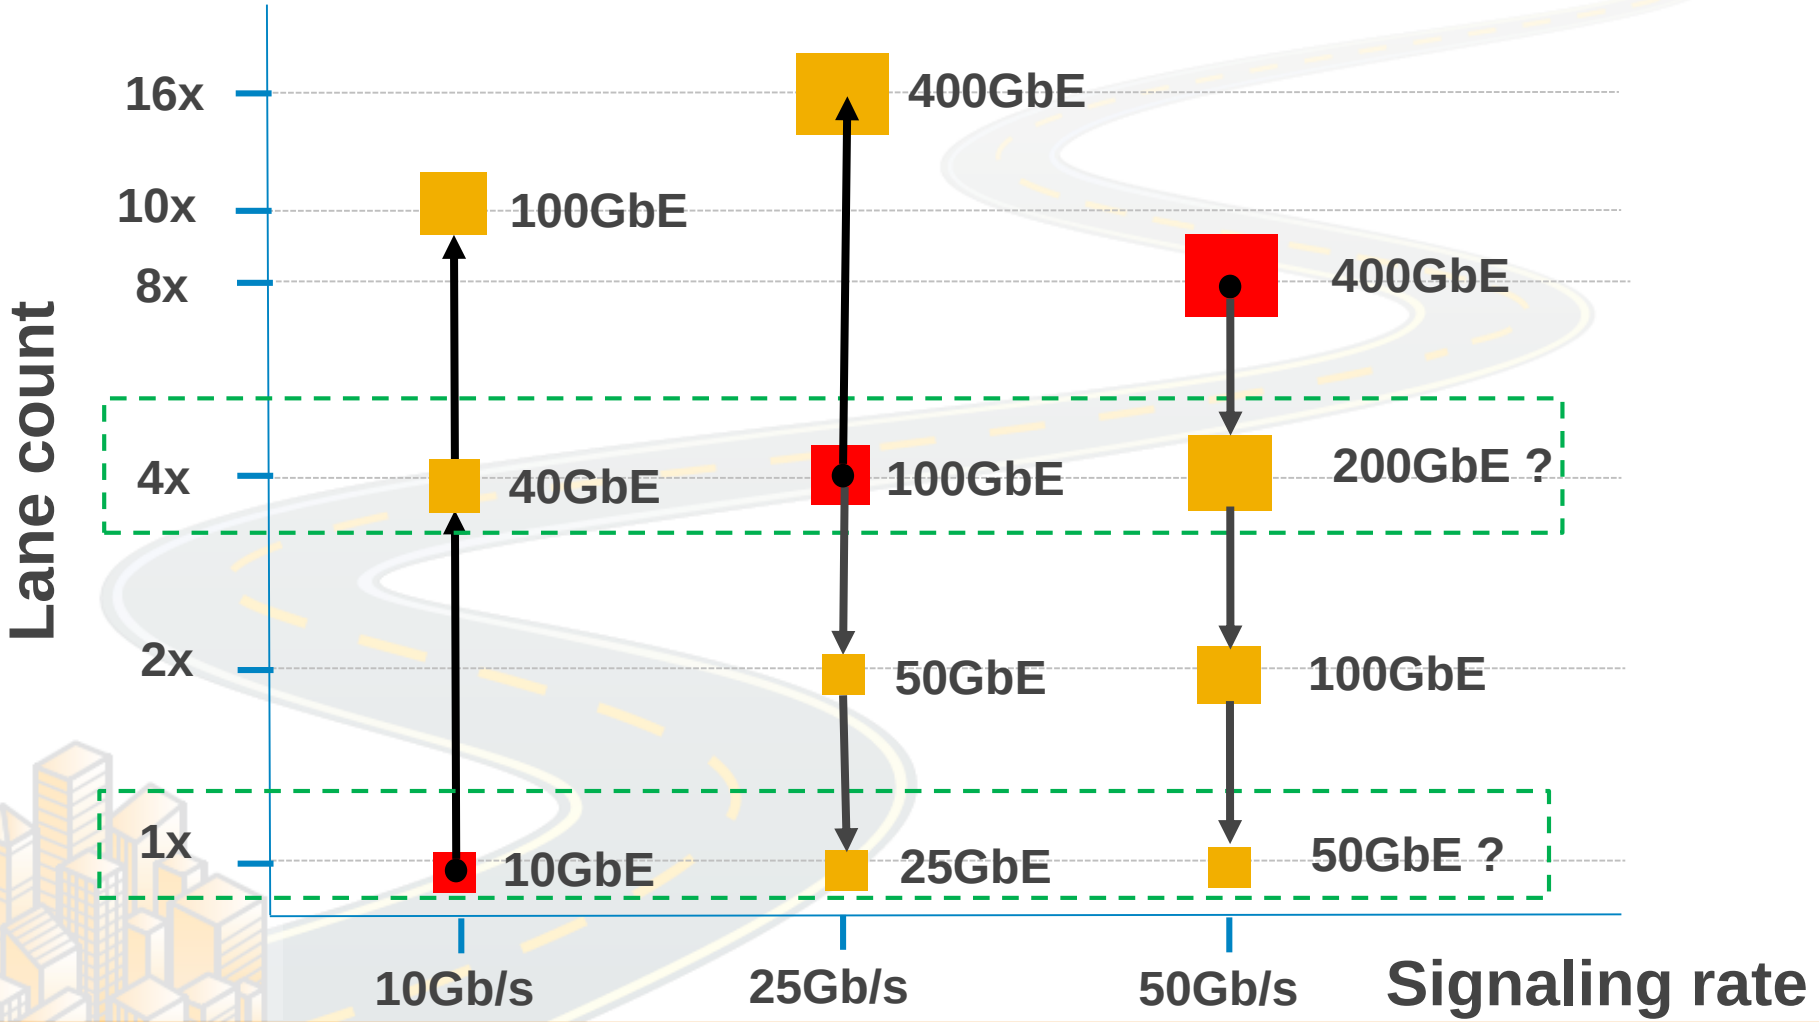

Ethernet ports using 10G SerDes

Data centers widely using 10GE servers, 40GE Network IO

- 128x10Gb/s switch ASIC

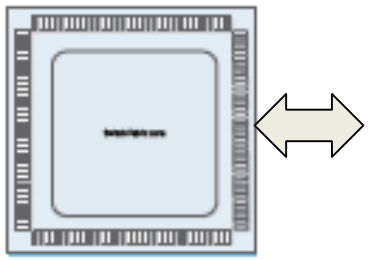

128x10GbE
32x40GbE
12x100GbE

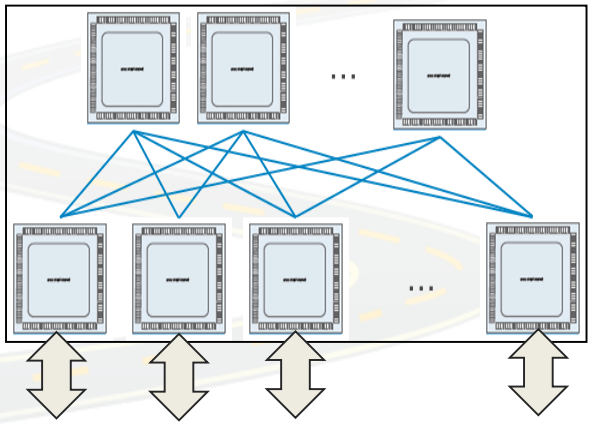

Large port count Spine switch
= N*N/2, where N is switch chip radix
N=32 → <= 512x40GE Spine switch
N=12 → <= 72x100GE Spine switch

- E.g. TOR configuration
  - 96x10GE + 8x40GE

- High port count of 40GE better suited for DC scale-out

Ethernet ports using **25**G SerDes

Data centers planning for **25**GE servers, **100**GE Network IO

- 128x**25**Gb/s switch ASIC

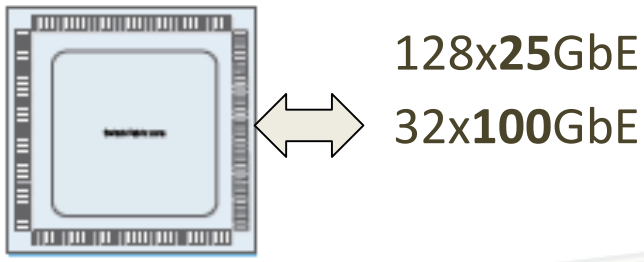

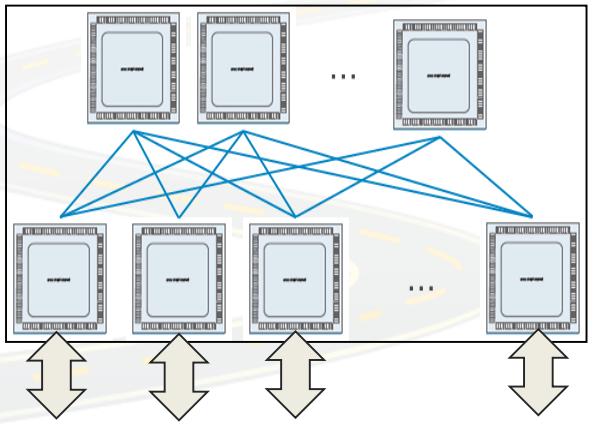

Large port count Spine switch

= N*N/2, where N is switch chip radix

N = 32 → <= 512x**100**GE Spine switch

- E.g. TOR configuration
  - 96x**25**GE + 8x**100**GE

- 100GE (4x25G) now matches 40GE in ability to scale

QSFP optics (40GE, 100GE)

- Data center modules need to support various media types, and reach

	MMF	SMF
Duplex 	• 100m	• 2km • 10km • 40km
Parallel 	• 100m • 300m	• 500m

- QSFP+ evolved to do just that
- QSFP28 following suit
- 4x lanes enabling compact designs
- IEEE and MSA specs.
- XLPPI, CAUI4 interfaces
- Breakout provides backward compatibility
  - E.g. 4x10GbE

Putting it all together

- E.g. 10/40GE Mega Data center
  - 288 x 40GE Spine switch
  - 64 Spine switches
  - 96 x 10GE Servers / Rack
  - 8 x 40GE ToR Uplinks
  - # Racks total ~ 2304
  - # 10GE Servers ~ 221,184
- Same scale possible with 25GbE servers, 100GE networking
- 100GE on same fiber optic cables as 40GE is possible

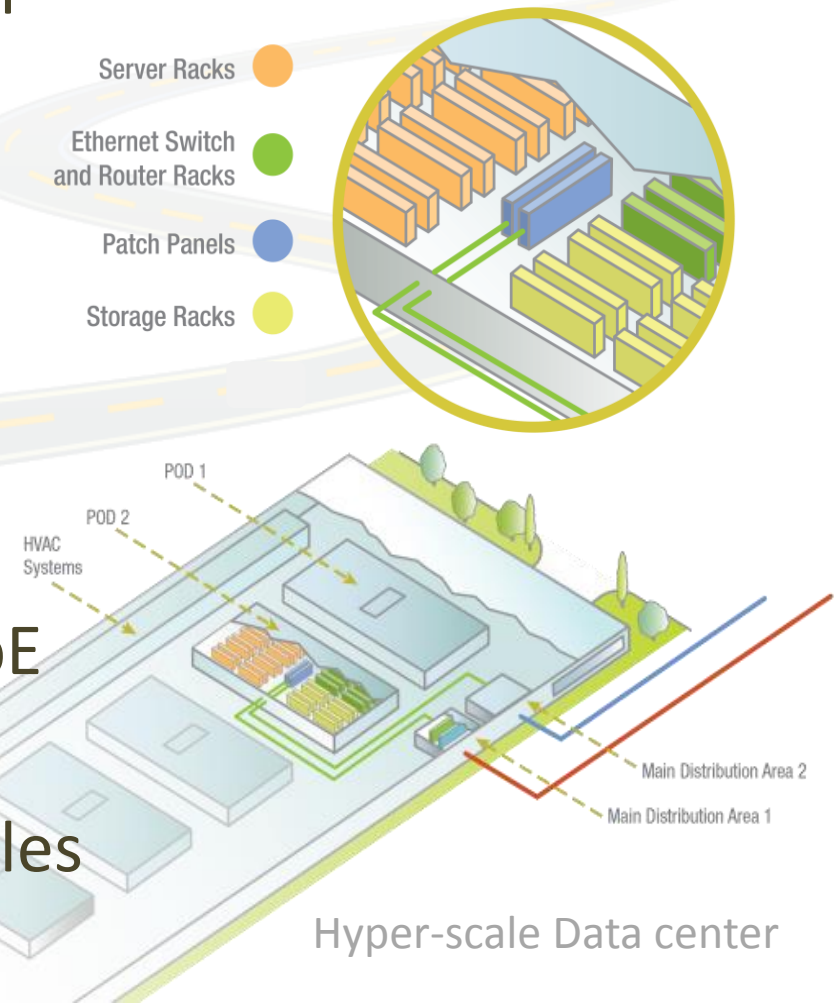

Hyper-scale Data center

Evolution using 50G SerDes

Next-gen switch ASIC

N x 50Gb/s SerDes chip

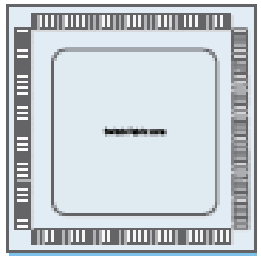

Radix

E.g. N = 128

128 x 40/50GbE

64 x 100GbE

32 x 200GbE

16 x 400GbE

Speed

- 50GbE Server I/O
  - Single-lane Speed > 25GE
- 200GbE Network I/O
  - Four-lane Speed > 100GE
  - Balances Radix v. Speed
- 200GE on same fiber optic cables as 100GE possible
- 4x50GE breakout possible
- DC scalability – same as 25/100GE, 10/40GE designs

200GE QSFP feasibility ?

- 50Gb/s optical lanes for SMF, MMF : Yes
- 8x Parallel and Duplex fibers: Yes
- 50Gb/s electrical specifications: Yes
- Electrical Connector : Yes
- Twin-ax for 50Gb/s lanes: Yes
- FEC striped over 4-lanes : Yes
- Power, Space, Integration ? Investigate.
  - Same questions as QSFP28 … will get solved over time
- For optical engineers – 200GbE allows continued use of Quad designs from 40/100GbE. Boring but doable ☺

The Ethernet Roadmap

Plan 50/200GbE standardization.

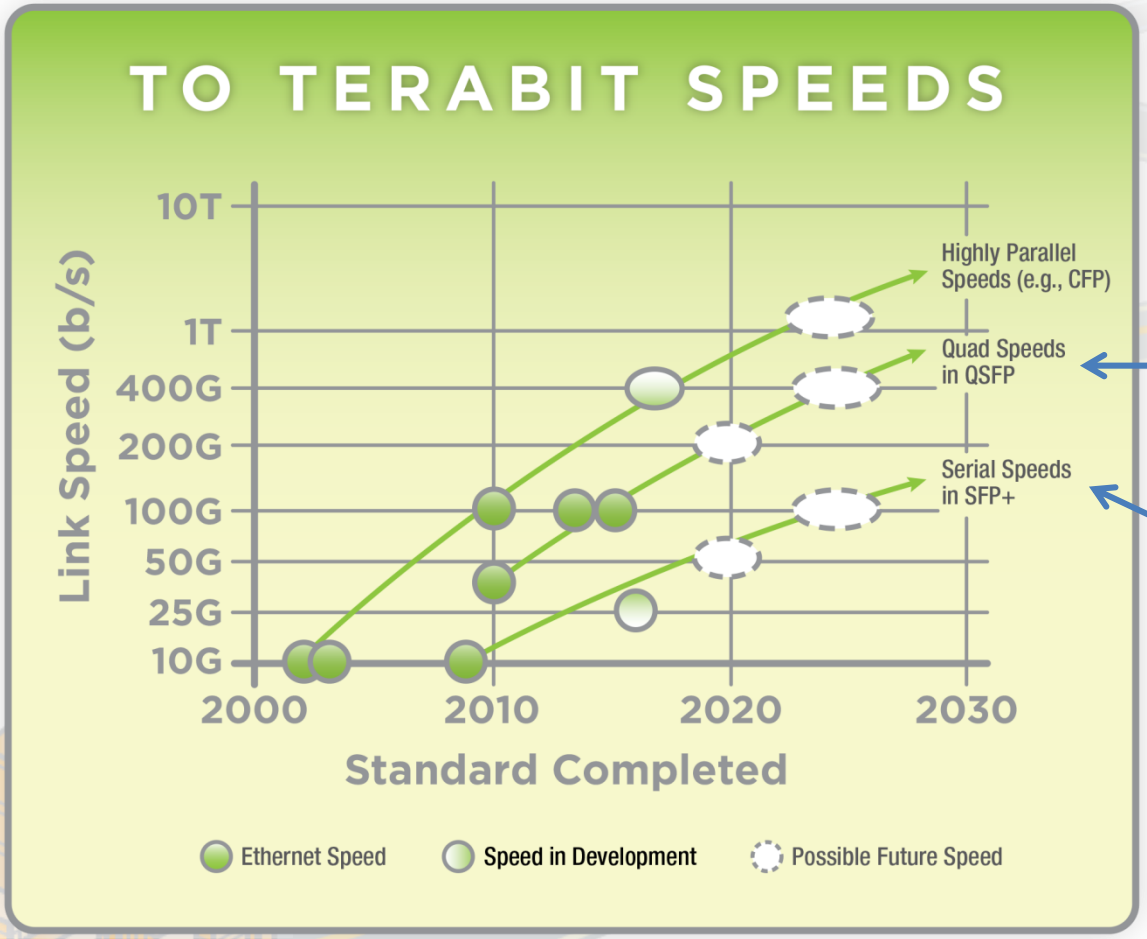

QSFP

400G >2020

200G - ~2019?

100G - 2015

40G - 2010

SFP

100G >2020

50G - ~2019?

25G - 2016

10G - 2009

Questions and Answers

Thank You!

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

ethernet alliance