

July 30, 2024



Marvell Introduces Breakthrough Structera CXL Product Line to Address Server Memory Bandwidth and Capacity Challenges in Cloud Data Centers

- *New product line is the industry's first to integrate Arm® Neoverse® processors, inline compression, and four memory channels.*
- *Structera™ A CXL near-memory accelerator family is optimized to address high-bandwidth memory applications such as deep learning and ML/AI on general-purpose servers.*
- *Structera™ X CXL memory-expansion controller family is the first to support up to 12 DDR4 memory modules, enabling module recycling to reduce e-waste and maximize server memory capacity for applications such as in-memory databases.*

SANTA CLARA, Calif., July 30, 2024 /PRNewswire/ -- [Marvell Technology, Inc.](#) (NASDAQ: MRVL), a leader in data infrastructure semiconductor solutions, today launched the [Marvell® Structera™ product line of Compute Express Link® \(CXL®\)](#) devices that enable cloud data center operators to overcome memory performance and scaling challenges in general-purpose servers.

To address memory-intensive applications, data center operators add extra servers to get higher memory bandwidth and higher memory capacity. The compute capabilities from the added processors are typically not utilized for these applications, making the servers inefficient from cost and power perspectives. The CXL industry standard addresses this challenge by enabling new architectures that can efficiently add memory to general-purpose servers.

The Structera product line comprises two CXL device families that are optimized for different use cases. The Structera A CXL near-memory accelerators are a new category of devices that integrate server-class processor cores and multiple memory channels with CXL to address high-bandwidth memory applications such as deep learning recommendation models (DLRM) and machine learning. The Structera X CXL memory-expansion controllers enable terabytes of memory to be added to general-purpose servers and address high-capacity memory applications such as in-memory databases. The Structera CXL device families are the industry's first to support four memory channels, integrate inline compression and use 5nm manufacturing processes.

In addition to the Structera A and X standard products, Marvell develops custom CXL silicon for cloud operators that is optimized for their unique architectures and workloads.

A New CXL Use Case: ML/AI Acceleration

The Structera A accelerators integrate 16 Arm® Neoverse® V2 cores and are optimized to improve the performance of high memory-bandwidth applications such as DLRM. DLRM is characterized by both sparse and dense memory operations. Sparse operations may hit a memory wall, in which memory bandwidth is insufficient for the available compute. The Structera A 2504 accelerator, the first product in the family, supports up to 200 GB/sec memory bandwidth and 4TB of memory capacity¹.

For a DLRM server using a single 64-core processor, adding a single Structera A device would increase the number of compute cores by 25% (64 vs 80), aggregate memory bandwidth by 50% (400 GB/sec vs 600 GB/sec), total memory by up to 4TB, and memory bandwidth per core by 25% (6.25 GB/sec vs. 7.5 GB/sec). It would also improve memory bandwidth power efficiency from 1W to 0.83W per GB/sec. Adding two Structera A devices would increase compute cores by 50%, double aggregate memory bandwidth, increase memory capacity by up to 8TB, increase memory bandwidth per core by 50%, and improve memory bandwidth power efficiency to 0.75W² per GB/sec.

The Structera X product family, by contrast, is designed for expanding memory capacity per server to address high-capacity memory applications. The Structera X devices are the industry's first to support four channels of DDR4 and DDR5 memory and support up to three DDR4 dual in-line memory modules (DIMMs) per channel (3DPC) for maximizing memory capacity per controller and server node. The devices are also the first CXL memory-expansion products that can simultaneously support two server CPUs to optimize power, space and memory utilization.

The Structera X 2404 enables operators to recycle their DDR4 DIMMs from decommissioned general-purpose servers. Millions of functional DDR4 DIMMs are expected to become e-waste over the next few years as operators replace existing general-purpose servers with new ones utilizing DDR5³. Using "free" DDR4 DIMMs to expand the capacity in these servers reduces capex by thousands of dollars per general-purpose server⁴. Repurposing decommissioned DDR4 DIMMs also addresses data center sustainability goals.

Structera products are the first in the industry to incorporate hardware-based inline memory compression that adheres to the Google and Meta specifications that have been submitted to the Open Compute Project (OCP).

"Memory access is a multifaceted problem and cloud service providers differ substantially when it comes to their goals and deployment strategies for CXL," said Bob Wheeler, principal analyst at Wheeler's Network. "Marvell has developed a long-term vision for CXL that capitalizes on this diversity of uses and will encourage hyperscalers and others to adopt CXL to scale the so-called memory wall."

"Our new Structera CXL product line will be a game-changer in enabling optimal resource utilization and lowering energy consumption for scaling memory-intensive workloads in the cloud," said Raghieb Hussain, president of products and technologies at Marvell. "Marvell is delivering on the promise of CXL and our commitment to customers to help them solve their most challenging issues. CXL will continue to be a critical technology enabler for accelerated infrastructure across our compute, connectivity and storage portfolios."

Broad Ecosystem Support

Several companies are developing plans to support the new Structera CXL product line.

"As CXL matures, it unlocks opportunities for innovation through disaggregation, enabling transformative applications across the entire data center ecosystem," said Raghu Nambiar, corporate vice president, Data Center Ecosystems and Solutions, AMD. "Our collaboration with Marvell is at the heart of ensuring our commitment to this ecosystem; delivering the performance and energy efficiency needed to deploy and scale next generation infrastructure to power the most demanding business critical and AI applications."

"The new Marvell Structera CXL controller family combines Marvell's compression technology and SoC expertise with the performance and efficiency benefits of Neoverse V2 in a way that exemplifies the pace of innovation only made possible by Arm Neoverse," said Mohamed Awad, senior vice president and general manager, Infrastructure Line of Business, Arm. "It's exactly this type of optimized processing that reduces energy consumption and enables the high-performance infrastructure that AI demands."

Debendra Das Sharma, Chief I/O Architect at Intel said, "Intel is delighted that Marvell is contributing to the evolving CXL ecosystem, including the utilization of memory expanders for capacity and bandwidth expansion solutions. These solutions enable the reuse of DDR4 DIMMs for memory capacity expansion while also increasing the memory bandwidth available to the CPU cores. With support for CXL 2.0, Intel Xeon 6.0 product family is at the forefront of these developments and appreciates the collaboration with Marvell to bring these solutions to market."

"Marvell is an important CXL ecosystem partner to Micron, and our work together will help advance the adoption of memory expansion in the industry. Our collaboration on memory interoperability will enable flexible and scalable memory resources to match increasing processor core counts and data-intensive workloads," said Vijay Nain, senior director of CXL Product Management at Micron. "By driving higher memory capacity and lower latency, data center operators can benefit from reduced capital and operating expenses."

Product Features

Marvell Structera A 2504

- PCIe 5.0/CXL2.0 x16
- Four channels of DRAM DDR5 6400 MT/sec
- Two memory DIMMs per channel (2DPC) support
- 16 Arm Neoverse V2 cores at 3.2 GHz
- Inline LZ4 compression and decompression
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot

Marvell Structera X 2404

- PCIe 5.0/CXL2.0 1 x16 or 2 x8
- Four channels of DRAM DDR4 3200 MT/sec
- Three memory DIMMs per channel (3DPC) support
- Inline LZ4 compression and decompression
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot

Marvell Structera X 2504

- PCIe 5.0/CXL2.0 1 x16 or 2 x8
- Four channels of DRAM DDR5 6400 MT/sec
- Two memory DIMMs per channel (2DPC) support
- Inline LZ4 compression and decompression
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot

Availability

Structera A near-memory accelerators, Structera X memory-expansion controllers and custom CXL silicon will sample in the fourth quarter of 2024.

About Marvell

To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for over 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

1. Based on 512GB DIMMs.
2. 64-core processor assumed to operate at 400W and 400GB/sec.
3. Marvell and industry analyst estimates.
4. The hypothetical assumes 4TB of memory added through 32 128GB DIMMs. Pricing per 128GB DDR5 DIMMs is estimated at \$280, the approximate market price in June 2024.

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:

Kim Markle
pr@marvell.com



View original content to download multimedia:<https://www.prnewswire.com/news-releases/marvell-introduces-breakthrough-structera-cxl-product-line-to-address-server-memory-bandwidth-and-capacity-challenges-in-cloud-data-centers-302209550.html>

SOURCE Marvell