EthernetRoadmap 2026-Side2-Press-3.pdf 1 11/17/25 12:02 PM

ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)

Artificial Intelligence and Machine Learning (ML) are driving the roadmap extending Ethernet speeds to 3.2T and beyond.

Ethernet for AI

The interconnect architecture within AI-driven data centers is rapidly evolving, incorporating advanced copper and fiber solutions to meet AI's soaring bandwidth demands while also reducing environmental impacts. Ethernet's progression towards higher speed interfaces, the widening variety of interconnect options, and advancements in power efficiency are ensuring that Ethernet can meet the needs of AI/ML workloads.

Ethernet has become the de facto standard for scale-out AI networking and is rapidly becoming the technology of choice for scale-up as well. It offers high bandwidth, low latency, and energy efficiency, all of which are needed to interconnect XPU clusters of up to, and beyond hundreds of thousands of XPUs in AI training and inference clusters.

Various international industry consortia are adopting Ethernet to support scale-out and scale-up networks for AI/ML. Of note are the Ultra Ethernet Consortium (UEC) Specification 1.0 defining a new network stack for AI, and the Ultra Accelerator Link (UALink) 1.0 Specification, which adopts the Ethernet PHY. Also, the Open Compute Project has launched the Ethernet Scale-Up Networking (ESUN) and Scale-Up Ethernet Transport SUE-T projects. ESUN is developing specifications for L2/L3 framing, header efficiency, lossless networking, and error recovery, while SUE-T is focusing on transport solutions. These and other advancements reinforce Ethernet as the foundation for end-to-end AI networking.

LATEST INTERFACES AND NOMENCLATURE

	Backplane	Twinax Cable	15-40m(OT) Single Twisted Pair	>100m (OT) Single Twisted Pair	100m (IT) Twisted Pair (2/4 Pair)	MMF	500m SMF	2km SMF	10km SMF	20km SMF	30 km SMF	40km SMF	80km SMF	Electrical Interface	Pluggable Module
10BASE–	T1S		T1S	T1L	T										
100BASE–			T1	T1L	T										
1000BASE–			T1		T										SFP
2.5GBASE–	KX		T1		T										SFP
5GBASE–	KR		T1		T										SFP
10GBASE–			T1		T	SR			LR BR10-D/U	BR20-D/U		ER BR40-D/U			SFP
25GBASE–	KR1 KR	CR1 CR/CR-S	T1		T (30m)	SR			LR EPON BR10-D/U	EPON BR20-D/U		ER BR40-D/U		25GAUI	SFP
40GBASE–	KR4	CR4			T (30m)	SR4/eSR4	PSM4	FR	LR4			ER4		XLAUI XLPPI	QSFP
50GBASE–	KR2 KR	CR2 CR				SR		FR	LR EPON BR10-D/U	EPON BR20-D/U		ER BR40-D/U		LAUI-2/50GAUI-2 50GAUI-1	SFP/QSFP
100GBASE–	KR4 KR2 KR1	CR10 CR4 CR2 CR1				SR10 SR4 SR2 VR1/SR1	PSM4 DR	CWDM4 FR1	LR4/4WDM-10 LR1	4WDM-20 LR1-20	ER1-30	ER4/4WDM-40 ER1-40	ZR	CAUI-10/CPPI CAUI-4/100GAUI-4 100GAUI-2 100GAUI-1-S/L	SFP/SFP-DD QSFP/QSFP-DD OSFP
						VR1/SR1	DR	FR1						EEI-100G-RTLR-1-S/L	
100G–							DR1-LPO							LEI-100G-PAM4-1	
200GBASE–	KR4 KR2 KR1	CR4 CR2 CR1				SR4 VR2/SR2	DR4 DR1	FR4 DR1-2	LR4			ER4		200GAUI-4 200GAUI-2-S/L 200GAUI-1	QSFP/QSFP-DD SFP-DD
						VR2/SR2								EEI-200G-RTLR-2-S/L	
200G–							DR2-LPO							LEI-200G-PAM4-2	
400GBASE–	KR4 KR2	CR4 CR2				SR16 SR8/SR4.2 VR4/SR4	DR4 DR2	FR8 FR4 DR4-2 DR2-2	LR8 LR4-6/LR4-10		ER4-30	ER8	400ZR	400GAUI-16 400GAUI-8 400GAUI-4-S/L 400GAUI-2	QSFP/QSFP-DD OSFP
						VR4/SR4	DR4	DR4-2 FR4						EEI-400G-RTLR-4-S/L	
400G–							DR4-LPO FR4-LPO							LEI-400G-PAM4-4	
800GBASE–	ETC-KR8/KR8 KR4	ETC-CR8/CR8 CR4				VR8/SR8 VR4.2/SR4.2	FR4-500 DR8 DR4	FR4 DR8-2 DR4-2	LR4 LR1	ER1-20		ER1	800ZR-A 800ZR-B 800ZR-C	800GAUI-8-S/L 800GAUI-4	QSFP-DD OSFP/OSFP-XD
						VR8/SR8	DR8	DR8-2						EEI-800G-RTLR-8-S/L	
800G–							DR8-LPO							LEI-800G-PAM4-8	
1.6TBASE–	KR8	CR8				VR8.2/ SR8.2	DR8	DR8-2						1.6TAUI-16-S/L 1.6TAUI-8	QSFP-DD OSFP/OSFP-XD

Gray Text = IEEE Standard Red Text = In Task Force Green Text = In Study Group Blue Text = Non-IEEE standard but complies to IEEE electrical interfaces

Orange Text = LPO MSA specification in early stages of standardization, not compliant with IEEE electrical interfaces

Dark Blue Text = OIF specifications, a new electrical interface, an aternative to IEEE 802.3 AUI for 500m and 2km PMDs, providing interoperability over fiber.

Warning! The Ethernet landscape is evolving rapidly – technologies listed here are subject to change.

ENERGY EFFICIENCY IN THE AI WORLD

- Blackstone estimates a **40% increase in electricity demand** in the United States over the next decade.[1]
- Gartner estimates the power required for data centers to run incremental AI-optimized servers will reach 500 terawatt-hours (TWh) per year in 2027, which is **2.6 times the level in 2023.**[2]
- The largest data center market globally is in northern Virginia, and the local utility, Dominion Energy, expects **power demand to grow by about 85% over the next 15 years,** with data center demand quadrupling.[3]
- SemiAnalysis forecasts Global Data Center Critical IT power demand will surge from 49 Gigawatts (GW) in 2023 to 96 GW by 2026, of which **AI will consume ~40 GW.**[4]
- Goldman Sachs forecasts **global power demand** from data centers **will increase 50% by 2027** and by as much as **165% by end of the decade** (vs 2023).[5]

Provision of energy to, and removing heat from, AI data centers is becoming a controlling limit. Data centers will account for about ~2% global electricity use in 2025 and their power usage is expected to double to more than 1,000 TWh by 2030 driven by GenAI.[3]

Ethernet is not the biggest power consumer in the data center, but it is material. Any watt used on the network is a watt not used on the main workload. It's expected that the Ethernet Industry will keep driving down the picojoules per bit with new technologies.

New PHY technologies, copper and optical interconnect advancements, and intelligent workload-aware traffic management are helping optimize energy use. Additionally, collaboration with AI-driven power management is emerging to further reduce energy waste. As Ethernet scales to 1.6T and beyond, balancing performance and energy footprint will be critical in supporting this global technology evolution.

1 "Blackstone (BX) Q2 2024 Earnings Call Transcript." The Motley Fool. July 18, 2024.
2 "Gartner Predicts Power Shortages Will Restrict 40% of AI Data Centers By 2027." Gartner. Nov 12, 2024.
3 "As GenAI Asks for More Power, Data Centers Seek More Reliable, Cleaner Energy Solutions." Deloitte. Nov 19, 2024.
4 "AI Data Center Energy Dilemma - Race for AI Data Center Space." SemiAnalysis. Mar 13, 2024.
5 "AI To Drive 165% Increase In Data Center Power Demand by 2030." Goldman Sachs. Feb 2025.

INTERCONNECT TECHNOLOGIES

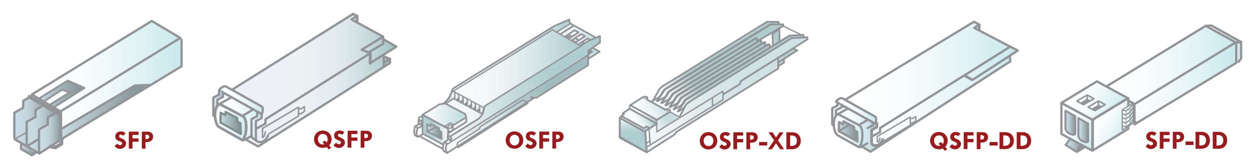

PLUGGABLE MODULES

Linear Pluggable Optics (LPO) and Linear Receive Optics (LRO)

The current high-speed optical market is dominated by retimed optics, but there is rapidly growing interest in linear-based solutions for optical modules which dramatically reduce the module power consumption. Linear Pluggable Optics (LPO) and Linear Receive Optics (LRO) are emerging techniques which remove all/some of the retiming circuitry found in traditional optics.

These implementations utilize common pluggable form factors of QSFP, QSFP-DD, and OSFP and are primarily targeted at 400GbE and higher markets. A fully linear optical module can operate at around half of the power of a similar retimed device. LRO is a half-retimed solution which achieves some of the power reduction while providing a higher quality transmitted optical signal.

CABLE TECHNOLOGIES

Active Electrical Cable (AEC) – Integrated retimer electronics for signal enhancement
Active Copper Cable (ACC) – Integrated redriver electronics for signal boosting
Active Optical Cable (AOC) – Integrated optical transceivers for low-power, high-speed connectivity

Both AECs and ACCs are active cables providing data transmission over copper cables in applications where standard direct attach cable lengths are insufficient. ACCs provide basic signal boosting for increased cable reach in cost-sensitive applications, whereas AECs offer enhanced signal regeneration capabilities suitable for even longer distances.

AOCs integrate fiber optics and embedded transceivers, providing high-bandwidth, low-latency, and low-power connectivity for short- to mediu m-range interconnects in high-speed Ethernet applications.

OPTICAL EVOLUTION

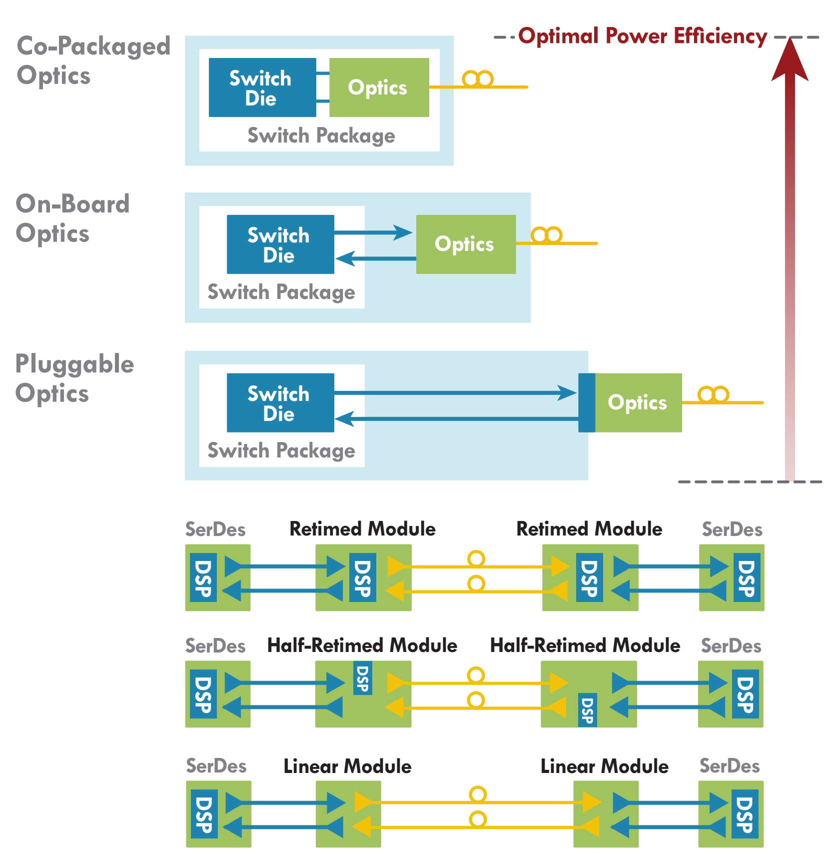

The ever-increasing demand for power efficiency in data centers is driving the transition to new levels of system and optics integration, such as **Co-Packaged Optics (CPO)**, **Co-Packaged Copper (CPC)**, and **LPO/LRO**. As data centers deploy higher and higher link speeds, the power consumption of SERDES links and the optical modules increase significantly. The need for reduced-power optical solutions is fueling innovation and creativity in this market.

To meet diverse deployment needs, **retimed, half-retimed, and linear optical modules** each offer varying levels of signal processing and power efficiency to optimize performance across different network architectures.

ETHERNET SPEEDS

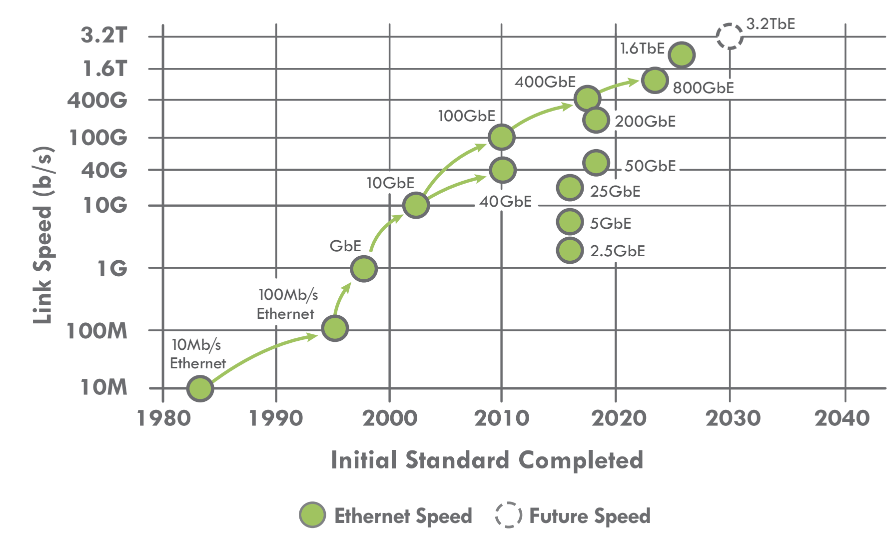

FATTER PIPES

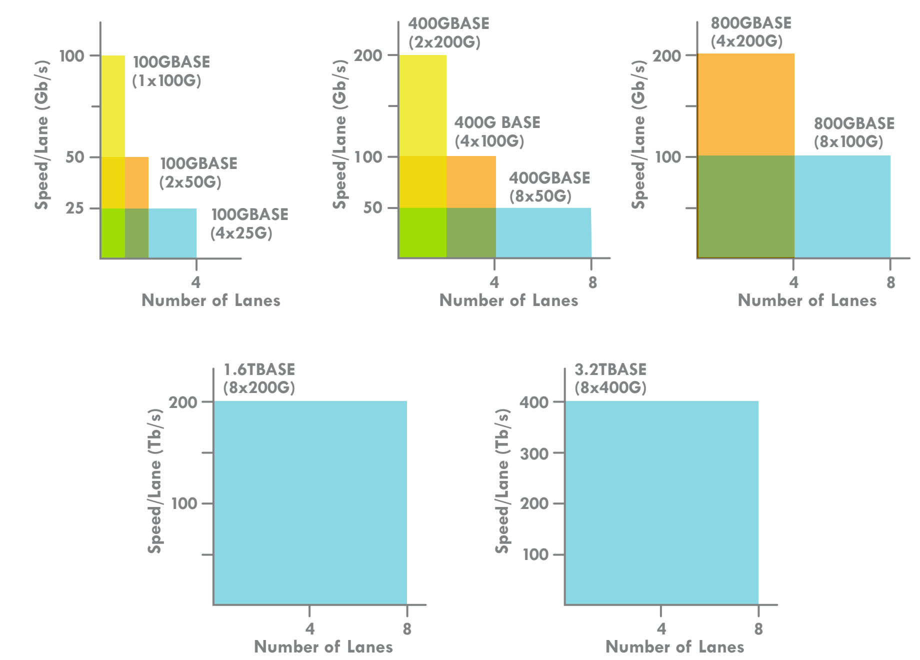

Total throughput (data rate) may be achieved in three general ways:

1 Aggregating multiple lanes **2 Increasing the per lane bit rate**
3 Increasing the bits transferred per sample (Baud)

This chart shows how multiple lanes can be used to generate similar speeds. The per lane speed times the number of lanes determines the total link speed.

PATH TO SINGLE LANE

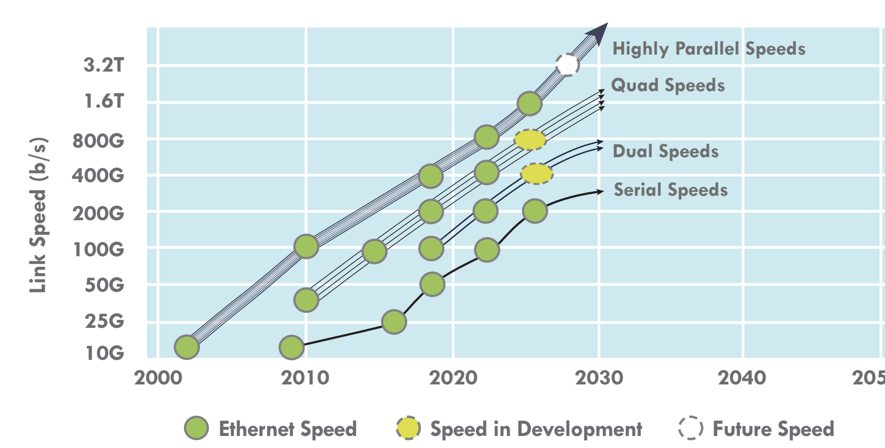

SIGNALING METHODS

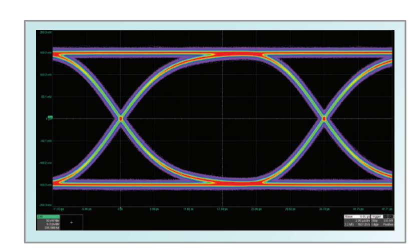

NRZ

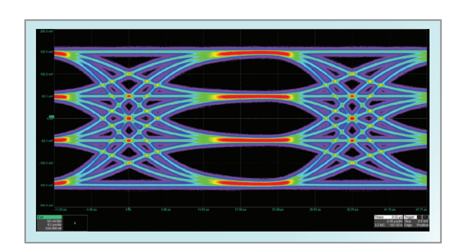

PAM4

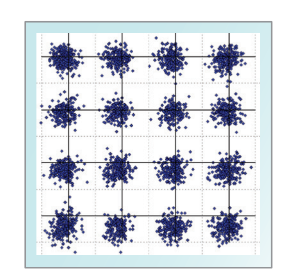

Coherent

Signaling Method Transitions:

- Non-Return-to-Zero (NRZ) used for 25Gb/s per lane and below
- Four level Pulse-Amplitude Modulation (PAM4) for 50Gb/s per lane
- Coherent signaling (both in-phase and quadrature modulation) for 100Gb/s per lane and above.