

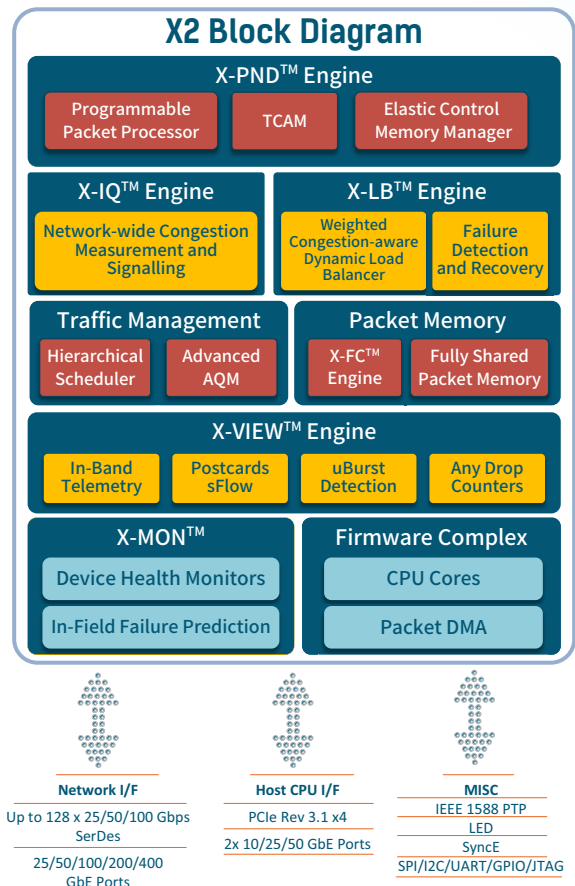


X2 PRODUCT BRIEF

X-Switch: Programmable Ethernet Switches

HIGHLIGHTS

- 12.8 Tbps full-duplex, non-blocking switching optimized for data centers, cloud, and AI networks.
- 100 Gbps PAM4 LR SerDes supporting retimer-less designs.
- Application Optimized Switching architecture powered by a programmable packet header processor enabling high performance, low latency and low power.
- X-PND™: Elastic resource allocation of control tables maximizing adaptation to application needs.
- Fully shared packet buffer for incast burst handling.
- Low latency with cut-through — 700ns first-bit-in-first-bit-out for typical data center applications.
- X-IQ™: UEC-ready programmable congestion measurement and signalling to facilitate dynamic load balancing and fast failure detection and recovery.
- X-LB™: UEC-ready weighted congestion-aware dynamic load balancing for reducing tail latency.
- X-FC™: Incast tolerance mechanism with HoL Blocking prevention.
- X-VIEW™: Telemetry suite for network trouble-shooting and optimization.
- X-MON™: In-field device health monitor for improved reliability.
- Integrated Firmware complex.
- Low power: less than 200W.
- X-SDK and SAI-SONiC Integration.



OVERVIEW

The X2 is the second generation chipset in the X-Switch family of programmable Ethernet switches. Supporting up to 12.8 Tbps full-duplex throughput, the devices are optimized for interconnecting AI/ML, storage and compute clusters in data centers.

These devices incorporate industry leading 100G PAM4 LR SerDes, enabling the designs with in-rack DAC connectivity for the 100G ecosystem without the need for re-timers, while future-proofed for an 800G optics ecosystem. The devices support port speeds of 25, 50, 100, 200, and 400 GbE.

The combination of a programmable packet processor and the X-PND™ engine, an elastic resource allocation mechanism, allow for flexible sharing of control tables. X2 also supports application-specific optimization, future emerging protocols and vendor-specific features.

Additionally, a comprehensive set of device features allows for optimizing AI, storage, and compute workloads.

Packet discard in incast scenarios is minimized using a fully shared packet buffer and X-FC™: a combination of queuing system and a fine-grained flow control protocol that extends the TOR switch(es) packet buffer to the Spine switch(es) while preventing HoL Blocking. High availability is maintained using a hardware-based link failure detection and recovery mechanism. RDMA workload performance is optimized using PFC and cut-through mechanisms.

X2 supports Ultra Ethernet Consortium (UEC) emerging congestion-sensing and load-distribution protocols with its native X-IQ™ and X-LB™ engines. X-IQ™ offers flexible metrics to measure congestion such as traffic rate, latency, queue-fill, and packet drops. The device also supports a variety of congestion signalling protocols, such as FCN, BTS, and dropped packet trimming (PT), all with flexibly assigned headers and metadata. X-LB™ dynamically adapts to network congestion indications, such as those supported by the X-IQ™. X-LB™ load balancing policy concurrently supports flow, flowlet and packet spraying, all according to per-flow specific requirements for order preservation.

An advanced telemetry engine, X-VIEW™, fuels network troubleshooting and optimization. X-VIEW™ includes in-band telemetry, any-cause verbose mirroring, black hole detection and localization, real time statistics histograms, SerDes link quality meters, and uBurst detection.

On-chip health analytics (X-MON™) increase in-field reliability and enable fleet maintenance.

The X2 consumes less than 200W of power enabling cost and power optimized systems that drive lower CAPEX and OPEX.

The X2 is complemented by extensive software packages: X-SDK and a comprehensive tool-chain, together simplifying integration with SONiC and other proprietary NOS's.

TARGET APPLICATIONS

The X2 is optimized for interconnecting AI/ML, storage and compute clusters within a data center's ecosystem. It enables building compact 1RU switches with large port densities of up to sixteen ports of 800G OSFP.



X2 - 16 OSFP Ports

PART ORDERING INFORMATION

The X2 product family is offered in different variations all sharing the same software and feature-set.

All products are fully inter-operable. For additional product information, contact sales@xsightlabs.com.

Part Order Number	Maximum Throughput	Network Facing SerDes	Port Configuration Examples
XLX2A128A	12.8 Tbps	128 x 100/50/25 Gbps	<ul style="list-style-type: none">• 128 x 25/50/100 GbE• 64 x 200 GbE• 32 x 400 GbE
XLX2A096A	9.6 Tbps	128 x 100/50/25 Gbps	<ul style="list-style-type: none">• 64 x 100 GbE and 64 x 50 GbE• 64 x 100 GbE and 8 x 400 GbE• 32 x 200 GbE and 8 x 400 GbE
XLX2A064A	6.4 Tbps	128 x 50/25 Gbps	<ul style="list-style-type: none">• 64 x 100 GbE• 32 x 200 GbE• 16 x 400 GbE
XLX2A032A	3.2 Tbps	128 x 100/50/25 Gbps	<ul style="list-style-type: none">• 128 x 25 GbE• 64 x 50 GbE• 32 x 100 GbE• 64 x 25 GbE and 16 x 100GbE

COPYRIGHT © 2017- 2024 XSIGHT LABS LTD. ALL RIGHTS RESERVED. THE INTELLECTUAL AND TECHNICAL CONCEPTS AND INFORMATION ("INFORMATION") CONTAINED HEREIN ARE PROPRIETARY OF XSIGHT LABS LTD. ("XSIGHT") AND ITS LICENSORS AND MAY BE PROTECTED OR GOVERNED BY VARIOUS PATENTS, PATENTS IN PROCESS, TRADEMARKS, TRADEMARKS IN PROCESS, AND ARE PROTECTED BY TRADE SECRET OR COPYRIGHT LAW. DISSEMINATION, REPRODUCTION OR OTHER UNAUTHORIZED USE OF THIS INFORMATION IS STRICTLY FORBIDDEN UNLESS PRIOR WRITTEN PERMISSION IS OBTAINED FROM XSIGHT. YOU MAY ONLY COPY, MODIFY, DISTRIBUTE, DISPLAY, LICENSE, OR SELL THE CONTENT IF YOU ARE GRANTED EXPLICIT PERMISSION SUCH AS WITHIN A WRITTEN AGREEMENT WITH XSIGHT OR OTHER WRITTEN TERMS BY XSIGHT THAT ACCOMPANY THE CONTENT. ALL INFORMATION CONTAINED HEREIN IS, AND REMAINS THE SOLE PROPERTY OF, XSIGHT AND ITS LICENSORS (AS APPLICABLE). XSIGHT RESERVES THE RIGHT TO MAKE CHANGES TO THE INFORMATION.