



Ambarella Launches X7, Its First Standalone AI Accelerator, to Add Physical AI to Any Host Processor

September 15, 2026 at 9:00 AM EDT

New edge AI accelerator brings Ambarella's third-generation CVflow® engine to Arm and x86 host systems

Leading OEMs in several key verticals are currently creating designs with X7

SANTA CLARA, Calif., Sept. 15, 2026 (GLOBE NEWSWIRE) -- Ambarella, Inc. (NASDAQ: AMBA), an edge AI semiconductor company, today announced the X7 edge AI accelerator, its first dedicated CVflow® AI accelerator chip. X7 serves as an AI co-processor for Arm and x86 host processors, adding or expanding AI capabilities in a system without a redesign of the underlying platform.

Existing installations of cameras, gateways or industrial controllers often have limited AI capabilities, or none at all, but may have years of operating life remaining. Owners and operators of these older deployments are now requesting AI features that were unavailable when the hardware was built. Millions of systems working in the physical world fall into this category. X7 is intended for OEMs and system builders responsible for that existing installed base, and for teams validating new AI features before committing resources toward a next-generation redesign. For physical AI, a category describing machines that perceive, reason and act in the physical world, X7 allows these machines to gain intelligence over their lifetimes.

"Many of our customers' systems in the field were designed before the AI workloads their end users now expect were possible," said Fermi Wang, President and CEO of Ambarella. "X7 brings the same third-generation CVflow engine that runs in more than 50 million of our AI SoCs to platforms our customers have already deployed at scale. Our integrated SoCs remain the strongest choice when a team is designing the entire system around AI, while X7 extends the same engine to every other scenario, whether the platform is already in the field or still being prototyped."

Power, thermal limits and system cost

A camera or an inspection system that once ran a single detection model now runs several simultaneously, often combining conventional computer vision, vision-language models and a natural-language interface. Whether AI features like these can successfully operate in an embedded product depends on inflexible factors such as power, thermal headroom, latency and system cost. In current customer designs, X7 operates within a power envelope between 2 and 5 watts, allowing the accelerator to run within a platform's existing constraints, maintain the enclosure size, and avoid unnecessary increases to the current bill of materials. X7 supports CNNs, vision transformers, multimodal transformers and hybrid networks using the same CVflow computer vision and deep learning processor Ambarella integrates into its SoCs.

How X7 integrates into a host system

X7 is a host-agnostic accelerator, connecting to the host over a single lane of PCIe Gen 3 or over USB 3.2 and appearing to the system as a standard PCIe device. The host can also be one of Ambarella's own SoCs, so a customer shipping an Ambarella-based product today has a path to improved AI performance for a given workload without a board redesign.

Software, tools and model support

X7 is supported by the Cooper™ Developer Platform, the same environment that serves every Ambarella AI SoC, including the compiler, quantization and profiling tools, the C++ and Python runtime APIs, and the scheduling and memory management that multi-model pipelines require. Models from the Cooper Model Garden in the Ambarella Developer Zone are designed to run on the X7 accelerator as well, with the same documentation, sample code and engineering support Ambarella provides to its SoC customers.

Because X7 shares the CVflow architecture with the overall Ambarella SoC portfolio, a model or perception pipeline built for one of the company's SoCs is designed to carry forward to X7 without requiring a software rewrite, following the same path used when a product line moves to a next-generation integrated SoC. The CVflow architecture is now in its third generation. Ambarella has been building neural processing into its silicon since 2015, and its AI processors number more than 50 million operating in the field today.

XCalibur reference design

XCalibur is Ambarella's reference design and development platform for X7. A standard 2280 M.2 card with LPDDR5 memory and PCIe throughput above 840 MB/s, XCalibur can process several video streams and AI models simultaneously. Ambarella intends to make the complete XCalibur design files available to customers, enabling an OEM or ODM to move from evaluation to a production module on the same design.

Markets, channels and availability

Target applications include retail and physical security, industrial inspection and manufacturing, mobile robots and drones, intelligent transportation and critical infrastructure, and enterprise edge systems such as kiosks, point-of-sale terminals, gateways and on-premise AI appliances. X7 also opens markets for new designs built around a host processor from another supplier, as well as designs built on an Ambarella SoC that require additional AI performance. The X7 product is not specific to cameras and serves a wide variety of digital and physical data modalities, which makes it suitable for edge-infrastructure inference systems and for endpoints such as robots.

Ambarella is bringing X7 to market through the company's direct sales channels as well as the indirect channels established via its relationships with Capgemini and Macnica announced on September 3, 2026. Capgemini works with large enterprises on complex deployments, while Macnica covers small and mid-sized customers, independent software vendors, OEMs, ODMs and systems integrators across the Americas and EMEA.

X7 is sampling now, and XCalibur evaluation kits are available to qualified customers. Ambarella will discuss X7 in further detail at the AI Infrastructure Summit, September 15-17 at the Santa Clara Convention Center.

About Ambarella

With an installed base of more than 50 million AI SoC units, Ambarella's products are utilized in a wide variety of physical edge AI applications, spanning edge endpoint and edge infrastructure use cases including physical security, vehicle safety, telematics, autonomy, portable video, aerial drones, and other emerging robotic applications. Building on this footprint, Ambarella offers a full-stack edge AI platform, from highly optimized silicon and programmable software to AI agentic frameworks that coordinate perception, decision-making and control across devices. Ambarella's low-power systems-on-chip (SoCs) integrate proprietary and highly efficient perception and deep learning neural network AI accelerators, enabling electronic systems to become more productive with partial or complete levels of machine autonomy. For more information, please visit www.ambarella.com.

Media Contact

Jonathan Miller

jmiller@ambarella.com

