

July 25, 2024



Marvell Teralynx 10 51.2T Ethernet Switch Enters Volume Production for Global AI Cloud Deployments

- *Unique architecture delivers high bandwidth, low latency, low power and programmability in a single platform for optimized cloud AI deployments.*
- *Seamless support for open network operating system software enables fast design and deployment for ODMs, OEMs and hyperscalers.*
- *Multiple customers designing with Teralynx[®] 10 Ethernet switch device to scale accelerated infrastructure.*

SANTA CLARA, Calif. , July 25, 2024 /PRNewswire/ -- [Marvell Technology, Inc.](#) (NASDAQ: MRVL), a leader in data infrastructure semiconductor solutions, today announced that the Marvell[®] [Teralynx[®] 10](#) Ethernet switch device is in volume production with customer deployments underway. The Teralynx 10 switch is a low power, programmable 51.2 Tbps Ethernet device with the industry's lowest latency¹, delivering leading performance for training, inference, general-purpose compute and other workloads to scale accelerated infrastructure in cloud data centers.

The Teralynx 10 switch device enters production as the growth in the size of AI clusters and the number of AI data centers is driving increased demand for high-bandwidth connectivity. Simultaneously, the industry is undergoing a shift from the exclusive use of proprietary network operating systems (NOS) to deployment of open network platforms, such as Linux Foundation's SONiC and SAI (switch abstraction interface), that enable operators to rapidly deploy switching from multiple manufacturers, speed up development timelines and optimize for their specific use cases.

Shipments of 51.2 Tbps switches are expected to skyrocket from approximately 77,000 in 2024 to 1.8 million in 2028, a 120% compound annual growth rate (CAGR), according to 650 Group. 650 Group's co-founder Alan Weckel said, "51.2T marks a watershed moment in data center switching. Cloud service providers will rapidly gravitate toward this switch capacity to build the next generation of infrastructure, placing a priority on multi-vendor diversity at the system and silicon levels."

The Teralynx 10 device is based on a clean-sheet switch architecture that simultaneously delivers high bandwidth, ultra-low latency, low power, high radix and zero-latency-loss programmability. This enables a single switching platform design to be used across multiple points of the network without compromising functionality or performance. Whether used in top-of-rack (ToR), leaf, spine, AI cluster, or network edge, the unified platform offers the attributes required for advanced AI networks and enables in-field feature updates to support new protocols.

The Teralynx 10 switch software platform is designed to deliver on the emerging transition to open networking for large scale data centers. Increasingly, operators are moving from traditional proprietary network operating systems to the open-source SONiC platform, similar to how Linux disrupted proprietary operating systems for server compute. An open-source approach enables operators to normalize their networking system software across multiple manufacturers and accelerate development timelines. From a business perspective, this approach allows for a broader ecosystem of switch manufacturers with increased silicon diversity and supply chain stability for cloud network operators. Operators are increasingly mandating open network software, and the Marvell Teralynx 10 switch platform provides a comprehensive solution with an industry-accepted programming model that can run on Marvell switch solutions as well as other compatible hardware. (See today's related announcement on SONiC.)

"AI deployments require a switch solution which is simultaneously high bandwidth, low power, low latency, and future proofed for evolving requirements. Marvell has delivered the most complete AI Ethernet switch solution available to the industry today, and we are pleased to deliver a production-ready solution for customers' expanding 51.2 Tbps deployments," said Nick Kucharewski, senior vice president and GM, Network Switching Business Unit, at Marvell. "This is the right product at the right time, as the market turns to innovative solutions from customer-focused suppliers like Marvell to satisfy their explosive demand for high bandwidth switching in AI data center buildouts."

Teralynx 10 Features and Benefits

The Teralynx 10 Ethernet switch offers leadership performance and high scale for the most advanced cloud and AI workloads:

- **Lowest latency:** Demonstrates 51.2 terabit-per-second Ethernet throughput with latency as low as 500ns, and sub-600 nanoseconds latency across all packet sizes (See video [here](#)). Low latency is essential for meeting the demands of AI, ML and distributed workloads, and it directly impacts job completion time and algorithmic efficiency. The low latency Teralynx switch gives cloud operators the ability to reduce operating expenses and increase their capacity for performing revenue-generating activity.
- **Industry-best radix:** 512 switching radix enables operators to reduce the number of switch tiers in large clusters, yielding dramatically lower power and total cost of ownership (TCO).
- **Low power consumption:** The switch consumes 1 watt per 100 gigabits-per-second of bandwidth.
- **Programmable:** The industry's only switch architecture that is fully programmable with no impact to packet processing capacity or latency. The Teralynx 10 device can be applied in multiple use cases at 51.2 Tbps. This flexibility gives data center operators investment protection as networking technologies evolve to handle new protocols.

Expansive Partner Ecosystem

Switch systems based on the Teralynx 10 device are available from major OEMs, ODMs, and ISVs to speed adoption, while providing greater choice and ability for customers to optimize for their unique environment.

"The proliferation of AI and machine learning in data centers represents a transformative shift in the industry, driving unprecedented advancements in computational power, data processing capabilities, and operational efficiency," said Gavin Cato, head of Portfolio Solutions and CTO, Connectivity and Cloud Solutions, Celestica. "Celestica congratulates

Marvell on the Teralynx 10 51.2T release and looks forward to continued collaborations in leading-edge Ethernet solutions to support bringing advanced AI capabilities to the market."

"We have been collaborating with Marvell to provide highly innovative and comprehensive solutions for our customers and we are very excited to add Teralynx 10 to our portfolio," said David Tsai, vice president and GM, Wistron Neweb Corporation, Networking Business Group.

"Keysight's industry-leading test solutions, like the AresONE-M 800GE platform, were instrumental in accurately benchmarking the Teralynx 10 low-latency capabilities. This precise measurement of data transfer through the switching fabric is critical in assessing optimal job completion times for complex AI training algorithms within massive data center clusters exceeding 100,000 GPUs. By collaborating with Marvell, we're empowering customers to confidently deploy the Teralynx 10 and achieve peak performance in their AI cloud infrastructure," said Ram Periakaruppan, vice president and general manager, Network Test & Security Solutions, Keysight.

"MultiLane is proud to extend its partnership with Marvell by introducing the Teralynx 10 platform into our Interoperability Labs in Fremont, Calif. As a leader in interoperability and compliance testing services, MultiLane continues to enhance its offerings to help data center operators streamline the adoption and deployment of high-performance switching fabrics," said Michel Haddad, general manager, Measurement Solutions Business Unit, MultiLane.

"Teledyne LeCroy Xena's leading test equipment enables the thorough performance benchmarking of the Teralynx 10 platform including testing of 800G Ethernet using the Z800 Freya platform," said Jacob Christensen, general manager, Teledyne LeCroy Xena.

Footnote 1: Based on testing by Keysight

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.

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 or achievements. Actual events or results 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:

Michael Kanellos

pr@marvell.com



View original content to download multimedia: <https://www.prnewswire.com/news-releases/marvell-teralynx-10-51-2t-ethernet-switch-enters-volume-production-for-global-ai-cloud-deployments-302206399.html>

SOURCE Marvell