




The SerDes Squeeze

Chips must talk by squeezing massive parallel data into a single, high-speed **serial stream** on one wire pair (Serializer/Deserializer).

The Speed-Up

10G → 25G → 56G → 112G → 224G

AI changed the tempo: AI training compute doubles every 5.2 months, cutting the traditional 2-3 year hardware upgrade cycle in half.



The Skin Effect:
At high frequencies, current violently crowds the surface. The center of the wire is completely wasted.

Dielectric Loss:
The surrounding insulation absorbs the high-frequency energy, turning it into heat.

At 224G, insertion loss reaches ~40 dB.
Only 1% of the original signal survives the trip.



WA!

The signal arrives smeared into mush.

Equalization (MLSE/Viterbi):
The receiver isn't just listening; it is actively fishing the signal out of the noise.



DSP

The Power Toll

This algorithmic fishing is exhausting. The DSP consumes roughly half the power (6-8W) of an entire 16W 800G optical module.

The Shrinking Radius

Copper never broke, it just got shorter. With every doubled generation, the economic radius of copper shrinks:


56G: 3 meters

112G: 2 meters

224G: 1 meter (or less)

Inside the rack, copper still reigns.
Example: NVIDIA's NVL72 rack uses 5,184 copper cables (3.2km total) because it requires zero optical conversion power and almost never fails.

NNGH!



Copper (Taxi Meter)

Pays by Distance. The further you go, the steeper the energy and signal loss.

Light (Toll Booth)

Pays by Port. You pay a fixed, high toll for optical conversion (lasers/DSP cost), but distance is nearly free (only 0.2 dB loss per kilometer).

The Crossover Slides Left:

As data rates double, copper's cost line gets steeper. The break-even point where light becomes necessary moves closer to the server every generation.

