Before 400G for Hyperscale Data Centers

Beyond 400G for
Hyperscale Data Centers
OFC 2019 Panel Th1A
7 March 2019
Chris Cole

FINISAR®

40G and 100G Shipment Data

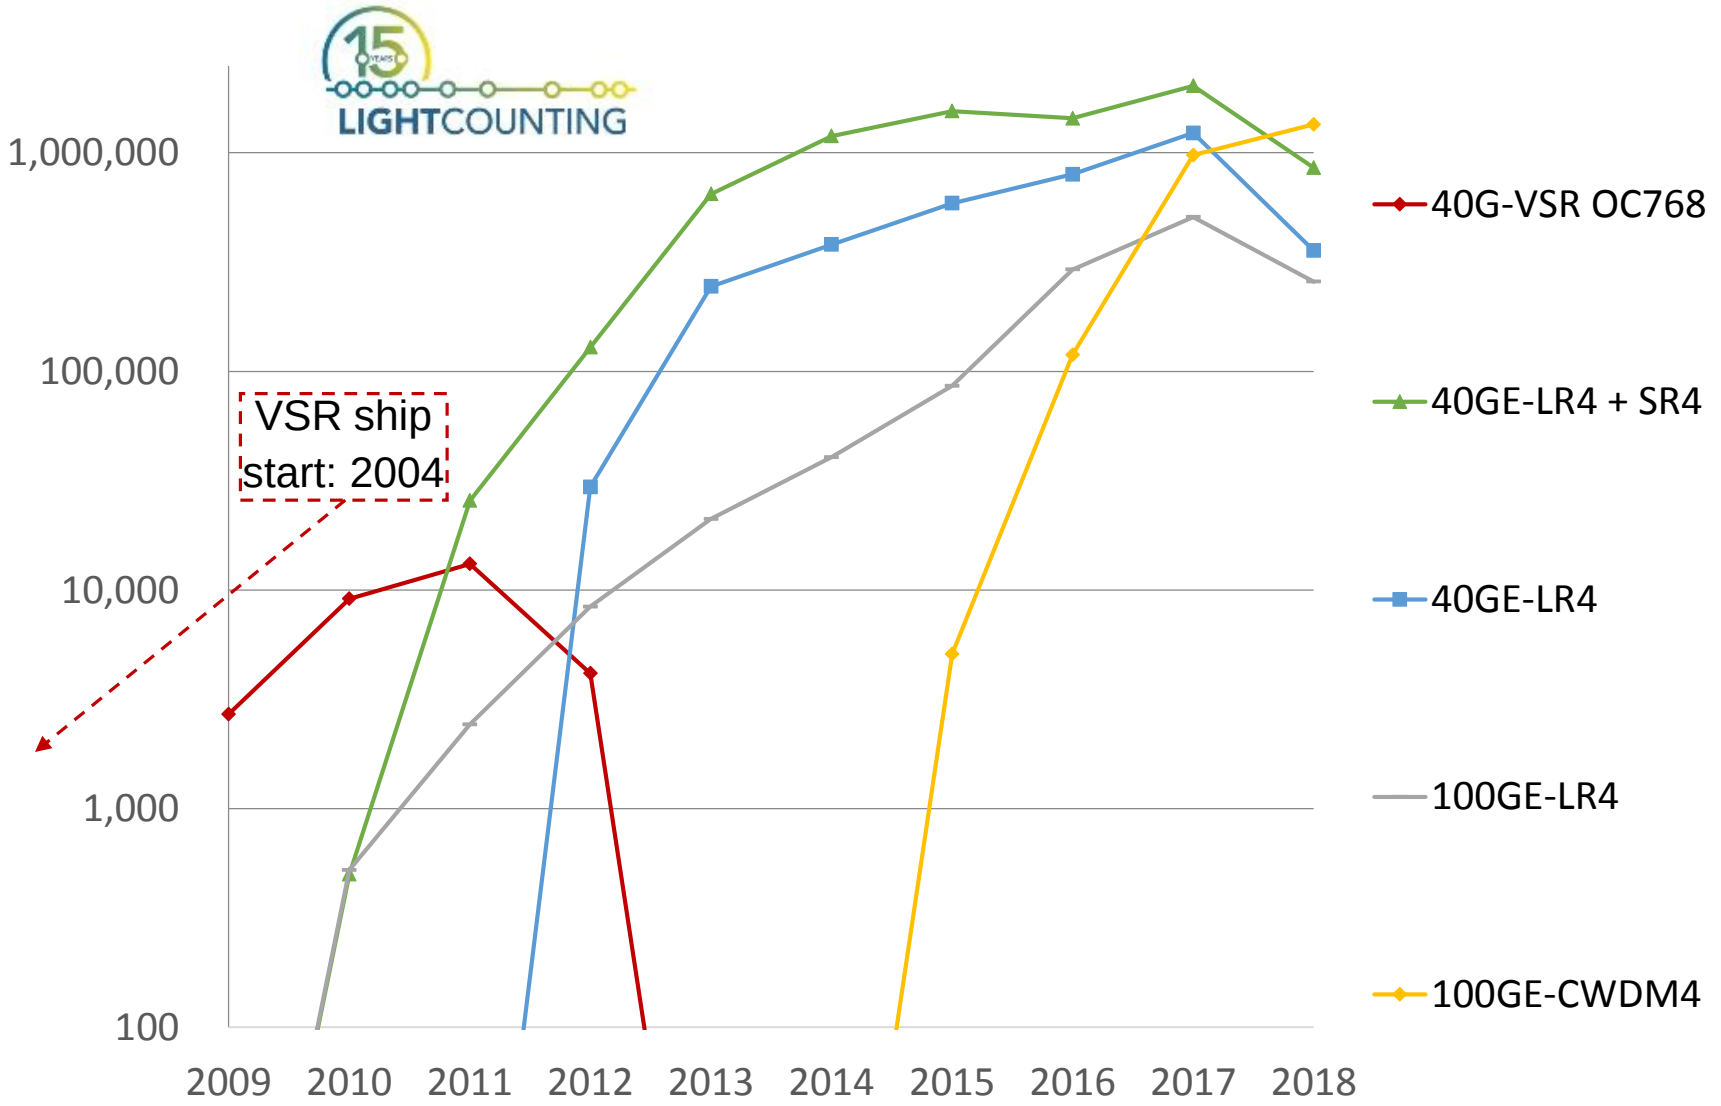

10G Shipment Data Milestones

- 1999: 10G Telecom start; OC192 300-pin
  - architecture: 622M ASIC I/O → 10G λ rate, 16:1 gearbox
  - BW intensive low-volume core/transport apps.
- 2002: 10GE Datacom start; 10GE LR, SR XENPAK, X2
  - architecture: 2.5G ASIC I/O → 10G λ rate, 4:1 gearbox
  - moderate ramp
- 2006: 10GE Datacom volume start; 10GE LR, SR SFP+
  - 480M switch ASIC: 48x radix
  - architecture: 10G ASIC I/O → 10G λ rate, no gearbox
  - steep ramp in datacenters
- 2007: 1st millionth 10GE port

40G Shipment Data Milestones

- 2004: 40G Telecom start; 40G VSR OC768 300-pin
  - architecture: 2.5G ASIC I/O → 40G λ rate, 16:1 gearbox
  - BW intensive low-volume core/transport apps.
- 2010: 40GE Datacom start; 40GE LR4, SR4 CFP
  - architecture: 10G ASIC I/O → 10G λ rate, no gearbox
  - moderate ramp
- 2011: 40GE Datacom volume start; 40GE LR4, SR4 QSFP+
  - 1.28G switch ASIC: 32x radix
  - architecture: 10G ASIC I/O → 10G λ rate, no gearbox
  - steep ramp in large datacenters
- 2014: 1st millionth 40GE port

100G Shipment Data Milestones

- 2010: 100G Telecom start; 100GE LR4 CPF
  - architecture: 10G ASIC I/O → 25G λ rate, 10:4 gearbox
  - BW intensive low-volume core/transport apps.
- 2014: 100G Datacom start: 100GE LR4 CFP2, CPAK
  - architecture 10G ASIC I/O → 25G I/O rate, 10:4 gearbox
  - moderate ramp
- 2015: 100G Datacom volume start; 100GE CWDM4, SR4, PSM4 QSFP28
  - 3.2T switch ASIC: 32x radix
  - architecture: 25G ASIC I/O → 25G λ rate, no gearbox
  - steep ramp in large datacenters
- 2017: 1st millionth 100GE port
- 2019: 100GE FR QSFP28, 4x100GE DR QSFP-DD
  - 12.8T switch ASIC: 128x radix

200G Shipment Extrapolation Milestones

- 2018: 200G Datacom start (see 2019)
- 2019: 200G Datacom volume start (2x200G FR4+ OSFP, 200GE FR4 QSFP56)
  - 12.8T switch ASIC: 64x radix
  - architecture: 50G ASIC I/O → 50G λ rate, no gearbox
  - steep ramp in large datacenters
- 2021: 1st millionth 200GE port
- 25.6T switch ASIC: 128x radix

400G Shipment Extrapolation Milestones

- 2019: 400G Telecom start; 400GE LR8 CFP8, 400GE LR8, FR4, LR4 OSFP, QSFP-DD
  - architecture: 50G ASIC I/O → 50G λ rate, no gearbox
  - architecture: 50G ASIC I/O → 100G λ rate, 2:1 gearbox
  - BW intensive low-volume core/transport apps.
- 2020: 400G Datacom start; 400GE FR4 OSFP, QSFP-DD
  - 12.8T switch ASIC: 32x radix
- 2021: 400G Datacom volume start; 2x400GE-FR4+ OSFP, 400GE-FR4+ QSFP112
  - 25.6T switch ASIC: 64x radix
  - architecture: 100G ASIC I/O → 100G λ rate, no gearbox
  - steep ramp into datacenters
- 2023: 1st millionth 400GE port

10G, 40G, 100G, 200G, 400G Ship Milestones

Gb/s	1	2	3	4	5	6	7	8	9	10	11
10	1999 TS			2002 DS				2006 DVS	2007 1st M		
40	2004 TS						2010 DS	2011 DVS			2014 1st M
100			2010 TS				2014 DS	2015 DVS		2017 1st M	
200							2018 DS	2019 DVS		2021 1st M	
400						2019 TS	2020 DS	2021 DVS		2023 1st M	

TS: Telecom Start

DS: Datacom Start

DVS: Datacom Volume Start

1st M: 1st Millionth xGE port

200G, 400G Forecast

specified bit rate	1st millionth pluggable Ethernet optical port with specified bit rate	1st millionth pluggable optical module with total specified bit rate
200Gb/s	2021 (early)	2021 (late)
400Gb/s	2023	2022 (early)

40GE & 100GE Standardization Background

- In July 2006 IEEE started studding follow-on rate to 10Gb/s
- 100Gb/s was the only serious contender
- In Jan. 2007 an intense 6 month debate started whether 40Gb/s is better for Datacenter

40GE, 100GE Standardization Background

- 100Gb/s pro arguments
  - 10x rate step minimizes deployment/operational cost by eliminating intermediate rate step
  - Investment focus on 25GB technology will lead to lower optics cost in the long-term
- 40Gb/s pro arguments
  - Mature 10GB technology ready for near-term, low-cost, low-risk, high-volume deployment
  - 1.28T switch ASIC radix:
    - 100G: 12x
    - 40G: 32x
- Both were adopted as IEEE 802.3 Ethernet rates

200GE, 400GE Standardization Background

- 400Gb/s pro arguments
  - 4x rate step minimizes deployment/operational cost by eliminating intermediate rate step
  - Investment focus on 100GB technology will lead to lower optics cost in the long-term
- 200Gb/s pro arguments
  - Mature 25GB technology ready for near-term, low-cost, low-risk, high-volume deployment
  - 12.8T switch ASIC radix:
    - 400G: 32x
    - 200G: 64x
    - 100G: 128x
- Both were adopted as IEEE 802.3 Ethernet rates

ECOC 2015 Cloud Datacenters Panel (C. Cole)

Ethernet Data Rates Standardized in IEEE

- Existing rate progression, Gb/s:
  - 1 → 10 → 40 → 100
- Resulting rate progression Gb/s:
  - 10 → 25 → 50 (& 40) → 100 → 200 → 400
- Rates after 400 (crystal ball,) Gb/s:
  - 800
  - 1600

7 March 2019 Prediction

Beyond 400G for Hyperscale Data Centers: 800G → 1600G

Before 400G for Hyperscale Data Centers

Thank You