January 6, 2025

Marvell Announces Breakthrough Co-Packaged Optics Architecture for Custom AI Accelerators

- *New Marvell AI accelerator (XPU) architecture enables higher bandwidth and longer reach scale-up fabric connections for custom AI servers.*
- *XPUs with integrated Co-Packaged Optics (CPO) enhance AI server performance by increasing XPU density from tens within a rack to hundreds across multiple racks.*
- *Marvell CPO leverages multiple generations of silicon photonics technology, which has been shipping for over eight years with more than 10 billion device hours of field operation.*

SANTA CLARA, Calif., Jan. 6, 2025 /PRNewswire/ -- Marvell Technology, Inc. (NASDAQ: MRVL), a leader in data infrastructure semiconductor solutions, today announced the advancement of its custom XPU architecture with co-packaged optics (CPO) technology. Building on its recently announced custom high-bandwidth memory (HBM) compute architecture, Marvell is now extending its custom silicon leadership by enabling customers to seamlessly integrate CPO into their next-generation custom XPUs and scale-up the size of their AI servers from tens of XPUs within a rack currently using copper interconnects to hundreds across multiple racks using CPO, enhancing AI server performance**.** The innovative architecture enables cloud hyperscalers to develop custom XPUs that achieve higher bandwidth density and deliver longer reach XPU-to-XPU connections within a single AI server – with optimal latency and power efficiency. The architecture is now available for Marvell customers' next-generation custom XPU designs.

The Marvell custom AI accelerator architecture combines XPU compute silicon, HBM and other chiplets with Marvell 3D SiPho Engines on the same substrate using high-speed SerDes, die-to-die interfaces and advanced packaging technologies. This approach eliminates the need for electrical signals to leave the XPU package into copper cables or across a printed circuit board. With integrated optics, connections between XPUs can achieve faster data transfer rates and distances that are 100X longer than electrical cabling. This enables scale-up connectivity within AI servers that spans multiple racks with optimal latency and power dissipation.

CPO technology integrates optical components directly within a single package, minimizing the electrical path length. This close coupling significantly reduces signal loss, enhances high-speed signal integrity, and minimizes latency. CPO enhances data throughput by leveraging high-bandwidth silicon photonics optical engines, which provide higher data transfer rates and are less susceptible to electromagnetic interference compared to traditional copper connections. This integration also improves power efficiency by reducing the need for high-power electrical drivers, repeaters and retimers. By enabling longer reach and higher density XPU-to-XPU connections, CPO technology facilitates the development of high-performance, high-capacity scale-up AI servers, optimizing both compute performance and power consumption for next-generation accelerated infrastructure.

First demonstrated at OFC 2024, the industry-first Marvell 3D SiPho Engine – which supports 200Gbps electrical and optical interfaces – is a fundamental building block for incorporating CPO into XPUs. The Marvell 6.4T 3D SiPho Engine is a highly integrated optical engine with 32 channels of 200G electrical and optical interfaces, hundreds of components such as modulators, photodetectors, modulator drivers, trans-impedance amplifiers, microcontrollers, and a host of other passive components in a single, unified device to deliver 2x the bandwidth, 2x the input/output bandwidth density, and 30% lower power per bit versus comparable devices with 100G electrical and optical interfaces. Multiple customers are evaluating the technology for integration into their next-generation solutions.

For more than eight years, Marvell has delivered silicon photonics technology for successive generations of high-performance, low power COLORZ® data center interconnect optical modules. This technology has been qualified and deployed in high-volume production by numerous leading hyperscalers to satisfy their growing data center to data center bandwidth needs. Marvell silicon photonics devices have logged over 10 billion field hours.

Marvell has been a pioneer in transforming interconnect technology to improve the performance, scalability, and economics of accelerated infrastructure. The Marvell interconnect portfolio includes high-performance SerDes and die-to-die technology IP for high-performance communication within custom XPUs, PCIe retimers for efficient short-reach connections between CPUs and XPUs on the same board, breakthrough CXL devices for overcoming memory challenges, Active Electrical Cable and Active Optical Cable digital signal processors for short-reach connections within a rack, an expanding range of PAM optical DSPs for rack-to-rack connections inside data centers and coherent DSPs and data center interconnect modules for linking data centers separated by thousands of kilometers.

"The Marvell custom AI accelerator with CPO architecture enables cloud hyperscalers to develop custom XPUs that will significantly increase the density and performance of their AI servers," said Will Chu, senior vice president and general manager of the Custom, Compute and Storage Group at Marvell. "Integrating optics directly into XPUs takes custom accelerated infrastructure to the next level of scale and optimization that hyperscalers must deliver to satisfy the growing demands of AI applications."

"AI scale-up servers require connectivity with higher signaling speeds and longer distances to support unprecedented XPU cluster sizes," said Nick Kucharewski, senior vice president and general manager of the Network Switching Business Unit at Marvell. "Integrating co-packaged optics into custom XPUs is the logical next step to scale performance with higher interconnect bandwidths and longer reach."

"Silicon photonics is vital for scaling accelerated infrastructure connectivity to address increasing bandwidth demands, interconnect distances, power consumption, and total cost of ownership," said Radha Nagarajan, senior vice president and chief technology officer of Optical Platforms at Marvell. "Since 2017, Marvell has pioneered the delivery of high-volume silicon photonics devices to top hyperscalers and leveraged this expertise to create a cutting-edge CPO architecture for the killer CPO use case of custom XPU connectivity."

"Cloud hyperscalers will integrate CPO technology into their next-generation custom XPUs and scale-up servers to meet the escalating performance demands of AI. We forecast that CPO will grow from less than 50 thousand port shipments today to over 18 million CPO ports by 2029 with most of the ports being deployed for connections within servers," said Vlad Kozlov, founder and CEO, LightCounting. "With its experience in optical technology and

custom XPUs, Marvell is ideally positioned to enable hyperscalers to unlock the potential of CPO and make it an integral part of their infrastructure."

Marvell and the M logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.

This press release contains forward-looking statements within the meaning of the federal securities laws that involve risks and uncertainties. Forward-looking statements include, without limitation, any statement that may predict, forecast, indicate or imply future events, results or achievements. Actual events, results or achievements may differ materially from those contemplated in this press release. Forward-looking statements are only predictions and are subject to risks, uncertainties and assumptions that are difficult to predict, including those described in the "Risk Factors" section of our Annual Reports on Form 10-K, Quarterly Reports on Form 10-Q and other documents filed by us from time to time with the SEC. Forward-looking statements speak only as of the date they are made. Readers are cautioned not to put undue reliance on forward-looking statements, and no person assumes any obligation to update or revise any such forward-looking statements, whether as a result of new information, future events or otherwise.

For further information, contact:
pr@marvell.com



View original content to download multimedia:https://www.prnewswire.com/news-releases/marvell-announces-breakthrough-co-packaged-optics-architecture-for-custom-ai-accelerators-302342765.html

SOURCE Marvell