Xsight Accelerate Program

Tell us what your project needs.
We have a solution ready.

Leverage our suite of ready-to-deploy deliverables to reduce the time-to-market of your project. Built on open architecture, our tried-and-tested solutions will power your innovation and scalability.

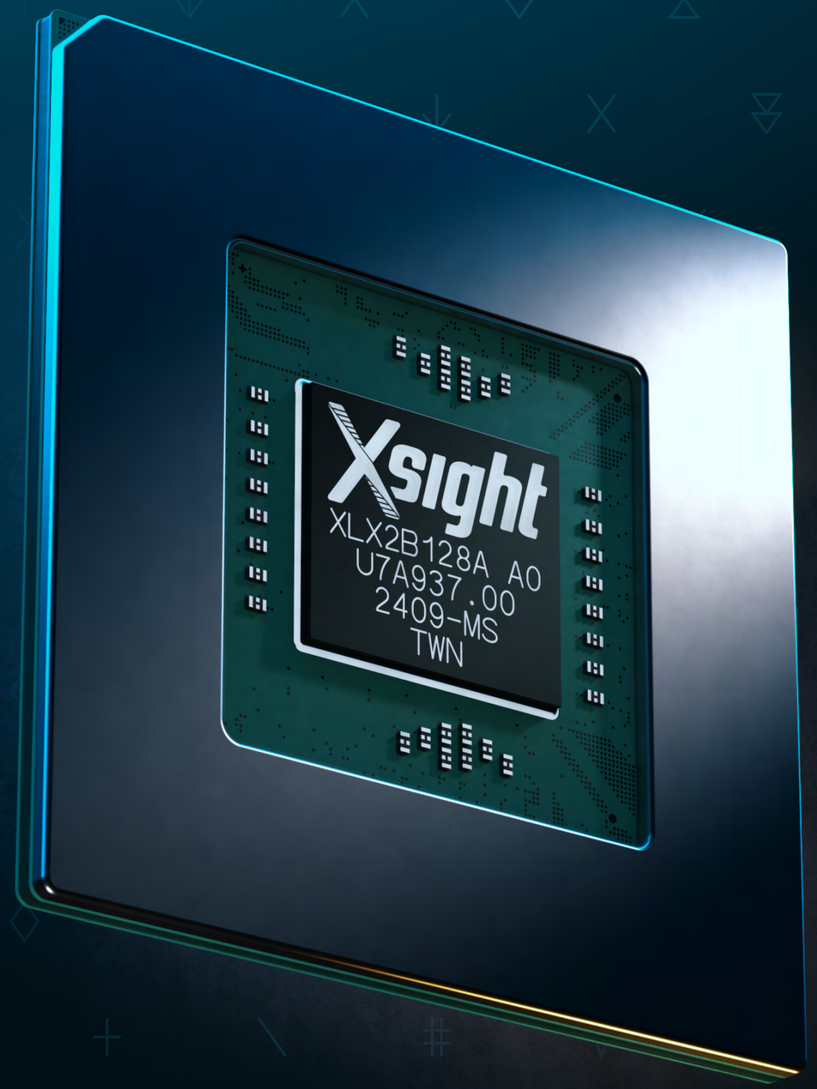

Contents

01 About Xsight Labs

04 What is the Accelerate Program?

06 Solution: Storage

09 Solution: Gateway and Network Packet Broker



About Xsight Labs

Efficient by design

Xsight Labs is a fabless semiconductor company focused on developing end-to-end connectivity solutions for AI and cloud infrastructure.

The company helps organizations overcome modern data center, AI, and edge infrastructure challenges with programmable products specifically designed to adapt to evolving business and technology requirements.



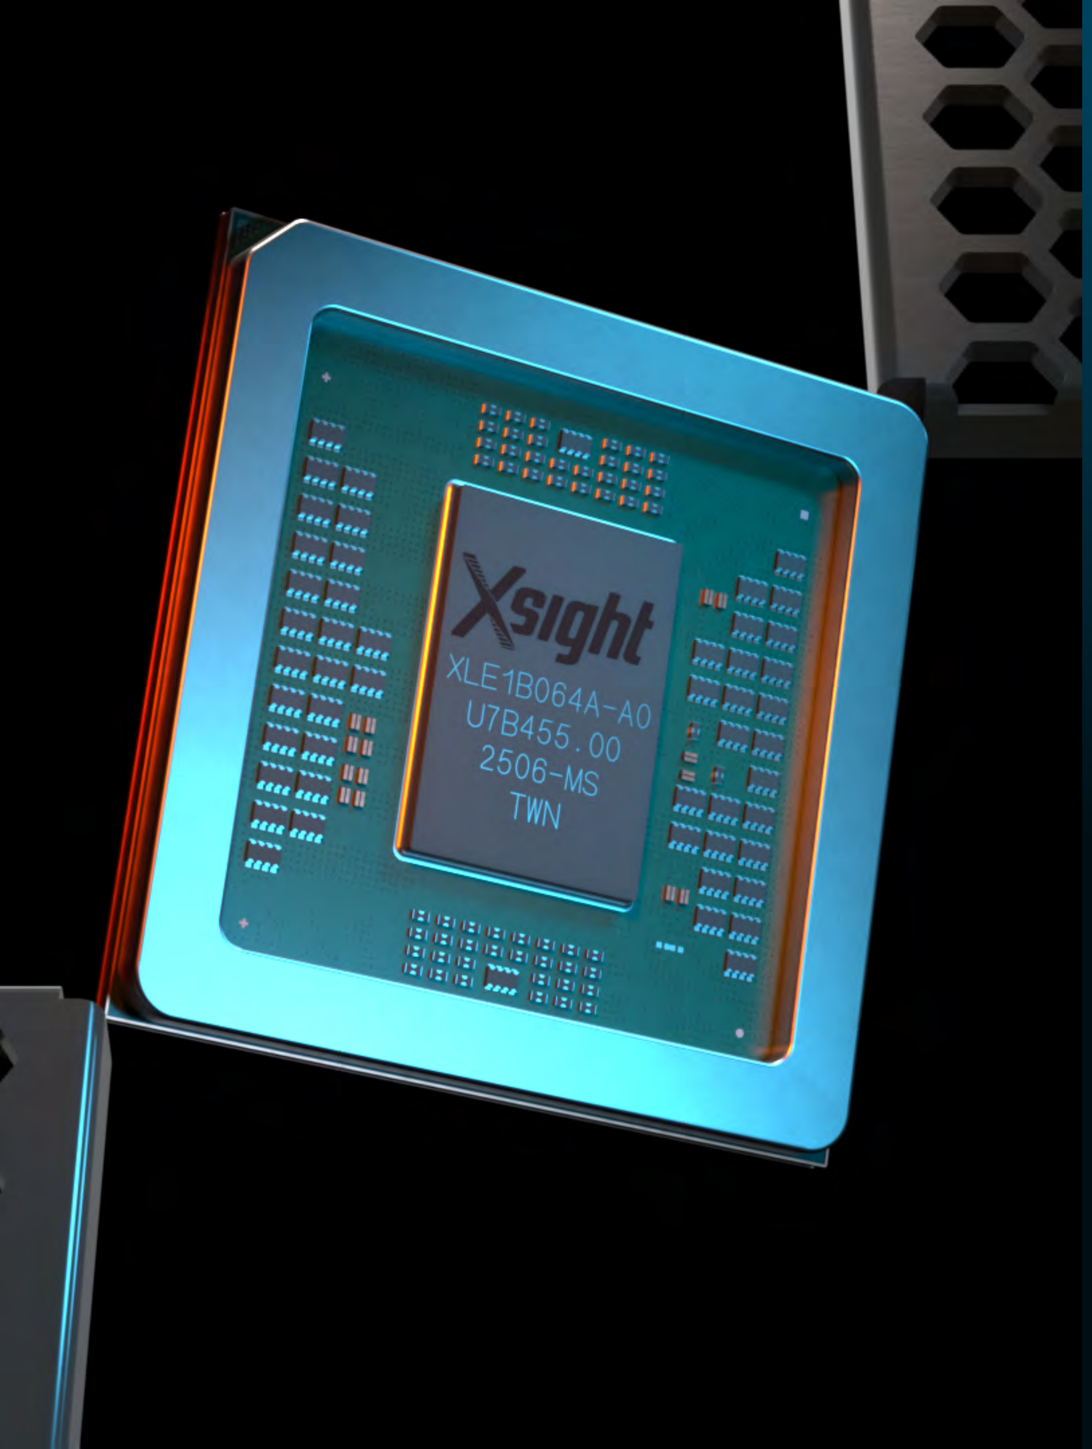

Adapted for AI

AI is transforming how businesses operate and innovate. Yet many legacy systems were built for consistency rather than adaptability and are struggling to keep up with growing AI and cloud requirements.

As demand for AI-ready applications continues to rise, organizations face significant operational and financial challenges that hinder transformation. Open, programmable architectures help avoid rigid, fixed-function environments that can limit innovation and increase complexity over time.

By removing traditional infrastructure constraints, Xsight Labs enables organizations to accelerate innovation and create a foundation for long-term growth.

Our vision and mission

To truly grow with confidence, partners need more than just silicon; they need a clear path to deployment. We help our partners build integrated, high-performance solutions on an open, programmable foundation. By combining silicon innovation with ecosystem collaboration, partners can accelerate development, reduce complexity, and bring differentiated infrastructure to market faster.



Open, Programmable Foundation
Build on flexible, standards-based infrastructure designed for AI, networking, and data-intensive workloads.



Performance at Scale
Deliver high-throughput, low-latency solutions with silicon and system design optimized for real-world workloads.



Built for Innovation
Co-develop and scale your solutions without the constraints of proprietary stacks. Whether you are optimizing a custom congestion algorithm or porting a complex security stack, our engineers work alongside yours to ensure line-rate performance.



Complete Partner Support
We help drive our customers' success with early access schemes, deep technical integration, co-development on open architectures, and co-marketing opportunities.

What is the Accelerate Program?

Xsight Accelerate delivers programmable, line-rate solutions across data center, AI, cloud, and edge environments. These solutions include ready-to-deploy deliverables including architectural blueprints, technical code, and integrated SONiC support.

The result: a significant reduction in time-to-market without compromising between speed, power, and security. With this foundation, organizations can stay agile and scale their infrastructure as demands change in an ever-shifting landscape.

This eBook covers our tried-and-tested solutions, all built on our proprietary technology and open architecture. But whether or not you find what you're looking for, talk to us. We'll gladly hear out your project's specific requirements and match it with a scalable solution.



The next-generation architecture powering our solutions

Xsight Labs is meeting modern data center challenges head-on by providing the critical silicon for next-generation environments. Xsight's X-Series switches and E-Series DPUs enable data center architects, AI infrastructure teams, and networking innovators to build and deploy intelligent infrastructure.

Our E-Series DPUs combine power efficiency with advanced compute capabilities. These software-designed DPUs enhance security and adaptability from cloud to core to edge, ensuring customers can handle critical workloads at scale.

X-Series switches are specifically designed for modern AI and cloud environments. Unlike older, fixed-function switch chips, the X-Series switch silicon is fully programmable, meaning that customers can implement custom behaviors and move data efficiently between AI, compute, and storage systems. This flexibility allows organizations to adapt to changing infrastructure requirements.

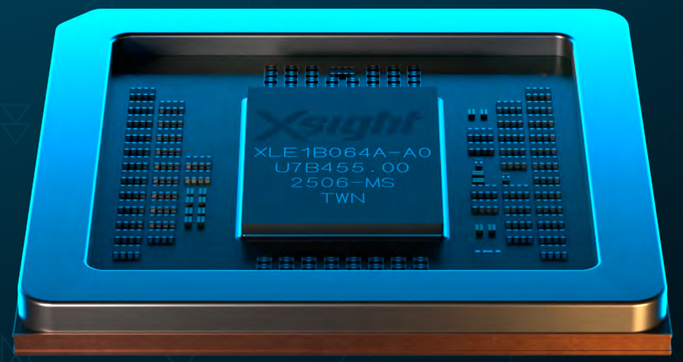

Xsight Labs' forward-thinking solutions provide the foundation organizations need to adapt, innovate, and scale with confidence.

Solution: Storage

Built for the AI era. Leverage programmable, line-rate storage at up to 800 Gbps with NVMe-oF acceleration, inline encryption, and advanced storage services powered by the E1 DPU.

Why Storage Matters

As data-intensive workloads continue to scale, storage infrastructure must keep pace. Traditional architectures place significant processing demands on host CPUs, while storage protection and data services add further overhead. Modern data centers need a different approach—one that delivers extreme performance without exceeding power or CPU limits.

Powered by the E1 DPU, Xsight's Accelerate Storage Solutions move demanding storage, networking, and security functions into dedicated, programmable hardware, helping organizations overcome traditional storage bottlenecks.

Key Features

800 Gbps Line-Rate Performance
Deliver line-rate stateful storage and network processing at up to 800 Gbps, with native NVMe-oF acceleration over NVMe/TCP and RoCEv2. Flexible fabric configurations support 2×400GbE, 4×200GbE, or 8×100GbE.

Offload Host CPU Processing
Move storage and network processing from the host CPU to the E1 DPU, freeing compute resources for applications and tenants.

Advanced Storage Services
Data services such as replication, data reduction, snapshots, clones, and EC can run directly on the E1 DPU.

Wire-Speed Security
Maintain transparent, wire-speed AES-XTS encryption with dedicated hardware acceleration. Hardware-accelerated PFC, ECN, and deep shared buffering support efficient, high-performance data fabric operation.

At a glance

Parameter	Description
Core Compute Architecture	Powered by high-efficiency on-path Arm processing elements (up to 64 N2 cores) unifying the infrastructure control and data plane layers.
Network Processing Capacity	Delivers line-rate stateful storage and network data pipeline bandwidth of up to 800 Gbps.
Fabric Interconnect Speeds	Turnkey support for multiple configurations: 2x400GbE, 4x200GbE, or 8x100GbE network fabrics.
Network Fabric Protocols	Native low-overhead, line-rate NVMe-over-Fabrics (NVMe-oF) offload via NVMe/TCP or RoCEv2.
Host Interface Link	Natively interfaces with host server nodes over a high-performance PCIe Gen5 x16 data path.
Direct Physical Storage Scale	Supports up to eight direct-attached PCIe Gen5 x4 NVMe SSDs on the E1-Server providing target functionality.
SATA Expansion Architecture	Scales up to 16 total SATA SSDs via four expansion cards running with downgraded PCIe Gen3 x2 link margins.
Hardware Acceleration Offloads	Dedicated on-path execution engines for block-level storage crypto (AES-XTS), zero-loss data fabric routing (PFC, ECN), and deep shared buffering logic.

Scale performance, not complexity.

Xsight combines the programmable E-Series DPU architecture with open, standards-based software interfaces to deliver storage acceleration without locking organizations into a fixed architecture. By integrating storage processing, networking, security, and data services into a configurable platform, Accelerate Storage Solutions can adapt to evolving infrastructure requirements while reducing host overhead.

Flexible deployment architectures support disaggregated E1-Server storage, local RAM-disk emulation, and local SATA SSD expansion—giving organizations the flexibility to tailor storage infrastructure to different workload and deployment requirements.

The result? A future-proof foundation capable of scaling alongside next-generation enterprise, cloud, and AI storage workloads.

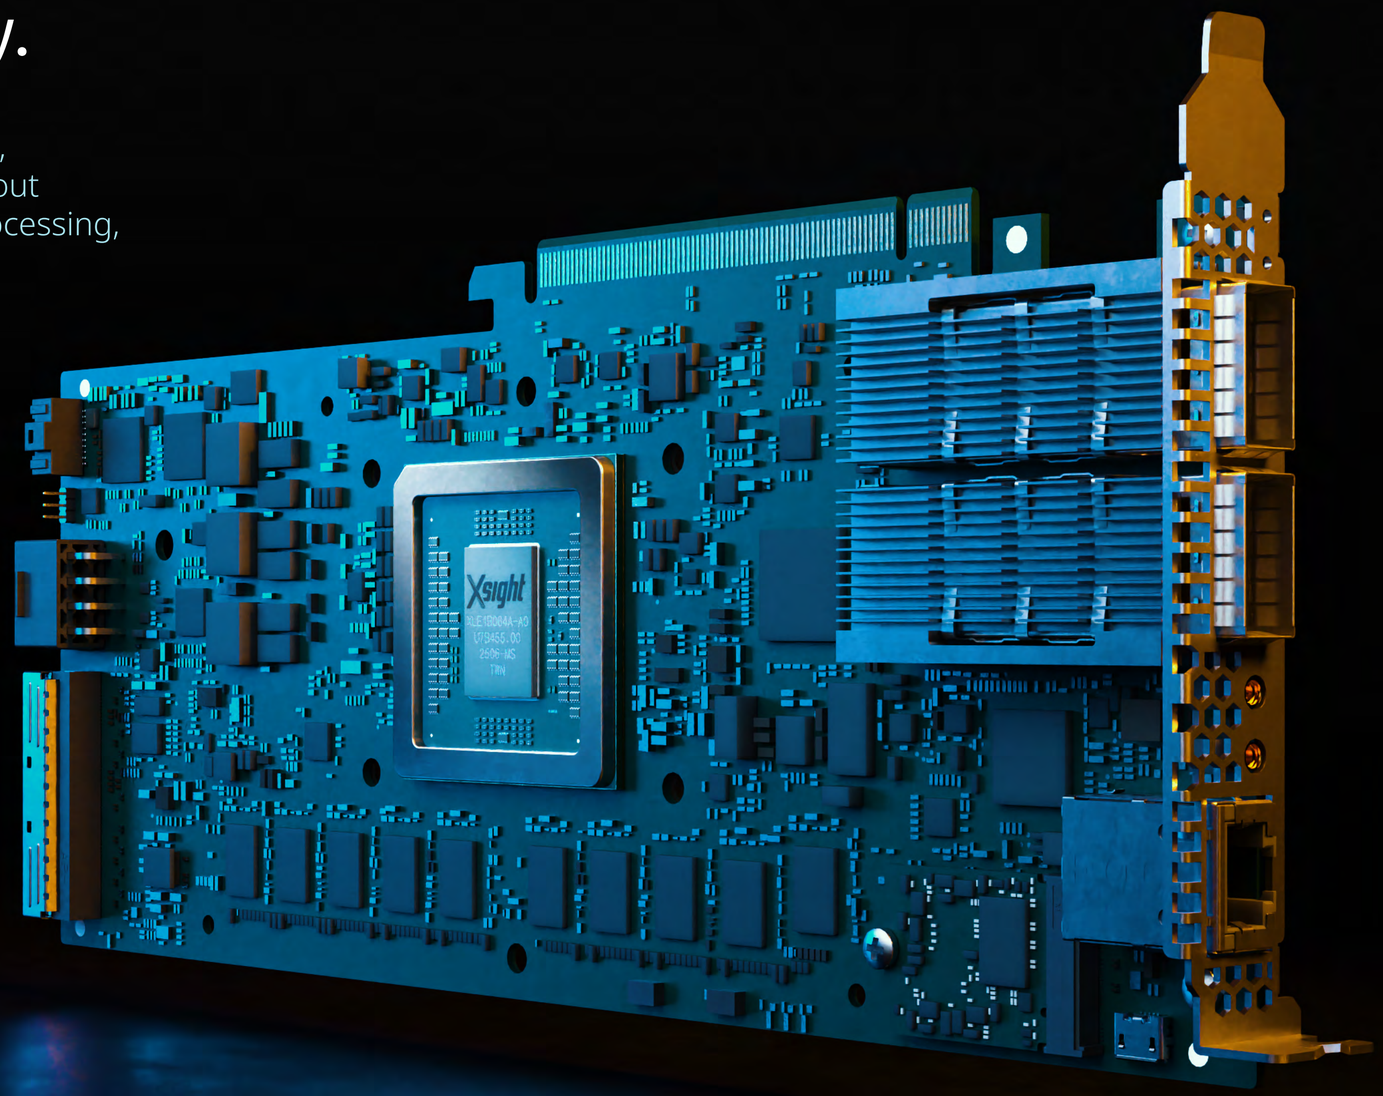

Solution: Gateway and Network Packet Broker

Built around our programmable 12.8 Tbps X2 Ethernet switch, this solution supports three primary deployment models spanning packet brokering and visibility, gateway functions, and programmable filtering and telemetry.

Get hardware-level speeds without locking packet-processing behavior

As a programmable switch, X2 supports the building of gateways and packet brokers directly to the silicon. This enables brokering, filtering, and gateway logic to run as software on the packet-forwarding engine at full line rate in under 200W.

Need new capabilities for your visibility fabric? Write it, load it, and run it. X2 gives you the performance to manage high traffic volumes while maintaining the flexibility you need to allow your packet-processing behavior to evolve for future requirements.

Key Features

Line-rate programmable filtering
The XISA parser walks up to 256 bytes of headers and generates lookup keys in software. Filter, classify, and steer at 12.8 Tbps with no recirculation penalties.

Tool-port aggregation and steering
128 x 100G PAM4 SerDes support 25 to 800 GbE port speeds. More than 1M table entries hold filtering and forwarding rules, 8x the scale of incumbent 12.8T silicon.

X-VIEW telemetry suite
In-band telemetry, sFlow, postcards, microburst detection, black hole localization, and any-cause verbose mirroring give the visibility fabric its own visibility.

Elastic table and buffer memory
X-PND allocates fully shared memory between control tables and packet buffer, so a filter-heavy broker profile and a buffer-heavy gateway profile run on the same chip.

At a glance

Parameter	Description
Switch Silicon Architecture	Programmable packet-forwarding engine with two stages: a Parser that walks up to 256 bytes of headers and a Match-Action Processor, executing in parallel rather than a serial pipeline
Network Processing Capacity	12.8 Tbps full-duplex, non-blocking; SKUs at 3.2, 6.4, 9.6, and 12.8 Tbps with identical software
Port Interface Speeds	25/50/100/200/400 GbE across 128 x 100G PAM4 LR SerDes; up to 16 x 800G OSFP or 32 x 400G
Forwarding Latency	450 to 700ns first-bit-in to first-bit-out, cut-through
Table Scale	More than 1M entries for filtering, steering, and forwarding rules
Traffic Management	ECMP, WECMP, LAG, MLAG, flowlet-granularity dynamic load balancing; WRED, ECN, and PFC congestion handling
Telemetry and Visibility	X-VIEW suite: in-band telemetry, postcards, sFlow, microburst detection, black hole localization, any-cause mirroring, real-time statistics histograms, SerDes link quality meters
Data Plane Programmability	XISA open instruction set (published on GitHub); P4 compiler development enabled

Parameter	Description
Software Stack	Native SONiC/SAI via XSAI; X-SDK with XSW, XHAL, and XBM behavioral model
Power Envelope	Under 200W at 12.8T, under 140W at 6.4T, under 90W at 3.2T
Timing and Synchronization	IEEE 1588 PTP and Synchronous Ethernet

One programmable switch. Multiple layers of flexibility.

X2 can be deployed across three key use cases: packet brokering and visibility, gateway services, and programmable filtering and telemetry. For packet-brokering use cases, there are a further three configuration options to choose from.

Deployment models

Packet broker and visibility fabric
Aggregate monitored traffic from production links, filter it in the XISA data plane, and distribute it to analysis tools at line rate. Port configurations up to 128 x 100GbE, 32 x 400GbE, or 16 x 800G OSFP

Gateway functions
Run cloud and carrier gateway services, including Carrier NAT and load balancing, on X2-based platforms instead of racks of x86 appliances. Pair X2 with E1 DPUs for stateful L4-L7 services.

Programmable filtering and telemetry via XISA
Build brokering features the silicon vendor never shipped: proprietary header parsing, custom tunnel termination, in-band telemetry insertion, and packet tracing, all written against a published instruction set.

Platform options for Packet Brokering

Fixed-function broker appliance
This incumbent approach leverages a fixed-pipeline switch ASIC or x86 appliance to deliver today's filtering feature set. New encapsulations, telemetry formats, or steering behaviors wait for the vendor's next hardware generation, and x86 datapaths are limited to 100G and 200G connectivity.

X2 programmable broker
An X2-based open networking platform, such as the Edgecore Universal ToR or Interface Masters Tahoe 3832 EX, runs brokering logic in XISA software at 12.8 Tbps line rate. Filtering behavior changes by loading new code, not new hardware.

Enhanced switch appliance (X2 + E1)
The Interface Masters Tahoe 3828 EXA adds two E1 800G DPUs to the X2 switch, combining 11.2T of L2-L4 brokering with 1.6T of stateful L4-L7 processing in 1RU at under 475W. One platform covers packet brokering, gateway services, and advanced instrumentation telemetry.

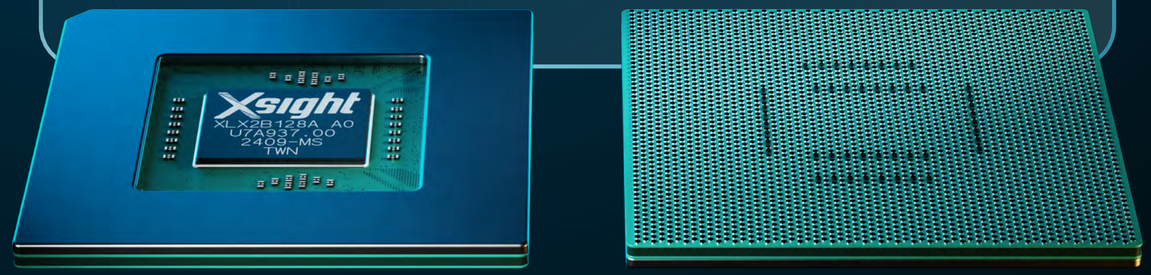



Get in touch today:
sales@xsightlabs.com
www.xsightlabs.com