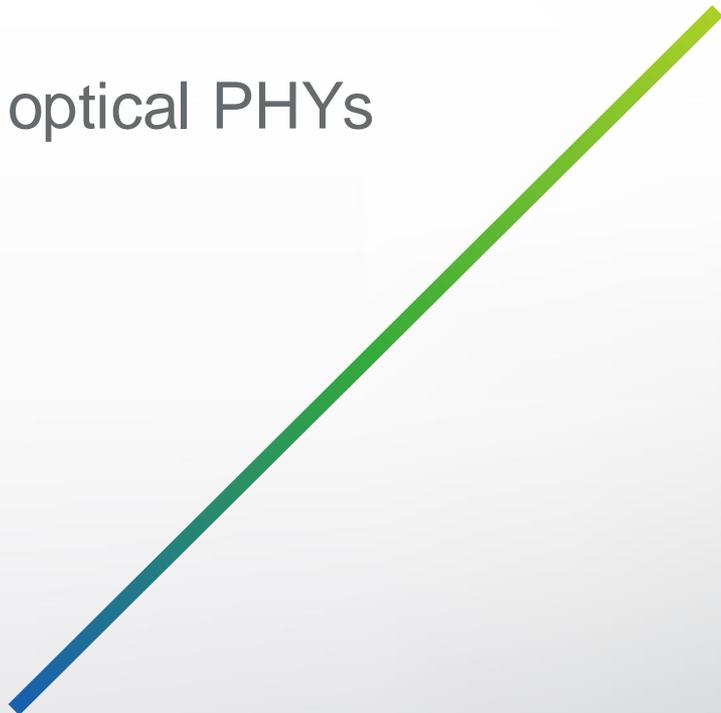


Broad Market Potential & Economic Feasibility:

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

Authors:

Justin Abbott (Lumentum)
Brian Welch (Luxtera)
Hai-Feng Liu (Intel)
David Lewis (Lumentum)



Supporters

- ...

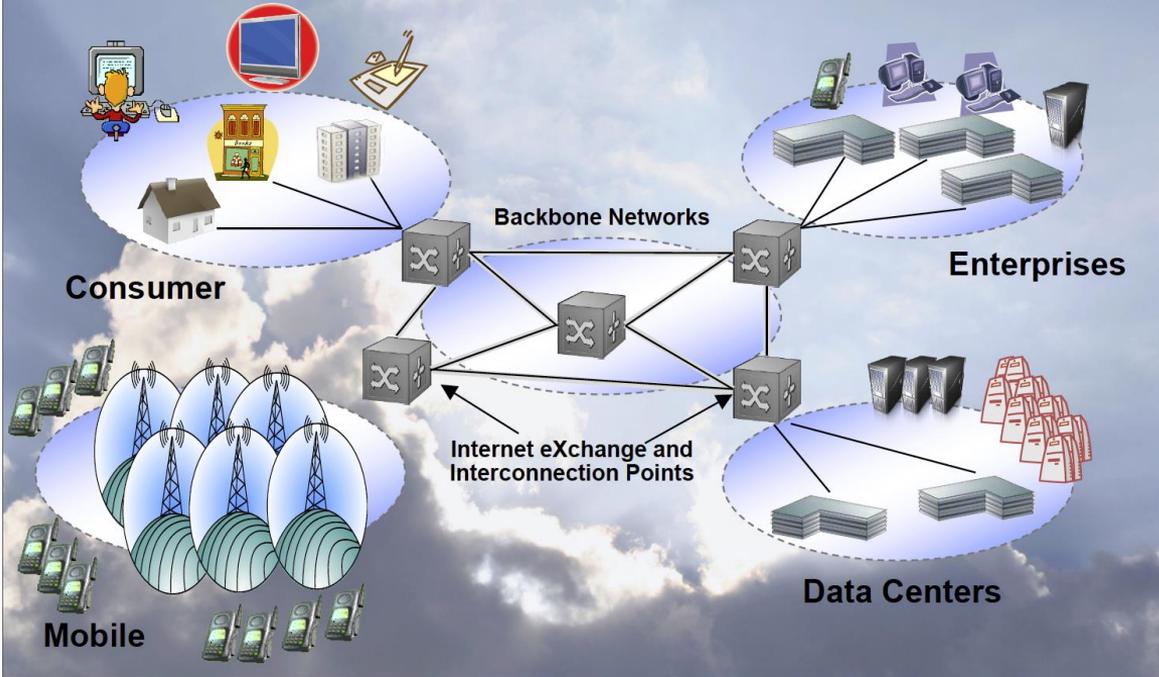
- ...

Proposed PHY Objectives

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

2 km & 10 km optics dominate throughout SMF ecosystem

The Ethernet Eco-System Today



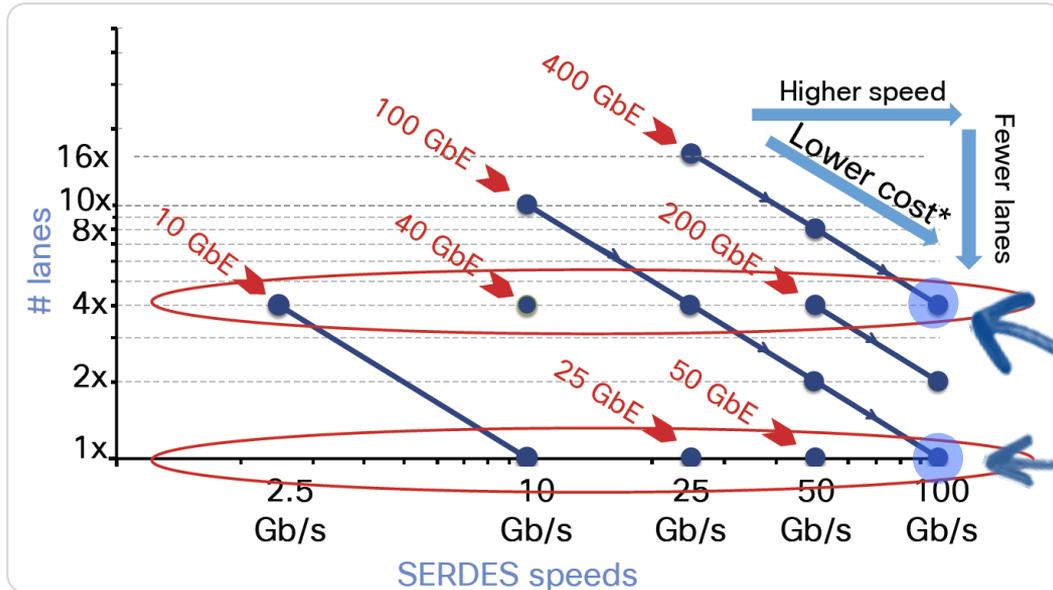
- Key Application
 - Telecom Client Interconnect
 - Enterprise
 - Hyperscale
- Advance in technologies allows for lane reductions to improve cost savings
- Aggressive 100G cost reduction is occurring now
 - Will continue as the market searches for the lowest cost solution
 - Key focus on < 10km reaches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

March 19, 2013

400 Gigabit Ethernet Call-For-Interest Consensus, V1.0
Orlando, FL, USA

6

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

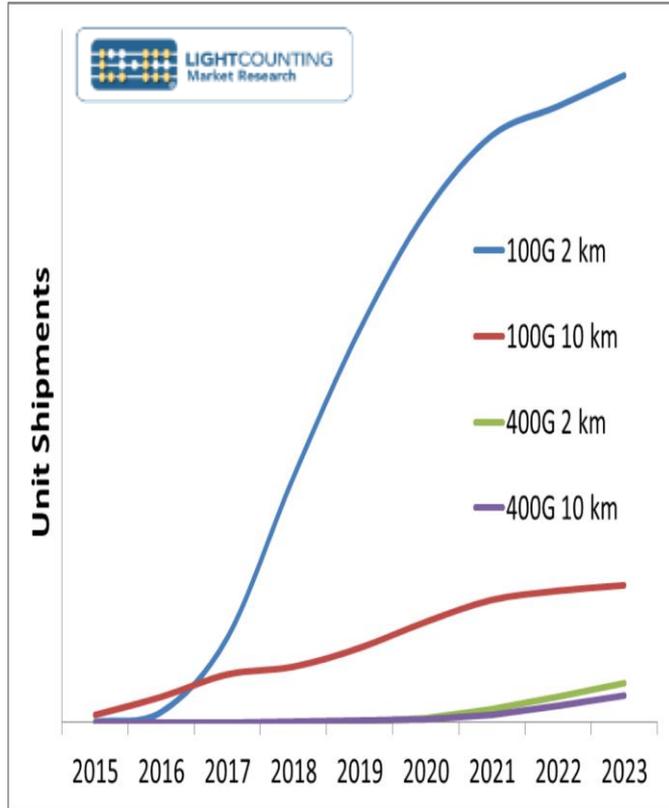
100G / 400G Lane Reduction

Lane reduction to lower cost and improved density

400G: 8 lanes → 4 lanes

100G: 4 lanes → 1 lane

Ethernet Market



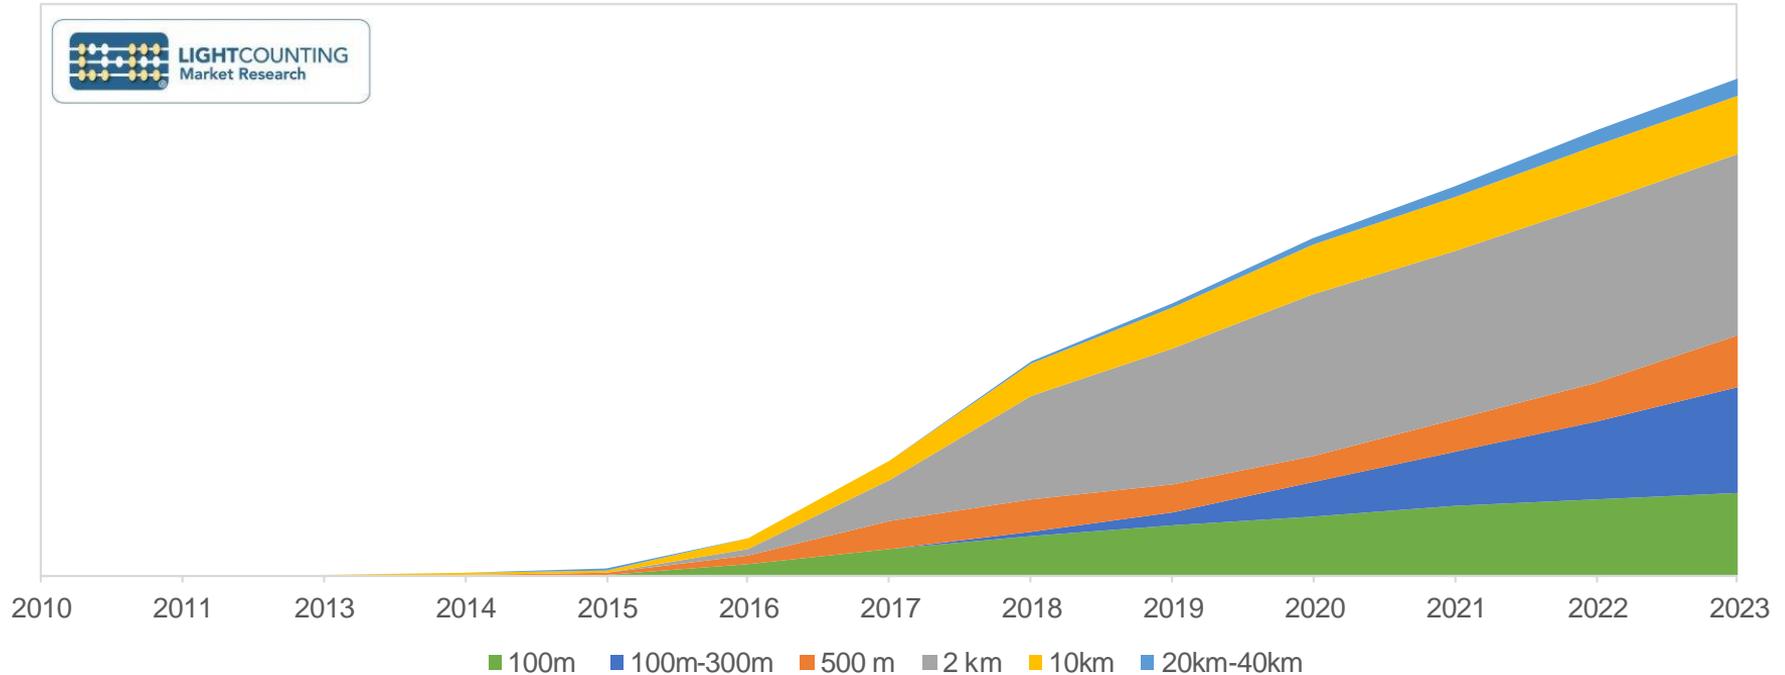
Courtesy Dale Murray, Light Counting

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 ASICs 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



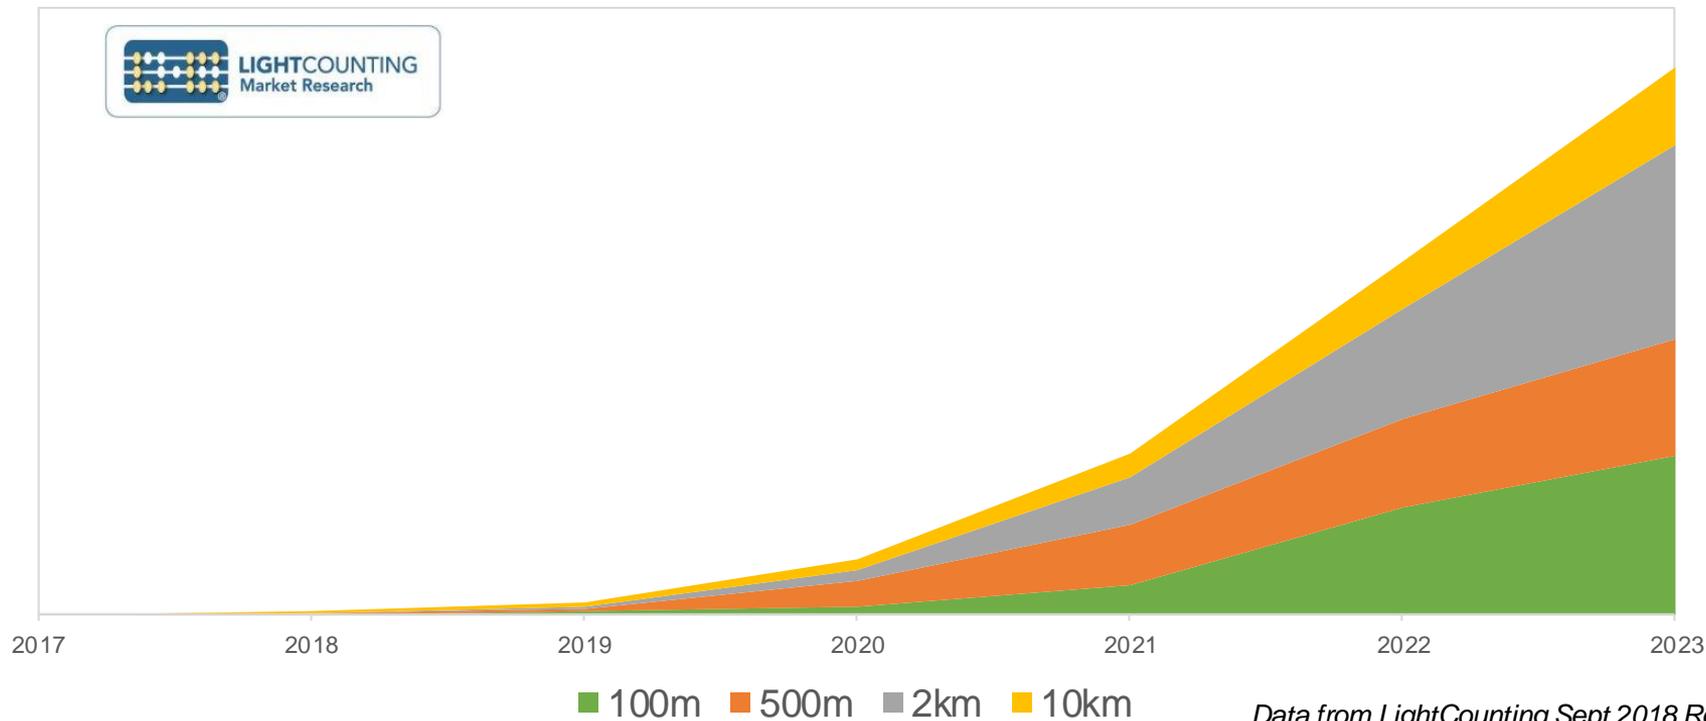
Data from LightCounting Sept 2018 Report

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



Application: Telecom / Enterprise

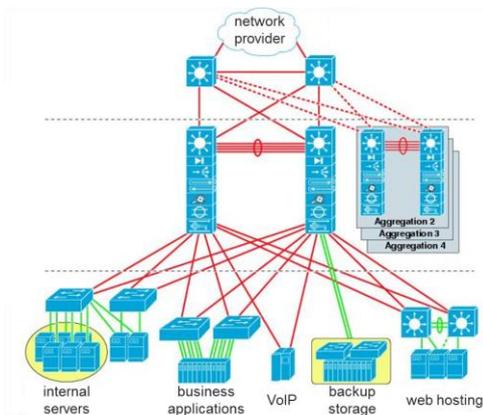
Telecom Client Interconnect

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

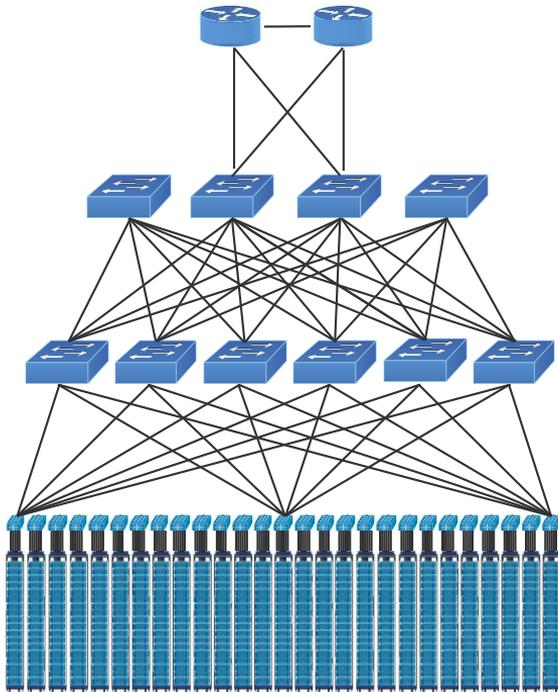
Enterprise



- 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



Gateway



Spine



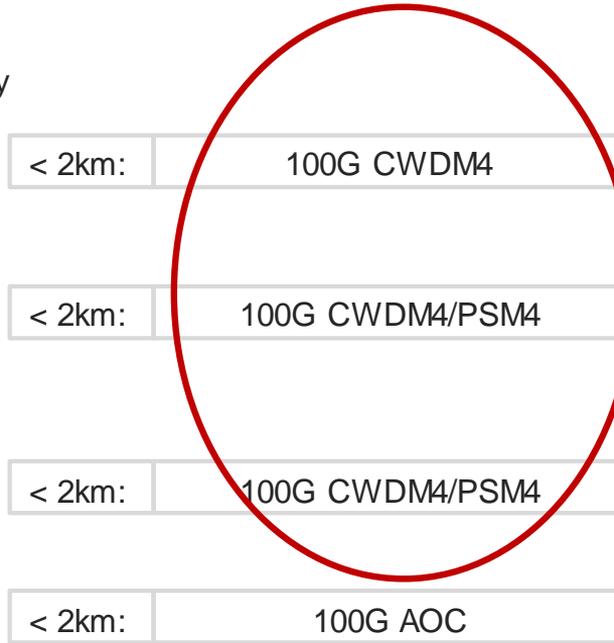
Leaf



ToR



Server &
Storage



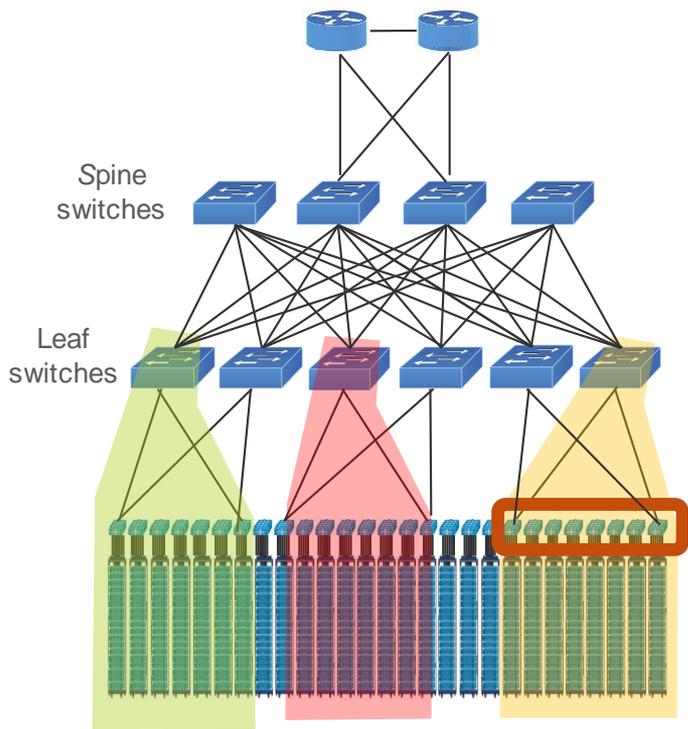
Low cost higher capacity
Interconnected needed

Subject of this
Project

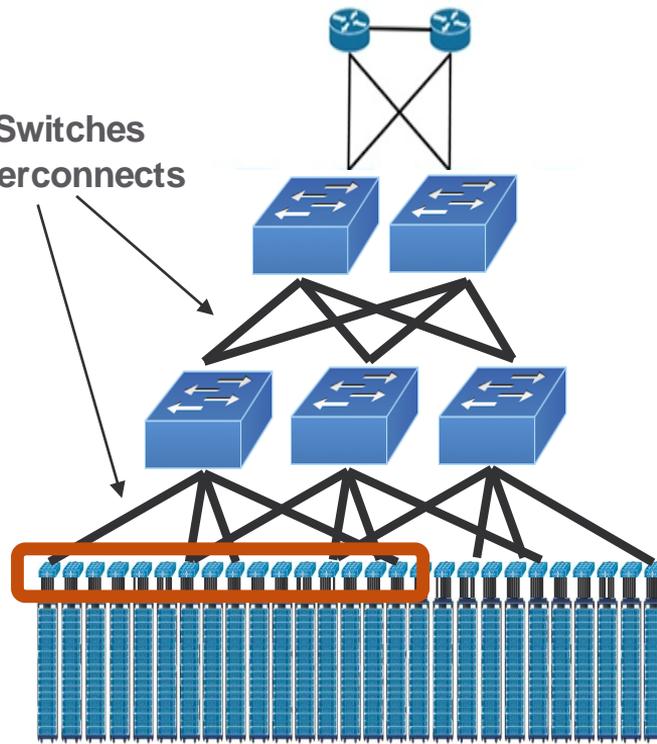
100G 2km x1 lane
400GbE 2km x4 lane

Value of 400G: Evolution to Flatter and Wider Networks

Larger Capacity Switches with Higher Capacity Ports

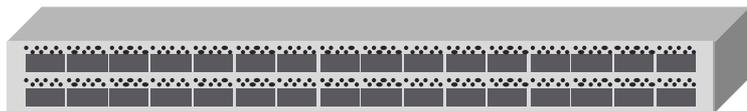
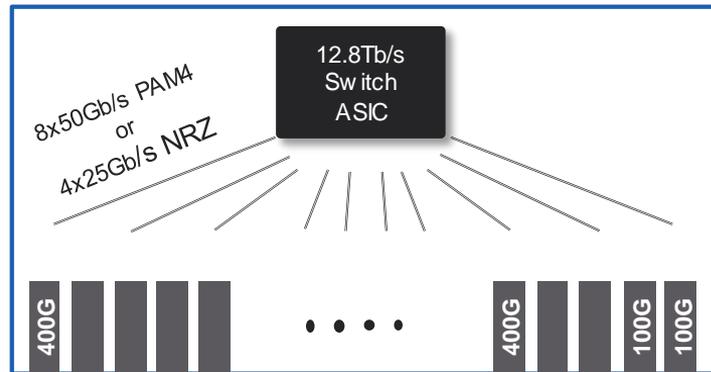
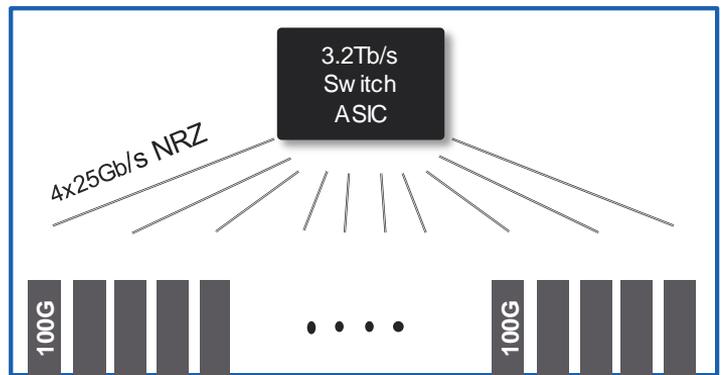


Larger, Fewer Switches
Fewer Optical Interconnects



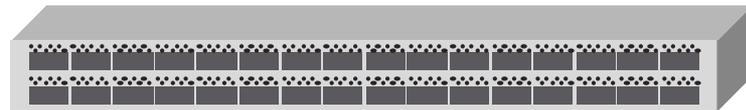
Network savings enabled by higher capacity interconnect links (e.g. 400G)

Ethernet Ecosystem to Support 400G and 100G



32 x 100GbE

3.2Tb/s sw itch ASIC available today from multiple vendors



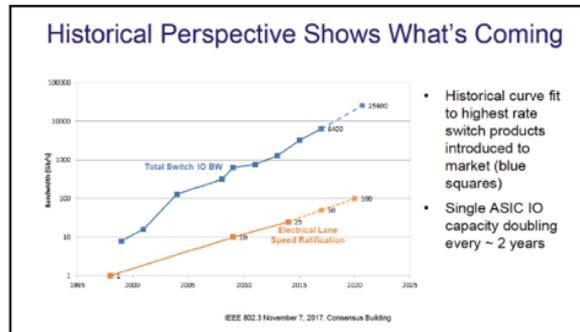
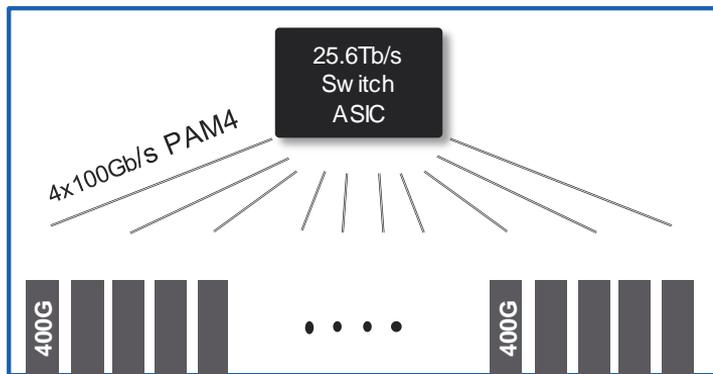
32 x 400GbE
With support for 100G

12.8Tb/s Sw itch ASIC's sampling today from multiple vendors

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

Supporting ecosystem exists for 400G

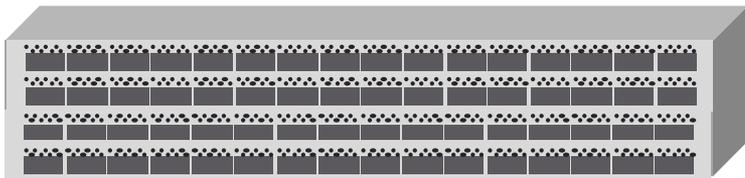
Matching ASIC IO to Module IO



IEEE P802.3ck's CFI:

http://www.ieee802.org/3/cfi/1117_3/CFI_03_1117.pdf

- 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



64 x 400GbE

Relative cost reduced when Module IO and ASIO IO are aligned

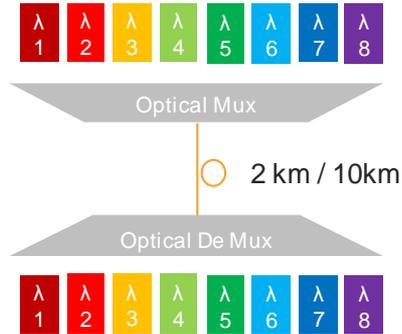
Economic Feasibility – Relative Cost

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

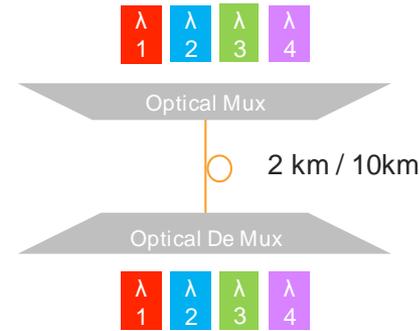
Transition from 4λ to 1λ results in significant relative cost reduction

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

Thank You