X1 PRODUCT BRIEF

Highly Integrated Ethernet Switches

- 12.8/25.6 Tbps full-duplex, non-blocking switching capability optimized for HPC, AI/ML, and data center fabric
- Flexible port speed configurations of 100, 200, or 400 GbE
- 100 Gbps LR PAM4 SerDes supporting backplane, DAC and 800 Gbps optical transceivers without the need for retimers
- Application Optimized Switching architecture powered by a programmable packet header processor enables high performance and minimizes latency and power
- X-PND™: Elastic resource allocation of control tables and packet memory, maximizing adaptation to application needs
- Fully shared packet buffer for incast burst handling
- X-IQ™: 64K queues and XFC™, a fine-grain, hardware-based flow control protocol, minimizing FCT
- Low power: Under 300W* for the 25.6 Tbps device and under 200W* for the 12.8 Tbps devices
- X-VIEW™: State-of-the-art telemetry for network troubleshooting and optimization
- X-MON™: In-field device health monitor for improved reliability

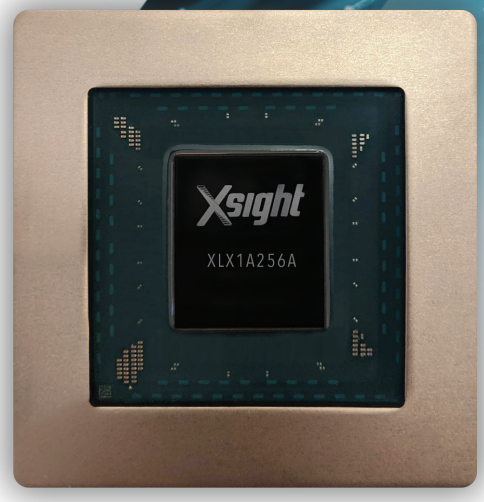

OVERVIEW

X1 is a family of highly integrated switching ASIC devices supporting 25.6 and 12.8 Tbps full-duplex throughput, optimized for interconnecting AI/ML, storage and compute clusters in data centers.

The X1 family of devices incorporates industry leading 100G LR PAM4 SerDes, enabling the design of in-rack DAC connectivity for the 100G ecosystem without the need for retimers, while future-proofed for 800G optics ecosystem. The X1 supports flexible port configurations using 100, 200, and 400 GbE speeds for port densities of up to 256 x 100 GbE, 128 x 200 GbE, or 64 x 400 GbE.

The devices support X-IQ™: 64K queues and XFC™, a fine-grained channelized flow control protocol accompanied by a fully shared buffer memory.

X-IQ™ minimizes Flow Completion Time (FCT) by maximizing burst tolerance, reducing packet loss rate, and minimizing fate-sharing and Head-of-Line (HoL) blocking among traffic flows.

The devices incorporate a programmable packet processor and X-PND™, an elastic resource allocation engine that allows flexible sharing of control tables and packet memory. The combination of both enables application-specific optimization, thus increasing performance while minimizing latency and power.

The X1 family of devices also supports X-VIEW™, an advanced telemetry engine for network troubleshooting and optimization. This includes in-band telemetry, any-cause verbose mirroring, black hole detection and localization, real-time statistics histograms, and uBurst detection.

On-chip health analytics (X-MON™) increase in-field reliability and reduce network failures within the devices.

Furthermore, these high-density devices minimize power consumption to less than 300W* for the 25.6 Tbps device and less than 200W* for the 12.8 Tbps device while running typical data center workloads.

*Typical Tier-1 data center customer use case.

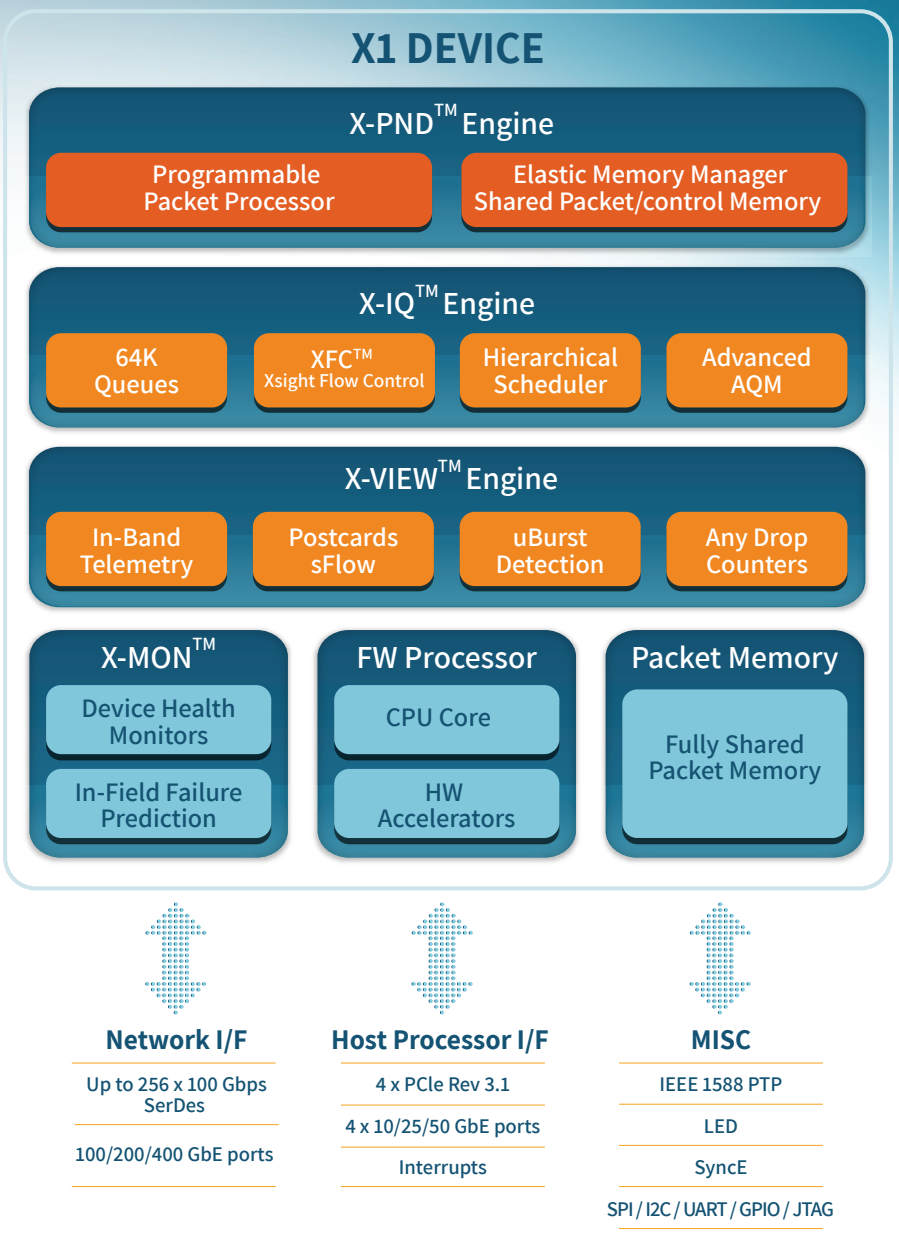

TARGET APPLICATIONS

The X1 family of devices 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 32 x 800G QSFP-DD/OSFP.

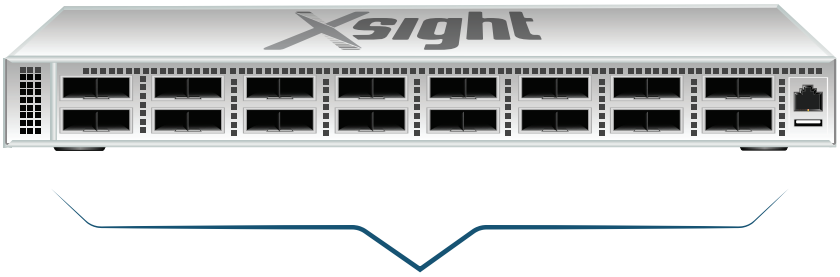

32 x 800G QSFP-DD/OSFP

FEATURES

Port Configuration

- 256 (X1-256), or 128 (X1-128) SerDes @ 100 Gbps
  - + Supports LR, DAC, and AEC without the need for retimers
  - + 800 Gbps optical transceivers
- Every group of 8 SerDes can be configured to 2 x 400, 4 x 200 or 8x100/50/25/10 GbE

Programmable Packet Processing and Parsing

- Support for application-specific processing to reduce latency and power
- Parallel hardware acceleration engines for increased fabric performance and reduced processing latency
- 256B parsing/packet editing capability

X-PND™ Engine

- Elastic resource allocation using a shared memory for packet buffer and control tables
- Flexible application specific memory partitioning

Packet Forwarding

- IPv4/6 Unicast routing
- IPv6/MPLS Segment routing

X-IQ™ Engine: Incast/Burst Tolerance

- Fully shared packet buffer memory minimizing packet loss rate
- 64K queues
- XFC™: Fine grained channelized flow control protocol
- Reduced flow completion time (FCT) with non-blocking minimal fate-sharing operation
- Low protocol handshake overhead

Traffic Management

- Scheduling and shaping
  - + 64K queues per device operating at full performance
  - + Three level scheduling/shaping hierarchy: physical port, virtual port, service and queue
- Advanced AQM
  - + Three drop precedence levels per queue
  - + RFC 7567 (WRED), Dynamic Buffer Allocation (DBA)

Load Balancing

- ECMP, WECMP, LAG and MLAG support
- Dynamic load balancing in flowlet granularity to optimize multi-path utilization and reduce hot-spots
- Smart hash generator
  - + Flexible selection of packet header fields and meta-data
  - + Supports a family of highly configurable hash functions to prevent hash polarization effects
  - + Consistent/resilient hashing
- Load imbalance detection using network telemetry

Congestion Notification and Lossless Networks

- ECN marking per RFC 3168, RFC 8311, RFC 8257 (DCTCP) and phantom queues
- Congestion notification message (CNM) generation
  - + Infiniband annex 17 RoCEv2 CNP
  - + IEEE 802.1Qcz congestion isolation
- IEEE 802.1Qbb PFC trigger/response
- IEEE 802.3x link level flow control

Control Plane Protection (CoPP)

- Four 10/25/50 GbE management ports, in addition to the network facing ports
- Up to 256 queues per management port supporting hierarchical scheduling and shaping
- Supports programmable packet classifiers and filters to flexibly associate packets with mirroring/ trap/discard profiles

Security

- TCAM based Ingress/Egress ACL
- Variable rule lengths supporting 5-tuples IPv4/6

NVMe Over Fabrics

- NVMe over RoCEv2/TCP forwarding with ECMP support
- Fine-grain flow classification using Infiniband BTH and NVMe Capsule fields in addition to IP/TCP/UDP 5-tuples
- TCP ACK prioritization to reduce latency, minimize retransmissions and increase throughput

X-VIEW™: Traffic Analytics and Telemetry

- Verbose ingress/egress mirroring
  - + Multiple target analyzers selected by the mirroring event
  - + SPAN, RSPAN, and ERSPAN types I, II, III with rich datasets added in INT/iOAM like headers
  - + Rich triggering events including: sFlow, discard, congestion and latency all associated with relevant information such as timestamp, drop reason, queue-ID, congestion level, inserted delay
- P4 INT/ IETF IOAM
  - + Supports Source, Transit and Sink nodes with up to 256B header editing
- Black holes and loop detection/ localization based on TTL and routing mismatches
- Switch health reports with hardware based export
  - + Variety of statistic counters including latency/queue size histograms, ingress/egress port utilization and discards
  - + Rich event collection in log queues. Every event is associated with a timestamp, 5-tuples and relevant dataset
  - + Events are exported using a DMA engine

uBurst/Burst Detection

- uBurst/Burst detection based on queue fill level and first order dynamics
- Report: uBurst start time, duration, size and victim flows

Clock Synchronization

- IEEE 1588 Precision Time Protocol: 1-step and 2-step PTP, Boundary and Transparent Clock operation
- ITU-T G.8261 Synchronous Ethernet

Advanced Visibility and Debug

- Discarded packets can be redirected to management processor with drop code
- Mirror to Host CPU on congestion
- Firmware based analysis and filtering

X-MON™: Xsight's Monitoring, Analysis & Predictive Maintenance Manager

- Performance degradation monitoring through the device lifetime from chip production, via system production and in-field use
- Detection and recovery from in-field transient and persistent errors

Firmware Processor

- Targets value added features that require high event rate processing in real time
- Accesses the entire device address space and application CPU DRAM
- Interrupts delivery mechanism to react to various system events
- Any-to-any DMA
- Hardware-based messaging and synchronization between the firmware processing core, data-plane processors and the application CPU

Xsight Software Development Kit (X-SDK)

- X-SDK delivers control APIs for the data center fabric, HPC, AI and feature set
- Multi-layered X-SDK design enables smooth integration with different NOS types
- X1 SW embraces open networking with OCP SAI and SONiC integration

PART ORDERING INFORMATION

The X1 product family is offered in 3 different variations, all of which share the same software and feature set and are are fully interoperable.
For additional product information, contact sales@xsightlabs.com

Part Order Number	Maximum Throughput	Network Facing SerDes	Port Configuration Examples
XLX1A256A	25.6 Tbps	256 x 100 Gbps	• 256 x 25/50/100 GbE • 128 x 200 GbE • 64 x 400 GbE
XLX1A128A	12.8 Tbps	128 x 100 Gbps	• 128 x 25/50/100 GbE • 64 x 200 GbE • 32 x 400 GbE
XLX1A128B	12.8 Tbps	256 x 50 Gbps	• 256 x 25/50 GbE • 128 x 100 GbE • 64 x 200 GbE • 32 x 400 GbE