AI FRAME GRABBER — EDGE AI · INDUSTRIAL

AI Frame Grabber-4PE

Edge AI hardware built for fleet deployment.

4-camera PoE+ GigE Vision capture and dual-NPU AI inference on a single PCIe card — with secure boot, full-disk encryption, A/B OTA, and container-ready NPU access built in.

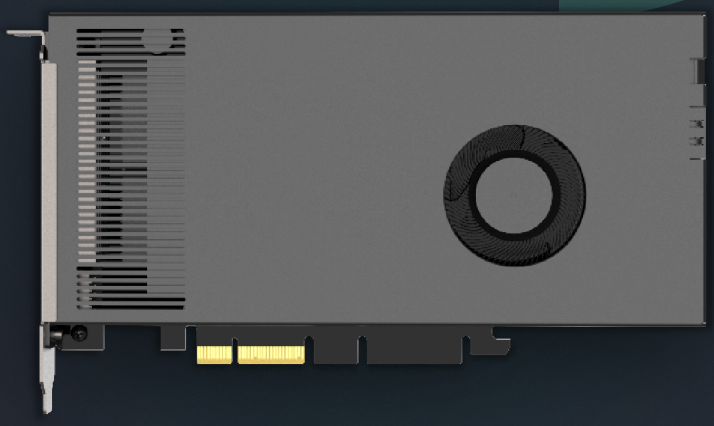

CORE HIGHLIGHTS

Four reasons 4PE ships.

4-Camera GigE Vision Acquisition

Synchronized capture across 4 PoE+ industrial cameras. Real-time multi-camera inspection on one card.

Dual-NPU Edge AI Compute

RK3588 + RK1820 inference on-card. 16 GB RAM, 128 GB eMMC for your apps and models. No external GPU.

Production-Grade Foundations for Edge AI Fleets

Secure boot, full-disk encryption, A/B OTA partitions, container-ready NPU access — pre-integrated in EVVR OS. Ready for your fleet platform.

Industrial Reliability

24/7 continuous inspection workloads. −20 °C to +65 °C operating range. Windows and Linux supported.

PERFORMANCE

AI compute on the card.

End-to-end measurements on production 4PE hardware. No marketing TOPS.

ResNet-50 v1.5 — Industry CNN Baseline

P99 latency **7.59** ms • Offline throughput **971** SPS

• @ 4 cameras × 30 fps **~8** passes / frame

Foundation Models — FP16, no quantization

MODEL	LATENCY	THROUGHPUT
DINOv3-ViT-S/16 (224)	22 ms	290 fps
DINOv2-ViT-S/14 (224)	27 ms	268 fps
DINOv3-ViT-B/16 (224)	60 ms	114 fps

Vision-language and edge-RAG pipelines on-card. No cloud round-trip.

16.8 W **Total board power** under sustained ResNet-50 Offline load — including 4× PoE+ control, 4× 2.5 GbE, dual-NPU, on-board DDR.

Internal benchmark on production 4PE hardware. LoadGen v5.1, MLPerf Inference Edge methodology. CNN: INT8 w8a8. Foundation models: FP16. End-to-end measurements include PCIe, DMA, and host-side overhead. **Not an MLCommons submission.**

TECHNICAL SPECIFICATIONS

The card, spec by spec.

PRODUCT MODEL	4PE
CPU	Rockchip RK3588
AI ACCELERATOR	RK1820
MEMORY	16 GB
STORAGE	128 GB eMMC
BUS INTERFACE	PCIe x4
POE	4× IEEE 802.3at PoE+ (25.5 W / port · 50 W total)
ETHERNET	4× 2.5 GbE
I/O	1× CAN · 1× GbE · 1× USB 2.0 · 1× USB-C · 1× DisplayPort
SERIAL	1× RS-232 · 1× Isolated RS-485
SUPPORTED CAMERAS	GigE Vision industrial cameras
MAX CAMERAS	4
SOFTWARE PLATFORM	EVVR OS — automation with deep-learning support
OS	Windows · Linux
POWER	12 V DC
DIMENSIONS	227.7 × 111 mm (W × H)
OPERATING TEMPERATURE	-20 °C to +65 °C
STORAGE TEMPERATURE	-20 °C to +65 °C
HUMIDITY	10–90 %, non-condensing

WHERE 4PE FITS

Built for teams shipping edge AI.

USE CASES

Smart-factory line upgrades • Multi-camera AOI and defect detection • Edge AI deployment (federated learning, branch-site inference)

INDUSTRIES

3C electronics • Automotive electronics • Semiconductor • Industrial equipment

Why 4PE

- **Deployment foundations built in.** Secure boot, full-disk encryption, A/B OTA partitions, container-ready NPU access — leverage out of the box, or integrate with your existing fleet platform.
- **Dual-OS support.** Windows and Linux, both first-class.
- **16 GB RAM + 128 GB on-board eMMC.** Headroom for your own applications, model variants, and dev tooling running alongside the inference pipeline. No external NVMe needed.
- **Full industrial I/O on the card.** CAN bus, isolated RS-485, USB-C, USB 2.0, DisplayPort, GbE.
- **Industrial reliability.** −20 °C to +65 °C operating range. PCIe x4 drops into existing IPCs.

Let's talk.

evvr.io/4pe · sales@evvr.io