Whitepaper:
Where Interconnect Meets AI Innovation

RF transmission over plastic cable (e-Tube) solves critical speed, cost, reach and weight challenges for terabit cable interconnect in AI/ML GPU Cluster buildouts

Introduction

"e-Tube" or radio frequency (RF) transmission over plastic cable, is an innovative technology that solves critical speed, cost, distance and weight challenges for high-speed interconnects often impacting today's hyperscale cloud datacenters, AI/ML large language model GPU buildouts, and high-performance computing clusters. As will be discussed, the technology is highly scalable to support speeds in excess of 3.2Tb/s and targeted for short interconnects (<10m) where cost, weight and power are essential.

Figure 1 shows a typical datacenter architecture with servers connected within a rack via a top of rack (TOR) switch. The racks are then interconnected through an additional layer of switches to allow access to all nodes in the system. Figure 2 shows an example of the new AI architectures which require a significant increase in mathematical computation via GPUs or custom devices. The interconnect density between CPUs and the GPU themselves are increasing substantially in the AI architecture. These internal connections are usually short spanning less than 5m.

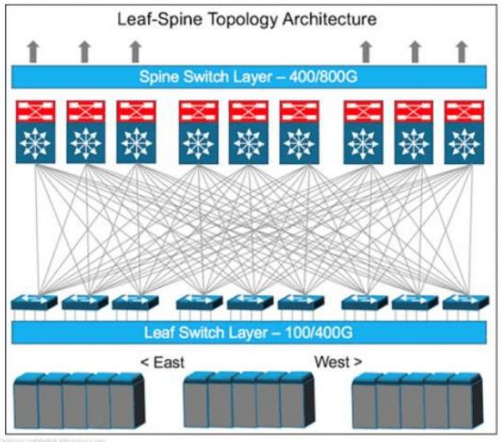

Figure 1: Data Center Leaf-Spin Architecture

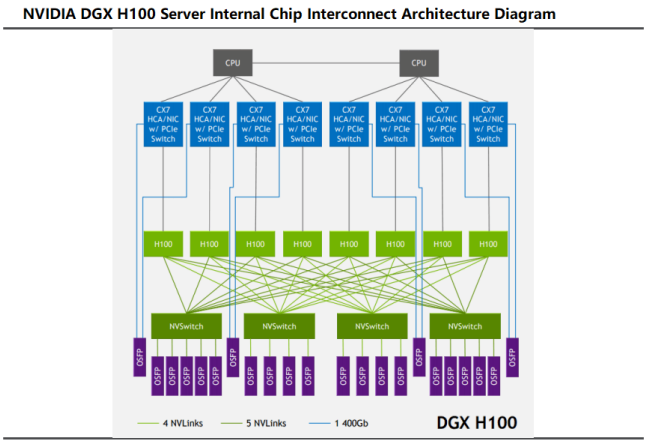

Figure 2: AI GPU & CPU Architecture (NVIDIA)

Interconnect speed and latency are imperative to the buildouts and are the driving force behind accelerating the electrical and optical data rates standardized by the market through IEEE and OIF.

Traditional data centers and communications markets use copper and optical interconnect technologies. Both have their advantages and disadvantages.

Copper is the interconnect of choice for short reach applications due to its low cost, simplicity, and high reliability. The limitation with copper is that, due to skin effects, the channel gets more challenging as you increase frequency and/or decrease the thickness as shown in Figure 3. This effect is seen on the PCB traces on a system linecard, across a backplane and on copper cables. In the cable example, a higher gauge which delivers a thinner and lighter cable with a better bend radius, means much higher insertion loss which translates to shorter distance that a given PHY will be able to drive.

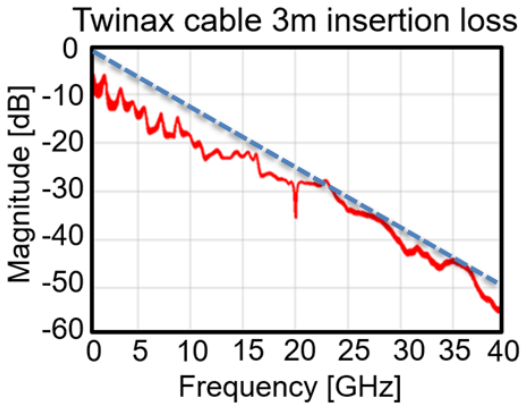

Figure 3: Twinax cable 3m insertion loss

Optical interconnects can provide distances beyond 10km but are usually more complex, power hungry and expensive due the electrical to optical conversions that require additional components such as optical DSP, transimpedance amplifiers (TIAs), laser drivers, and lasers. Typically, copper is used until it can no longer meet the cable reach and/or weight requirements primarily because it is 4-10x lower cost and 2-4x lower power than optical.

Process technology, modulation techniques and DSP equalizer topologies have improved enabling the industry to deliver higher speed PHY devices capable of driving copper traces, connectors and cables for distances that cover up to 5-7 meters. When the PHY on a host ASIC (CPUs/GPUs/switches) alone has not been capable, integrating retimers and/or equalizers into cables (active cables) has delivered sufficient distance without having to resort to optics. As shown in Figure 4, with hundreds of millions of interconnects across boards, backplanes, and top of rack, the incentive to maintain the cost and

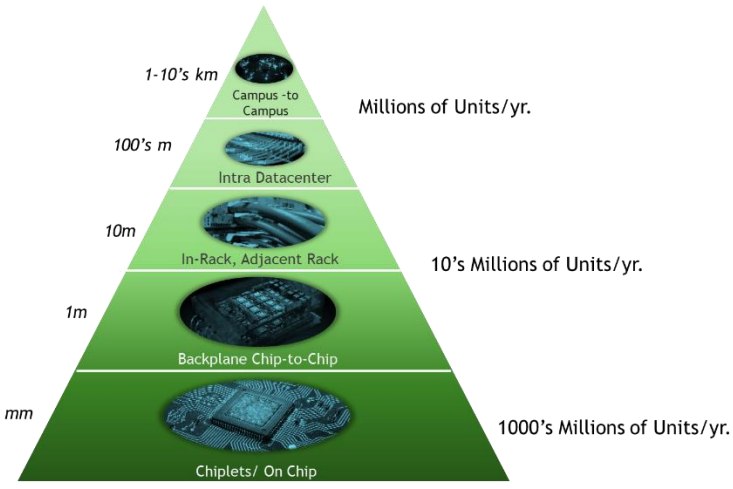

power benefits of copper is unmatched. But at 100Gb/s SerDes speed and above, the PHYs are struggling to deliver the distance with a reasonable copper gauge (weight and size) and at 200Gb/s, even the transitions across the PC board are limited and active electrical cables with retimers will only support short cable reach. Co-packaged optics is one solution that is gaining popularity but while the density and distance supported is impressive, it still suffers from the cost, heat, power, and reliability challenges inherent in optical solutions. While it may provide a reasonable solution for the mid-layer switches in the data center, co-packaged optics is still cost and power prohibitive for the high volume in-rack and adjacent rack and backplane applications.

What is e-Tube

e-Tube is an interconnect technology in which mmWave RF signal is transmitted and received via a plastic dielectric core. A thin metal cladding on top and bottom sides of the core directs the RF signal while preventing cross-talk from other cores in the same cable bundle.

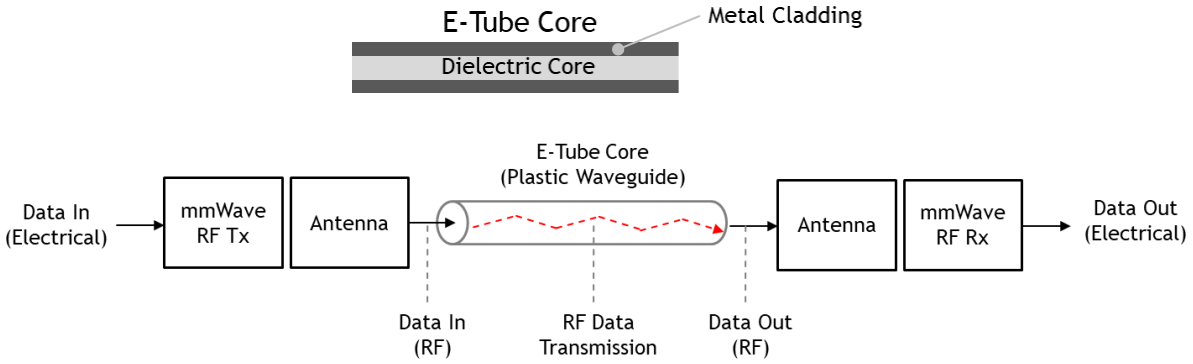

Figure 5: E-Tube Architecture

In the e-Tube architecture shown in Figure 5, a transmitter IC upconverts a baseband signal into a higher band millimeter wave RF frequency. The RF signal is then launched into the e-Tube core using an on-chip Microstrip-to-Waveguide Transition (MWT). The e-Tube core connects the RF transmitter and receiver chips as shown in Figure 6. The e-Tube core to board connector is implemented using a 90-degree hollow waveguide array arrangement, with each dielectric core inserted horizontally into the corresponding opening slot. The connector is designed to couple the signal vertically to the RF transmitter and receiver ICs, which increases the area efficiency of the e-Tube link. Such 90-degree redirection achieves low-profile board-edge connection of the e-Tube. On the receiver side, the e-Tube signal is down-converted by an RF receiver IC to retrieve the baseband signal.

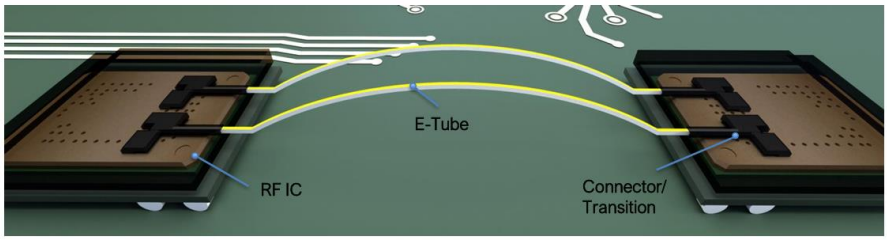

Figure 6: E-Tube Board Connectivity

e-Tube has a similar cost structure to copper. But unlike copper, the e-Tube material has a flat frequency response and fixed bandwidth across the band that does not change as the data rates change. There is no need to compensate for the high-frequency loss like copper that reduces the cable length and/or requires heavier cables. This also means that the e-Tube used for 100Gb/s SERDES can be used for 200Gb/s and future 400Gb/s, making e-Tube scalable for many interconnect generations. The primary use cases for the e-Tube technology in data center infrastructure is for short reach (<7m) cables used in in-rack, adjacent rack, backplanes (typically <2m) and intra-board "flyover" interconnects. It has other potential applications for autonomous vehicles, aerospace, or industrial control where interconnect bandwidth is increasing significantly with cable length, and weight being critical.

e-Tube Architecture

The architecture in Figure 7 outlines an e-Tube architecture that supports 224Gbps over each core using dual carrier frequencies RF transmit and receive to deliver 112Gbps across each carrier.

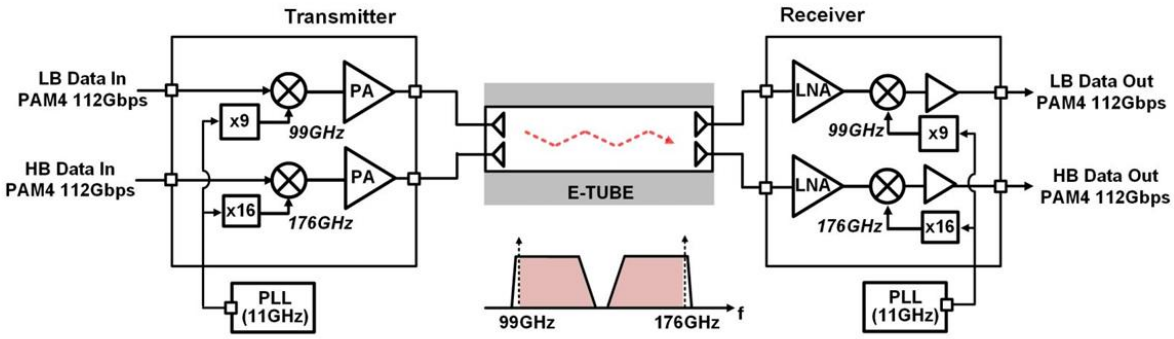

Figure 7: e-Tube Dual Carrier Architecture for 224Gbps per Core

In this implementation, the cutoff frequency of Low Band (LB) and High Band (HB) are at 99GHz and 176GHz, respectively, and each band has 32GHz bandwidth. A Tx chip includes low band (LB) and high band (HB) RF Tx, and each Tx up-converts the 112Gbps PAM4 input data into the corresponding carrier frequency band. As shown in the drawings in Figure 8,

the two Microstrip-to-Waveguide Transitions (MWT) for LB and HB are implemented on the package substrate and transfer the LB and HB Tx outputs from the chip to a waveguide-Tube core. A Chip-to-core connector (referred to as a deflector) is a duplexer with a Y-junction waveguide that combines the LB and HB Tx output into a single e-Tube core.

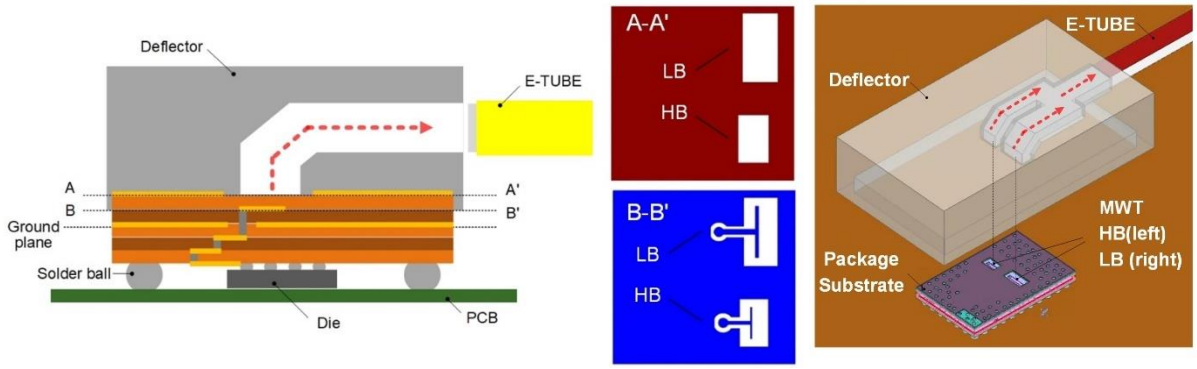

Figure 8: e-Tube Chip to Core Connector

The signal propagates through the e-Tube core and the deflector, MWTs, and the RF Receiver splits and down-converts the LB and HB signals. The LB and HB signals utilize two different carrier frequencies which are generated by the corresponding frequency multiplier while receiving the external 11GHz PLL output.

The electrical characteristics for an 800G (8x112Gbps) cable are shown in Figure 9. The RF ICs are designed to support link budget for a cable distance of up to 7m. Also note that the latency which includes both in the electronics and flight time in the cable is low and the power per module is only ~5.0W.

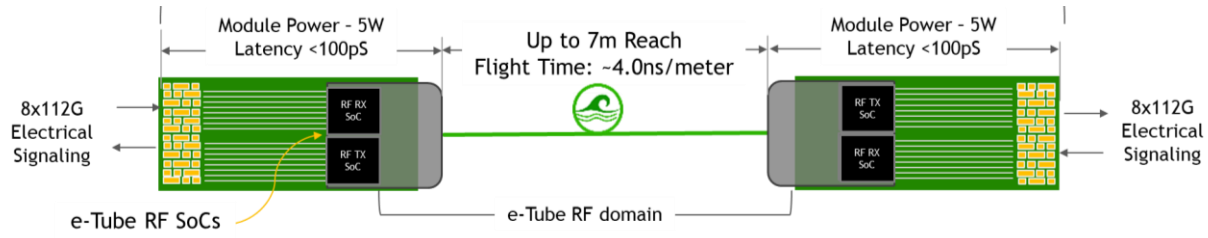

Figure 9: 800G e-Tube Cable

The e-Tube cable will support MSA defined form factors such as OSFP, OSFP-XD, and QSFP-DD. For 800G, the cable will contain 8 e-Tube cores with a total diameter of 8,1mm. This equates to a diameter and bend radius similar cable built with 32AWG copper wiring. Figure 10 shows a comparison of the 800G dual carrier e-Tube solution compared to an 800G copper AEC with 112Gbps per pair.

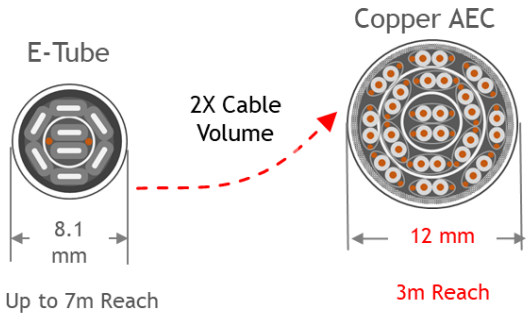

Figure 10: E-Tube Core vs Copper AEC

Figure 11 below shows the eye diagram of the LB and HB RX signals at 112Gb/s PAM-4 with a PRBS pattern of 2^{15}-1 over a 3m e-Tube channel. A BER of less than 10^{-10} is estimated from the results of the eye diagram. The group delay (latency of the link without the channel) of RF Tx/Rx with the channel is 80ps which only comes from the RF and the analog circuits.

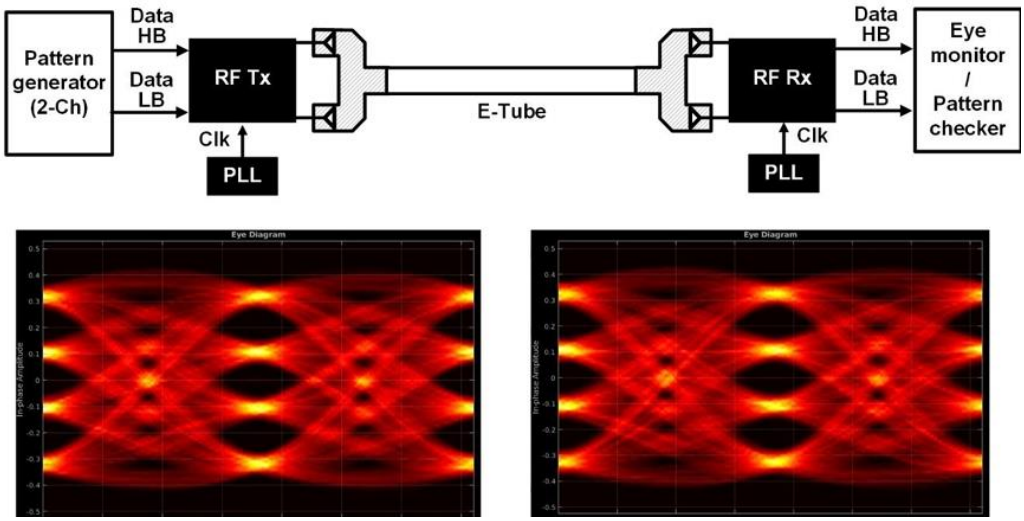

Figure 11: e-Tube Eye Diagram at 112Gbps PAM-4

Scalability of e-Tube beyond 800Gbps

Because of its flat frequency response and fixed frequency across the band, e-Tube is highly scalable. The same dual carrier architecture will be extended to support 224Gb/s PAM4 per carrier at 130GHz (LB) and 260GHz (HB) with 60GHz bandwidth per carrier, enabling the same e-Tube core to support 448Gb/s. Figure 12 shows the solutions for 800Gbps and 1.6Tbps using current serdes data rates of 112G and 224G from the electrical domain and a combination of carrier frequency, channel bandwidth and modulation in the RF domain. Even higher data rates up to 3.2Tbps and beyond can be supported using

increased carrier frequency and advanced multiplexing and/or modulation schemes while maintaining the same base e-Tube construction.

	Electrical Domain			E-Tube RF Domain				
Cable Generation	SERDES Data Rate	Modulation	Coax Pairs (TX & RX)	Carrier Frequency	Channel Bandwidth	Modulation	RF Data Rate	E-Tube Cores (TX & RX)
800G	112G	PAM4	16 (8 Tx, 8 Rx)	99GHz/176GHz	30GHz/30GHz	ASK	224G	8 (4 Tx, 4 Rx)
1.6T	224G	PAM4	16 (8 Tx, 8 Rx)	130GHz/260GHz	60GHz/60GHz	ASK	448G	8 (4 Tx, 4 Rx)

Figure 12: Electrical & e-Tube RF Characteristics and Scalability

e-Tube Technology Implemented in standard MSA Form Factors

The e-Tube cable will be deployed in MSA defined module form factors including QSFP-DD, OSFP and OSFP-XD. The RF IC, connectors, antenna and cable are interconnected inside the module form factor. Figure 13 outlines the implementation of the e-Tube technology in standard module form factors.

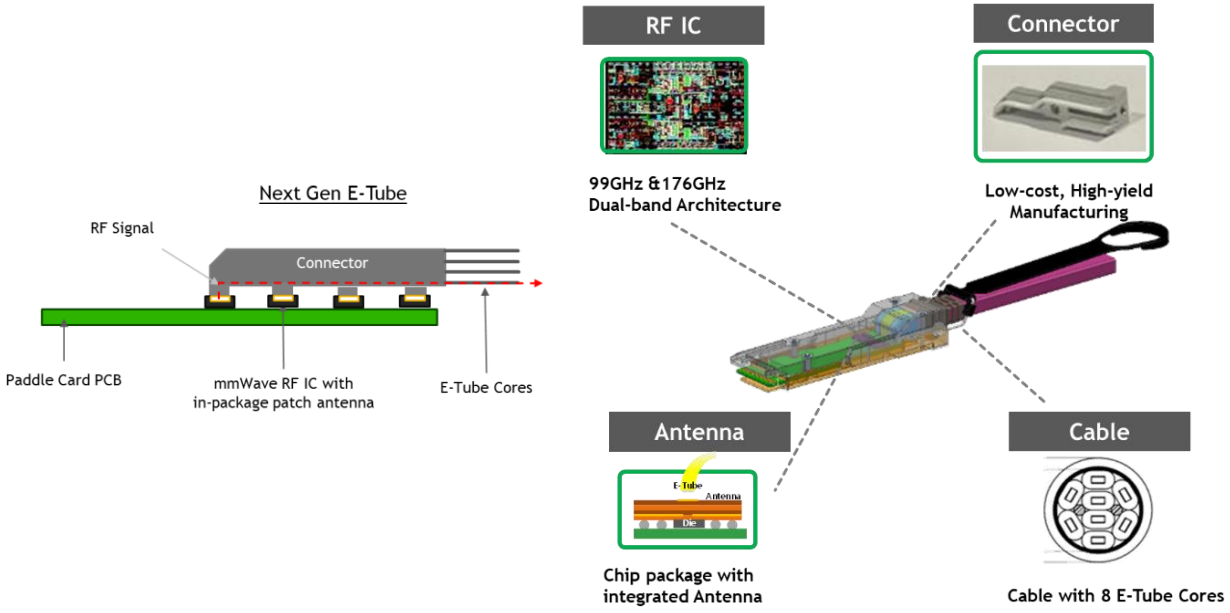

Figure 13: e-Tube Implementation in Standard MSA Form Factors

Comparison of Short and Medium Reach MSA Cable Options

The picture below outlines the standard types of module interconnect used today for short-to mid-range interconnect. The lowest cost option is always the direct attach cable (DAC) which has no active components and relies only on an external PHY to drive the cable. The active copper cable (ACC) uses an equalizer (typically analog) to help improve the signal integrity of the link. These components are generally smaller and lower power than a full DSP-based solution but with less improvement in distance and/or thickness. At high speeds, the Active Electric Cable (AEC) typically includes a full equalizer CDR with DSP to support longer and/or thinner cables. Finally, for long reach capability up to 2km there is the active optical cable (AOC). The AOC must include an optical DSP and all the optical components required to support the optical link including photo diode and TIA for receive and a laser driver and laser for transmit. The DSP and optical assemblies raise the power and cost significantly. There is also a new optical version called linear drive optics (LDO), a derivative of the AOC, that lowers power by eliminating the power-hungry DSP and leverages the board level PHY along with a small level of equalization available from the optical solution.

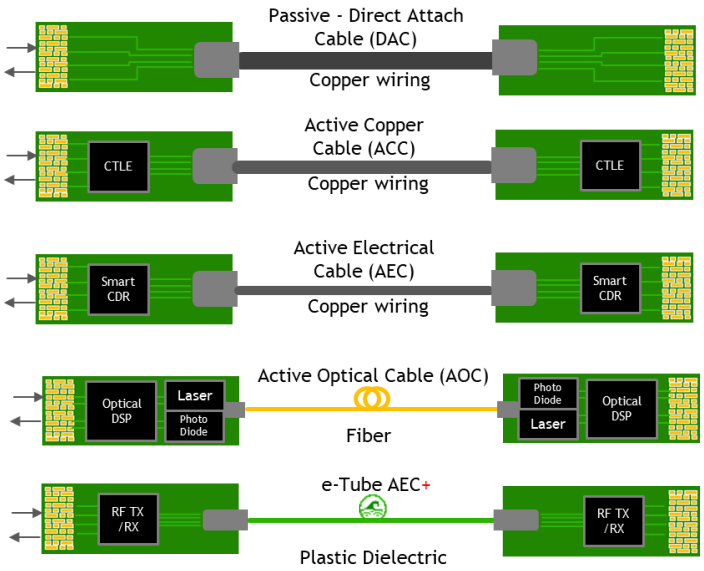

Figure 14: Standard Module Interconnects

As shown in Figure 15, e-Tube provides fundamental energy savings of >50-70% compared to the optical AOC and LPO solutions as indicated in the energy efficiency graph below with economics similar to that of the copper ecosystem.

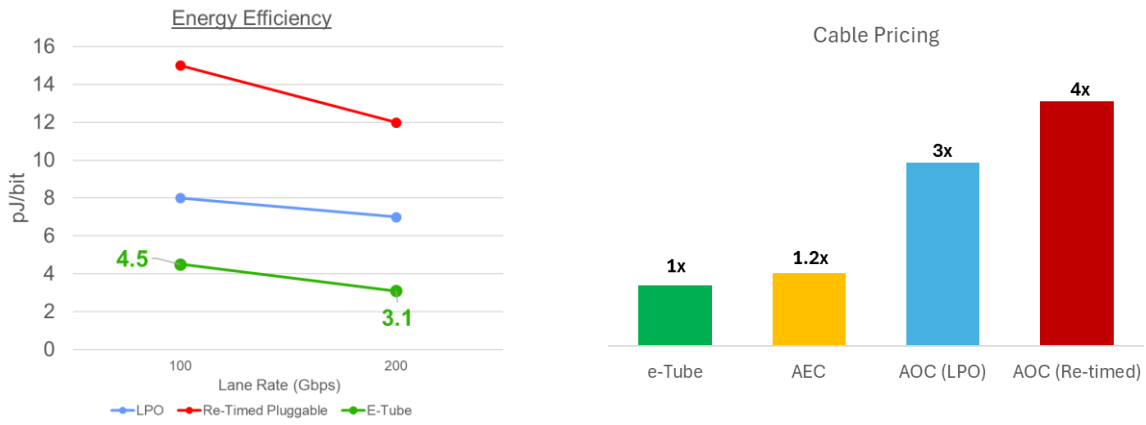

Figure 15: e-Tube Energy Efficiency and Price Comparisons

Figure 16 summarizes the fundamental trade-offs between e-Tube and the other solutions for in-rack and adjacent rack applications based on critical attributes. Power is an issue for these large data centers so as data rates increase and copper becomes more limited, adding a massive increase in power for optical solutions is not practical.

Attribute	e-Tube	AEC Retimed	AOC LPO	AOC Retimed
Power Consumption	🟢	🔴	🔴	🔴
Cost	🟢	🟢	🔴	🔴
Module Latency	🟢	🔴	🟢	🔴
Cable Reach	🟢	🔴	🟢	🟢
Mechanical Form Factor	🟢	🔴	🟢	🟢
Serviceability/Interoperability	🟢	🔴	🟢	🟢

Figure 16: e-Tube Comparisons

Conclusion

e-Tube solves fundamental copper and optical interconnect challenges in AI/ML data center applications. The scalable technology provides the cable reach, density, weight and bend radius for in-rack and adjacent rack applications and is dramatically less expensive and more power-efficient than available optical solutions, while maintaining the cost and power benefits of copper.