25 Gb/s Ethernet Over Single Mode Fiber Call For Interest Consensus Presentation

IEEE 802.3

David Lewis, Lumentum

Kohichi Tamura, Oclaro

Peter Jones, Cisco

Dallas, TX

Nov 9th-12th, 2015

Supporters

Dale Murray, Light Counting

Rick Rabinovich, ALE

Alan Tipper, Semtech

Bharat Tailor, Semtech

Peter Jones, Cisco

Mark Nowell, Cisco

Kohichi Tamura, Oclaro

Kiyo Hiramoto, Oclaro

Raymond Xie, Huawei

Scott Kipp, Brocade

John D'Ambrosia, Independent

Dave Ofelt, Juniper

Pirooz Tooyserkani, Cisco

Chad Jones, Cisco

Joel Goergen, Cisco

Amrik Bains, Cisco

Mark Gustlin, Xilinx

Gary Nicholl, Cisco

Jonathan King, Finisar

Mike Dudek, Qlogic

Yong Kim, Broadcom

Matt Brown, APM

Jerry Pepper, Ixia

Thananya Baldwin, Ixia

George Zimmerman, CME Consulting, Inc. / Commscope

Richard Mei, Commscope

Jacky Chang, HPE

Daryl Inniss, Ovum

Brian Teipen, Adva Optical

Kapil Shrikhande, Dell

Kent Lusted, Intel

Jan Filip, Maxim

CFI Objectives

- To gauge the interest in studying single mode fiber PMD(s) for 25 Gb/s Ethernet

- We do not need to:
  - Fully explore the problem
  - Debate strengths and weaknesses of solutions
  - Choose a solution
  - Create a PAR or 5 Criteria
  - Create a standard

- Anyone in the room may vote or speak

Overview: Motivation

- Address 25GE links longer than 100m.
- Develop cost optimized 25GE SMF PMD(s).
- Allow other markets (e.g., Enterprise, Metro) to adopt 25GE.
- Initial applications?

  Enterprise campus; Metro network access
- Similar initiatives?

  Mobile Front-haul – 25G for Common Public Radio Interface (CPRI)

What Are We Talking About?

- Application spaces that could move to 25Gb/s lanes (1X or 4X) over SMF.
- 25GE SMF provides optimized single lane switch/router connectivity
- Enable 25GE to move from DC to campus and beyond.

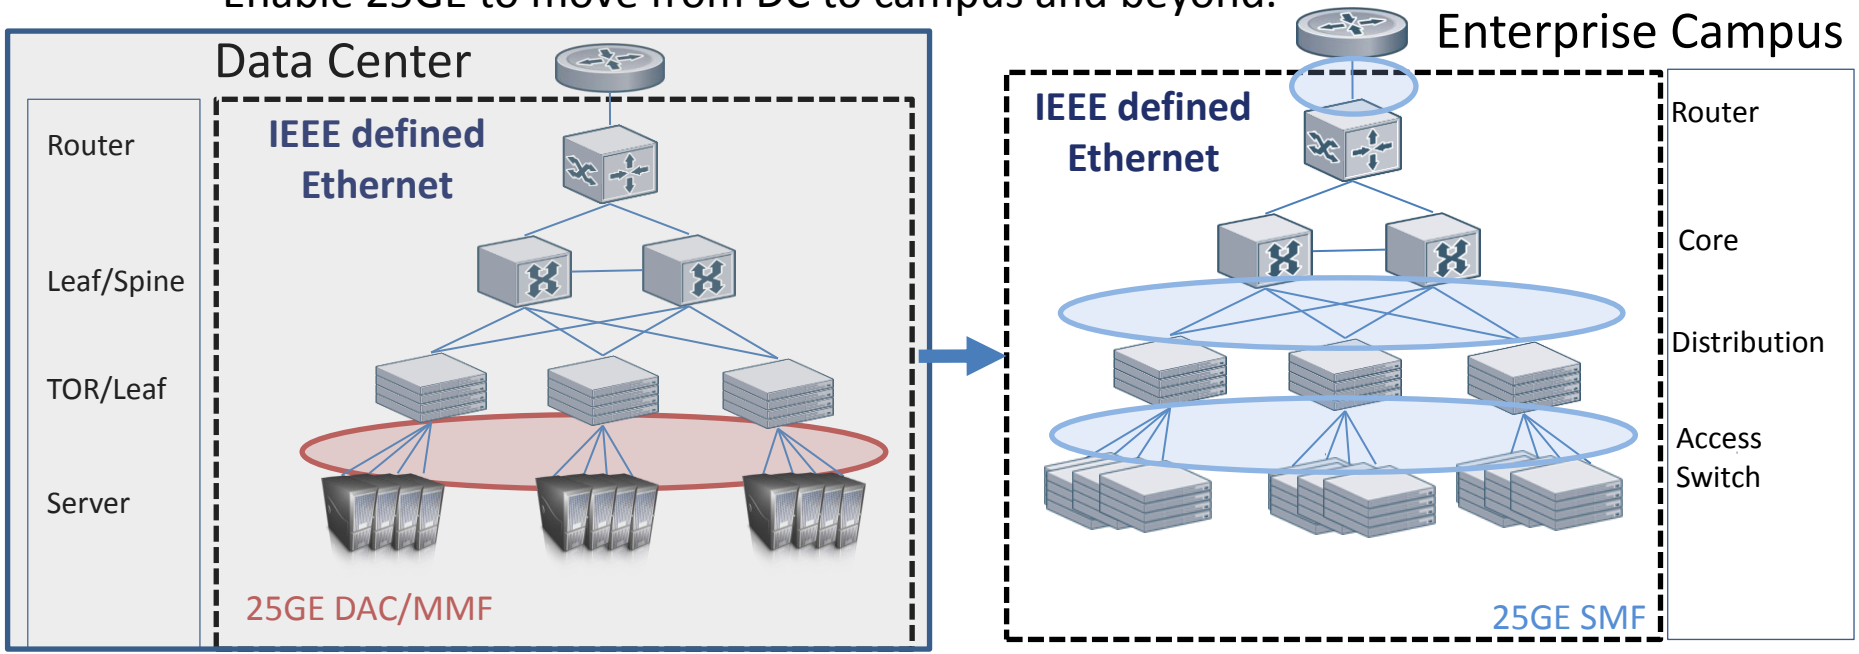

Agenda

- Overview Discussion
  - 25GE SMF PMDs – David Lewis, Lumentum
- Presentations
  - 25GE SMF Market Drivers – Peter Jones, Cisco
  - 25GE SMF Technical Feasibility – Kohichi Tamura, Oclaro
  - 25GE SMF: Why Now? – David Lewis, Lumentum
- Straw Polls

Market Drivers

25GE SMF PMD

Peter Jones, Cisco

Ethernet Evolution

Leading edge markets (e.g., Cloud DC, SP) drive speeds

- Initial adoption: 10G ~2004; 40G ~2012; 100G ~2013; 25G ~2016

Other markets (e.g., Enterprise DC/Campus) more cost sensitive

- 1/4-lane solutions enable cost reductions and volume adoption
- Key question - “What do I need & what can I afford?”

Ethernet market is wide & varied (and getting more so).

- Today we have active projects from 100Mb/s to 400Gb/s
- *Ethernet “speciation” in action.*

Speciation: an evolutionary process that leads to the formation of a new species

1GE/10GE: MMF to SMF Transition

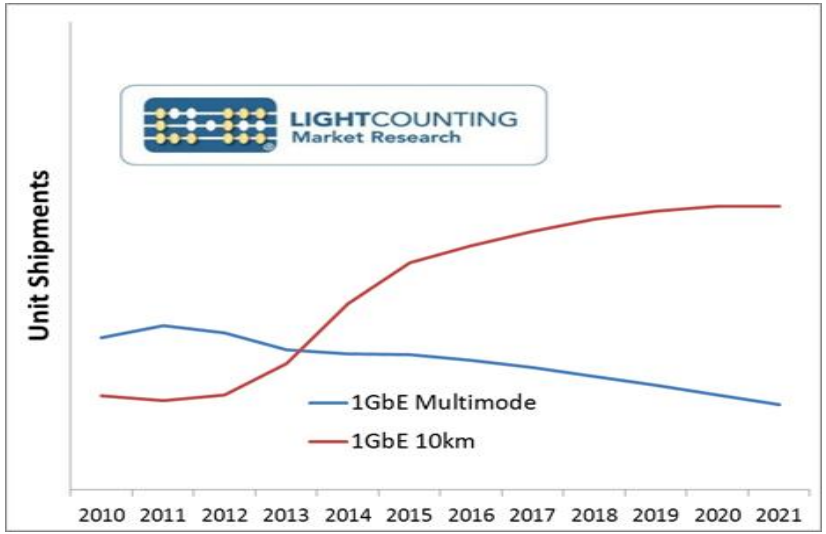

- SMF surpasses MMF volume:
  - 1GE 2013, 10GE ~2018
- Key form factors – SFP, SFP+

- Increasing deployment of longer 1GE/10GE runs drives SMF

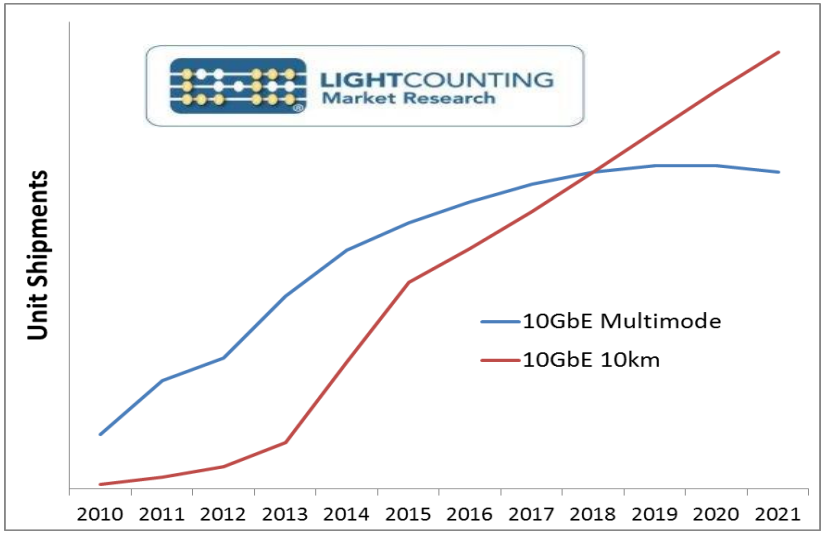

LightCounting "High Speed Datacenter Optical Interconnects Report" (June'15)

25GE: MMF to SMF Transition

- Like 10GE, 25GE SMF starts slower than MMF
  - 10GE MMF reach is 400m (OM4)
  - 25GE MMF reach is 100m
  - 25GE SMF transition could be faster than 10GE
- Key form factor – SFP28

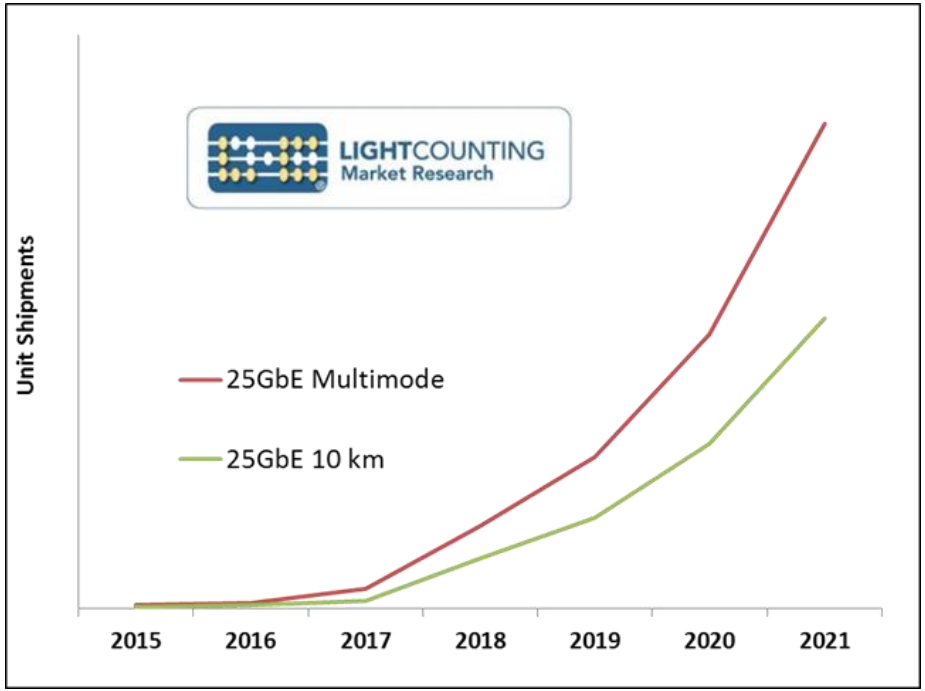

LightCounting <will check>

25G Module Forecast

Ethernet + CPRI

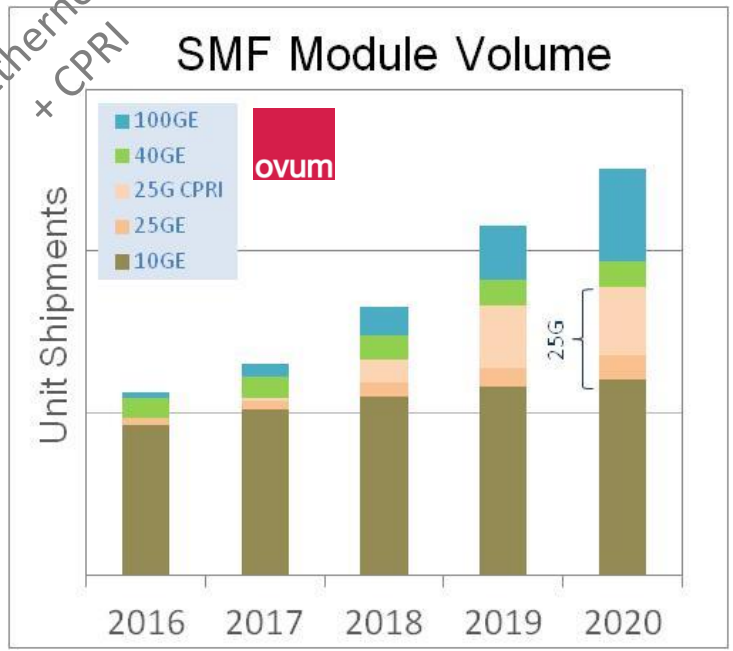

Ethernet Only

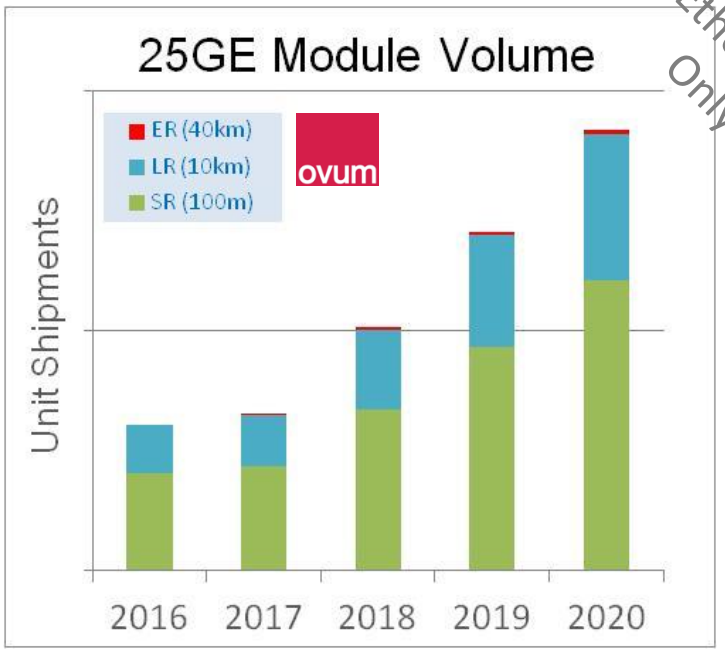

- 25G SMF volume to approximates 100G SMF through 2020
- 25G SMF is mostly LR. Around 3% of SMF is ER in 2020

Reference: Ovum "Total OC Forecast Spreadsheet: 2014-2020" (August 2015)

“Stop me if you’ve heard this before.” Part 1

25Gb/s Ethernet Connectivity

- Enables similar topology as 40Gb/s & 10Gb/s
  - Single 25Gb/s SFP28 port implementation or Quad 25Gb/s QSFP28 breakout implementation possible
  - Maximizes ports and bandwidth in ToR switch faceplate
  - Dense rack server
  - Within rack, less than 3m typical length

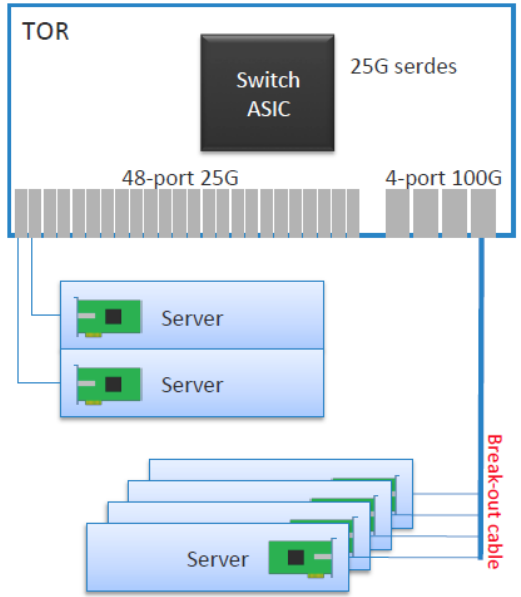

13

IEEE 802.3 Call For Interest – 25Gb/s Ethernet over a single lane for server interconnect – July 2014 San Diego

From 25GE CFI: moving to 25Gb/s optimizes TOR switch design

Other switches (e.g., Campus distribution) can benefit from the same optimization.

25GE CFI: www.ieee802.org/3/cfi/0714_1/CFI_01_0714.pdf

“Stop me if you’ve heard this before.” Part 2

25Gb/s I/O Efficiency

- Switch ASIC Connectivity limited by serdes I/O
- 25Gb/s lane maximizes bandwidth/pin and switch fabric capability vs. older generation
- Single Lane port maximizes server connectivity available in single ASIC
- 25Gb/s port optimizes both port count and total bandwidth for server interconnect

For a 128 lane switch:

Port Speed (Gbps)	Lane Speed	Lanes / port	Usable ports	Total BW (Gbps)
10	10	1	128	1280
25	25	1	128	3200
40	10	4	32	1280
40	20	2	64	2560
100	25	4	32	3200

Using 25Gb/s ports maximizes connectivity and bandwidth.

14

IEEE 802.3 Call For Interest – 25Gb/s Ethernet over a single lane for server interconnect – July 2014 San Diego

From 25GE CFI: moving to 25Gb/s optimizes ASIC price performance

Other switches (e.g., Campus access/distribution) can benefit from the same optimization.

25GE CFI: www.ieee802.org/3/cfi/0714_1/CFI_01_0714.pdf

Enterprise Switching – Rate Evolution

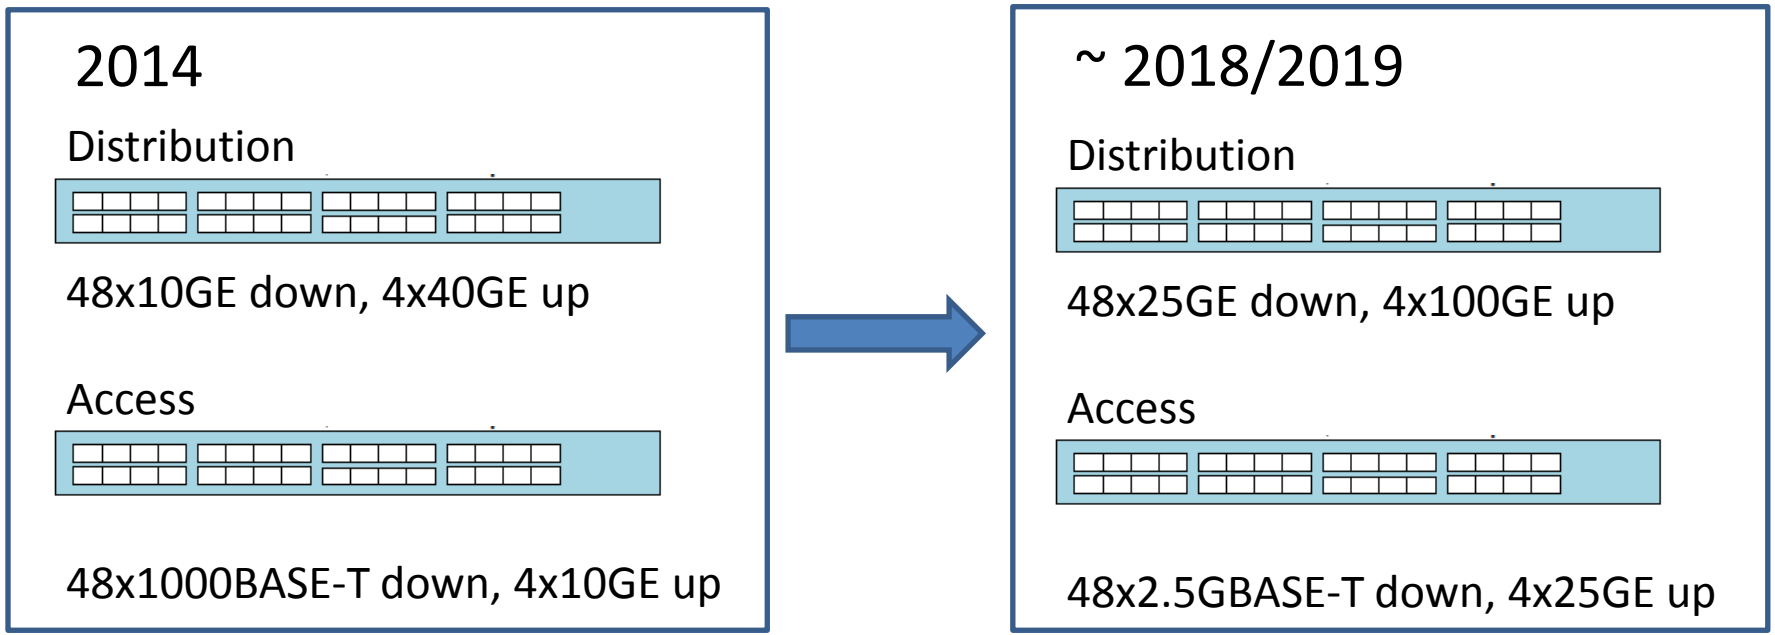

- 2.5X rate, same network model
- Fits nicely with 802.3bz 2.5G/5GBASE-T

Metro-Access – Rate Evolution

- Core moving to 100GE (4 x 25G)
- Simplify system by eliminating redundant gear-boxing & retiming
- Lower cost, lower power

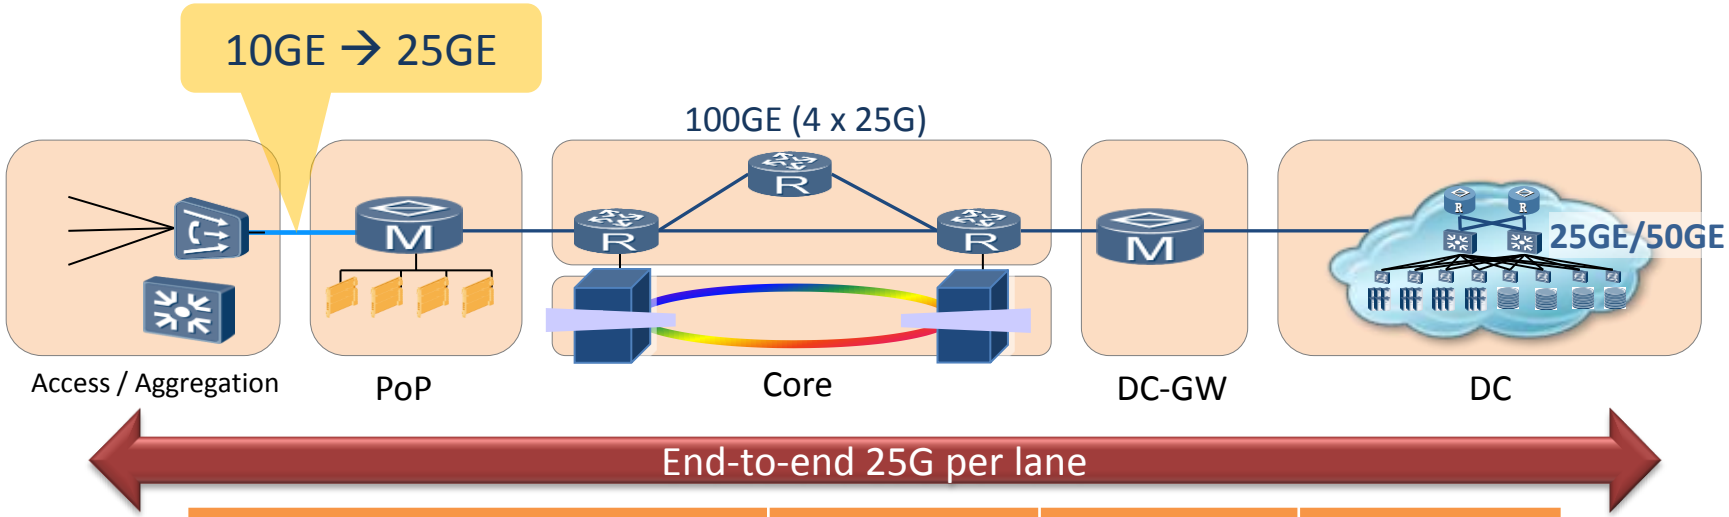

Current Interconnect Distance	10km	40km	80km
Fixed Broadband Network	60-70%	20%	<5%
Mobile Broadband Network	45%	45%	10%

Source: Huawei

Market Drivers Summary

- 25GE ecosystem is missing a story for > 100 meters
- Enterprise & Metro markets (among others) need longer runs
- 25GE is the natural successor to 10GE
- Leveraging 25G lane rates with 25GE/100GE just makes sense