ROBOFLOWOS ON NVIDIA JETSON

RF-THOR128

Flagship edge AI appliance for high stream counts and on-device vision-language models



8+ camera streams | Edge VLM inference | 4× 25GbE

RF-THOR128 is Roboflow's most powerful edge appliance. Built on the NVIDIA Jetson T5000 (Thor) module with 128 GB of memory, it runs Roboflow OS to deploy and manage computer vision models and Workflows across many camera streams at once, with the headroom to run vision-language models and other large AI models directly on the production line. A QSFP28 port provides four 25GbE links for high-bandwidth camera aggregation.

System specifications

Compute module	**NVIDIA Jetson Thor T5000**
Storage	**1 TB NVMe SSD**
Networking	**1× 5GbE (RJ45)** **1× QSFP28 (4× 25GbE)**
Operating system	**Roboflow OS**
Runtime	**NVIDIA JetPack 7.0**
Recommended workload	**8+ concurrent camera streams, edge VLM inference**

Compute module

Per NVIDIA published module specifications

AI performance	**Up to 2070 TFLOPS (FP4)**
GPU	**NVIDIA Blackwell, 2560 CUDA cores, 96 5th-gen Tensor Cores**
CPU	**14-core Arm Neoverse-V3AE 64-bit**
Memory	**128 GB 256-bit LPDDR5X, 273 GB/s**
Module power	**40 W – 130 W**

Physical & environmental

Power input	**19–36 V DC, terminal block**
Operating temperature	**−10 to 35 °C (0.7 m/s airflow)**
Size (W × D × H)	**195 × 200 × 71.5 mm**
Weight	**2 kg**
Certifications	**CE, FCC, CCC, KCC, BSMI, UL**

Network interfaces

Port	Qty	Typical use
5GbE (RJ45)	**1**	Uplink to plant network
25GbE (QSFP28)	**4**	High-bandwidth cameras or switch uplink, via one QSFP28 port

Generative AI at the edge: RF-THOR128 has the memory and compute to run vision-language models on-device alongside detection and segmentation Workflows, without sending images to the cloud.