

Synapse Core-X9

Enterprise AI Acceleration Engine

High-performance PCIe Gen 5.0 inference engine for real-time edge intelligence and deep learning workloads.

STATUS: PRODUCTION

REV Q3-2026

PART: SYN-X9-64G

750 TOPS

INT8 TENSOR COMPUTE

64 GB

LPDDR5X (680 GB/s)

75 W

MAX ACTIVE TDP

< 1.2 ms

SUB-MILLISECOND LATENCY


PRODUCT OVERVIEW

The **Synapse Core-X9** is a fourth-generation domain-specific neural processing card designed for ultra-low latency inference in enterprise server and edge gateway environments. Built on a 4nm architecture, Core-X9 delivers unprecedented compute density for Vision Transformers, LLMs, and multi-modal neural networks.

KEY ARCHITECTURAL FEATURES

- ✓ **32-Core Matrix Engine:** Dedicated tensor cores optimized for transformer attention and dense matrix multiply.
- ✓ **Dynamic Precision Scaler:** Hardware support for FP16, BF16, INT8, and quantized INT4 execution.
- ✓ **Secure Vault Enclave:** On-chip AES-256 hardware root-of-trust, encrypted memory, and secure boot.
- ✓ **Zero-Copy Driver SDK:** Seamless integration with PyTorch, ONNX Runtime, and TensorRT pipelines.
- ✓ **UltraLink Interconnect:** Dual 100Gbps chip-to-chip links for multi-card cluster scaling without host bottlenecks.

HARDWARE BLOCK DIAGRAM



TECHNICAL SPECIFICATIONS

COMPUTE & PERFORMANCE

Peak Performance	750 INT8 TOPS / 375 BF16 TFLOPS
Tensor Execution	32 Matrix Pipeline Units
On-Chip Cache	128 MB High-Bandwidth SRAM
Precision Support	FP16, BF16, INT8, INT4, TF32

MEMORY & INTERFACE

Memory Capacity	64 GB Unified LPDDR5X
Memory Bandwidth	680 GB/s (256-bit bus)
Host Interface	PCIe 5.0 x16 (Up to 64 GB/s)
Chip-to-Chip	Dual UltraLink 100 Gbps ports

ELECTRICAL & THERMAL

Power Consumption	75 W Active Peak / 8 W Idle
Form Factor	Standard PCIe Single-Slot FHHL
Thermal Solution	Passive Heatsink / Active Fan Option
Operating Temp	-40°C to +85°C (Industrial)

SOFTWARE & COMPLIANCE

Software Stack	Synapse Neural Studio v4.2
Frameworks	PyTorch, TensorFlow, ONNX, C++
Certifications	CE, FCC Class A, RoHS, ISO 9001

TARGET WORKLOADS & APPLICATIONS

- 👁️ **Vision Analytics:** Real-time 4K multi-stream object tracking and video segmentation.
- 💬 **LLM Edge Serving:** Low-latency token generation for 7B–13B parameter models.
- 🔒 **Automotive & Robotics:** High-reliability sensor fusion for autonomous vehicle systems.

ORDERING INFORMATION

Part Number	Memory	Grade & Thermal
SYN-X9-64G-I	64 GB LPDDR5X	Industrial (-40°C to +85°C), Passive
SYN-X9-64G-C	64 GB LPDDR5X	Commercial (0°C to +70°C), Active Fan
SYN-X9-32G-C	32 GB LPDDR5X	Commercial (0°C to +70°C), Active Fan

SYNAPSE TECHNOLOGIES INC.

450 Innovation Parkway, Suite 800, Tech District

✉️ sales@synapse-tech.example 🌐 www.synapse-tech.example

📞 +1 (800) 555-0199 © 2026 Synapse Technologies. All rights reserved.