

Marvell® Teralynx® 10 Data Center Ethernet Switch

The world’s lowest-latency programmable switch family with advanced telemetry – 25.6 Tbps through 51.2 Tbps

Overview

Operators are building data centers at a rapid pace for public, private and hybrid cloud use cases. Ongoing adoption of cloud, machine learning, mobile, video, anything-as-a-service (XaaS), edge computing and big data analytics, is driving increased traffic inside data center networks. Architectural shifts such as microservices-based applications and the move to distributed storage using flash/NVMe over fabric are further increasing data center Ethernet traffic. Beyond the need for bandwidth, critical data center requirements include deeper, actionable analytics, low latency, programmability and greater power efficiency. These factors demand an innovative and focused approach to scalable, high performance networking silicon.

Marvell® has engineered the Teralynx® Ethernet switch silicon family from the ground up to provide the most highly optimized data centers network switches. Featuring industry-leading low latency, unmatched visibility through fine-grained telemetry, and rich programmability, the Teralynx family includes a broad range of data center network switches with full software/SDK compatibility across the entire family.

Marvell Teralynx 10 switches are the world’s highest performance switch chips, with up to 51.2 Tbps throughput and industry-leading low latency, analytics and programmability. It is ideal for data center fabrics as well as AI/ML and HPC environments. These switches integrate Marvell’s proven, robust 112G LR SerDes, supporting up to 64 ports of 800GbE or

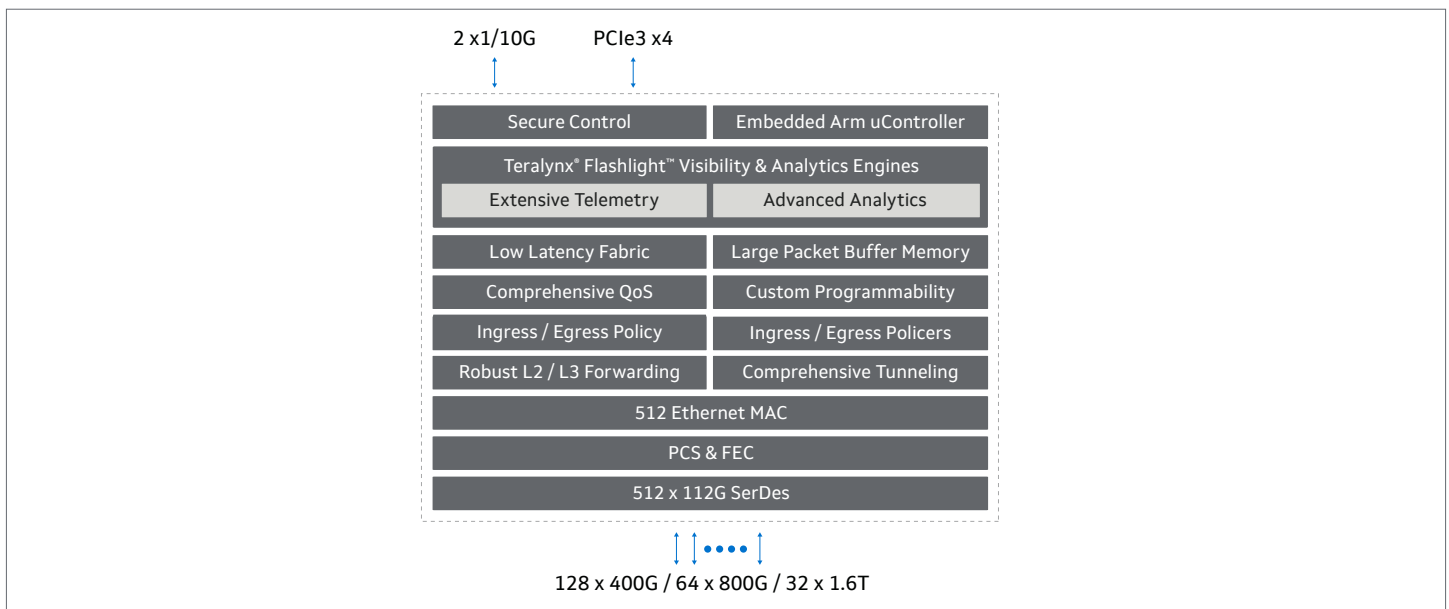
128 ports of 400GbE. Teralynx 10 enables high radix switch architectures that can collapse data center switch tiers, dramatically reducing cost and power.

Marvell Teralynx 10 switches have the following key building blocks:

- **Programmable feature-rich forwarding pipeline:** A programmable forwarding pipeline architecture that delivers a comprehensive data center feature set with deterministic low latency, and ability to support new forwarding protocols through permutable flex-forwarding.
- **Highly scalable fabric:** A scalable and low-latency switch and buffer fabric architecture that delivers comprehensive QoS, superior traffic management and large packet buffers.
- **Advanced Teralynx® Flashlight™ telemetry & analytics:** A ground-up telemetry architecture that delivers extensive real-time visibility and actionable granular network analytics to troubleshoot and resolve network issues quickly helping customers increase network resiliency and uptime.

Software for Teralynx 10 switches is fully compatible with Teralynx 7 to help customers transition easily, reduce development costs and accelerate time to market.

Block Diagram



Key Features

Features	Benefits
<ul style="list-style-type: none">Switch capacity up to 51.2 Tbps with large buffers	<ul style="list-style-type: none">Enables customers to deploy fewer network switches and tiers dramatically reducing cost, power, latency, while significantly simplifying management
<ul style="list-style-type: none">Up to 512 100G SerDes that support 25G, 50G and 100G I/O speed with proven, robust interoperability	<ul style="list-style-type: none">Lowers the cost per bit and enables higher scale I/O such as 400 GbE and 800 GbE while maintaining backward compatibility with 25G NRZOEM & ODM switches can provide high port count 100G, 200G, 400G and 800G connectivity using 50G/100G PAM4 with support for 25/50/100/200/400 and 800 GbE connectivity
<ul style="list-style-type: none">Industry-leading low latency in cut-through & store-and-forward modes	<ul style="list-style-type: none">Reduces time in networking and shortens job completion time, especially for distributed workloads like AI/ML, HPC and distributed databases
<ul style="list-style-type: none">Teralynx® Flashlight™ visibility and telemetry, including support for P4 in-band network telemetry (INT) along with critical extensions	<ul style="list-style-type: none">Enable predictive, faster and more accurate issue resolution, higher automation, and self-aware, and self-healing networks
<ul style="list-style-type: none">Line-rate permutable flex-forwarding to accommodate future networking protocols with software upgrades	<ul style="list-style-type: none">Enables support for custom and new packet forwarding protocols without requiring ASIC spins
<ul style="list-style-type: none">Power-efficient architecture	<ul style="list-style-type: none">Facilitates overall data center efficiency and lower total cost of ownership
<ul style="list-style-type: none">Comprehensive IP forwarding, including congestion aware routingHighly scalable/flexible layer 2 and 3 tables for IPv4, IPv6 and hybrid networksTunneling protocols including IP-in-IP, GRE, MPLS, VXLAN and Geneve	<ul style="list-style-type: none">Optimized packet forwarding enables higher network utilization and prevents congestionSame switch can be deployed in multiple different use cases, saving cost
<ul style="list-style-type: none">Advanced QoS/traffic management feature set including DCB and RoCE	<ul style="list-style-type: none">Enables easy deployment of applications with stringent network requirementsDelivers higher performance for demanding applications like AI/ML and storage
<ul style="list-style-type: none">PTP (IEEE 1588v2) and SyncE support	<ul style="list-style-type: none">Provides timing and synchronization for applications that are sensitive to timing accuracy and precision (e.g., trading, online gaming, telecom)

Customer Deployment Scenarios

- Data center networking infrastructure for the cloud, AI/ML, HPC, high-end enterprise and service providers
- Next-generation switches for ToR, leaf, spine and data center interconnect

Software Support

- Common software development kit (SDK) across entire product line
- Support for open APIs, including OCP SAI (Switch Abstraction Interface)

Switch System Examples

- **51.2T**
 - 1U: 32 x 1.6T (OSFP-XD)
 - 2U: 64 x 800G (OSFP or QSFP-DD800)
- **25.6T**
 - 2U: 64 x 400G (QSFP-DD or OSFP)
 - 1U: 32 x 800G (OSFP or QSFP-DD800)
 - 4U: 128 x 200G (QSFP56)

Ordering Information

Part Number	Capacity (Tbps)	# of SerDes @ Gbps	25/50G ports	100G ports	200G ports	400G ports	400G ports
98TX9180	51.2T	512 @ 25/50/100	512	512	256	128	64
98TX9160	25.6T	256 @ 25/50/100	256	256	128	64	32



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 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.

Copyright © 2023 Marvell. All rights reserved. Marvell and the Marvell 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.