200G/wavelength MMF optical PHYs

Call For Interest (CFI) consensus meeting presentation

17 July 2025

IEEE 802.3 NEA Meeting

Today’s Panel

- Speakers
  - Mabud Choudhury, Lightera
  - Earl Parsons, CommScope

- Additional Panelists
  - Guangcan Mi, Huawei
  - Ernest Muhigana, Lumentum
  - Ramana Murty, Broadcom
  - Roberto Rodes, Coherent

Contributors

Vipul Bhatt, Coherent
Jose Castro, Panduit
Jerry Chen, Alibaba Cloud
Weiqiang Cheng, CMCC
Mabud Choudhury, Lightera
Dipak Chudasama, Trumpf Photonic Components
John D'Ambrosia, Futurewei, U.S. Subsidiary of Huawei
Vince Ferretti, Corning
Ali Ghiasi, Ghiasi Quantum LLC
Chris Kocot, Coherent
Angela Lambert, Corning
Hao Liu, China Telecom
Flavio Marques, Lightera
Jeff Maki, Juniper Networks
Vladimir Kozlov, LightCounting
Guangcan Mi, Huawei
Ernest Muhigana, Lumentum
Ramana Murty, Broadcom
Chengguang Pang, CMCC
Earl Parsons, CommScope
Matthew Peters, Lumentum
David Piehler, Dell Technologies
Roberto Rodes, Coherent
Xia Sheng, China Telecom
Hans Spruit, Trumpf Photonic Components
I-Hsing Tan, Broadcom
Yi Tang, Cisco
Craig Thompson, NVIDIA
Howard Trieu, Lightera
Haojie Wang, CMCC
Alan Weckel, 650 Group
Yu (Helen) Xu, Huawei
Zhiping Yao, Alibaba Cloud

CFI objectives

- To measure the interest in addressing:
    - 200G/wavelength MMF optical PHYs
- 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
- RESPECT … give it, get it

Agenda

- **Introduction**
- **Motivation**
- **Market Drivers**
- **Technical Feasibility**
- **Why Now?**
- **Straw Polls**

Motivation

- The introduction of artificial intelligence (AI) networks has led to increased deployment of low cost, low power, short reach optical links in back-end networks
- These back-end links are an addition to front end networks for server-attachment
- This proposed study group will look at short reach (TBD) MMF PHYs using 200G per wavelength to match emerging 200G SerDes
- The motivation is to leverage multimode technology including advanced packaging and volume manufacturing in sensing applications to address the ongoing cost and power consumption pressures on optical interconnects in the web-scale and AI datacenter market
- Adding 200G/lane capabilities enables higher port densities and lower cost per bit

What are we talking about?

What are AI Networks?

General Purpose vs. Scale-Up versus Scale-Out (UEC) Networks

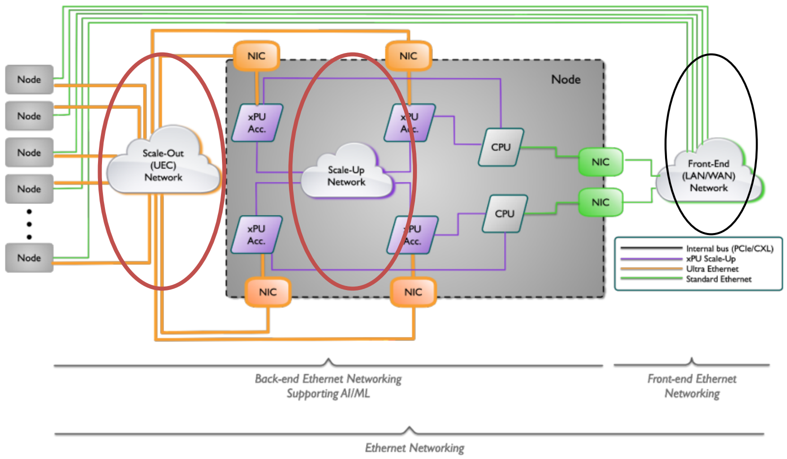

Source: Ultra Ethernet Consortium

Used with permission from J Metz, UEC

- The author is aware that there are different representations of different implementations of AI Networks.
- The key takeaway is there are three types of networks for AI:
  - Front-end / traditional Ethernet
  - Back-end networks
    - Scale-up
    - Scale-out

dambrosia_nea_01a_2501

Applications for early adoption of short-reach 200G PMDs include Scale-Out & Scale-Up Networks in AI clusters. Applications can also include the Front-End Network/traditional Ethernet.

Market Drivers

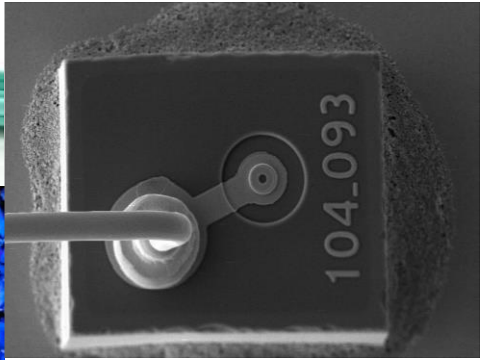

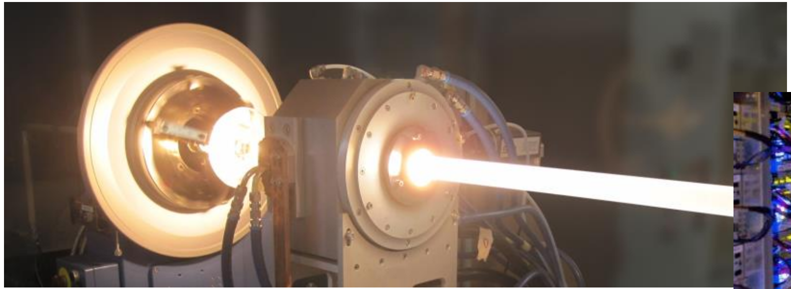

Market Drivers - Overview

- There is a growing demand for low cost, low power fiber links with 200 Gb/s signaling and with reaches much shorter than the shortest links defined in P802.3dj
- These short reach links can be constructed using VCSELs and multimode fiber
- Opportunities to use VCSELs and multimode fiber in scale-up

Ethernet Switch: Data Centers

from dambrosia_nea_01a_2501

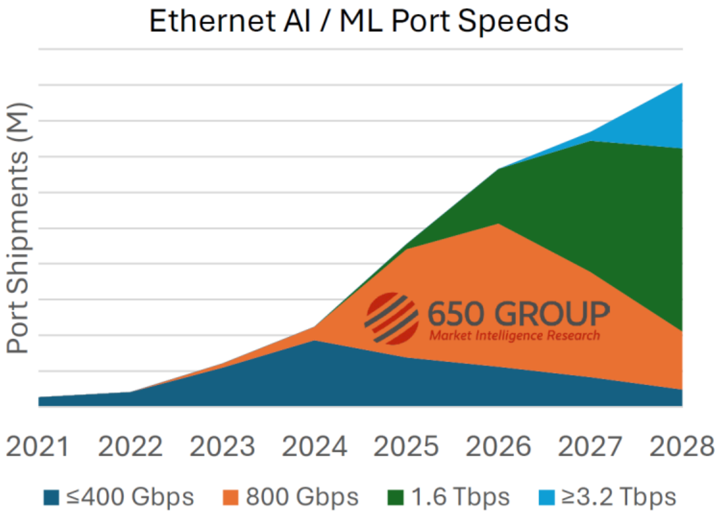

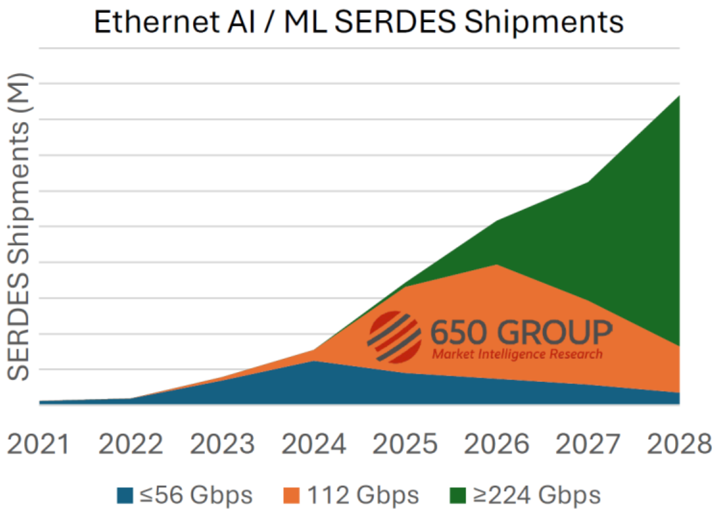

Data Source: Provided by and used with permission by Alan Weckel, 650 Group.

MMF for Large-Scale AI: Efficiency and Resilience

Jose Castro, Panduit

- MMF offers lower cost, power consumption, and better contamination resilience than SMF, making it the ideal choice for the first layer of the scale-out network.
- Depending on power and cooling, large numbers of accelerators connect to Leaf switches over distances under 30 m—expected to decrease as GPU density per rack increases.
- OM4 multimode fiber provides ≥150 GHz of modal bandwidth at 30 meters (corresponding to an EMB of ≥4700 MHz·km at 850 nm).
  - Depending on the VCSEL spectrum, low chromatic dispersion can also be achieved.

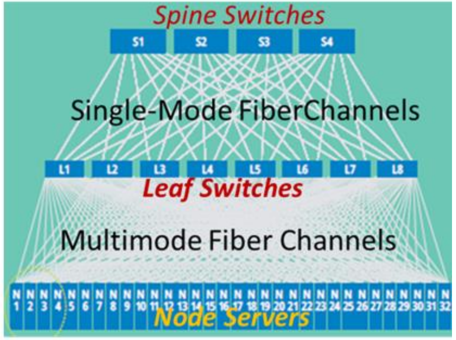

Logical Topology of an AI POD

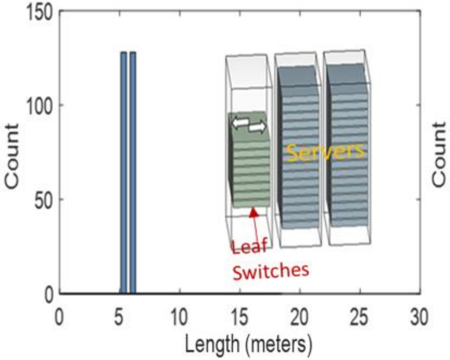

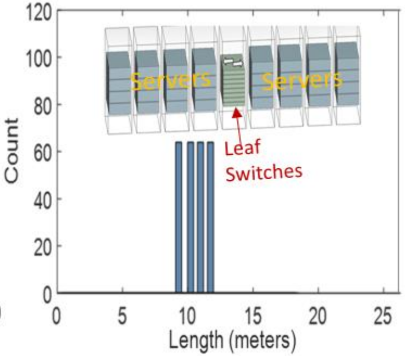

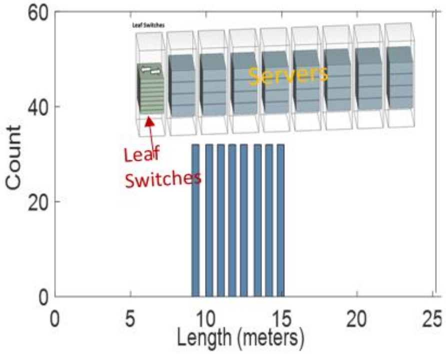

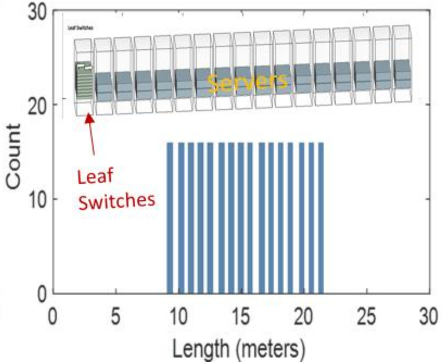

AI scale-out networks are dominated by very short links (< 50 m)

Earl Parsons, CommScope

- Latency requirements keep links between GPUs and first layer switching short

Link lengths from example AI cluster deployed by cloud company in 2024.

Chart represents thousands of multimode fiber links.

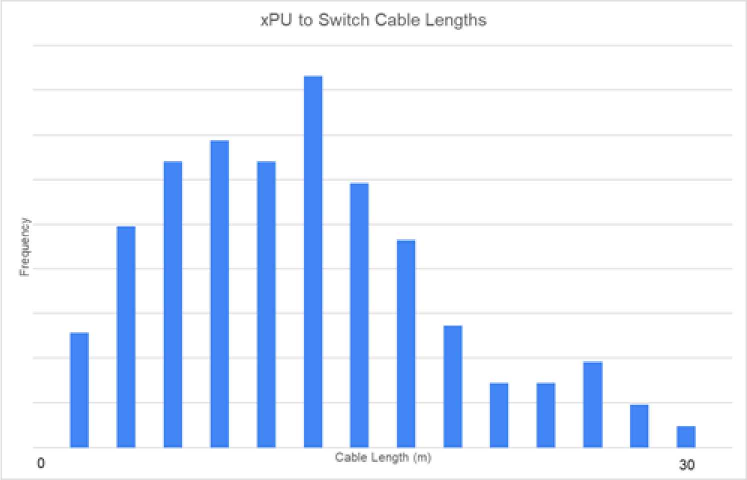

OM4 System Length – AI Intrabuilding Cable

Vince Ferretti, Corning and Angela Lambert, Corning

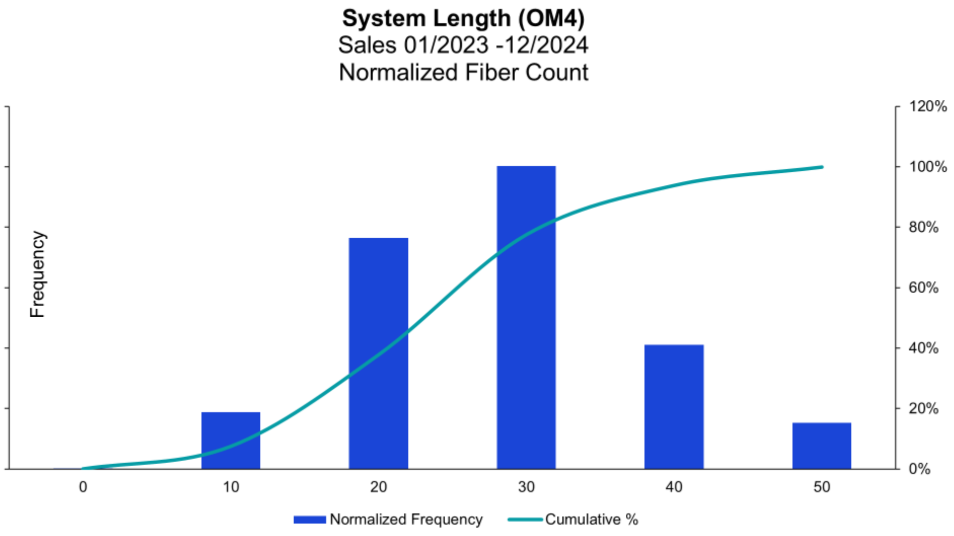

- 30m OM4 reach covers 78% of data center links
- 50m OM4 reach covers 100% of data center links
- Average OM4: ~28m

CORNING

AI scale-out MMF cable length

Howard Trieu, Lightera and Mabud Choudhury, Lightera

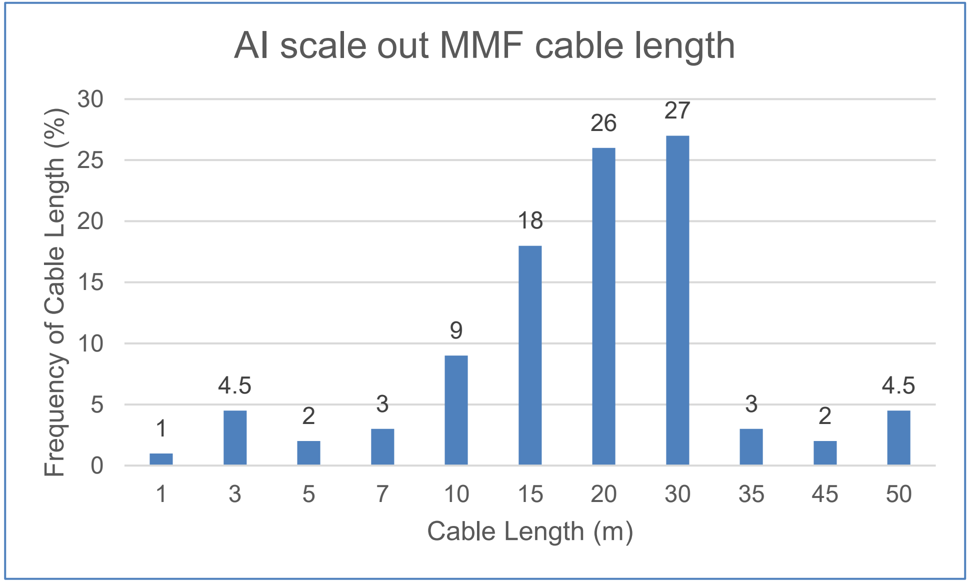

- AI scale out, xPU to switch, MMF cable length
- 90.5% ≤ 30 m
- 100% ≤ 50 m
- Chart data represents many thousands of MMF links
- Timeframe: 2024
- OM3/OM4 (predominantly OM4)

Multimode Trunk Cables Distribution, LatAm - 2024

Flavio Marques, Lightera

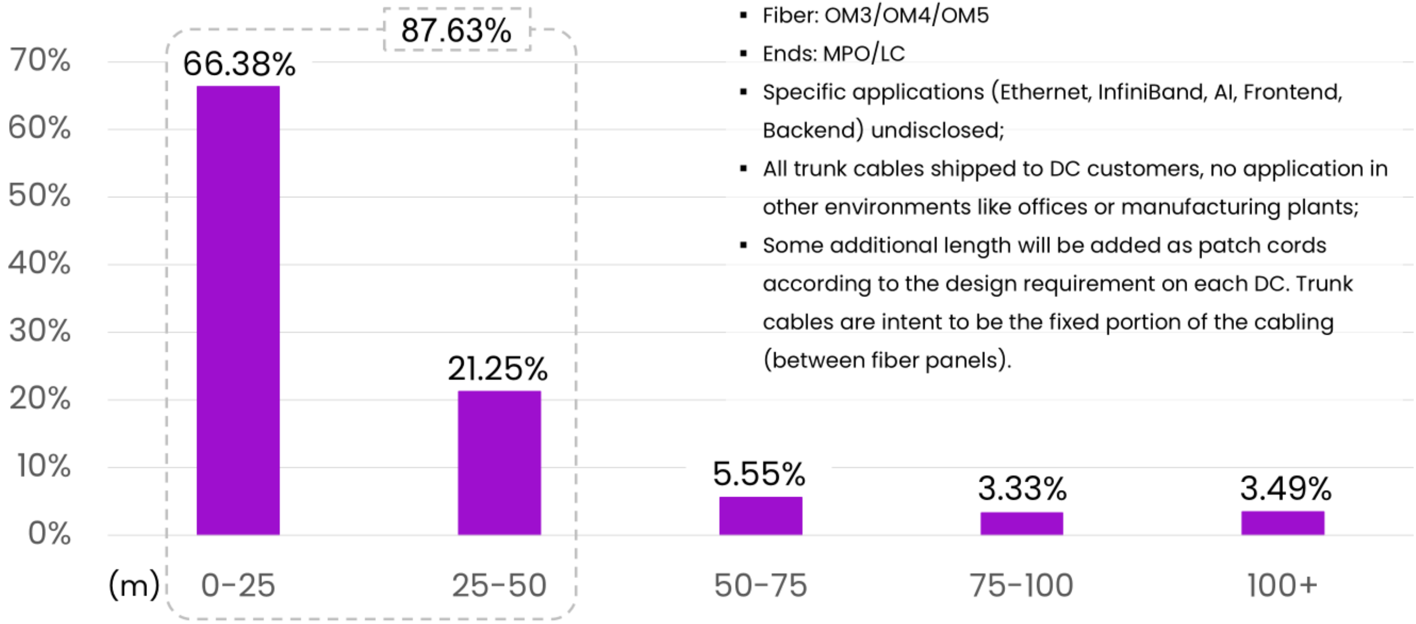

NOTES:

- Fiber: OM3/OM4/OM5
- Ends: MPO/LC
- Specific applications (Ethernet, InfiniBand, AI, Frontend, Backend) undisclosed;
- All trunk cables shipped to DC customers, no application in other environments like offices or manufacturing plants;
- Some additional length will be added as patch cords according to the design requirement on each DC. Trunk cables are intent to be the fixed portion of the cabling (between fiber panels).

Lightera

Why not use AOCs?

Inherent limitations of AOCs

- AOCs require on-site installation
  - Must route transceiver ends thru pathways
  - Longer AOCs hinder deployment speed
  - Risk of damaging fiber during installation
  - **Some OEMs do not offer AOCs for 100G+ lanes**
- AOCs with breakouts even more difficult
  - Breakout involves routing multiple transceiver ends
  - Endpoint location diversity becomes challenging
- AOCs require rip and replace with each generation
  - 18-month refresh cycles
  - Structured cabling persists for multiple generations

Case for 200G VCSEL Technology in AI Systems

Craig Thompson, NVIDIA

Optical Interconnect in AI systems - the numbers:

- **Proportion of Optical Interconnect in AI systems becomes dominated by scale-up fabric.**

- **Multiple factors driving NVIDIA towards optics for scale-up:**
  - **GPU BW continues to increase by factor of 2x every 2 years:**
    - **2.4Tb/s -> 3.6Tb/s -> 7.2Tb/s >> 7.2Tb/s**
  - **GPU domain sizes increasing by factor of 2-4x every 2 years:**
    - **8x -> 32x -> 72x >>72**
  - **GPU-L1 Cable length increasing >2m to 30m+**
    - **Cu reach limited at higher lane speeds**
  - **GPU-L1 Cable density becoming connector and cable gauge limited**

- **VCSEL/MM links are a viable option for bridging the gap between copper and longer-reach single mode optics.**

- **Supportive of a short reach multimode standard for 1x200G and 4x200G.**

Case for 200G VCSEL Technology in AI Systems

Craig Thompson, NVIDIA

Optical Interconnect in AI systems - the numbers:

TODAY:

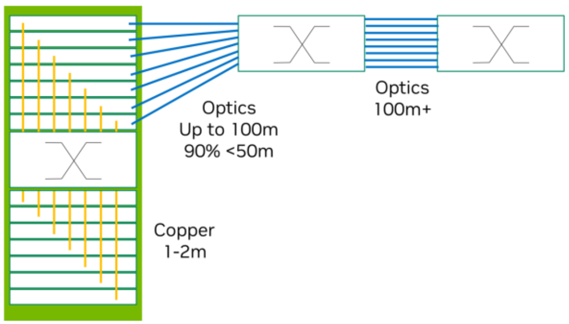

FUTURE:

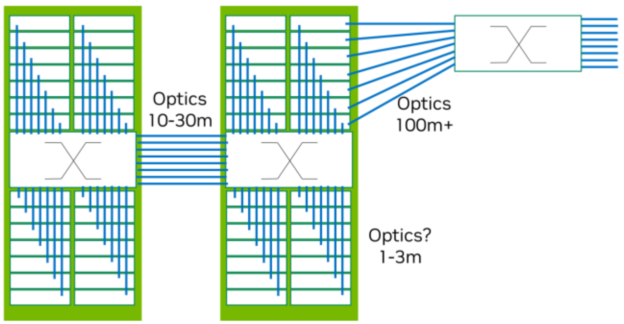

- **Majority of GPU-GPU links <2m Cu today:**
  - **Scale-up fabric ~9x BW of scale-out network**
  - **~5:1 copper vs optics**
- **Cu used for cost and reliability**
- **Cu limited by reach and density (connectors)**

- **Majority of GPU-GPU links will remain <50m**
- **Min 50% of scale-up links must >2m**
- **Need low-cost, reliable, power-efficient optics**
- **Must hit targets:**
  - **30m distance**
  - **Approaching copper-level reliability**
  - **<5pJ/bit (not including PHY)**
  - **1Tbps/mm**

If optics can be applied to 100% scale-up, volume grows >20x over next 5 years

200G/L SR use case and requirement

Jerry Chen, Alibaba Cloud and Zhiping Yao, Alibaba Cloud

- Use case
  - Parallel for breakout applications
  - AOC and Transceiver for servers to switch connections
  - Transceivers for switch to switch connections
- Distance
  - 50 meters required for transceivers; covers 60% connections
  - 30 meters is currently a space for AOCs at Alibaba
  - 70 or 100 meters desired for transceivers with inner-FEC enabled
- Cost & Power
  - Cost<50% of DR, Power<80% of DR
- Reliability
  - FIT<100, AFR(Annual failure Rate)<0.1%

200G VCSEL Optics

Jeffery Maki, Juniper Networks

200G-Lane VR/SR TOR-NIC Interconnects

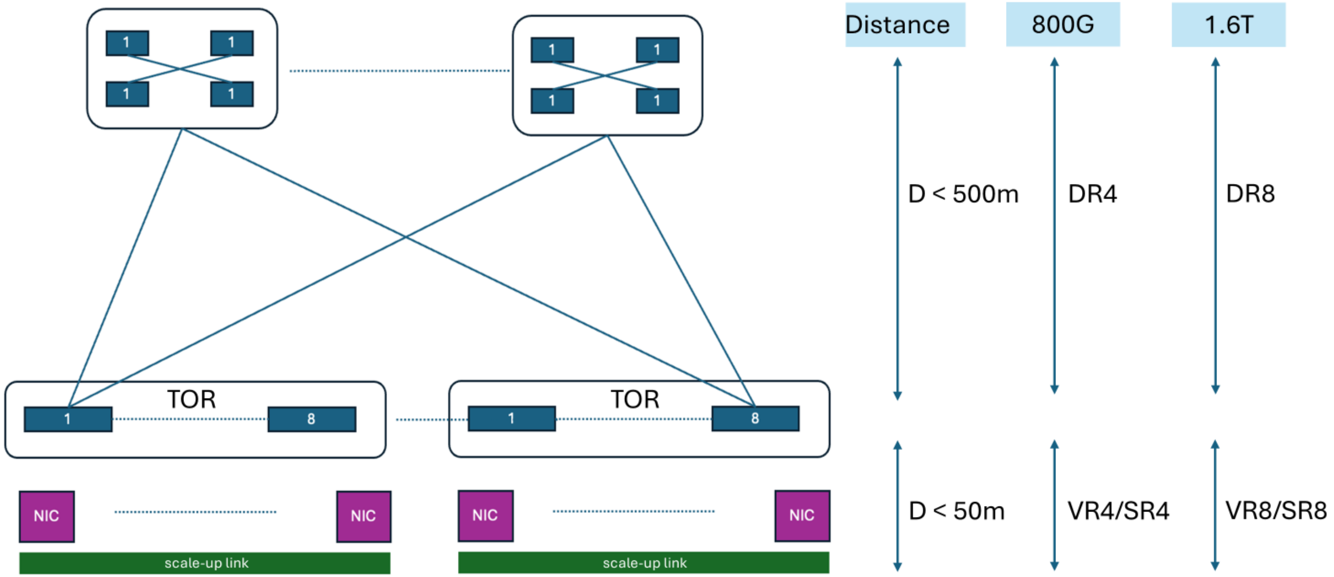

200G VCSEL Optics

Jeffery Maki, Juniper Networks

Need for 200G-Lane VR/SR Optics in AI DC

Higher Bandwidth Density

- 200G/lane module can double the bandwidth capacity within the same physical space, allowing for more efficient use of data center rack space. This is particularly needed for AI data centers.

Improved Power Efficiency

- 200G/lane has fewer lanes used to achieve the same data rates as previous 100G/lane generation, which leads to better power efficiency & reduced cooling requirements

Simplified Network Design

- The need for gearboxes is reduced, simplifying the design of optical modules
- Uniform electrical and optical lane rate leads to less complex designs

Performance

- 200G/lane better provides for adoption of 1.6T Ethernet that enables lower latency switching, which is key for AI data centers
- 200G/lane technology helps drive NIC capacity

Juniper Public

Opportunity for 200G optics using multimode fiber

David Piehler, Dell Technologies

- The opportunity for lower-cost, lower-power, moderate-distance 200G optics using multimode fiber exists in legacy datacenter applications but its usage will be overwhelmed by its use in scale-up and scale-out AI fabrics. This contribution focuses on these applications.
  - The demand for 200G optics begins at lengths where 200G active and passive copper solutions likely fail which is about 2-3 m for scale-up and 5-6 m (?) for scale-out. I believe > 80% of demand will be for < 30 m distances.
  - Traditional optical form factors (*e.g.* OSFP) will be used in scale-out. Scale-up will introduce new form factors which could be anything from very-high density LPO to highly-integrated CPO, with a continuum of other viable potential form factors in-between.
  - The architectures for AI fabric scale-up and scale-out tend toward connecting every device-to-every device or “breakout everywhere.” Heavy use of channel bundling at the physical layer (via WDM) or layer 1 (via high-speed MAC domains) make little sense. No new high-speed MACs need to be created.
  - All-to-all AI fabrics consists of hundreds to thousands (and beyond) of individual links in both scale-up and scale-out. AOCs make little sense, due to high breakout requirements. Also, connectivity management is out-sourced to a fiber shuffle. The fiber shuffle allows AI fabrics to be built all at once, instead of as a process of adding single point-to-point links one after another.

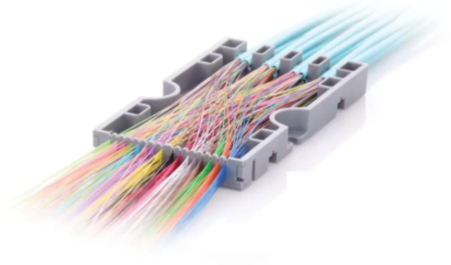

DELL Technologies

200G/L Optical Link based on MMF has a broad market use

Guangcan Mi, Huawei and Yu (Helen) Xu, Huawei

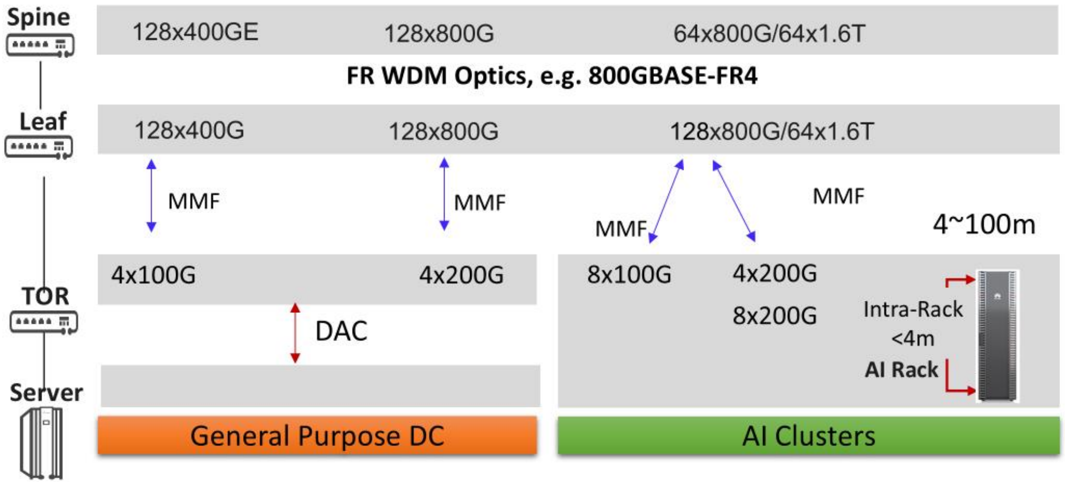

Data Center use cases

- In both General Purpose DC and AI Clusters
- Evolving from past generations e.g. 100G/L
- Reuse established fiber installation is preferred
- Cost effective is key decision factor in both cases
- In AI cluster, 200G/L optics will see opportunity in intra-rack reach, if overall cost and power makes sense

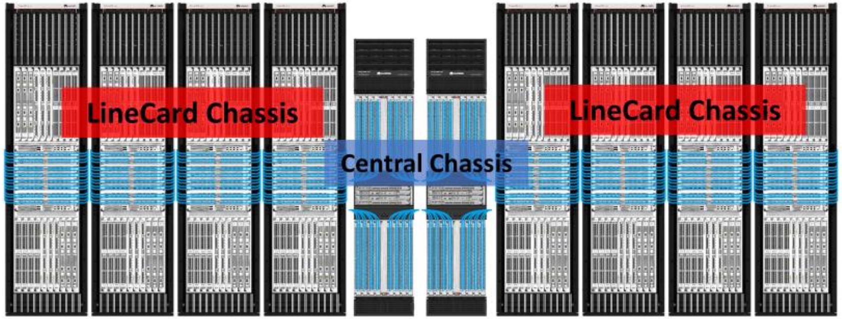

Carrier Network---Router Chassis

- **Very High Density** (bps/mm^2) on front panel of Central chassis
- **Low heat footprint** (W/mm^2) is key factor in system design, therefore, low power is top priority for optical module.
- Use high density form factor
- Maintaining MMF link is preferred in system upgrade

200G/L MMF optics for 800GE/1.6TE

Haojie Wang(CMCC), Weiqiang Cheng(CMCC), Chengguang Pang(CMCC)

- CMCC's data centers and AI clusters extensively utilize short-reach optics, representing more than 50% of deployed modules.
- For leaf-to-spine connections, 400G FR4 is used; server-to-leaf connections use 400G SR8 with a convergence ratio of 1:1. SR8 accounts for about 50% of deployments.
- The current infrastructure predominantly operates at 400G rates, phased transitions to 800G or leapfrog directly to 1.6T are planned for subsequent stages.
- From the roadmap for electrical and optical interface evolution, NICs are presently based on 56G SerDes, which will subsequently upgrade to 112G SerDes. The corresponding optical interfaces will transition to 400GE QSFP112 or 800GE SR8. For 224G SerDes, the adoption of 200G/L VR/SR optics for 800GE or 1.6TE will be positive depending on the industry maturity of VCSEL and oDSP.

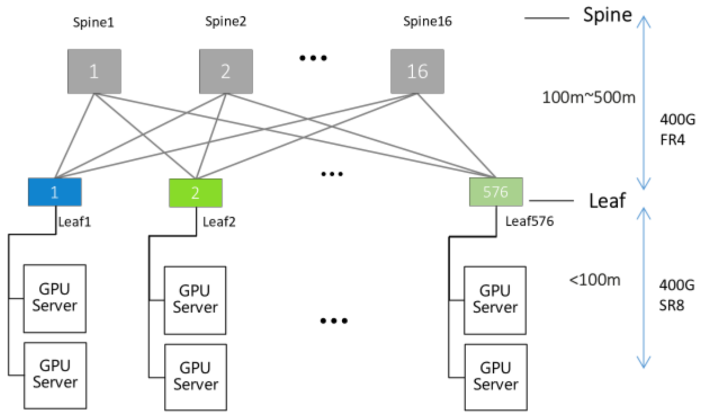

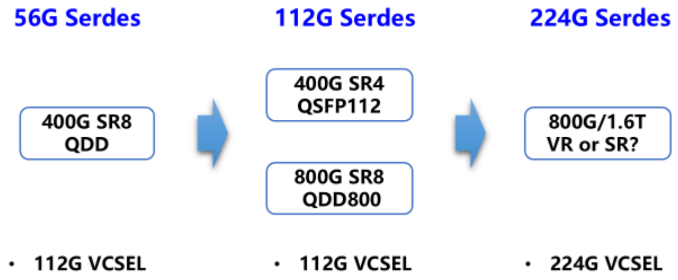

200G/L MMF use case, Cloud & AI data centers

Hao Liu, China Telecom and Xia Sheng, China Telecom

China Telecom's Interconnect Solution for Next Generation Data Center

- **Cloud & AI Data Center are Cost Sensitive & Power Consumption Sensitive**
- **MMF Solution is the best choice and technical feasible**

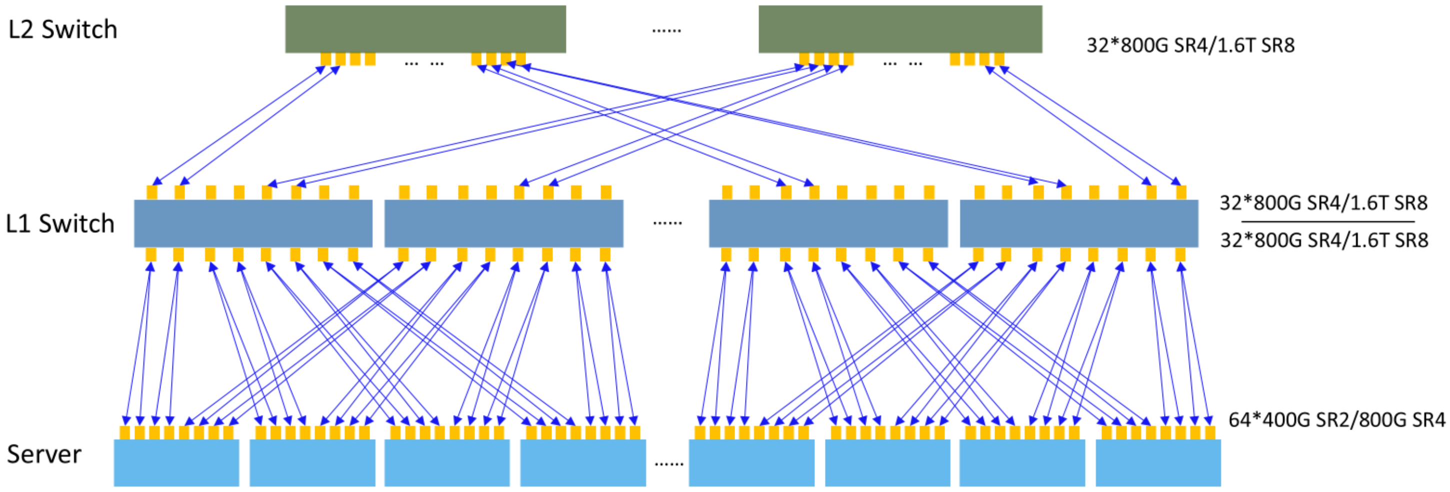

Technical Feasibility

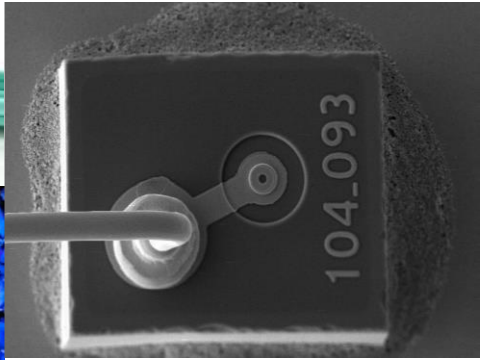

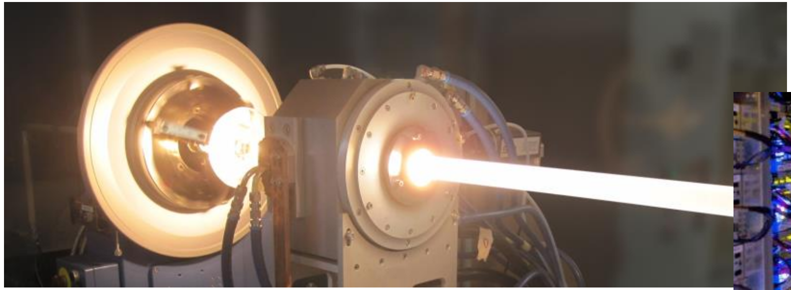

802.3 Architecture View

Guangcan Mi, Huawei, Yi Tang, Cisco, and Ramana Murty, Broadcom

- Multimode PMDs will fit into the broader 200G per lane optical PMDs with no compatibility issues with existing host designs
- Study PMD + FEC to explore reach. Additional data encoding (FEC layer) will be discussed during the study group phase, allowing for a thorough examination of tradeoffs.
- Adopt PCS/PMA developed in 802.3dj
- Leverage the Inter-Sublayer Link Training (ILT) defined by 802.3dj for PAM4 IMDD optical links (Annex 178B)

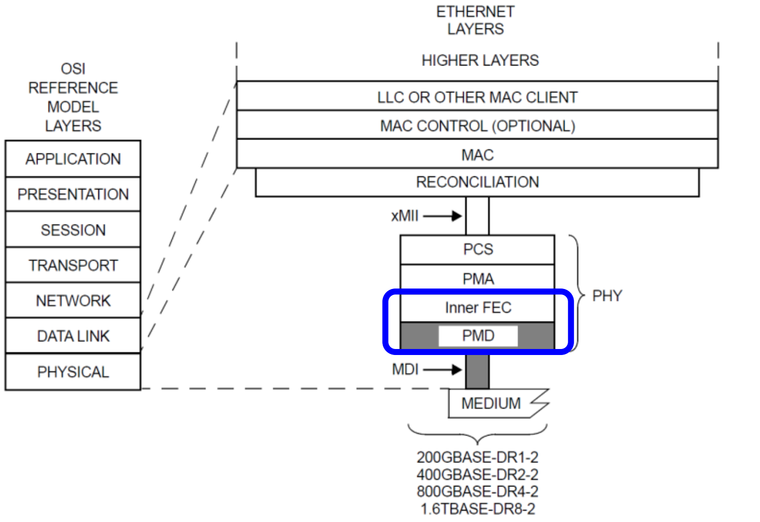

FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
xMII = GENERIC MEDIA INDEPENDENT INTERFACE

Technical Feasibility: 200G Multimode Fiber Link

Ramana Murty, Broadcom and I-Hsing Tan, Broadcom

Development of VCSEL for 200G per lane multimode link is in progress:

- There is substantial increase in device bandwidth over 100G VCSEL
- Adaptive equalization enables demonstration of an open eye
- Links of 30 m and 50 m have been demonstrated
- 850 nm wavelength extends the use of OM3 and OM4 multimode fibers to the next generation

ER	2.7 dB
T	Room

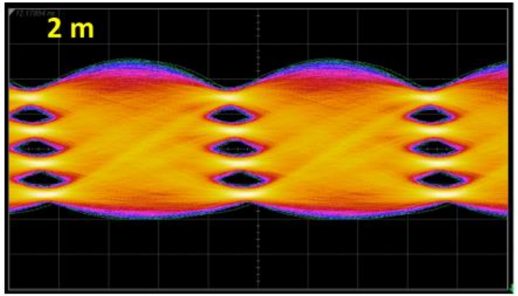

Symbol rate	106.25 GBd
Pattern	PRBS15Q
SIRC filter	53.1 GHz
DCA FFE	25-tap

50 m Link

Tmax = 7 operating over 1 hr

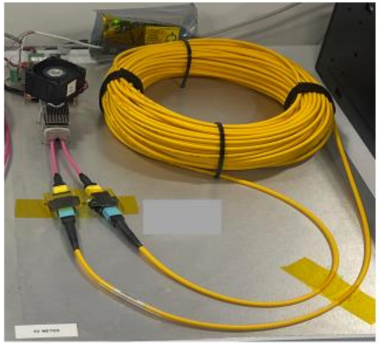

Symbol rate	106.25 GBd
Pattern	PRBS31Q
Temperature	Room
EMB	> 6500 MHz·km

200G VCSEL - Technical Feasibility

Chris Kocot, Coherent and Roberto Rodes, Coherent

2025 OFC Demo

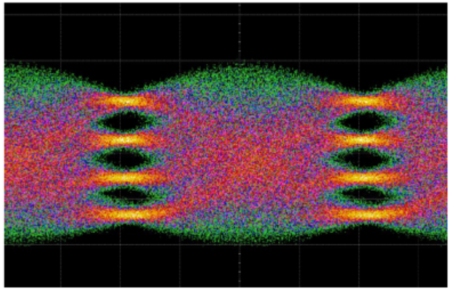

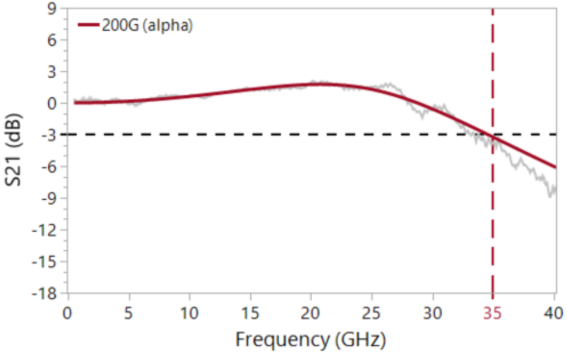

- 850nm VCSEL
- VCSEL 3dB bandwidth of 35 GHz
- Transceiver includes DSP+ laser driver + VCSEL
- 212.5 Gb/s data transmission
- PAM-4 eye over 30m OM4 fiber

Wavelength >850nm also possible:

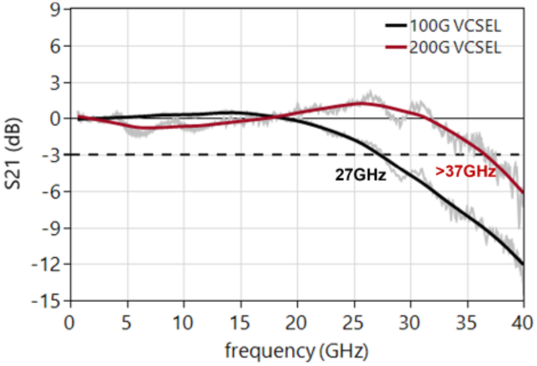

940nm VCSEL 3dB bandwidth >37 GHz

COHERENT

200G/Lane MMF Technical Feasibility

Yu (Helen) Xu, Huawei and Guangcan Mi, Huawei

Eye Diagram

Simulated

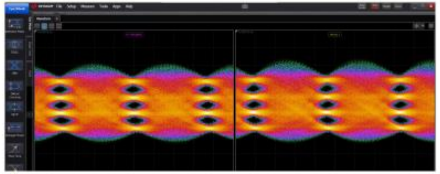

- Tx bandwidth: 52GHz@3dB (with pre-emphasis.)
- Rx bandwidth: 35GHz@3dB
- Wavelength: 850nm

Tested @ 850nm

Vendor A

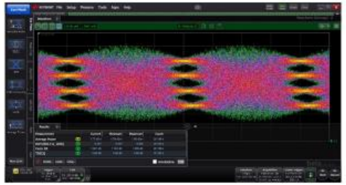

Vendor B

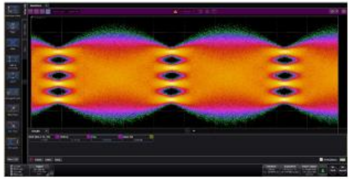

B2B BER vs OMA

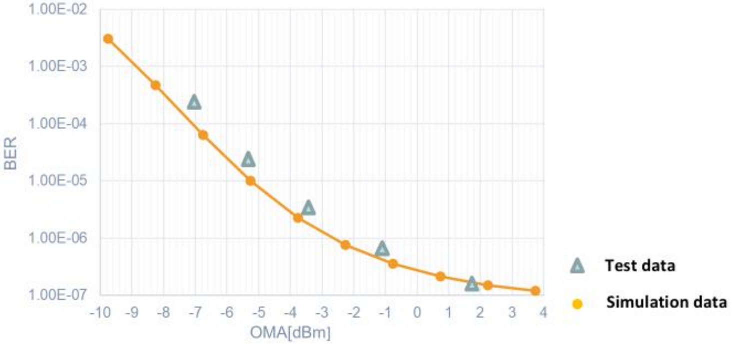

HUAWEI

200G VCSEL | S21 and high temperature LIV measurements

Ernest Muhigana, Lumentum and Matthew Peters, Lumentum

- 1st design iteration with 4.5um oxide aperture at <u>1060nm</u> → ~33GHz 3dB BW
  - DC testing of prototype wafer: devices can operate to 85°C substrate temperature without roll-over to >11mA

Measured 3dB Bandwidth

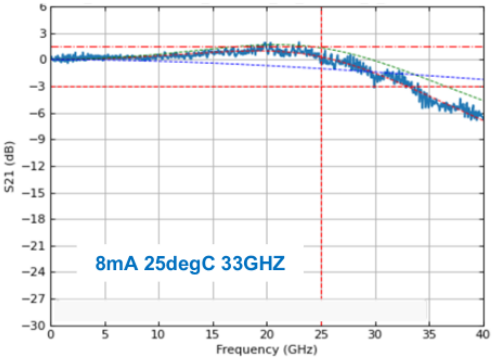

Measured DC LI over Temperature

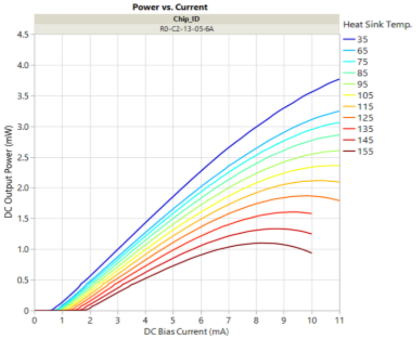

Measured DC IV over Temperature

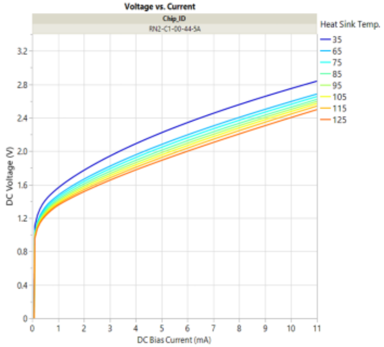

- 2nd iteration w/ updated design in wafer processing → Target 40GHz 3dB BW
- Leverage existing <u>B</u>ottom <u>S</u>ide-<u>E</u>mitting VCSEL technology, used in sensing applications, allowing for dense (flip-chip) arrays and volume proven 2.5/3D packaged structures

LUMENTUM

200G VCSEL/PD feasibility

Hans Spruit, Trumpf Photonic Components and Dipak Chudasama, Trumpf Photonic Components

Photodiodes for 200G

Spectral bandwidth

- Responsivity from 840 nm to about 1600 nm wavelengths
- Based on InGaAs absorption layer

Reliability

- No fails so far after 3000 hours of testing at 140°C -2V and -4V

HF performance

- Capacitance < 100 fF
- Measured ~ 25 GHz f3dBe at 20 μm diameter
- > 30 GHz bandwidth by tuning diameter and intrinsic layer thickness

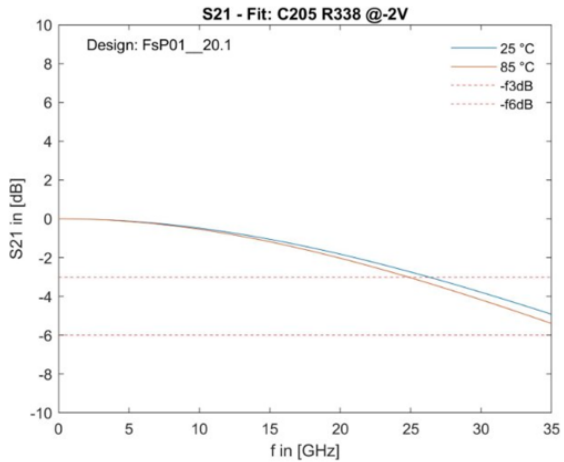

Photodiode for longer wavelengths and high data rate using InGaAs lattice matched to InP

TRUMPF

Re-use of Receiver Components

- Except for the low-cost photodiode, receiver components for 200G optics are not unique to multimode receivers
- Transimpedance amplifiers, downstream clock recovery and signal processing circuits will be the same functions as used in longwave receivers (802.3dj implementations, for example)
- In fact, such re-use will lead to further improvements in economies of scale

Why Now?

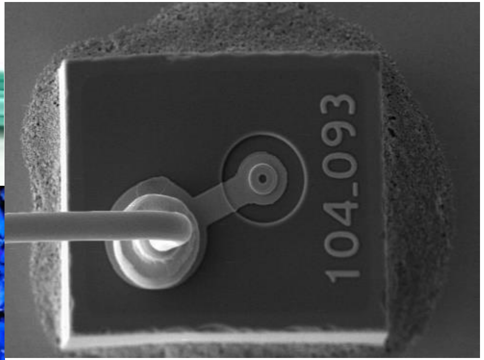

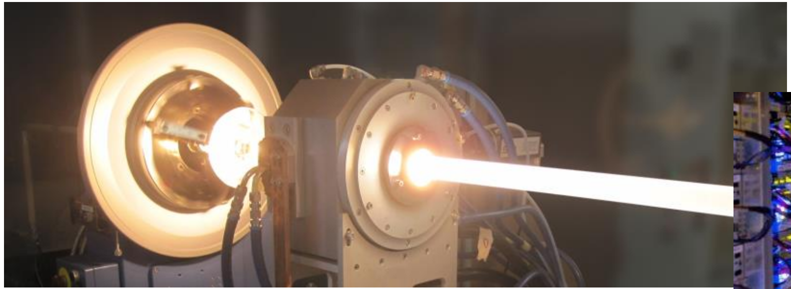

IEEE 802.3 MM standards have set the foundation for short reach optical links that are cost-optimized solutions for broad market adoption

Publication/CFI dates

- **IEEE Std 802.3df™-2024 – 800G-SR8**
  - Published Mar 2024. CFI Oct 2020
- **IEEE Std 802.3db™-2022 – 400G-SR4**
  - Published Dec 2022. CFI Nov 2019
- **IEEE Std 802.3cm™-2020 – 400G-SR8**
  - Published Mar 2020. CFI Nov 2017
- **IEEE Std 802.3cd™-2018 – 200G-SR4**
  - Published Feb 2019. CFI Nov 2015
- **IEEE Std 802.3bm™-2015 – 100G-SR4**
  - Published March 2015. CFI Jul 2011
- **IEEE Std 802.3ba™-2010 – 40G-SR4**
  - Published Jun 2010. CFI Jul 2006

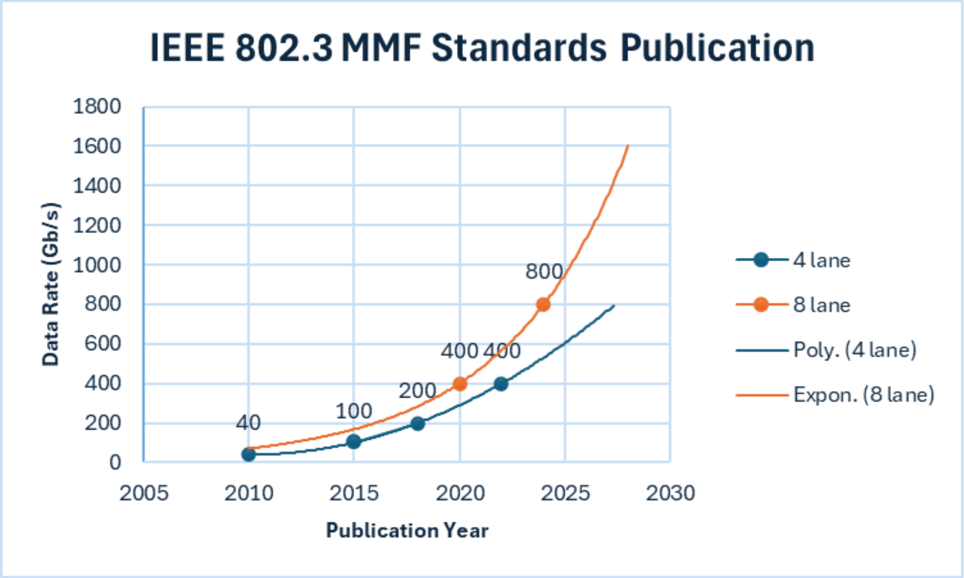

Three generations of SR4/SR8 transceivers

VCSELs are not giving up! SR transceivers continue to deliver the lowest cost and power.

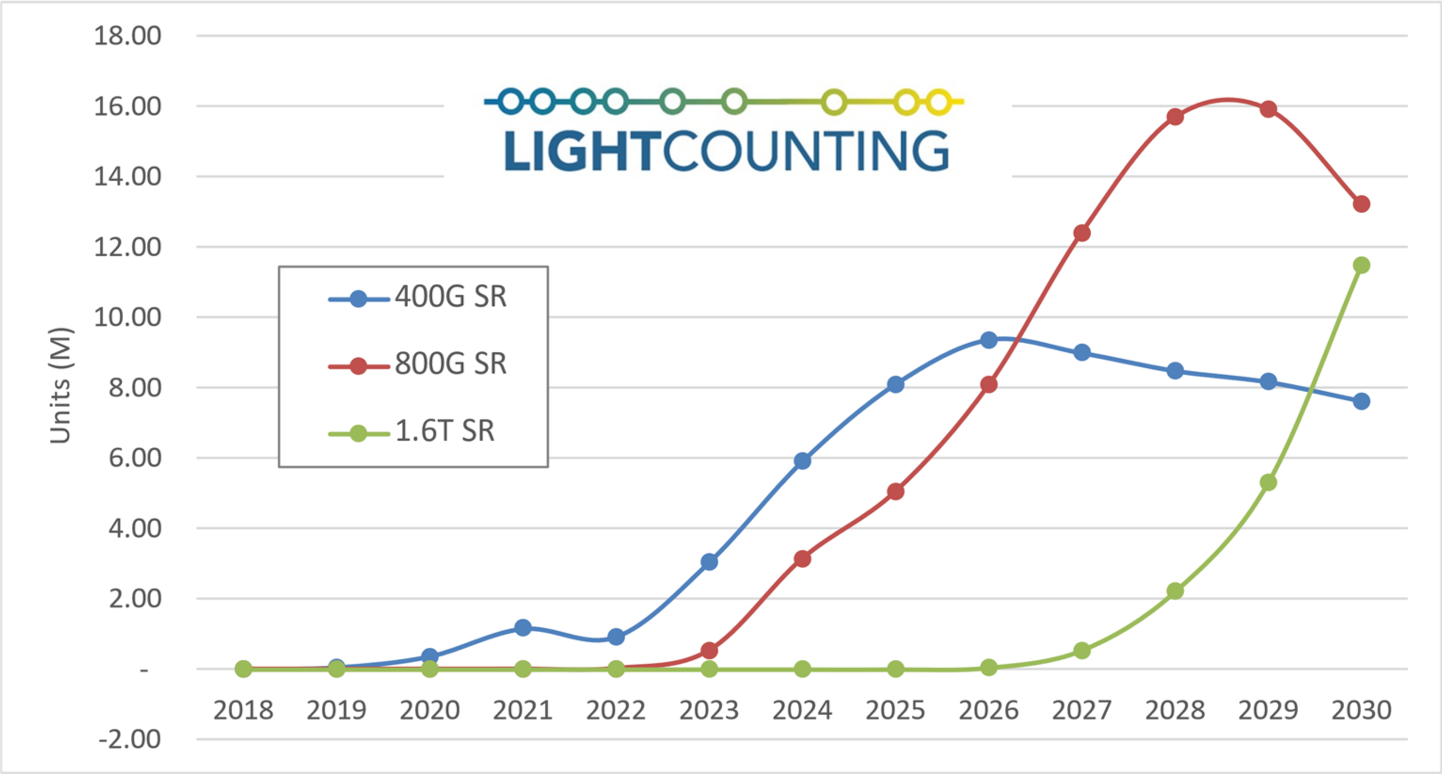

Data Source (Feb 2025): Provided and used with permission from Vlad Kozlov, LightCounting

CFI Market Drivers & Technical Feasibility relative to P802.3dj Objectives

Based on CFI market drivers/use cases and technical feasibility, if CFI and SG are approved, SG might consider targeting objectives for 1, 2, 4, 8 lanes of 200G
– aligning with and leveraging IEEE P802.3dj objectives and developments

IEEE P802.3dj Objectives

Ethernet Rate	Signaling Rate	AUI	Backplane	Cu Cable	SMF 500m	SMF 2km	SMF 10km	SMF 20km	SMF 40km
200 Gb/s	200 Gb/s	200GAUI-1 C2C C2M	200GBASE-KR1	200GBASE-CR1	200GBASE-DR1	200GBASE-DR1-2			
400 Gb/s	200 Gb/s	400GAUI-2 C2C C2M	400GBASE-KR2	400GBASE-CR2	400GBASE-DR2	400GBASE-DR2-2			
800 Gb/s	200 Gb/s	800GAUI-4 C2C C2M	800GBASE-KR4	800GBASE-CR4	1.800GBASE-DR4 2.800GBASE-FR4-500	1. 800GBASE-DR4-2 2. 800GBASE-FR4	800GBASE-LR4		
	800 Gb/s						800GBASE-LR1	800GBASE-ER1-20	800GBASE-ER1
1.6 Tb/s	100 Gb/s	1.6TAUI-16 C2C C2M							
	200 Gb/s	1.6TAUI-8 C2C C2M	1.6TBASE-KR8	1.6TBASE-CR8	1.6TBASE-DR8	1.6TBASE-DR8-2			

12 March 2025 Overview - IEEE 802.3 "Ethernet for AI" Assessment 1

Source: John D'Ambrosia, IEEE P802.3dj Chair

Summary

- Multiple industry experts are identifying a market need for 200G/lane short reach optical links for AI back-end compute clusters and for front-end/traditional Ethernet networks. These networks prioritize power and cost for short reach applications. All key value propositions for VCSEL-MMF links.
- The technology for 200 Gb/s per wavelength VCSEL-MMF links has reached a level that suggests the time is right to study an interoperable Ethernet specification that will have broad market adoption.
- Seek to initiate a Study Group to explore the potential market requirements and feasibility of, and to develop a PAR and CSD for 200 Gb/s per wavelength MMF optical PHYs.
- Leverage developments in P802.3dj in that effort.

Proposed Study Group Chartering Motion

- Approve the formation of a Study Group to explore the potential market requirements and feasibility of addressing AI/Data Center networks, and to develop a PAR and CSD for 200 Gb/s per wavelength MMF optical PHYs

Supporters

Vipul Bhatt, Coherent
Matt Brown, Alphawave Semi
Jose Castro, Panduit
Connie Chang-Hasnain, Berxel Photonics
Jerry Chen, Alibaba Cloud
Chan Chen, AOI/Independent
Weiqiang Cheng, CMCC
Mabud Choudhury, Lightera
Dipak Chudasama, Trumpf Photonic Components
John D'Ambrosia, Futurewei, U.S. Subsidiary of Huawei
Piers Dawe, Nvidia
Mike Dudek, Marvell
Ahmad El-Chayeb, Keysight
Vince Ferretti, Corning
Wanchao Gao, Accelink Technology
Ali Ghiasi, Ghiasi Quantum LLC

Xiaoming Han, Vertilite
Xiang He, Huawei
Robert Hu, HG Genuine
Zhaoyang Hu, Crealights Technology
Tom Issenhuth, Huawei
Kenneth Jackson, Sumitomo Electric Industries
John Johnson, Broadcom
Mark Kimber, Semtec
Toshiharu Kiuchi, Sony Semiconductor Solutions
Beth Kochuparambil, Cisco
Chris Kocot, Coherent
Vladimir Kozlov, LightCounting
Daniel Kuchta, Nvidia
Angela Lambert, Corning
Ryan Latchman, Macom

Supporters

Jon Lewis – Dell Technologies
Jing Li, YOFC
Hao Liu, China Telecom
Hai-Feng Liu – HG Genuine
Kent Lusted, Synopsys
Jeff Maki, Juniper Networks
David Malicoat, Senko Advanced Components
Flavio Marques, Lightera
John Marshall, AMD
Marco Mascitto, Nokia
J Metz, AMD
Guangcan Mi, Huawei
Tom Mitcheltree, US Conec
Andy Moorwood – Keysight Technologies
Jianwei Mu, Hisense Broadband
Ernest Muhigana, Lumentum

Ramana Murty, Broadcom
Ray Nering – Cisco
Tiger Ninomiya, Accelink Technology
Mark Nowell, Cisco
David Ofelt, Juniper Networks
Thomas Palkert, Samtec-Macom
Chengguang Pang, CMCC
Earl Parsons, CommScope
Matthew Peters, Lumentum
David Piehler, Dell Technologies
Roberto Rodes, Coherent
Toshiaka Sakai – Socionext
Xia Sheng, China Telecom
Rames Sivakolundu – Cisco Systems Inc.
Yung Sung Son – Optomind
Hans Spruit, Trumpf Photonic Components

Supporters

Peter Stassar, Huawei
Min Sun, Tencent
Tomoo Takahara – 1FINITY
I-Hsing Tan, Broadcom
Yi Tang, Cisco
Craig Thompson, Nvidia
Marek Tlalka, Macom
Pirooz Tooyserkani – Cisco
Howard Trieu, Lightera
Emma Wan, Baidu
Haojie Wang, CMCC
Alan Weckel, 650 Group
James Withey, Fluke
Yu (Helen) Xu, Huawei
Jin Xu, YOFC
Lu Xuu – Huawei
Zhiping Yao, Alibaba Cloud
Rang-Chen (Ryan) Yu, Terahop
Al Yuen, Picojool

Straw Polls

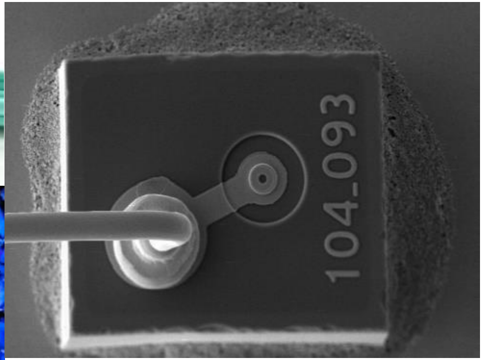

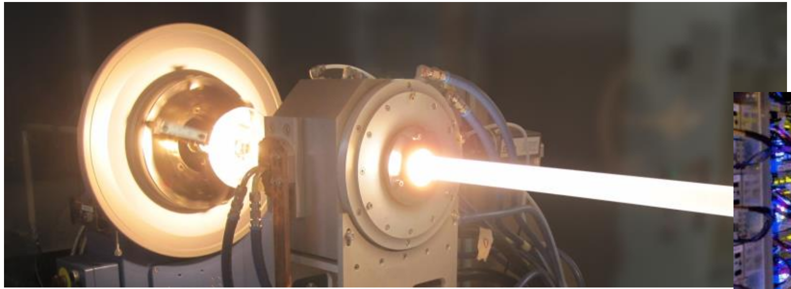

Straw Poll Questions

1. Should a study group be formed to develop a PAR, CSD responses, and objectives for “200G/wavelength MMF optical PHYs”?
   - ❑ Yes: 63
   - ❑ No: 0
   - ❑ Abstain: 2

2. If formed, will you participate in this Study Group?
   - ❑ Yes: 51 individuals, 36 affiliations

Straw Poll Questions

1. Should a study group be formed to develop a PAR, CSD responses, and objectives for “200G/wavelength MMF optical PHYs”?
   - ❑ Yes: 63
   - ❑ No: 0
   - ❑ Abstain: 2

2. If formed, will you participate in this Study Group?
   - ❑ Yes: 51 individuals, 36 affiliations

Next Steps

- Make the call for interest during the IEEE 802.3 Opening Plenary meeting on Monday, 28th July
- The vote to determine if a Study Group will be formed will take place at the IEEE 802.3 Closing Plenary meeting on Thursday, 31st July
- If approved, request formation of “200G/wavelength MMF optical PHYs” Study Group by IEEE 802 EC
- If approved,
  - Teleconference(s) in August to start the discussion (will post on NEA Ad hoc reflector)
  - Creation of Study Group page /reflector
  - First Study Group meeting [hybrid] anticipated for Sep 2025 Interim

Thank you!

Back Up

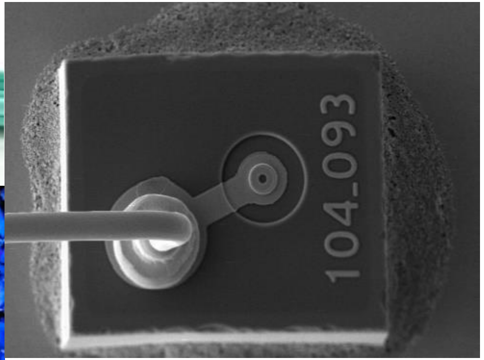

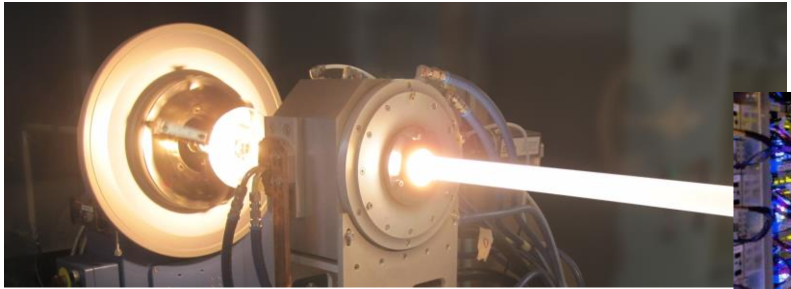

CFI Announcement: 200 Gb/s per wavelength MMF optical PHYs

Links comprising multimode fiber (MMF) cable and VCSEL-based transceivers have played a key role in implementing multiple generations of Ethernet data rates in data centers for short reach. Ethernet has a proven track record of reusing and leveraging technology to enable new cost-optimized solutions for broad market adoption in these short-reach applications. IEEE 802.3db and IEEE 802.3df Ethernet projects defined specifications for 100 Gb/s, 200 Gb/s, 400 Gb/s, and 800 Gb/s operation over MMF using 100 Gb/s signaling. These Ethernet standards have gained market adoption in high bandwidth, high growth artificial intelligence (AI) back-end networks, as well as front-end networks for server-attachment, due to their lower power and lower cost than other optical technologies and their longer reaches than copper technologies. The continual growth of bandwidth demand has driven the evolution of even higher Ethernet speeds, most recently with 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet using 200 Gb/s signaling specifications being developed by the P802.3dj project. Now, the technology for 200 Gb/s per wavelength VCSEL-MMF links has reached a level that suggests the time is right to study an interoperable Ethernet MMF specification that will have broad market adoption.

This is a call for interest to initiate a Study Group to explore the potential market requirements and feasibility of addressing AI/Data Center networks, and to develop a PAR and CSD for 200 Gb/s per wavelength MMF optical PHYs.