Arista Introduces Next-Generation 1.6Terabit Portfolio for AI Fabrics

2026-06-09

New 7060XE7 Series expands Etherlink™ for scale-out and scale-up applications

SANTA CLARA, Calif.--(BUSINESS WIRE)-- Arista Networks (NYSE: ANET), a leader in modern AI fabric networking, today announced the Arista **7060XE7 Series**, a new portfolio of 1.6T networking platforms designed specifically as the foundation for rack-scale AI infrastructure. As AI workloads scale from thousands to hundreds of thousands of XPUs, the network has evolved from a standalone layer into a critical backplane for a tightly-integrated AI supersystem.

The 7060XE7 Series represents Arista's transition from providing high-performance switches to delivering comprehensive rack-scale systems. By addressing the extreme density, power, and thermal efficiency requirements of the AI era, these platforms enable customers to build scale-up and scale-out AI fabrics optimized for air, liquid and hybrid-cooled environments, maximizing compute density per kilowatt of power.

As AI training and inference workloads evolve, the network must transition from a standalone layer to an elastic and integrated backplane. To support diversity of AI accelerators, Arista has engineered the 7060XE7 Series to serve as a tightly coupled rack-scale AI supersystem. This includes collaborating with our ecosystem of customers and technology partners to develop these large-scale AI fabrics.

"The AI era requires a shift in how we think about the network - it is no longer a standalone layer of infrastructure, but a tightly integrated component of the AI supersystem," said Tyson Lamoreaux, Senior Vice President, Cloud and AI Networking at Arista Networks. "With the 7060XE7 Series, we are delivering massive-scale 1.6T systems that combine world-class reliability and the differentiation of EOS with liquid cooling and low-power optics to help our customers build AI fabrics designed for maximum performance and power efficiency."

Delivering Rack-Scale AI Fabrics with Diverse AI Accelerators

The 7060XE7 Series delivers a broad portfolio of fixed switch platforms and configurable rack-scale systems, enabling flexible system counts per rack for vertical and horizontal AI workflows. These systems enable compute and rack optimized networks with a suite of **Arista EOS®** (Extensible Operating System) features required by AI fabrics. Each system provides low-latency and intelligent packet buffering to manage the intense microbursts typical of AI communication and collective patterns, ensuring consistent high-performance connectivity intra- and inter-rack for compute and memory semantics.

Arista AI Networks at Scale with Meta, Microsoft and Oracle

For more than a decade, Arista has worked closely with our Cloud and AI titan customers to support their ever-growing AI clusters.

"As AI workloads grow, the physical infrastructure must evolve to support higher density and power efficiency," said Gaya Nagarajan, Vice President of Infrastructure at Meta. "Arista's 1.6T platforms and liquid-cooled designs align with our focus on open, scalable AI fabrics that meet the requirements of next-generation training and inference."

"As AI infrastructure scales to support increasingly demanding workloads, high bandwidth, power-efficient networking becomes foundational," said Rani Borkar, President, Azure Hardware Systems and Infrastructure at Microsoft. "Our collaboration with Arista on the 1.6T Ethernet interface helps enable the next generation of AI clusters with greater interconnect capacity for Azure Maia, Microsoft's AI accelerator, and Fairwater, Microsoft's extreme-scale AI datacenters, while preserving operational simplicity across our infrastructure."

"Oracle Cloud Infrastructure requires networking that can scale with the demands of large AI training jobs," said Mahesh Thiagarajan, Executive Vice President, Oracle Cloud Infrastructure. "Arista Networks' 1.6T platforms provide the throughput, determinism, and stability needed for our RDMA-based AI fabrics, while Arista EOS delivers operational consistency and performance at scale across our global AI infrastructure."

The Latest Etherlink Member: Massive Scale and Throughput

The 7060XE7 Series utilizes the latest SerDes technology to provide high-speed bandwidth of 100 Terabits/sec, 1.6 terabits /sec throughput across ports and signal integrity across the rack. Support for Linear Pluggable Optics (LPO) reduces interconnect power consumption by approximately 60%, thereby lowering the total cost of ownership.

The portfolio includes distinct configurations to meet different data center needs:

- 7060XE7-64PS and 7060XE7-64PRS Rack Switches: Air-cooled systems providing support for both Integrated heat sink (IHS) and Riding heat sink (RHS) optics, giving deployment flexibility in a 4RU form factor.

- 7060XE7-64PRS-RV3-L: A specialized 2OU liquid-cooled platform for high-density clusters, utilizing 224G SerDes. This system utilizes DC power from the ORv3 rack and contains no internal fans, integrating with liquid-cooled XPU servers to maximize power efficiency.
- 7060XE7-128PE: Provides 128 800G ports in an air-cooled 4RU design, utilizing 100G SerDes, for environments requiring deployment flexibility and backward compatibility.

Silicon Partnerships with AMD and Broadcom

The transition to 1.6T networking is supported by an ecosystem of switch silicon and NIC providers. Arista is working closely with AMD on next-generation compute silicon and NICs to enable scale-out AI fabrics.

“AMD is committed to delivering open, standards-based AI infrastructure that gives customers choice across the full compute stack, from CPUs and GPUs to advanced networking technologies and unifying software solutions,” said Ravi Pendekanti, corporate vice president, Data Center Solutions Group, AMD. “As AI systems scale, the network is becoming a critical foundation for performance and efficiency. By working with partners like Arista to advance the Ethernet networking ecosystem, we are helping enable flexible, high-performance AI fabrics built on open standards.”

Arista’s 15-year partnership with Broadcom has culminated in an integrated system architecture that delivers best-of-breed and optimized hardware platforms. By unlocking the on-chip resources and the features of the Tomahawk® 6 silicon while maintaining software consistency with Arista EOS or Open NOS, the switches are transformed into a unified rack-scale radix for high-density compute for both air and liquid-cooled AI environments.

“Our collaboration with Arista on 1.6T networking showcases the power of co-innovating with industry-leading switching silicon to meet the relentless bandwidth and latency demands of next-generation AI workloads,” said Dr. Asad Khamisy, senior vice president and general manager, Core Switching Group, Broadcom. “By powering Arista’s 1.6T Ethernet solutions with Broadcom's industry-leading Tomahawk 6 102-Tbps Ethernet switches, we are expanding interconnect capacity for massive AI clusters, enabling optimized scale-up and scale-out connectivity while maintaining the robust, open standards the industry relies on.”

Arista EOS for AI Operational Excellence

The 7060XE7 Series is designed to support both Arista EOS and open Network Operating Systems, providing cloud titan customers with a choice of operational flexibility. The platforms support a rich suite of "Smart AI" software features designed for high-performance AI resilience, congestion management and versatility, including:

- Load Balancing: Dynamic Load Balancing (DLB) and Cluster Load Balancing (CLB) are used to manage job

completion times and optimize traffic distribution across the fabric.

- MRC Fabric Resiliency: The inclusion of full support for MRC (Multipath Reliable Connection) addresses the "failure amplifier" challenge, ensuring that a single link flap does not stall a massive training job. This is complemented by Link Layer Retry (LLR) to maintain performance stability at the physical layer.
- Congestion Management: Support for PFC-aware DLB, PFC-aware ECN, and telemetry provides granular visibility into network health and prevents head-of-line blocking.
- Feedback Mechanisms: Congestion Signaling (CSIG) and Fast CNP provide rapid hardware-level feedback to resolve congestion issues.
- Diagnostics and Platform Uniformity: Integration with NetDI™ (Network Diagnostics Infrastructure) and PDI (Platform-Dependent Infrastructure) enables these platforms to support diverse operating environments. NetDI provides a uniform way to access device telemetry and diagnostics, while PDI manages the low-level hardware specific to each platform, facilitating software portability across different hardware architectures.

Availability

- 7060XE7-64PS/PRS (64x 1.6T Air-Cooled): Q4 ‘26
- 7060XE7-64PRS-RV3-L (64x 1.6T Liquid-Cooled): Q1 ‘27
- 7060XE7-128PE (128x 800G Air-Cooled): Q1 ‘27

Additional Resources

Blog: Three Genius Ideas for AI Fabrics

White Paper: AI Networking

Register for our Webinar: June 25, 2026, “**A New Era of Rack-scale AI Fabrics**”

About Arista Networks

Arista Networks is an industry leader in data-driven, client-to-cloud networking for large AI, data center, campus, and routing environments. Its platforms deliver availability, agility, automation, analytics, and security through an advanced network operating stack. Arista is an S&P 500 member and was recognized on the 2026 Fortune 500 list of the largest United States corporations by revenue. For more information, visit **www.arista.com**.

 Additional information and resources can be found at **www.arista.com**. This press release contains forward-looking statements including, but not limited to, statements regarding the performance and capabilities of Arista’s products

and services. All statements other than statements of historical fact are statements that could be deemed forward-looking statements. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in the forward-looking statements, including rapid technological and market change, customer requirements, and industry standards, as well as other risks stated in our filings with the SEC available on Arista's website at **www.arista.com** and the SEC's website at **www.sec.gov**. Arista disclaims any obligation to publicly update or revise any forward-looking statement to reflect events that occur or circumstances that exist after the date on which they were made.

Media Contact:

Amanda Jaramillo
Corporate Communications
Tel: (408) 547-5798
amanda@arista.com

Investor Contact:

Investor Advocacy
Rudolph Araujo
Rod Hall
Tel: (408) 547-8080
ir@arista.com

Source: Arista Networks