


Open AI Factory

Own Your **Compute**. Own the Future of AI.




Production ready, validated Open AI Infrastructure for high-bandwidth, high-availability deployments.

Micas Networks and Broadcom’s validated open AI infrastructure solution integrates compute, networking, cooling, and optics into a deployable architecture. Customers can now achieve performance, availability, power efficiency, interoperability, and ownership of their AI Factory from the start.



Front View



Rear View

PRODUCT SPECIFICATIONS

Switch Model	M2-W6950-64X3-DR8-CP10-DC
Ports	64 x 1.6T MMC DR8 Optical
Console Port	2 x RJ45 Console Port
MGMT Port	2 x rj45 1000 BASE-T to CPU and BMC 1 x QSFP29 100G in-band management
Switching Capacity	102.4T
CPU	AMD Fire range 16 Core
System Memory	DDR5 (32GB, up to 64GB)
System Storage	1TB (Pluggable)
BMC Option	AST2520
Switch ASIC	BCM78919 (TH6 CPO)
Power Supplies	48V DC Input
Fans	N/A
OS	SONiC
Airflow Options	Front-to-rear
Max/Typical Power	Est. 3415W/ Est. 2557W (25°C@0m, 600B line-rate)
Consumption Dimensions (W x D x H)	Est. 3415W/ Est. 2557W (25°C@0m, 600B line-rate)
Weight (with all modules)	TBD

PERFORMANCE vs pluggable
≥40%


51.2T Co-Packaged Optics in a 4RU, 128 × 400G FR4 system.

SCALE & POWER SAVINGS
38.28% switch power reduction

Validated with Supermicro + AMD + Micas at production volume.

OPEN FABRIC
>1M port-hours · zero link flaps

A fabric operators own end to end. No proprietary lock-in.



AMD Instinct™ MI355X Platform

Optimized for heavy LLM scale, featuring the high-capacity HBM3E memory hierarchy and full ROCm software ecosystem.


8-GPU Architecture

288 GB
HBM3E / GPU

2.3 TB
Total Memory

8 TB/s
Bandwidth / GPU

ROCm
Open Ecosystem



AMD Pensando™ Pollara 400 AI NIC

Optimized for heavy LLM scale, featuring the high-capacity HBM3E memory hierarchy and full ROCm software ecosystem.

UEC-Ready NIC

400 Gbps
HBM3E / GPU

Intelligent
Load Balancing

Active
Congestion Mgmt