

Broad Market Potential & Economic Feasibility:

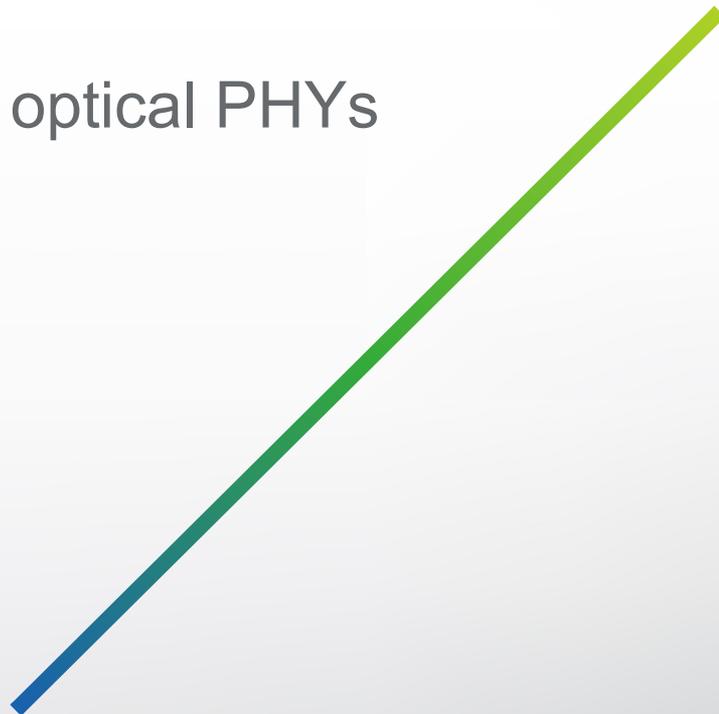
IEEE 802.3 100 Gb/s per lane optical PHYs
Study Group

Authors:

Justin Abbott (Lumentum)

Brian Welch (Luxtera)

Hai-Feng Liu (Intel)



Supporters

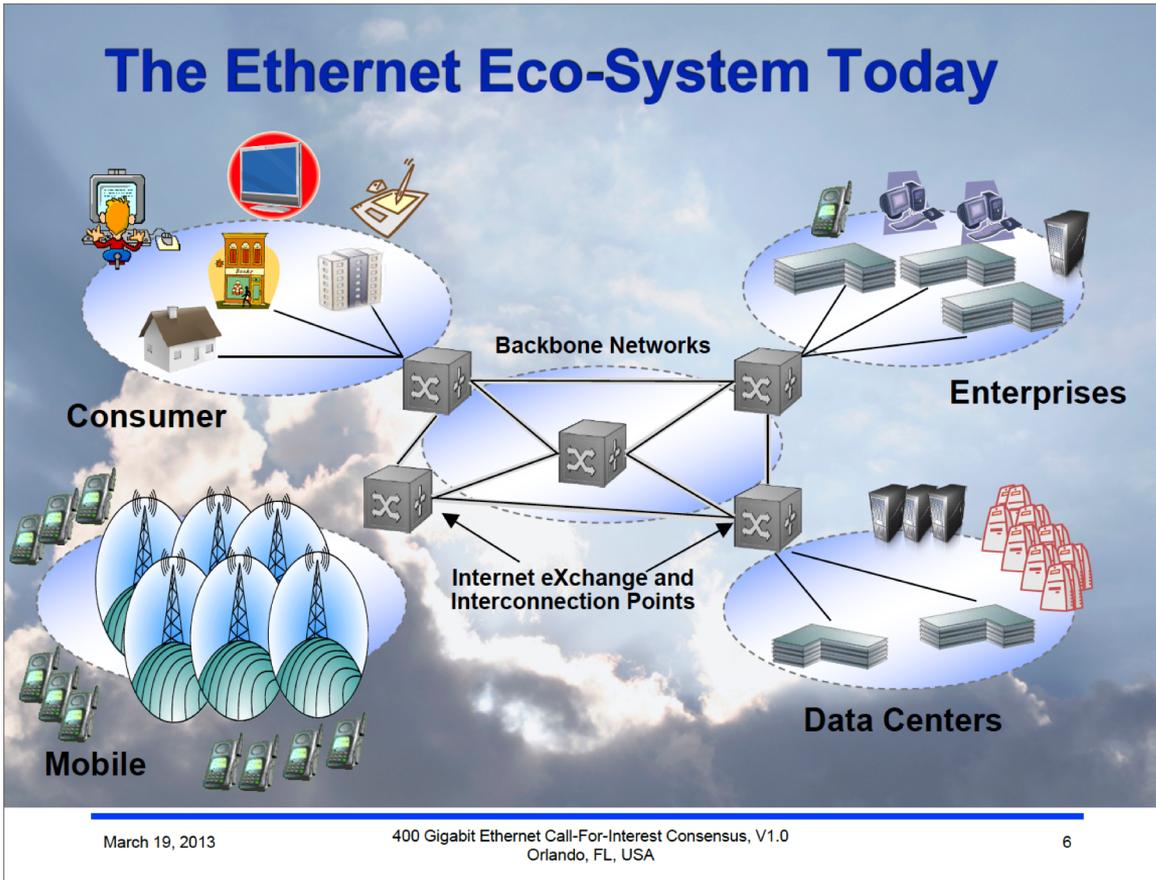
- ...

- ...

Proposed PHY Objectives

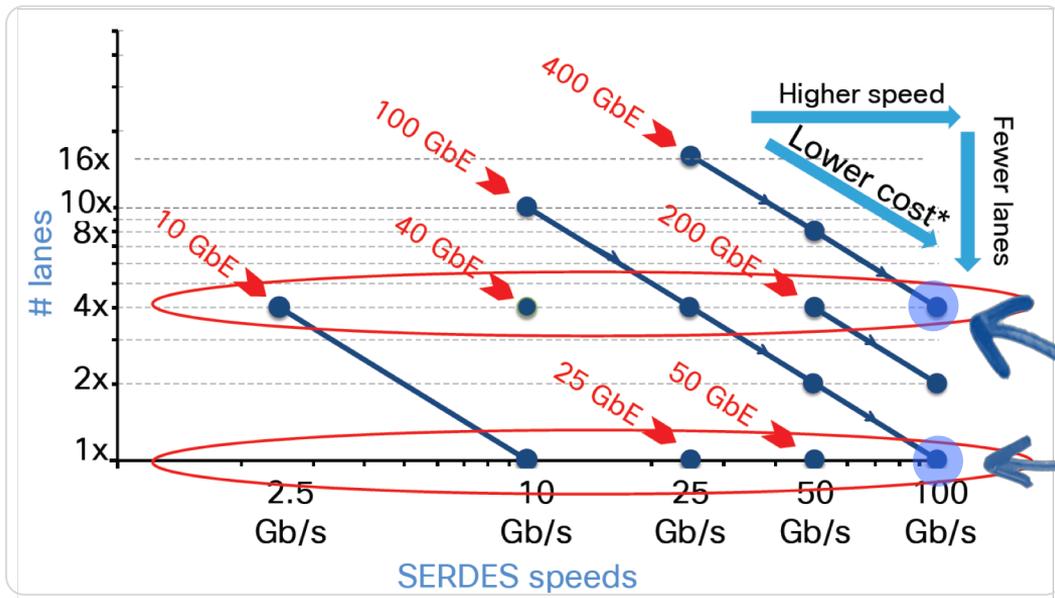
- Define single-lane 100 Gb/s PHYs for operation over
 - SMF with lengths up to at least 2 km
 - SMF with lengths up to at least 10 km
- Define four-lane 400 Gb/s PHYs for operation over
 - SMF with lengths up to at least 2 km
 - SMF with lengths up to at least 10 km

2 km & 10 km optics dominate throughout SMF ecosystem



- Key Application
 - Telecom Client Interconnect
 - Enterprise
 - Hyperscale
- Advance in technologies allows for lane reductions to improve cost savings
- Aggressive 100G cost reduction is occurring now
 - Will continue as the market searches for the lowest cost solution
 - Key focus on < 10km reaches
- 400G market adoption about to start as networks are exceed capacity and density that 100G and Nx100G can provide
 - New technologies enable lower cost solutions than are currently standardized in Ethernet

Ethernet Trend



http://www.ieee802.org/3/cfi/1118_1/CFI_01_1118.pdf

Trend

1x and 4x bus widths dominate volume adoption

Begin with multi lane, evolve towards single lane

Driven by new technologies that enable cost reductions and higher density

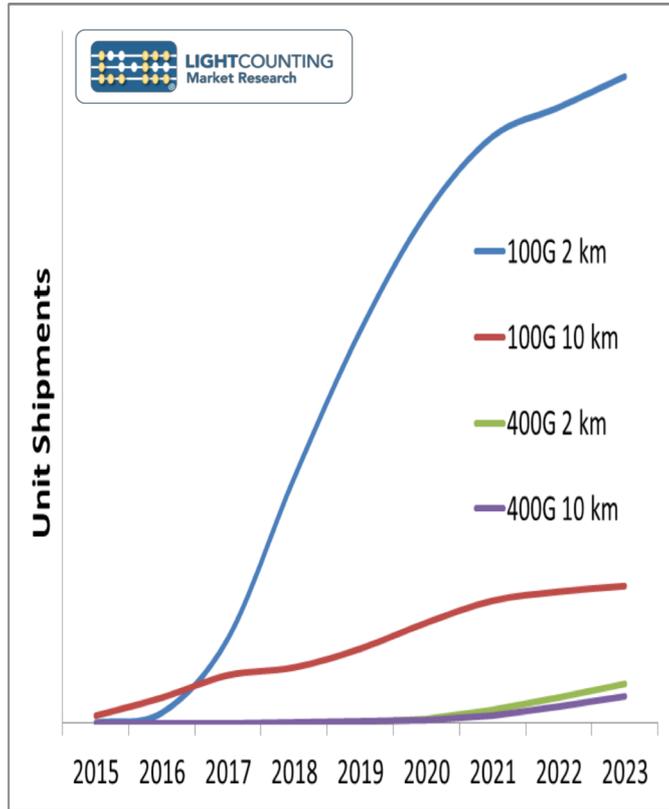
100G / 400G Lane Reduction

Lane reduction to lower cost and improved density

400G: 8 lanes → 4 lanes

100G: 4 lanes → 1 lane

Ethernet Market



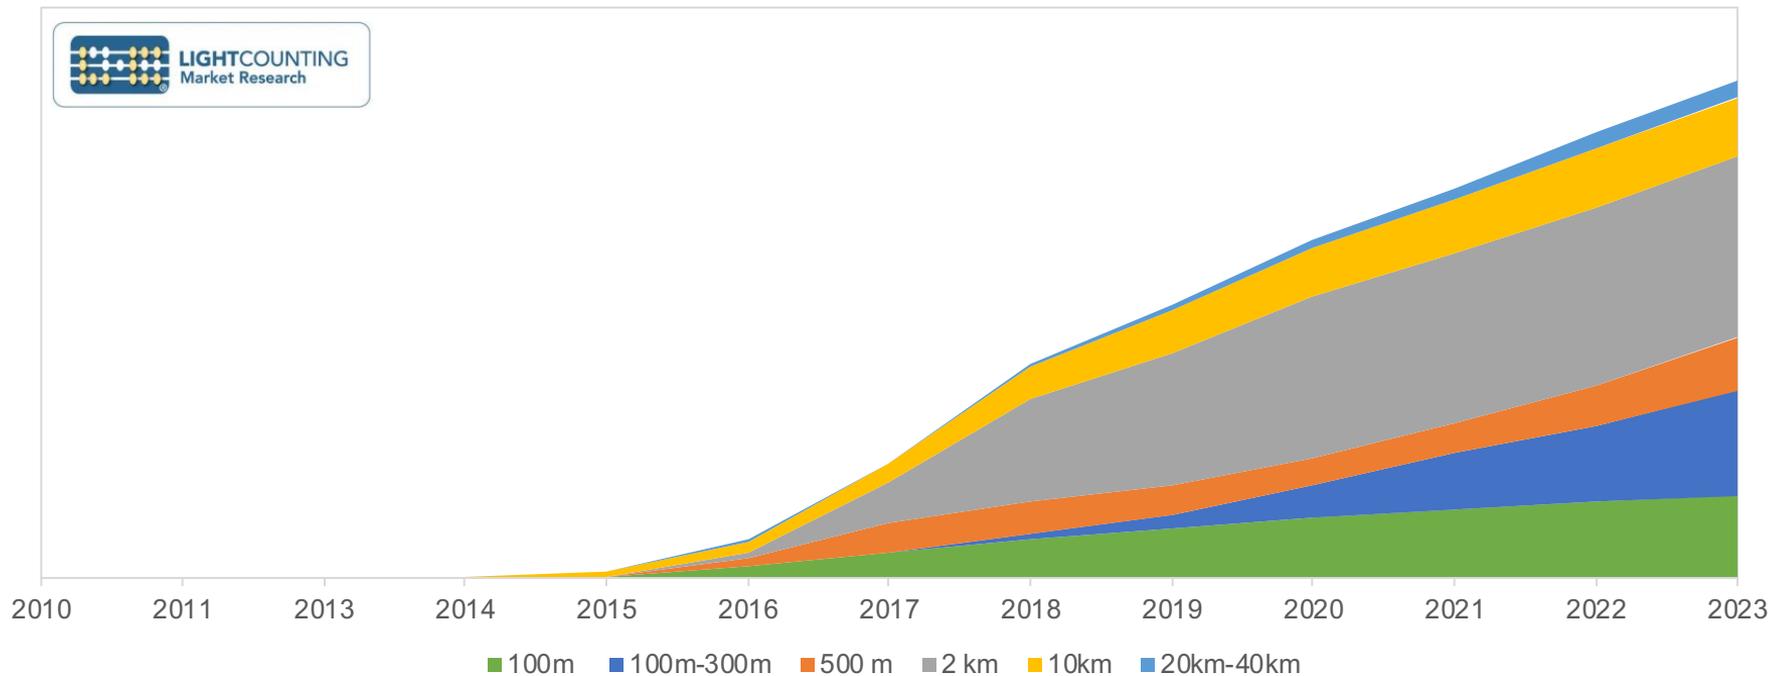
Courtesy Dale Murray, Light Counting

Comments

- Strong market traction for 100G 2km and 10km
 - 2km based on 100G CWDM4 MSA
 - 10km based on IEEE LR4
- 100G 2km volume ramp coincided with the availability of x4 lane CWDM4 MSA interfaces
- Aggressive 100G cost reduction continuing as volume increases
 - Will continue as the market searches for the lowest cost solution
 - Key focus on < 2km reaches
- 400G ecosystem beginning to emerge
 - 12.8TB switch ASIC's sampling
 - 32 port hardware switches demo'ing
 - 400G x4 lane 2km optical modules demo'ing
- 400G interconnects expected to surpass 4M in 2023
- Existing Ethernet standards do not cover targeted 400Gbps x4 optical interconnect

Broad Market Potential: 100G

Reach Share by Year



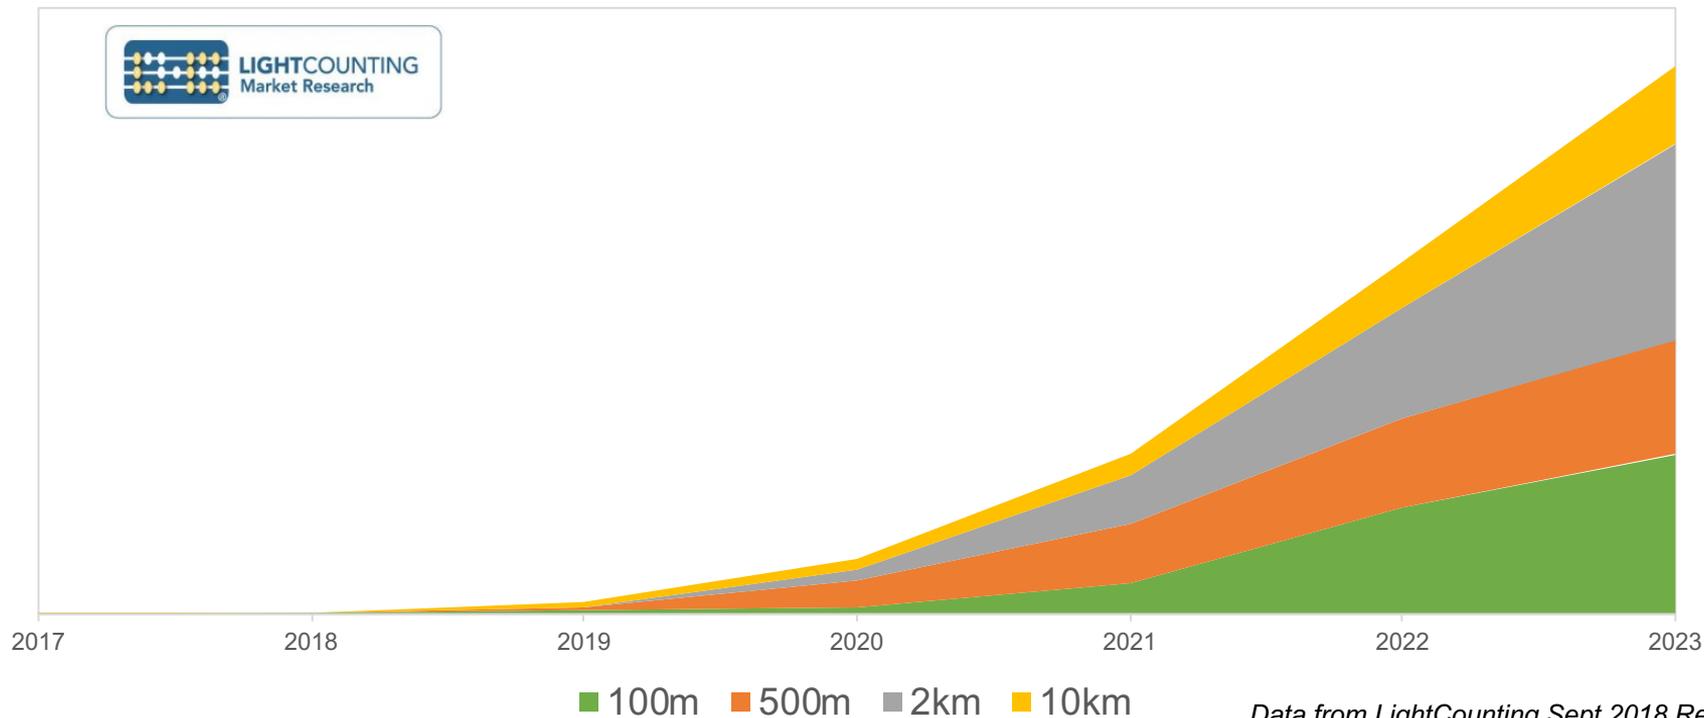
Data from LightCounting Sept 2018 Report

Broad Market Potential: 100G

- 2km and 10km reaches represent significant volume for optical interconnects
 - Through 2023 2km expected to be > 40% of all 100GE (by volume)
 - Through 2023 10km expected to almost 15% of all 100GE (by volume)
- Currently 100GE at 2km and 10km serviced by 25G/Lane optical solutions
 - 2km: 100G-CWDM4 (MSA)
 - 10km: 100GBase-LR4 (IEEE), 100G-4WDM10 (MSA)
- Transition to higher host rates (50G and 100G serdes) will favor faster lane rates
 - Currently two 100G solutions > 25G/lane, both shorter reach: 100GBase-SR2, 100GBase-DR

Broad Market Potential 400G

Reach Share by Year

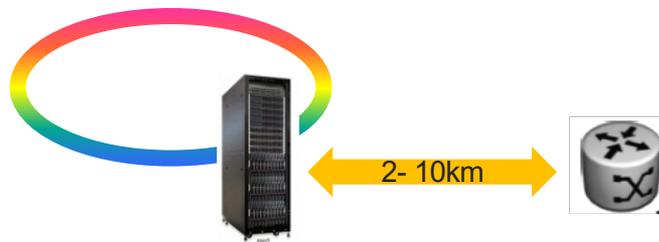


Data from LightCounting Sept 2018 Report

Application: Telecom / Enterprise

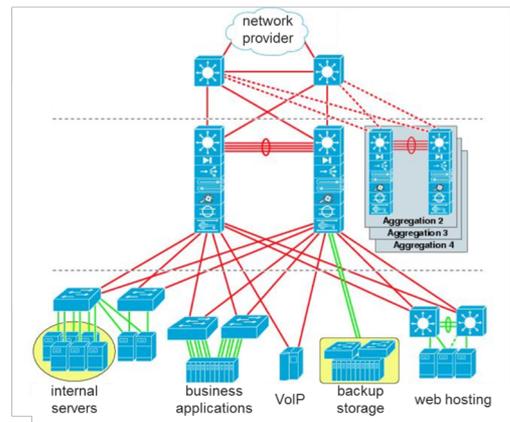
Telecom Client Interconnect

DWDM System
Carrying Ethernet Traffic



- Interface between different supplier and/or operator networks
 - need well-defined, measurable, IEEE interoperability standard
- < 2-10km interface: primarily 100G LR4 today
- 400G FR8/LR8 provides a capacity upgrade, but newer technology can enable relatively lower cost x4 lane interconnects

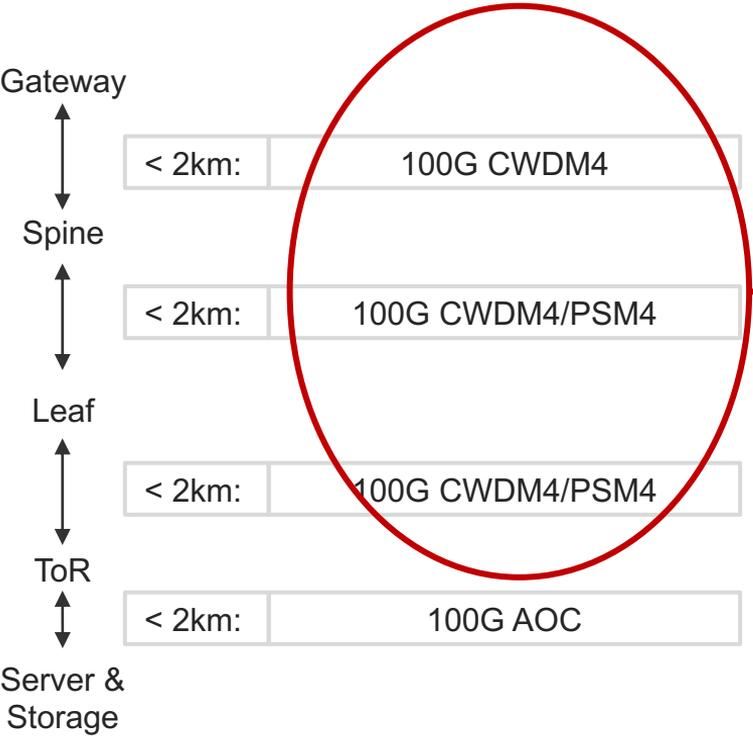
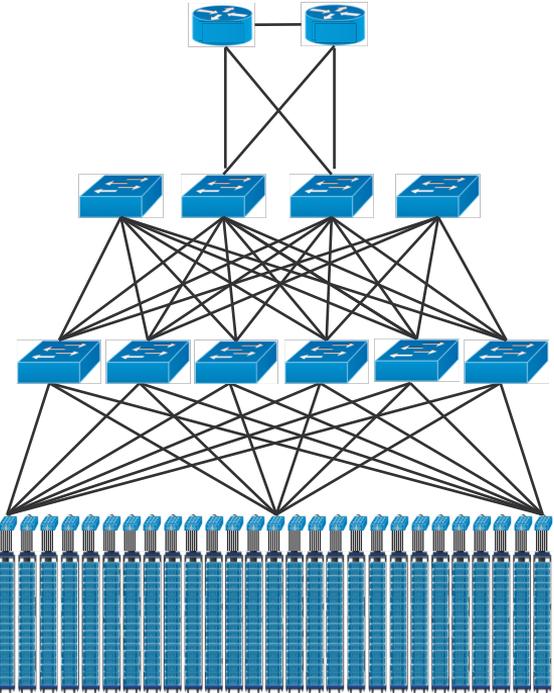
Enterprise



- Modest scale and large variations in scale
- Functions often not virtualized but use dedicated hardware
 - Requires standardized interfaces
- Primarily < 100m, but also requires <2km
 - 100G x1 lane provides relative lower cost / higher density upgrade path for <2km

Application: Hyperscale

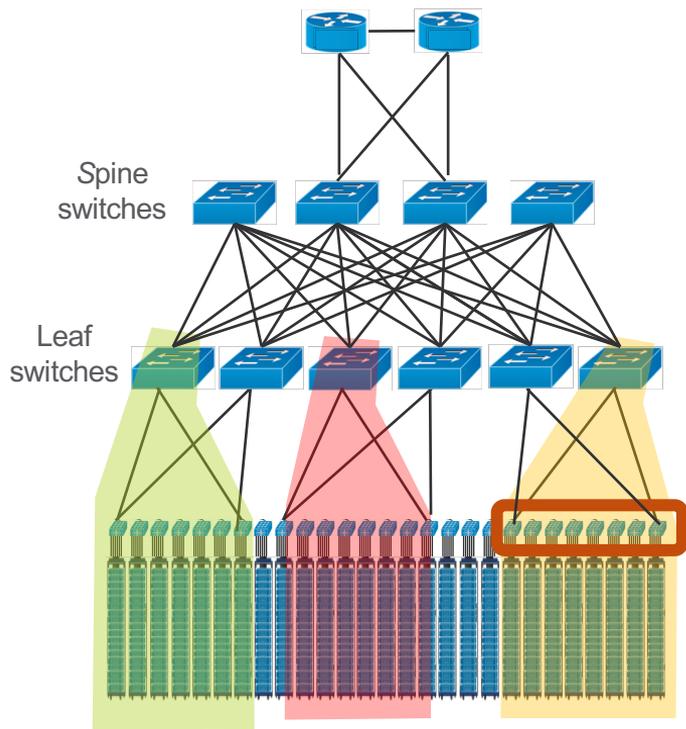
Example of Hyperscale Architecture



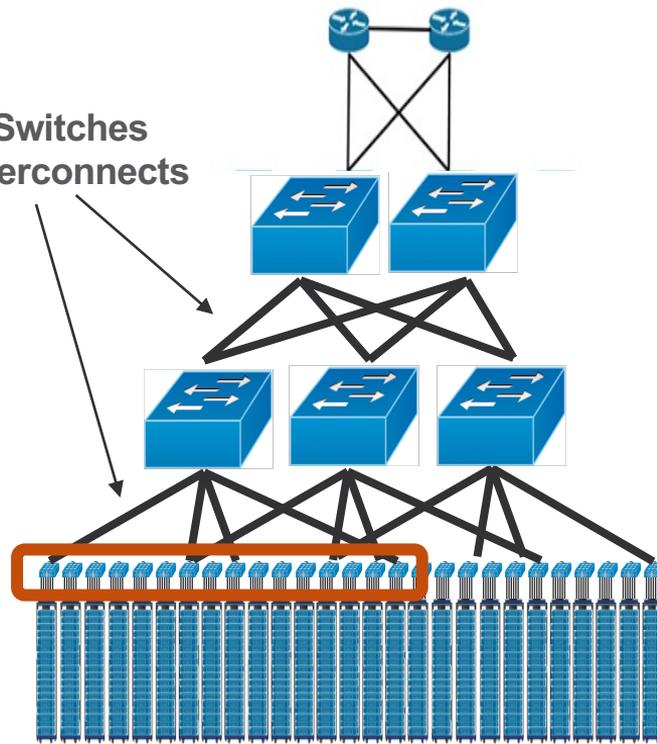
Low cost higher capacity
Interconnected needed
Subject of this
Project
100G 2km x1 lane
400GbE 2km x4 lane

Value of 400G: Evolution to Flatter and Wider Networks

Larger Capacity Switches with Higher Capacity Ports

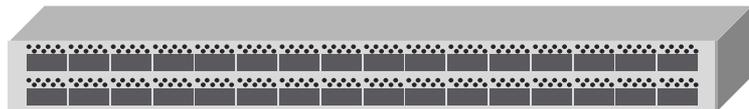
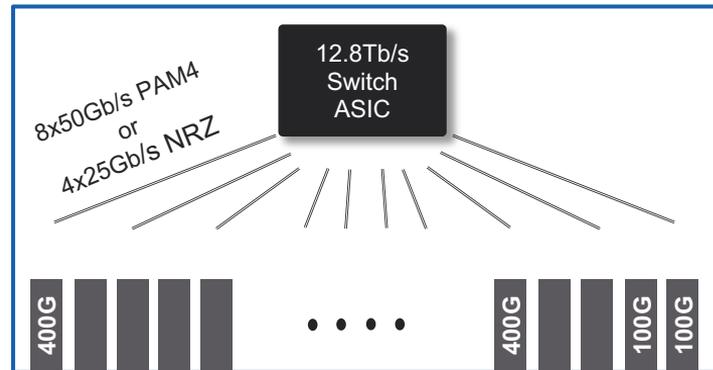
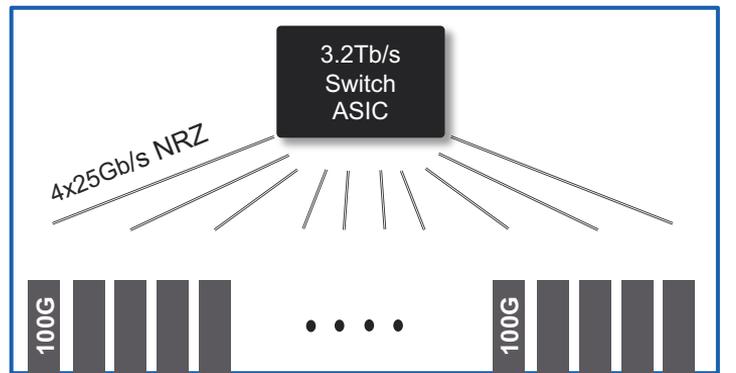


Larger, Fewer Switches
Fewer Optical Interconnects



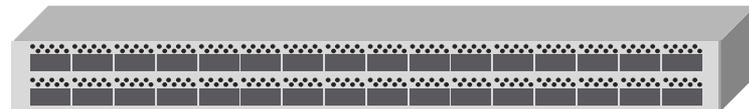
Network savings enabled by higher capacity interconnect links (e.g. 400G)

Ethernet Ecosystem to Support 400G and 100G



32 x 100GbE

3.2Tb/s switch ASIC available today from multiple vendors



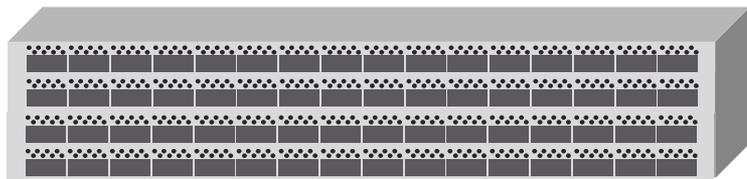
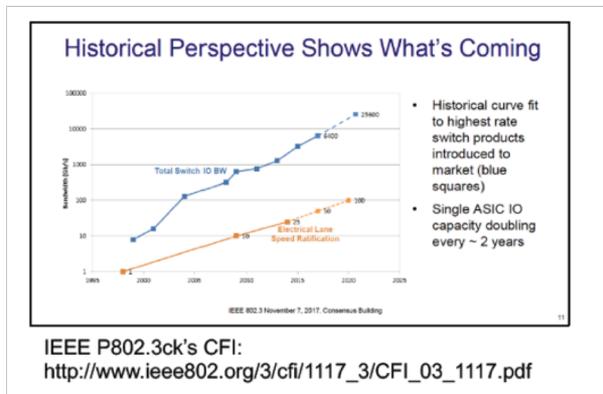
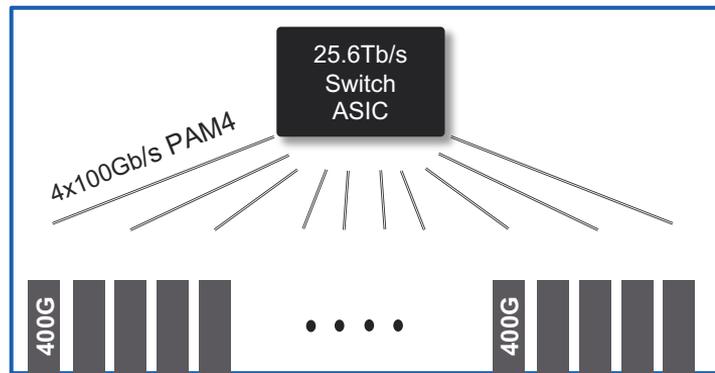
32 x 400GbE
With support for 100G

12.8Tb/s Switch ASIC's sampling today from multiple vendors

400G optics match switch silicon throughput in a 32x400G 1RU
Electrical interface also supports 100G

Supporting ecosystem exist for 400G

Matching ASIC IO to Module IO



64 x 400GbE

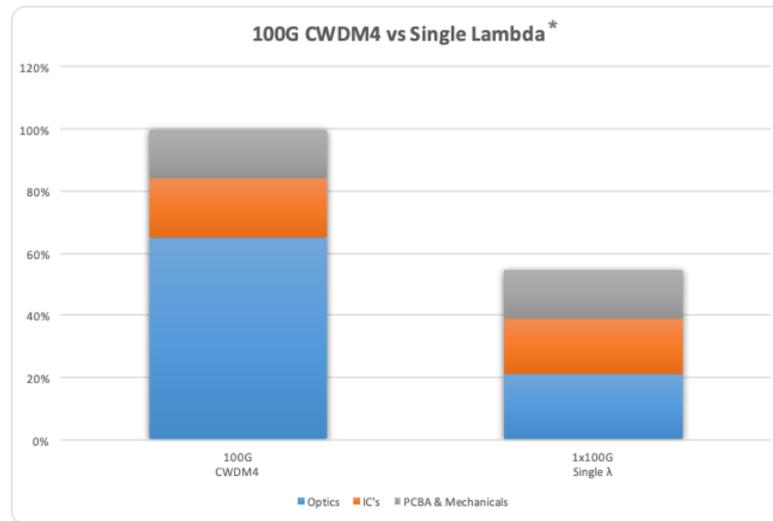
- ASIC IO is increasing
- Potential to leverage IEEE 802.3ck (100G electrical)
- Optical module is simplified when module IO matches ASIC IO
 - 100G → 100G re-timer can be used
 - Complicated gearbox / reverse gearbox not needed
 - 400G FR4 / LR8 would require a more complicated reverse gearbox

Relative cost reduced when Module IO and ASIO IO are aligned

Economic Feasibility - Relative Cost Comparisons (4λ vs 1λ)

- Optics
 - > 50% of the 4λ optical module cost is due to the optics
 - Significant savings are realized in moving from 4λ to 1λ
 - Assembly, test, alignment, yield
- IC's
 - Relative cost remains nearly constant in moving from 4λ to 1λ
 - Advanced nodes require significant initial investment, but can be amortized over the high volume 100G market
- PCBA & Mechanicals
 - Relative cost remains nearly constant in moving from 4λ to 1λ

Transition from 4λ to 1λ results in relative cost reduction > 40%



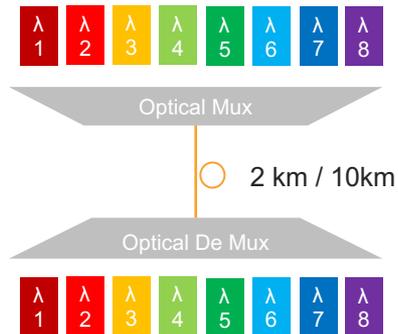
* Based on Lumentum data

Note

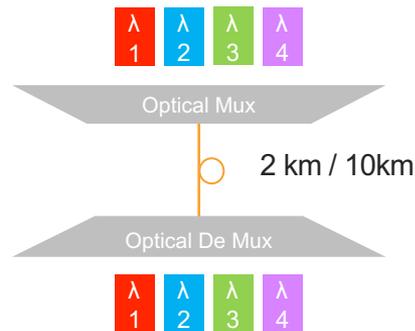
- CY2019 Relative Cost
- 4λ assumes CWDM4 implementation
- Comparison is at similar volumes

Reduced Optical Lane Count → Reduced Relative Cost

400GBASE FR8 / LR8



x2 Optical Lane Reduction



- As with 100G $4\lambda \rightarrow 1\lambda$, reducing the 400G optical lane count is expected to minimize the overall relative cost
- Further savings may be realized as moving from 8 optical lanes to 4 opens an option to relax the wavelength grid

Broad Market Potential

- **Each proposed IEEE 802 LMSC standard shall have broad market potential. At a minimum, address the following areas:**
 - **Broad sets of applicability**
 - **Multiple vendors and numerous users**
-
- Ethernet is widely deployed in telecom client interconnects and in switch-to-switch applications in hyperscale and enterprise data centers. The forecasted demand for 100 GbE and 400 GbE interconnects in these applications is approaching 12M modules per year by 2023 assuming continuous cost reduction from current solutions.
 - There is a significant opportunity to address this demand with 100 Gb/s per wavelength technologies that allow for reduced lane count, increased density, and lower cost 100 Gb/s and 400 Gb/s modules.
 - The call-for-interest for “100 Gb/s per lane optical PHYs” had 80 participants. Approximately 55 individuals affiliated with at least 43 companies indicated they would support the standardization process. It is anticipated that there will be sufficient participation to effectively complete the standardization process including individuals from end-users, equipment manufacturers and component suppliers.

Economic Feasibility

Each proposed IEEE 802 LMSC standard shall provide evidence of economic feasibility. Demonstrate as far as can reasonably be estimated, the economic feasibility of the proposed project for its intended applications. Among the areas that may be addressed in the cost for performance analysis are the following:

- a) **Balanced costs (infrastructure versus attached stations).**
- b) **Known cost factors.**
- c) **Consideration of installation costs.**
- d) **Consideration of operational costs (e.g., energy consumption).**
- e) **Other areas, as appropriate**

-
- The cost factors for Ethernet components are well known.
 - Reasonable cost for the resulting performance will be achieved in this project as established by prior experience in the development of:
 - Ethernet optical specifications based on 100 Gb/s per wavelength PMDs including 400GBASE-DR4 and 100GBASE-DR.
 - Non-Ethernet optical specifications for 2 km and 10 km SMF reaches at 100 Gb/s and 400 Gb/s.
 - In consideration of installation costs, the project is expected to use proven and familiar media consistent with industry deployments, namely single-mode fiber.
 - 100 Gb/s per wavelength PMDs reduce the number of optical transmit/receive components by up to 4 times compared to 25 Gb/s and 50 Gb/s per wavelength solutions which lowers overall network power consumption. For 100 Gb/s the need for optical mux/demux components is eliminated which further reduces cost and power consumption of these PMDs.
 - Network design, installation and maintenance costs are minimized by preserving network architecture, management and software.