

A large, stylized 'X' graphic is positioned on the left side of the slide. It has a 3D effect with a blue and white gradient on its upper left arm and a dark blue gradient on its lower right arm.

Empowering the Ecosystem for AI and Beyond

How Xsight Labs can help you bring differentiated
and efficient solutions to market faster


Xsight Labs is a fabless semiconductor company building end-to-end connectivity solutions for AI and hyperscale data centers. Our switch silicon and data processing units provide the bandwidth, efficiency, and programmability required by next-generation networks.

We take a fundamentally different approach to chip design.

Our open architecture empowers customers to customize network behavior without waiting for new hardware generations.

Our products deliver up to 40% lower power consumption

than incumbent solutions and are fully programmable, capable of supporting protocols that have yet to even be imagined.




What You Can Build with Xsight Labs


Organizations that partner with Xsight Labs gain the blueprint for a new generation of high-margin, intelligent infrastructure. By leveraging our X-series switches and E-series DPUs, partners move beyond selling commodity hardware to delivering specialized, high-value solutions for AI, cloud, and data-intensive environments.

Deployable Infrastructure Platforms

Partners use Xsight Labs silicon to build and brand several “ready-to-deploy” product categories:



Intelligent AI Storage and Networking: Partners like Hammerspace built an Open Flash Platform (OFP) that integrates E1 DPUs directly into storage arrays. This allows them to sell high-density storage solutions that offer 10x the density and 9x the power savings for AI training and inference.



Next-Gen Smart Switches and Appliances: Partners such as Edgecore Networks and Interface Masters build fully programmable 800G and 12.8T SDN switches. These are sold as branded enterprise networking gear that supports open ecosystems like SONiC and P4, allowing partners to capture market share in the growing white-box and open-networking markets.



AI-Driven Security Services: Cybersecurity partners integrate their software stacks onto our 800G DPU add-in cards (AIC) and 1RU edge servers. This allows them to sell a complete, air-gapped Zero Trust security appliance that handles real-time threat detection at line rate without taxing the customer's host PU.

ROI Made Simple

Xsight Labs's open, software-defined architecture changes the economic model for infrastructure partners:

Accelerated Time-to-Revenue:

Because our XISA architecture and programmable pipelines are software-defined, partners can develop and deploy new features like custom offloads or specific protocol support, in software. This allows them to bring new products to market months (or years) ahead of competitors waiting for custom hardware spins.

High-Value Vertical Customization:


Partners move away from price-per-port competition by creating tailored solutions for specific workloads. Whether it's optimizing network behavior to reduce GPU cluster congestion or designing efficient data pipelines for storage, partners can charge a premium for specialized performance.

Expansion into Neo-Cloud and Edge Markets:

Our power-efficient, single-die designs allow partners to build solutions for high-growth, space-constrained environments from Starlink-style satellite constellations to 6G Edge wireless infrastructure.

Open, Programmable Silicon:

Partners can accelerate innovation and differentiation with programmable silicon that adapts to specific products and markets. Open and customizable architectures (XISA, unified SDKs, Arm-based design) enable seamless integration and co-development.



Xsight X2 Switch Silicon: The Foundation for High-Margin Infrastructure

For partners building next-generation networking, the X2 is more than a switch, it is a programmable platform. By moving away from fixed-function silicon, partners can escape the commodity hardware trap and build high-value, differentiated solutions.

Revenue Opportunities for X2 Partners


Partners use Xsight Labs silicon to build and brand several “ready-to-deploy” product categories:

- **Monetize “Zero-Wait” Innovation:** Proprietary silicon forces partners to wait 2+ years for hardware refreshes to add new features. With X2’s XISA architecture, partners can sell software-based feature updates (like new congestion control for AI clusters) instantly, creating recurring revenue streams without new hardware spins.
- **Capture the AI Fabric Market:** Use X2 to build Ultra Ethernet Consortium (UEC)-ready fabrics. Partners can offer “Lossless GPU-to-GPU” networking solutions that command a premium in the AI market compared to standard leaf-spine switches.
- **Vertical-Specific Appliances:** Instead of a general-purpose switch, partners can build specialized appliances for High-Frequency Trading (leveraging sub-700ns latency) or Satellite/Edge Networking (leveraging the 40% lower power profile used by SpaceX Starlink).

Partner Insight: Because X2 is fully programmable at the assembly level, you aren’t just a reseller of a chip, you are the architect of the network’s behavior.

“Xsight’s X2 chip will form an integral part of the terabit routing needs of the next generation of Starlink satellites.”

Michael Nicolls, VP Starlink Engineering, SpaceX




Xsight E1 DPU:

Building the Cloud-on-a-Chip

The E1 DPU is the industry's first and only fully software-defined 800G DPU. For partners, the E1 represents a shift from "NIC offload" to "Server-Class Processing," placing 64 Arm Neoverse N2 cores directly in the data path.

How Partners Build and Profit with E1

- **Consolidated Infrastructure Sales:** Partners can build all-in-one appliances that converge networking, security, and storage into a single 1RU E1-Server. This reduces a customer's rack footprint and power draw while allowing the partner to sell a single, high-margin appliance.
- **Software-Defined Storage Gateways:** Use the SPDK and NVMe-oF support to build high-performance storage targets. Partners like Hammerspace use this to deliver global data environments where the E1 handles the massive data movement required for AI training at line rate.
- **Line-Rate Security Services:** Cybersecurity partners can port their existing Linux-based XDP/DPDK stacks to the E1 to offer 800G Zero Trust appliances. By running security entirely on the DPU, you can sell air-gapped security that doesn't steal CPU cycles from the customer's primary workload.



The Only Fast Path Commercial Advantage

Traditional DPUs have a slow path and a fast path, when a packet hits the slow path, performance drops.

- **Partner Benefit:** E1's architecture means your solutions deliver predictable performance with only a fast path. This allows you to offer stricter service level agreements (SLAs) to your customers, a key differentiator when working with Cloud Service Providers and Enterprises.


Expanded Form Factors and Integration

To support different partner business models, the E1 is available in three distinct integration points:

- **E1-Server (1RU System):** Best for Solution Providers who want a branded, turnkey appliance for Edge AI or Security.
- **E1-AIC (PCIe Add-in Card):** Best for Server OEMs looking to upgrade existing Gen5 server portfolios to 800G AI-ready nodes.
- **E1-SoC (Direct Silicon):** Best for Hyperscalers and Custom System Builders designing proprietary hardware from the ground up.

“The E1 is setting new standards of performance.”

Dong Wei Arm, Standards Architect and Fellow



Xsight Labs collaborators aren't just resellers; they are architects of next-generation infrastructure. By integrating X2 and E1 silicon, these partners have developed high-performance solutions that are shipping and scaling today.

Larch Networks

The Solution: Larch Networks has developed a hardened, carrier-grade SONiC distribution specifically optimized for Edgecore X2 switches. This solution provides partners with a "turnkey" open-networking OS that supports complex routing protocols and telemetry out of the box.

Hammerspace

The Solution: Hammerspace integrates the E1 DPU into their Open Flash Platform (OFP) to create a "flattened" AI storage architecture. This eliminates the need for intermediate storage servers, allowing GPUs to talk directly to flash memory. This integration delivers 9x power savings and 10x storage density, making it the gold standard for AI "Warm Storage."

Niagara Networks

The Solution: Niagara Networks utilizes Interface Masters' X2-based hardware to power their high-visibility Packet Brokers. By using X2's programmable pipelines, Niagara delivers line-rate TLS decryption and traffic intelligence, allowing SecOps teams to inspect 400G / 800G traffic without dropping a single packet.

RedFig Networks

The Solution: RedFig has validated its Network Analytics and Zero Trust security stacks on the E1-Server (1RU) and E1-AIC platforms. They provide a "Portal Accelerator" for managed service providers (MSPs) to deploy AI-driven threat hunting and real-time network forensics as a subscription service.

"Legacy storage is collapsing under the weight of AI. By fusing our orchestration software with Xsight's 800G DPU, we flatten the data path and turn every Linux device into shared storage. The result is an open architecture that scales linearly with performance, slashes cost, and feeds GPUs at the speed AI demands."

David Flynn, CEO of Hammerspace




Foundational Technology Partners

We build our solutions on trusted technology platforms that provide industry-leading capabilities.

Revenue Opportunities for X2 Partners

- **Arm:** Neoverse N2 cores power the E1's 64-core compute complex. Arm SystemReady Level 6 certification ensures off-the-shelf Linux compatibility, enabling partners to use standard tools like XDP, DPDK, and SPDK without proprietary SDKs.
- **Alphawave:** Provides the industry-leading 100G SerDes technology that enables the high-speed, low-latency connectivity required for 800G networking.
- **TSMC:** Our X2 and E1 silicon is manufactured on TSMC's advanced 5nm process, which is the secret sauce behind our 40% power reduction compared to legacy 7nm and 16nm incumbents.



Trusted Standards. Scalable Ecosystems.

Xsight Labs entire product line is aligned to open standards and industry initiatives that promote interoperability and innovation.

We've taken this approach for three important reasons: open standards remove vendor lock-in (meaning far greater interoperability), accelerate time to market, and give partners the freedom to build differentiated infrastructure solutions on a shared foundation.

Together, the standards here enable hyperscale-class performance, validated in real-world testing at 800G and beyond.

Ultra Ethernet Consortium

Member since November 2023. X2 includes UEC-ready features for AI cluster networking. Enables partners to design high-performance AI clusters with predictable latency and congestion control.

P4

Open ISA (XISA) enables P4 compiler development for fully programmable packet processing. Allows partners to create custom packet processing pipelines; for example, embedding security, telemetry, or workload-aware routing directly into the data path.

SONiC

E1 achieved the industry's first SONiC-DASH Hero testing at 800G, validating 14.25 million connections per second. It gives partners a production-ready, open network OS to build and validate solutions quickly, reducing dependency on proprietary network stacks.

Open Flash Platform

Co-founded with Hammerspace, Linux Community, and Los Alamos National Laboratory to define AI warm storage architecture. Creates a new tier of AI storage architecture, allowing partners to build solutions that bridge high-performance compute with scalable, cost-efficient data access.

What's the Opportunity for Infrastructure Partners?

Data centers have always been evolving. But 2022 was the year things shifted to AI. Since then, the rulebook has been rewritten.

Traditional workloads like web apps, data, and general low-intensity traffic needed data centers built around air-cooled CPUs, with an emphasis on virtualization and comparatively low-level compute demands.

Infrastructure today must support thousands of GPUs operating as a single system, handling high-throughput, low-latency AI workloads at scale. This shift is exposing the limits of traditional, fixed-function silicon which is where programmable infrastructure comes in.

14% growth ¹

CAGR data center sector growth through to 2030

**12% energy usage
increase year on year ²**

Over the last five years

**21% of the
world's energy ³**

Data centers could account for up to 21% of overall global energy demand by 2030



The Size of the Opportunity for Infrastructure Companies

Today, dedicated AI data centers are being built all over the world. The global data center sector will likely expand at a 14% CAGR through 2030. That means there is, and will continue to be, high demand for data center components.

Power is now one of the largest operational costs, with hyperscale facilities often facing energy bills running into millions of dollars per month.

Following the initial surge in model training, the landscape is increasingly dominated by inference. This transition reflects a market moving from development to delivery, where the primary demand lies in serving active GenAI chatbot interactions.

Limitations of Closed and Legacy Silicon

Closed, legacy silicon limits flexibility, locking partners into fixed features, slow upgrade cycles, and vendor-controlled roadmaps. This makes it harder to adapt to evolving AI workloads. Open, programmable architectures remove these constraints, enabling faster innovation, deeper customization, and more efficient, scalable solutions tailored to modern infrastructure demands.

An Open Future

Open source software (OSS), brings freedom. But it also reduces costs by letting you do more with existing appliances, avoids vendor lock-in, and increases operational agility which means you can change direction on your schedule.

Three Forces at Work

- 1 Energy consumption and performance per watt are critical
- 2 Vendor lock-in is limiting
- 3 Software-defined development is the way forward for infrastructure designers who want programmability and interoperability



Why Partner with Xsight Labs?

A trusted, open, partner-first innovation platform that helps you bring differentiated, efficient, and scalable solutions to market faster.

Partnering with Xsight Labs gives organizations the foundation to build, differentiate, and scale next-generation infrastructure solutions with confidence.

Create Differentiated Solutions, not Commodity Infrastructure

Xsight's open, software-defined architecture removes the constraints of traditional silicon, allowing you to design infrastructure that behaves exactly how your customers need it to. Instead of relying on fixed features, you can build tailored capabilities that set your solutions apart in competitive markets.

Accelerate Time to Market and Innovation

With programmable pipelines, open standards, and a shared SDK, your teams can develop, test, and deploy new features in software, without waiting for hardware refresh cycles. This enables faster iteration, quicker response to customer demands, and a more agile approach to product development.

Unlock New Revenue and Growth Opportunities

Xsight Labs supports partners with early access to product roadmaps, co-development opportunities, and joint go-to-market initiatives. By participating in a growing ecosystem of infrastructure providers and solution builders, you can expand your offerings, reach new customers, and create additional revenue streams.

Build on a Platform Designed for Long-Term Partnership

Beyond technology, Xsight provides the technical support, integration expertise, and enablement framework needed to design, deploy, and scale solutions with confidence. Through close collaboration and ongoing innovation, partners are equipped to evolve their offerings as new workloads, standards, and opportunities emerge.



Let's Build Next-Generation Infrastructure Together

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 Partner Innovation

Co-develop, integrate, and scale solutions within an open ecosystem without the constraints of proprietary stacks. Whether you are optimizing a custom congestion algorithm on X2 or porting a complex security stack to the E1's 64 Arm cores, our engineers work alongside yours to ensure line-rate performance.



Powerful Partner Benefits

Early access, deep technical integration and co-development on open architectures. Plus, co-marketing opportunities and access to customers and go-to-market channels.





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

¹ <https://www.jll.com/en-uk/insights/market-outlook/data-center-outlook>

² <https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai>

³ <https://mitsloan.mit.edu/ideas-made-to-matter/ai-has-high-data-center-energy-costs-there-are-solutions>

