

October 12, 2026



ADDING and REPLACING Marvell Unveils Second-generation Structera X, Industry-first Six-channel CXL 3.2 Memory Expanders for AI Inference

- **Expanding memory capacity for AI:** Second-generation Marvell® Structera™ X 3406 and 3506 CXL 3.2 memory expanders add six DDR4/DDR5 channels (50% more than the prior generation) and PCIe Gen 6 connectivity, supporting up to 9TB of DDR4 or 6TB of DDR5 per controller.
- **Addresses growing memory demand:** Multi-host sharing and inline LZ4 compression deliver up to 2–3x more effective capacity, with up to 40–50% faster time to first token at equivalent context length.
- **Preview and availability:** Demoing at the OCP Summit, the Structera X 3406 and 3506 are expected to sample in Q1 2027, extending the Structera X 2404 and 2504 architecture already shipping to hyperscalers.

SANTA CLARA, Calif.--(BUSINESS WIRE)-- In the "**Industry Ecosystem Support**" section, please note the addition of a new fourth paragraph, a quote from Yashar Bayani, director, Hardware Engineering at Meta.

This press release features multimedia. View the full release here:
<https://www.businesswire.com/news/home/20261012029557/en/>

Marvell Six-channel CXL 3.2 Memory Expander for AI Inference

The updated release reads:

MARVELL UNVEILS SECOND-GENERATION STRUCTERA X, INDUSTRY-FIRST SIX-CHANNEL CXL 3.2 MEMORY EXPANDERS FOR AI INFERENCE

- **Expanding memory capacity for AI:** Second-generation Marvell® Structera™ X 3406 and 3506 CXL 3.2 memory expanders add six DDR4/DDR5 channels (50% more than the prior generation) and PCIe Gen 6 connectivity, supporting up to 9TB of DDR4 or 6TB of DDR5 per controller.
- **Addresses growing memory demand:** Multi-host sharing and inline LZ4 compression deliver up to 2–3x more effective capacity, with up to 40–50% faster time to first token at equivalent context length.
- **Preview and availability:** Demoing at the OCP Summit, the Structera X 3406 and 3506 are expected to sample in Q1 2027, extending the Structera X 2404 and 2504 architecture already shipping to hyperscalers.

[Marvell Technology, Inc.](#) (NASDAQ: MRVL), a leader in data infrastructure semiconductor solutions, today announced the second-generation Marvell Structera X 3406 and Marvell Structera X 3506 CXL memory expanders, designed to help cloud, enterprise and AI

infrastructure providers build more efficient, scalable and higher-performance systems for the next era of AI and computing. Designed to address the evolving needs of leading hyperscale datacenters, Structera X addresses the most demanding data center workload requirements.

AI and ML workloads are rewriting the rules of data center design. Growing LLM sizes, expanding context windows and growing KV-cache requirements are driving unprecedented memory demands that traditional server architectures simply cannot meet. Memory is a critical constraint of AI performance and infrastructure efficiency.

Modern AI inference workloads, from LLM serving to deep learning recommendation models (DLRM), are memory bound. KV-cache capacity directly determines time to first token (TTFT), throughput, and XPU utilization. Structera X 3406/3506 addresses these bottlenecks with:

- **Industry-first six-channel CXL memory expansion:** Six DDR4/DDR5 memory channels deliver 50% more channels than the previous generation, increasing memory bandwidth and capacity so AI servers can sustain larger models and longer context windows - up to 18 DDR4 RDIMM slots supporting 9TB for Structera X 3406 and 12 DDR5 RDIMM slots supporting 6TB for Structera X 3506.
- **Industry-first four host connectivity:** Enables memory capacity sharing with up to four hosts and up to 256GB/s PCIe6.0 x16 of host bandwidth, improving memory utilization and efficiency across AI clusters.
- **5x higher LZ4 compression throughput from previous generation:** Up to 250GB/s compression/decompression throughput unlocks up to 2–3x more effective memory capacity, depending on workloads.
- **StructGuard™ enterprise-grade RAS:** Provides advanced error correction to enhance memory reliability and resilience against DRAM device failures.
- **Industry-first post-quantum security:** Combines post-quantum cryptography, Caliptra 2.0, CXL IDE and AES-256 memory encryption to provide advanced data protection.

For DLRM workloads, Structera X delivers up to 2x higher queries per second and 2–3x more effective embedding-table capacity per controller. For LLM serving, Structera X delivers up to 40–50% faster TTFT at equivalent context length, or up to 2x longer supportable context windows at equivalent latency compared to the previous generation.

Building on the success of the first-generation Structera X 2404/2504, the Structera X 3406/3506 leverages the proven software stack and ecosystem integrations enabling an accelerated ramp. The Marvell CXL ecosystem continues to expand across the full data center stack through close partnership with CPU and memory vendors, OEMs, ODMs and system integrators.

“CXL is starting to gain widespread adoption in AI systems since they reap strong benefits from its very high bandwidth and support of large memory sizes,” said Jim Handy, general director of Objective Analysis, and author of a CXL market report. “In addition, today’s memory shortage and subsequent high prices are driving IT managers to use CXL to re-deploy used DDR4 DIMMs into DDR5 systems for immediate cost savings. In the near future, these same managers will use CXL’s memory pools to reduce stranded memory and drive greater savings, and they will find controllers that offer compression, as this one does, to be even more interesting.”

By bringing together connectivity, memory and storage technologies, the expanded data center portfolio improves how compute and memory work across AI systems, increasing utilization, infrastructure efficiency and AI performance. These next-generation offerings increase access to memory resources, improving both memory and compute utilization, helping customers optimize AI infrastructure performance as workloads continue to grow in scale and complexity.

“Marvell is the only provider with a complete end-to-end CXL portfolio spanning memory expansion (Structera X), near-memory acceleration (Structera A), and PFMA (Photonic Fabric™ Memory Appliance), pod-scale CXL memory-sharing appliance, giving customers a complete portfolio from direct-attach to disaggregated memory pooling,” said Will Chu, executive vice president and general manager, Custom Cloud Solutions at Marvell. “We are delivering industry-leading CXL controllers, co-developed with leading hyperscalers to address the most demanding AI infrastructure requirements.”

Industry Ecosystem Support

“AMD is advancing an open ecosystem that gives enterprises the freedom to build infrastructure around their evolving workloads,” said Raghu Nambiar, corporate vice president, Data Center Ecosystems and Solutions, AMD. “CXL is an important part of that vision, enabling customers to expand memory capacity and use resources more efficiently as AI and data-intensive applications grow. Our collaboration with Marvell builds on the performance and scalability of AMD EPYC processors to help enterprises realize those benefits, bringing greater flexibility and long-term value to their data center investments.”

“As AI becomes more agentic, memory and compute need to scale together to meet growing infrastructure demands. Marvell’s latest Structera X controllers bring CXL 3.2, memory expansion and compression to the Arm ecosystem, giving hyperscaler customers more options to efficiently support complex AI workloads,” said Eddie Ramirez, vice president of go-to-market, Cloud AI, Arm.

“Marvell and Alibaba are collaborating to advance CXL memory solutions for the rapidly growing demands of AI infrastructure. As AI models and workloads continue to increase in scale, expanding memory capacity and bandwidth without compromising system performance and efficiency has become a critical requirement. CXL enables a more flexible approach to memory expansion and sharing, providing a path to scale memory resources across increasingly complex AI systems. Our collaboration with Marvell brings together its leadership in CXL technology with Alibaba’s experience in large-scale AI and cloud infrastructure to develop memory architectures designed for the next generation of AI data centers.”

“Meta is pleased to work with Marvell to advance CXL solutions that address the growing memory demands of compute infrastructure,” said Yashar Bayani, director, Hardware Engineering at Meta. “As systems scale, increasing memory capacity while improving infrastructure efficiency and cost becomes critical. CXL provides a flexible approach to expanding memory resources and optimizing infrastructure economics, and we look forward to working with the ecosystem, including Marvell to help enable the next generation of scalable infrastructure.”

“CXL is reshaping how memory is deployed in data centers, and our work with Marvell ensures customers can confidently deploy Structera with Samsung DDR memory for reliable, high-performance systems,” said Jangseok Choi (JS Choi), vice president of the

Memory Product Planning Team at Samsung Electronics. "The move to CXL 3.2 and six memory channels in Structera X 3406/3506 further strengthens the flexibility our joint customers need."

"Working with Marvell to validate Structera with SK hynix memory supports our shared goal of making memory expansion more accessible and flexible," said Uksong Kang, Head of Next Generation Product Planning and Enabling at SK hynix. "Structera X 3406/3506 gives customers the tools they need to build future-ready, disaggregated architectures with less friction and more choice."

Availability

Marvell Structera X 3406 and Structera X 3506 CXL memory expansion controllers are expected to sample in Q1 2027.

Visit Marvell Booth #B31 at the [2026 OCP Global Summit](#), Oct. 12-15 in San Jose, Calif., to learn more about the company's newest memory and storage offerings.

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 30 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 and carrier architectures transform—for the better.

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.

View source version on businesswire.com:

<https://www.businesswire.com/news/home/20261012029557/en/>

Media Contact:

Erica Coleman
pr@marvell.com

Source: Marvell Technology, Inc.