Photonic Fabric : Photonic Scale-Up Network for Accelerated Computing

Hot Interconnects 2025

Ravi Mahatme for team Celestial AI

celestial AI

Introducing Celestial AI

Celestial AI is creators of the Photonic Fabric™ Technology Platform – a revolutionary photonic interconnect platform

Ultra-Low Latency | Lowest Power | > Tbps Bandwidth

Photonic Fabric is the only optical interconnect technology that *can* be integrated in middle of a silicon die – bypassing the beachfront for IO

Leveraging Photonic Fabric, we build connectivity & system products to transform AI & HPC infrastructure

Scale-Up Connectivity

Shared, high-capacity, low-latency memory

Revolutionary Photonic Link for Accelerated Computing

Photonic Fabric Link (PFLink ™)

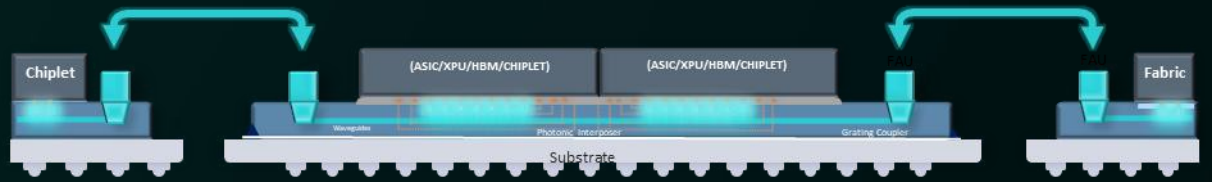

01. Compact, Thermally Stable Optical Modulator
    Chip-to-chip packaging with XPUs dissipating 1000s of watts

02. Multiple Packaging Options Tailored for Customer Applications
    Compatible with industry standard packaging flows

03. No DSP – Linear Drive Optics
    High SNR, low BER, close proximity of optics to electronics:
    Eliminates need for DSP

04. Full Stack E-O-E Link Optimization
    Protocol adaptive Network Convergence Layer (NCL)
    Full Electrical to Optical to Electrical (EOE) link management
    FEC, CRC, FLIT Replay

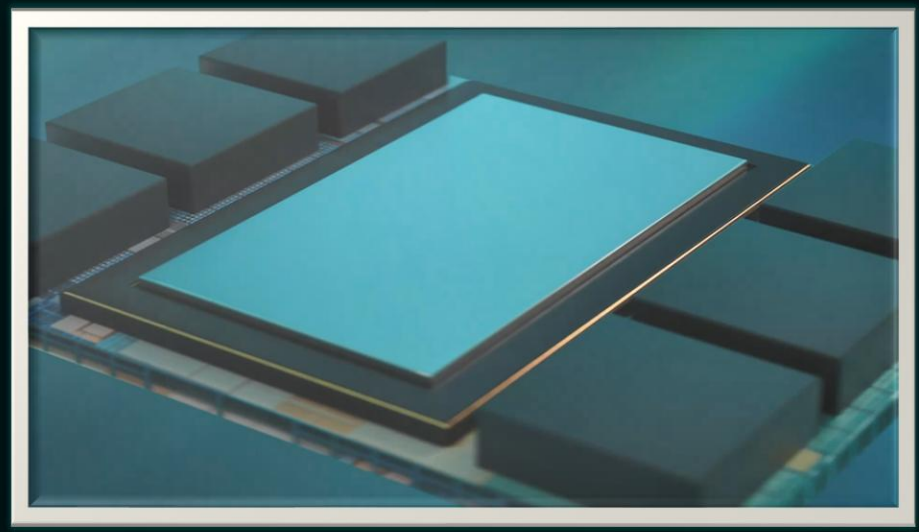

Celestial AI Product Portfolio

Connectivity (PFLink™) Products

1. PF Chiplets: Package level UAL compatible XPU/Switch connectivity (ULTRA ACCELERATOR LINK™)
2. PF IP: Silicon level XPU connectivity
   - AMS IP (E-O-E technology)
   - OMIB™ (Optical Multi-chip Interconnect Bridge)

   (Scale-In + Scale-Up)
3. PF-NIC: MCM[1] package for server level connectivity

(Scale-Up)

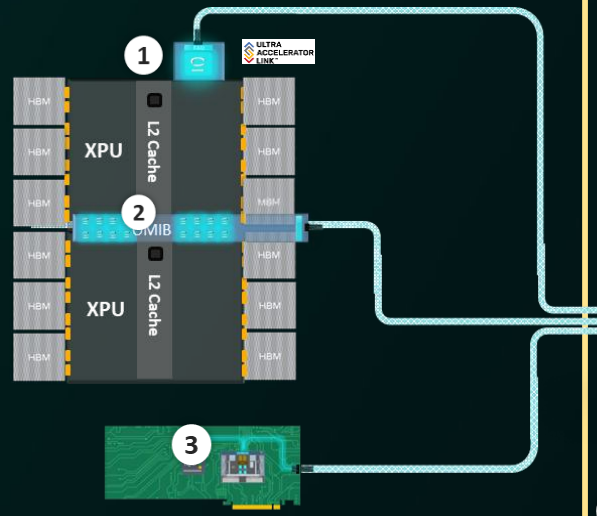

1. MCM = Multi-Chip Module
2. Available with UAL Switch Providers
3: Available for ODM partners

System Products

4. PF Module: MCM[1] package for scale-up fabric & in-network memory
5. PF Switches[2]: Rack-scale scale-up switch
   PF Appliances[3]: Rack-scale scale-up switch with in-network memory

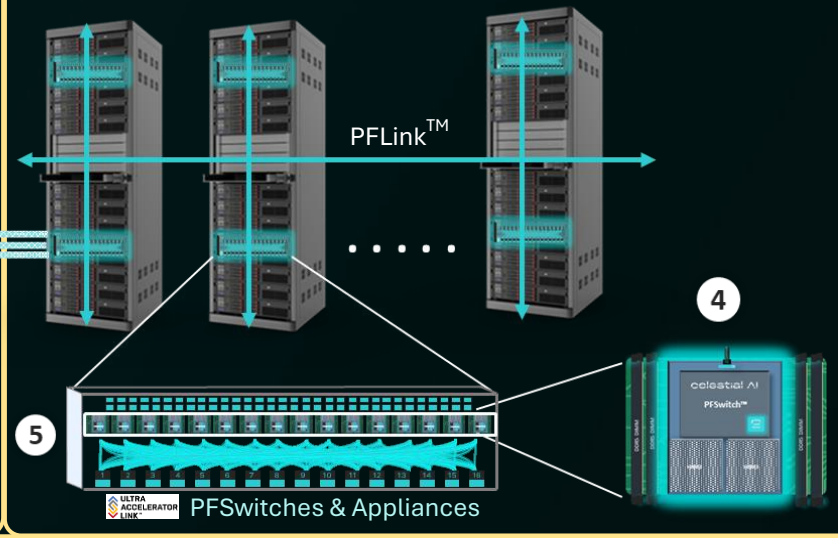

Software: Link, System & Fleet level for link management, telemetry, RAS & Predictive Insights

Photonic Fabric Optical Connectivity Chiplet

For XPU level Integration

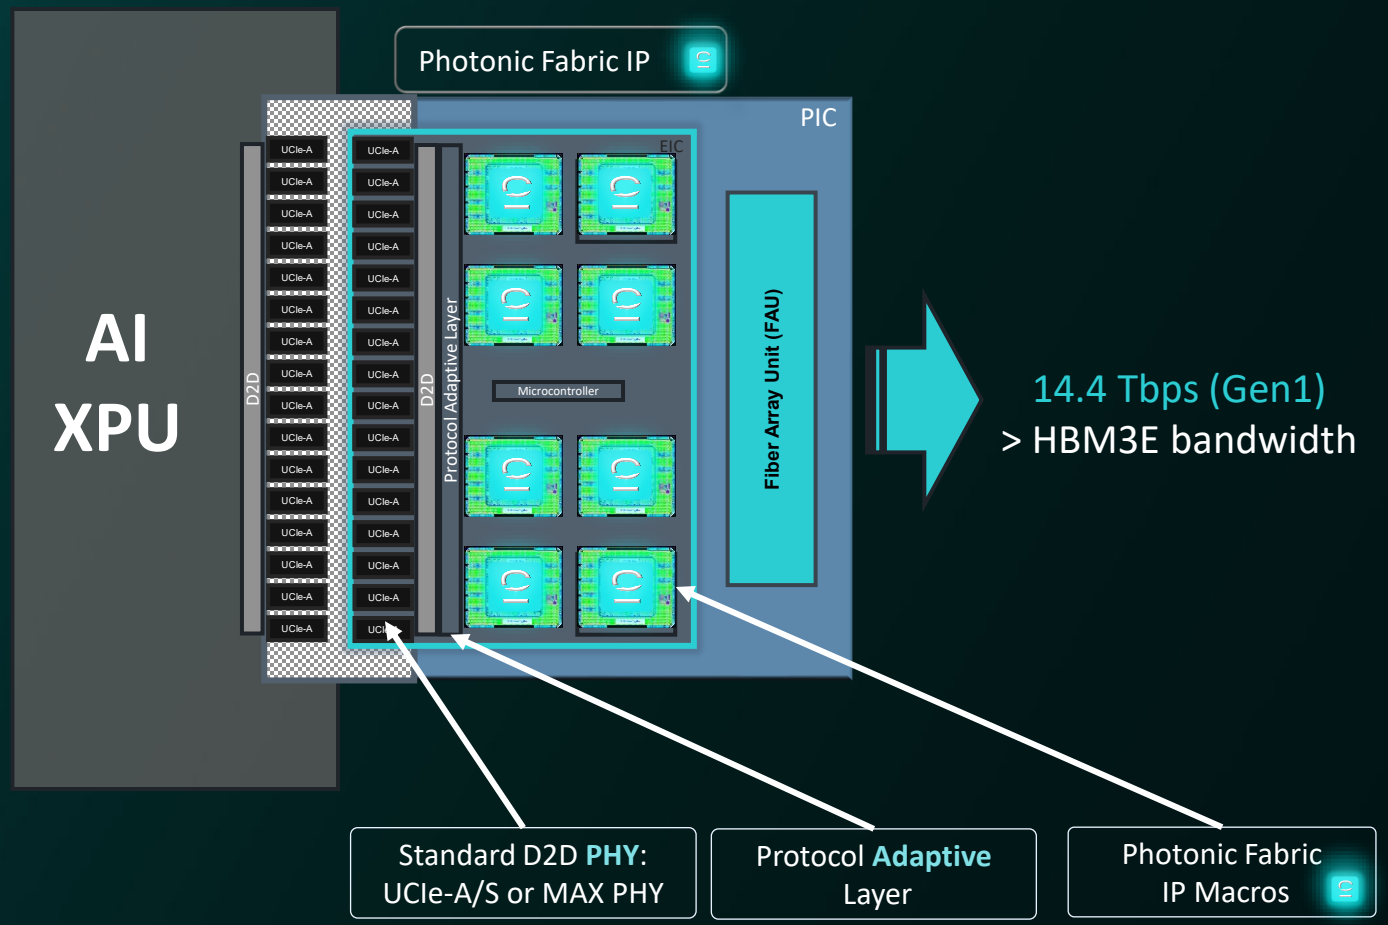

Protocol Adaptive:
AXI, HBM/DDR, UAL, CXL etc.

Electronics (EIC) TSMC 4/5nm	**AMS (Analog Mixed Signal) Macro**	Transmitter, Modulator Driver Analog Equalizer Receiver, TIA SerDes - Photonic Optimized
	OMAC	Link Mngt, FEC, Flit Gen, Retransmission, PCS
	NCL	Network Convergence Layer Protocol adaptive
Advanced Packaging		
Photonic Interposer or Bridge		Photonic Integrated Circuit (PIC) FAU (Fiber Array Unit) EAM: Compact Thermally Stable Modulators Photodiodes SiPh Waveguides Grating Coupler

ULTRA ACCELERATOR LINK™

Photonic Fabric™ Module & Switch/Appliance

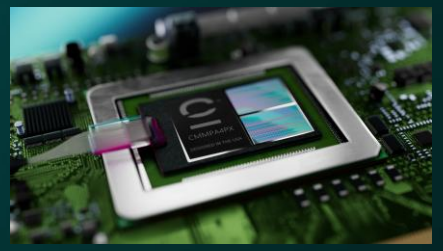

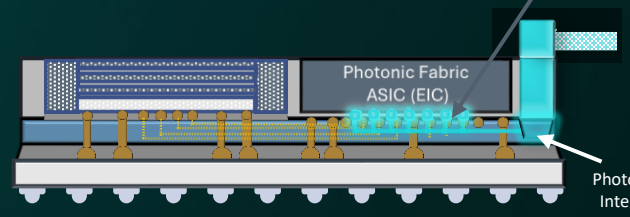

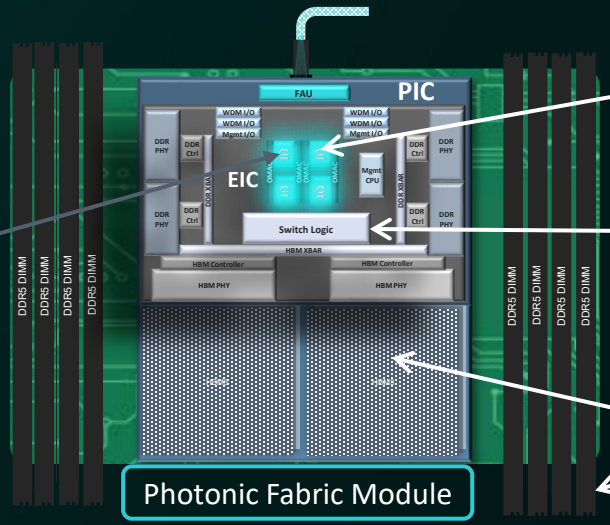

PFLink:
Interior Die Photonics Link

- No Beachfront Limitations:
- Frees space for memory & I/O

PFSwitch:
High bandwidth 8 Tbps Electronic Scale-Up Switch

- 16 ports x 448Gbps

In-Network Shared Memory

- 2 x 36GB HBM3e Stacks
  - 64 virtual channels
- 8 DDR5 channels → 2TB (256GB x8)
- Bandwidth per Module: 7.2Tb/s
- HBM operates as write-through cache for DDR
- Hardware Semaphores

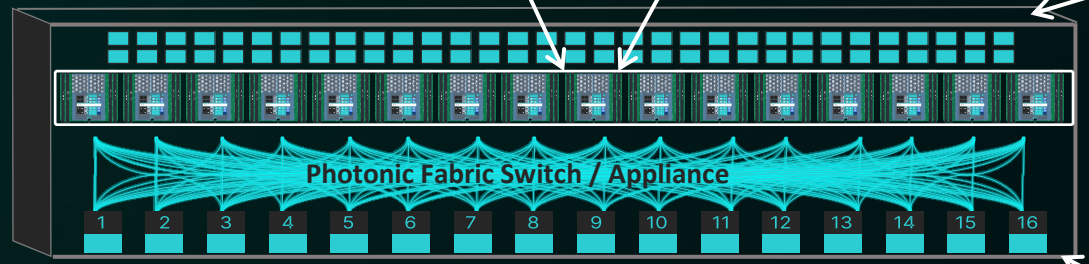

Photonic Fabric Appliance:
16 PFModules = PFSwitch with In-Network Memory

- 115 Tbps Switch
- 16 Radix Photonic I/O ports
- 16x Photonic Fabric Modules in a 2U Appliance
- 33TB Memory Capacity (1TB of HBM & 32TB of DDR)
- Bandwidth & Latency of HBM3E with Capacity & Cost of DDR5
- Built - in hardware semaphores to guarantee synchronization

XPU Integration with PF Appliance

XPU-to-XPU & XPU-Memory

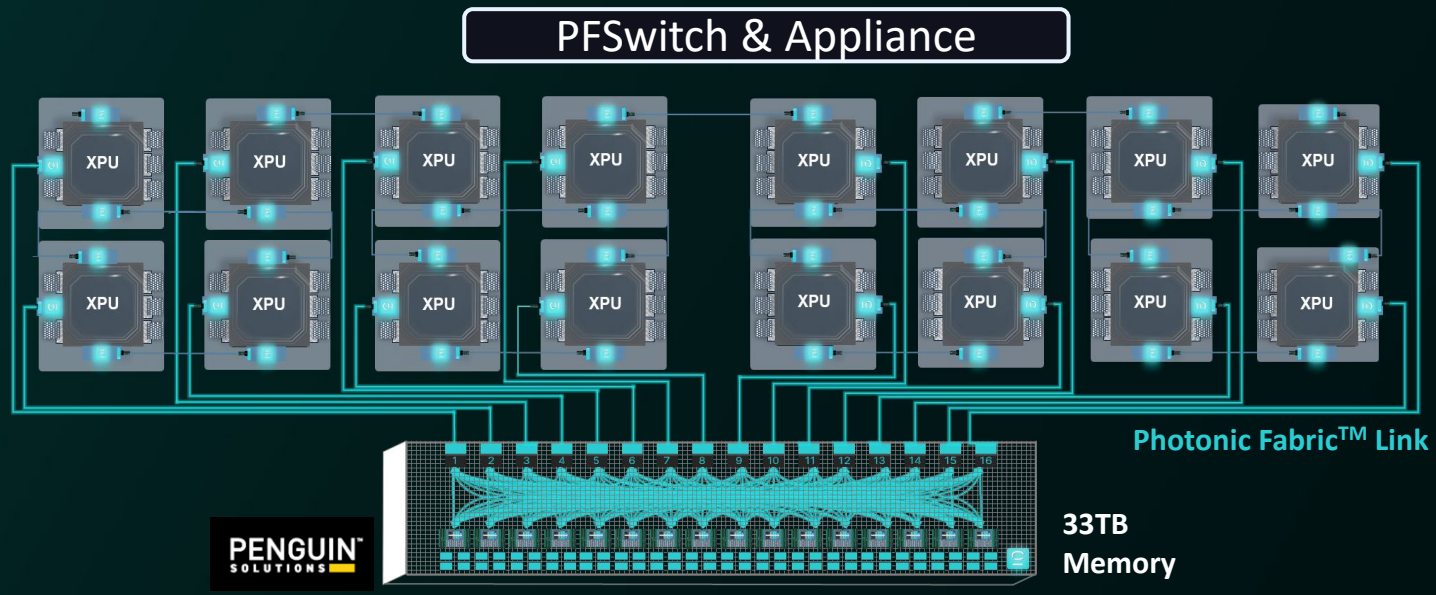

- ✦115 Tbps Switch for scale-up network
- ✦ All-to-All Connectivity for efficient Collective Comms
- ✦32TB Unified/Shared Memory Space
- ✦Broadcast & Reduce across all connected XPUs

Photonic Fabric™ Technology Platform Enables Cluster Scale AI Processing

PF-NIC : Add-In Card for Server Level Connectivity

- Server level connectivity
- Supports CXL 3.0 /2.0 capable server
- Offers low-latency access to PF Appliance (for shared memory)

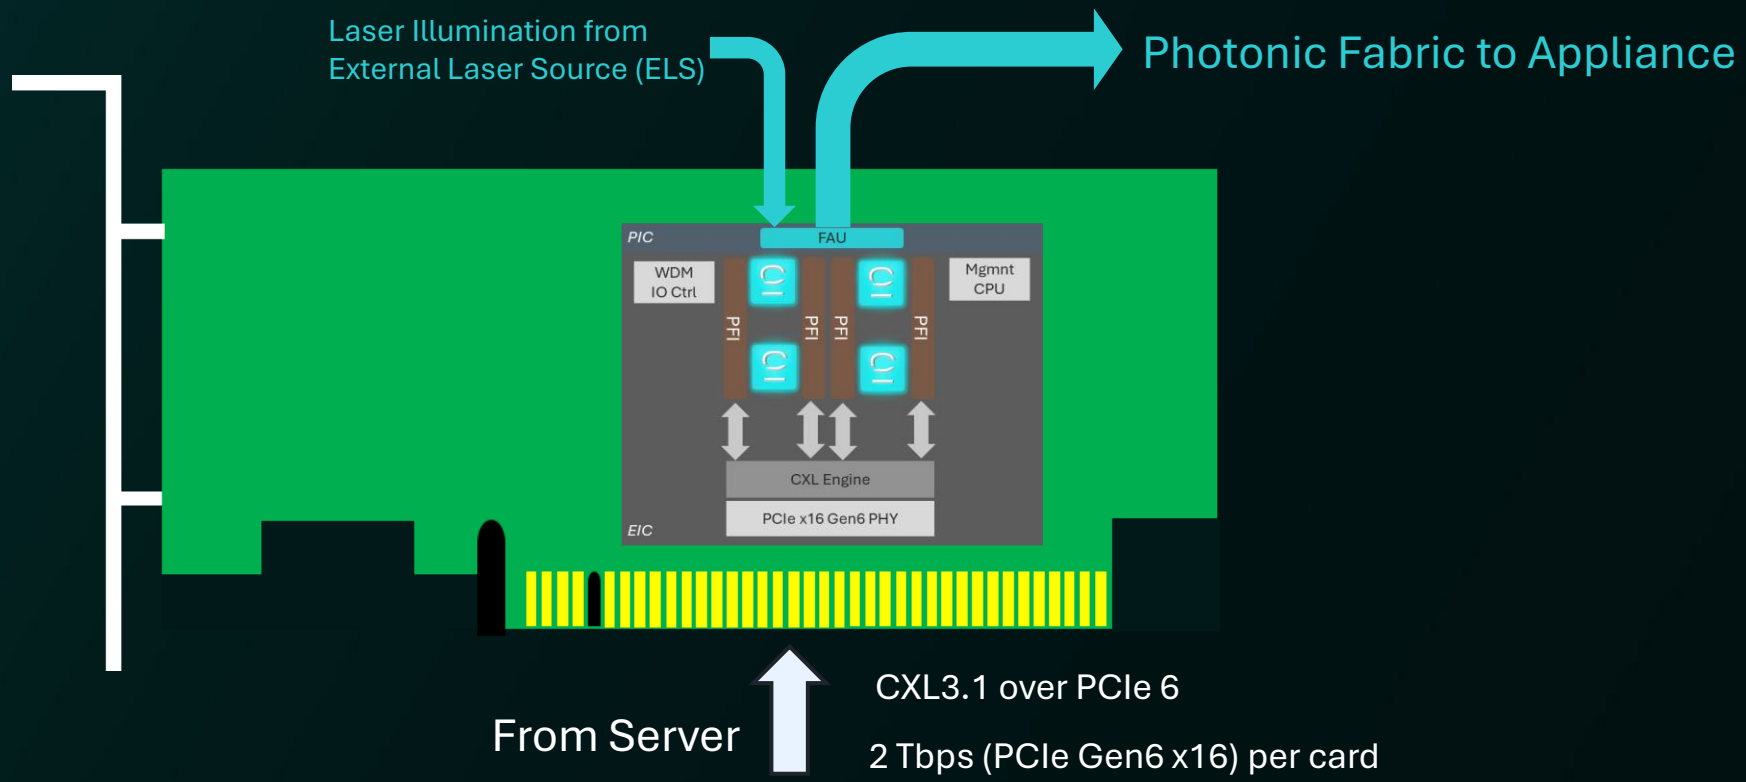

High Capacity, Low-Latency Shared Memory Across Servers

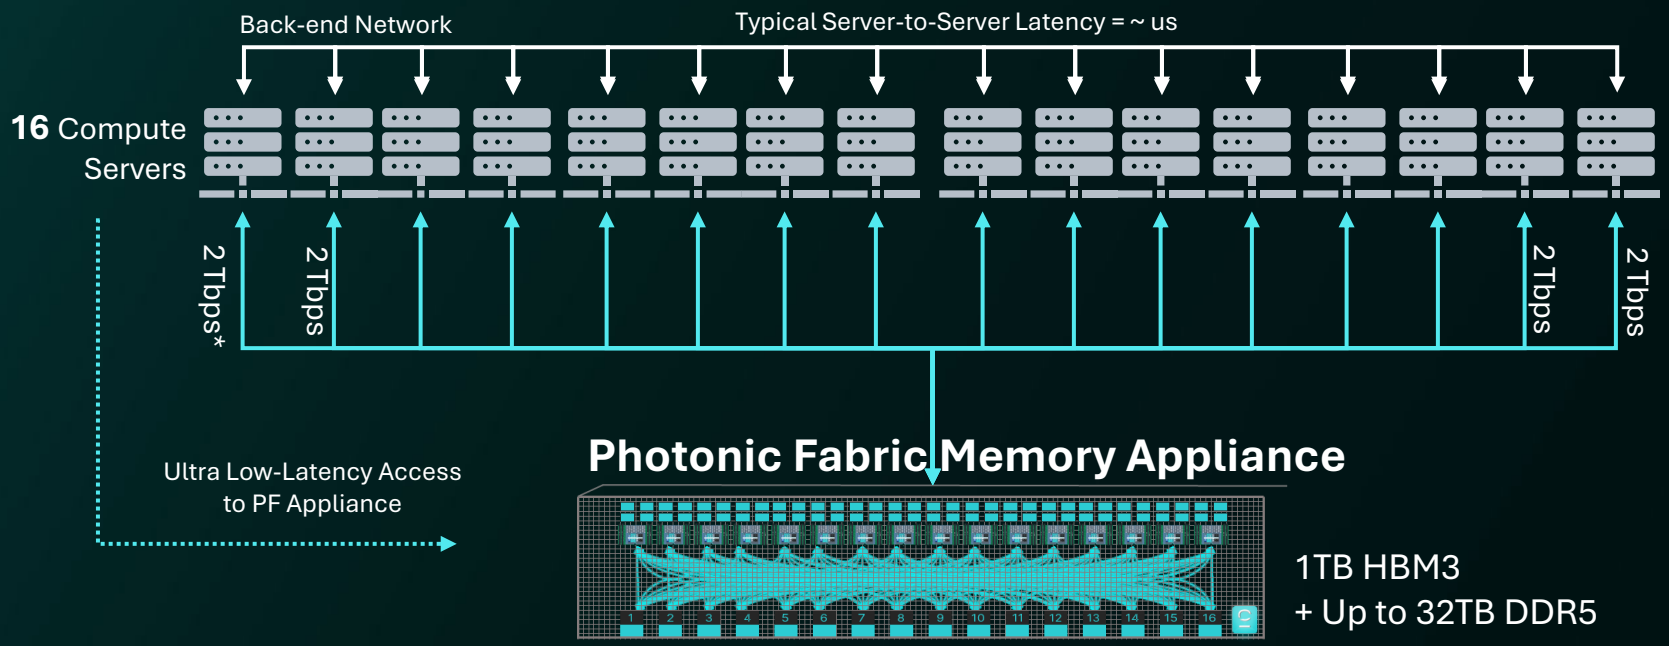

Typical topological configuration with 16 servers connected to PFMA

* 2 Tbps = PCIe Gen6 bandwidth

PFMA Delivers Transformative Benefits Beyond TCO Reduction

Reduce TCO

- Optimize resource utilization
- Lower infrastructure costs
- Decrease operational expenses

Reduce Sharding

Deterministic Tail Latencies (p99 and p99,9)

Improve Throughput/ Performance (QPS, TTFL, TPOT)

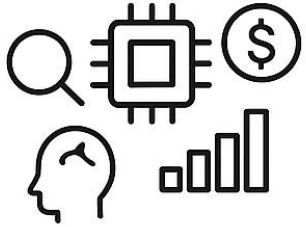

Enables new use-cases & revenue opportunities not possible today

Applications of Photonic Fabric Products

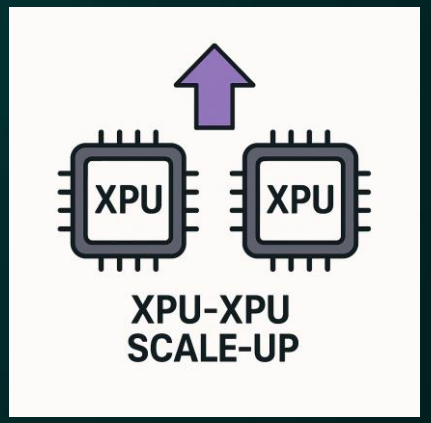

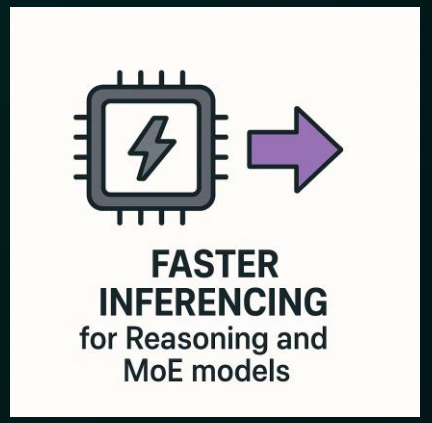

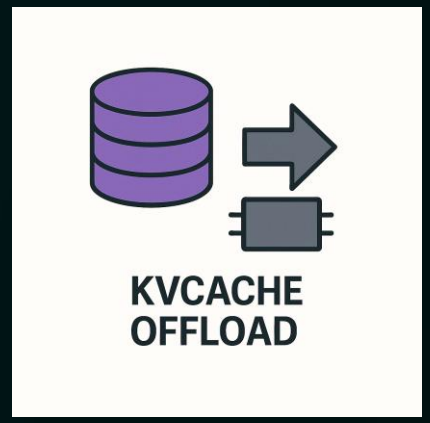

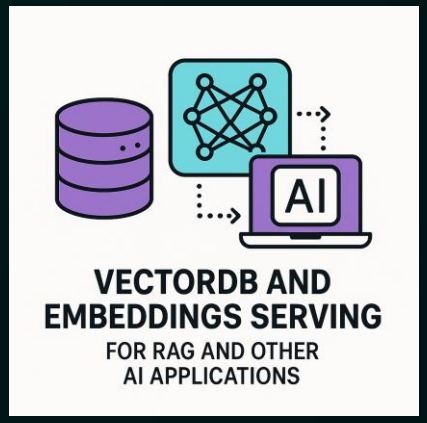

Celestial AI @ HotChips 2025

Aug 26th Aug 2025, 830am PST,
Memorial Auditorium, Stanford University, CA

Sign-Up: https://www.hotchips.org/

Conference Day 2: Tuesday, August 26, 2025

Time (PDT)	Title	Presenters
7:45AM-8:30AM	*Breakfast/Registration*	
8:30AM-10:30AM	*Optical* Chair: Borivoje Nikolic	
	Celestial AI Photonic Fabric Module (PF Module) - The world's first SoC with in-die Optical IO	Phil Winterbottom, Celestial AI

Phil Winterbottom
Chief Technology Officer
Celestial AI

Intelligence,
Illuminated ™