





The Transmit (Tx) Path

An internal laser turns your outgoing electrical bits into pulses of light.




The Receive (Rx) Path

A Photodetector (PD) captures faint incoming light, converting it back into microamps of electricity. A Transimpedance Amplifier (TIA) sits right next to the PD to amplify the weak signal to a usable voltage.



The Fiber Capillary



Send and receive share a single bidirectional fiber cable. They use two different wavelengths—one upstream, one downstream—so data doesn't collide.

Dozens of households take turns on this single fiber, merging at a 1:32 or 1:64 passive splitter before reaching the operator's central office.





Copper (Electricity)

- **Cost:** Cheap; always the technology that **arrives first**. Always the technology that arrives first.
- **Distance Limit:** Maxes out at roughly **3 meters** in a data center before the signal severely degrades.




Light (Fiber Optics)

- **Cost:** Expensive to deploy; **requires physical trenching or retrofitting** old buildings.
- **Distance Limit:** Total attenuation climbs incredibly slowly. The signal loss per kilometer is **tiny** and fixed.

Whether an operator bothers to retrofit an old apartment block depends entirely on who is paying. Once bandwidth demands rise beyond a few meters, copper hits a wall.




Inside an 800G module is a symmetric electrical-to-optical signal chain



The Scaling Problem

Training a frontier AI model requires wiring over 100,000 GPUs into a single, synchronized network. This spans thousands of racks across massive floors, requiring latency to stay impossibly low over hundreds of meters.

The Scale Multiplier



Top-of-Rack (<3m): Servers connect to the first switch via cheap Copper DAC.

Beyond the Rack: East-West Traffic flowing sideways between servers dominates the network. Spanning these distances requires light.

The Ratio: Depending on network topology, a single GPU demands 2 to 3+ high-speed optical modules (800G or 1.6T).


GPUs computing in parallel are like a row of wide-open mouths all eating bandwidth at once.



Analog design difficulty climbs exponentially as speed generates heat

The 1.6T Heat Problem: 1.6T is not sixteen 100G lanes; it is eight 200G lanes. Doubling the lane rate means redesigning every chip to handle immense heat. An 800G module draws ~16W, and the DSP alone consumes half of that.

QSFP-DD




Width: 18.35mm

Boasts the largest installed base due to backward compatibility with 200G/400G, but struggles to shed the heat of next-gen speeds.



OSFP



Width: 22.58mm

A slightly larger physical cage providing more surface area for heat dissipation. It is the preferred, mandatory choice for 1.6T high-power AI deployments.

Electricity that co-packaged optics is trying to save is entirely dictated by the exponential thermal cost of these DSPs. Speed generates heat, and heat dictates design.