




THE INTERCONNECT
FOUNDATION FOR
HIGH-PERFORMANCE AI
TRAINING AND
INFERENCE



TRAINING AND INFERENCE

Different computational requirements



PASSAGE PHOTONICS ACCELERATES ALL THREE.



3x

Reduction in
Time-to-Train

3x

Reduction in TTFT
for inference prefill


2x

Faster inference
decode



TRAINING


Train trillion parameter MoEs faster





INFERENCE PREFILL

Passage reduces TTFT



Lightmatter prefill simulator. Batch size 8192, 1024, and 64 for 1K, 8K, and 131K content length



INFERENCE DECODE

Passage lowers inter-token latency (ITL)

2x faster decode
vs Copper

Across short-to-long context lengths

ScaleScope — Lightmatter LLM workload simulator

Model: 10T MoE LLM · 340B active parameters · 1,024 experts

Hardware Assumptions: Optically-connected 16x128 XPU HyperX with 4.0 TB/s link BW, 0.5 us hop latency vs. copper-based 128-XPU racks with 4.0 TB/s scale-up BW at 1us inter-hop latency, 200 GB/s scale-out at 4 us scale-out latency



THE LASER AND FIBER WALL

Solved by BiDi and VLSP



BiDi

One fiber, Tx
and Rx paths
included.



Photonics has been viewed as a dongle for electrical SerDes

56G → 112G → 224G → 448G

ETHERNET SERDES ROADMAP



With the OCI MSA, photonics is in the driver seat.
DWDM and BiDi

Gen 1 200 Gbps

Gen 2 400 Gbps


Gen 3 800 Gbps

Gen 4 1600 Gbps



PHOTONICS

BiDi cuts wiring count in half





Q1 2027

Universal NPO optical engine. BiDi.



PASSAGE L20

6.4 Tbps per direction

12.8 Tbps aggregate

32 ports × 200 Gbps per lane

200 Gbps PAM4 SerDes · IEEE 802.3dj

2-λ BiDi 1311 / 1331 nm · DR-compatible

30 W TDP max



Available in the NVIDIA NVLink Fusion Ecosystem



BiDi

150 miles

of fiber saved

20,000

Fewer connectors
and failure points

16%

Savings in total
network spend

for a typical 512 XPU scale-up cluster, 32 Tbps / XPU





VLSP

Very Large
Scale Photonics




Conventional Discrete Laser Bottleneck Scaling and Supply Constrained






ACCURACY AND STABILITY NEEDED

The future of photonic links requires accurate, precise, and multi-lambda lasers.





CONVENTIONAL COMPONENTS.



DISCRETE InP LASERS

GUIDE IT'S A CHIP.



GUIDE™ VLSP LASERS



128+

Lasers per chip

0.1 FIT

Reliability and
self-healing

0.25 dB

Power Accuracy

64λ+

multi-color


300mm

CMOS SOI Wafers

3 GHz

Frequency Accuracy


GUIDE



CONVENTIONAL



Each laser requires many discrete components.



Optical components are integrated on chip.




PHOTONICS IN EVERY FORM


The 10-year roadmap for photonics in silicon today

Optical Engine

NPO
2027



CPO
2028



Interposer
2029+





**All of these platforms in our
validation datacenter and shipping to
customers today.**





OPEN SILICON PHOTONICS FOR AI SYSTEMS

**Growing ecosystem for CPO. XPU tray,
switch tray, fiber routing, rack design.
Via broad collaboration in the OCP.**



Celestica

CORNING

DELL Technologies

flex

FOXCONN

GUC

hyvesolutions

intel

JABIL

KEYSIGHT
TECHNOLOGIES

Qualcomm

QCT

wiwynn

With Lightmatter and others

