

Alcatel-Lucent OmniCompute 8100

AI Server

The Alcatel-Lucent OmniCompute 8100 (OC8100) AI Server is a high-performance, scalable, GPU-based server ideally suitable for AI/ML (Artificial Intelligence/Machine Learning) and HPC (High-Performance Computing) applications. The server is ideal for training large language models, automation, object classification, and recognition cases. The system is designed with eight AMD Instinct MI325X GPUs and dual AMD EPYC 9005/Turin series processors. The AMD MI325X GPU integrates 256 GB of HBM3 memory each.

Each MI325X offers a 16-lane PCIe Gen5 host interface and seven AMD Infinity Fabric™ links for full connectivity between eight GPUs. AMD Infinity Fabric technology delivers excellent I/O efficiency, scaling and communication within and between industry-standard accelerator module (OAM) device packages.

Key Features and Benefits

- Featuring eight AMD Instinct™ MI325X GPUs based on CDNA™ 3 architecture
- Dual 5th Gen AMD EYPC processors with advanced data center performance and efficiency
- 7 Infinity Fabric™ links with 128 GB/s for full connectivity between 8 GPUs in a ring (scale-up)
- 16 PCIe Gen5 with 128GB/s bandwidth for scale-out network
- Up to 1.3X AI performance and up to 2.4X HPC performance vs. competitive accelerators
- AMD ROCm™ software optimized for generative AI and HPC applications, easy for migrating existing code

Product Picture


O8100 AI Server front: final product may look slightly different.



O8100 AI Server rear-angled: final product may look slightly different.



O8100 AI Server rear-top: final product may look slightly different.



Product specifications

Feature	Description
Form Factor	<ul style="list-style-type: none"> • 8 RU
Compute Node	<ul style="list-style-type: none"> • CPU: AMD EPYC™ 9005/Turin series processor, 2 sockets EPYC 9575F 64C 3.3GHz 400W SP5 • Memory: 24 x 64GB RDIMM 2 RANK 5600 • Operating System: Linux • BIOS: AMI BIOS (32 MB SPI Flash ROM) • Front End Interface: 2 x BCM957608-P2200GQF00 Dual-Port 200GbE QSFP112 PCIe Ethernet NIC • Storage: 2 x 1.92TB U.2 2.5 NVME4 1DWPD (non SED)
GPU	<ul style="list-style-type: none"> • AMD MI325X 8 GPUs with AMD Infinity Fabric
Input/Output	<ul style="list-style-type: none"> • USB: 2 x USB3.0 ports • VGA: 1 x D-Sub 15-pin port • RJ-45: 1 x GbE port dedicated for IPMI • Buttons: UID button, power button, reset button • LEDs: <ul style="list-style-type: none"> • UID led • PWR led • Baseboard management controller • Indicator (over temperature warning indicator/fan and PSU fail)
Scale-Out Interface	<ul style="list-style-type: none"> • Remote Direct Memory Access (RDMA), RoCEv2 • 8 x LP slots for switch board (PCIe x16) • 8 x BCM957608-P1400GDF00 Single-Port 400G QSFP112-DD PCIe Ethernet NIC
Storage	<ul style="list-style-type: none"> • Front: <ul style="list-style-type: none"> • 6 x 7.68TB U.2 2.5 NVME4 1DWPD (non SED)
BMC	<ul style="list-style-type: none"> • AST2600
Power-Supply (PSU)	<ul style="list-style-type: none"> • 3300 W hot-swap RPSU, 80+ Titanium (4+2 redundancy)
Fans	<ul style="list-style-type: none"> • 15-hot-swappable fans
Dimensions	<ul style="list-style-type: none"> • W x D x H: 448 x 850 x 351 mm (17.63 x 33.46 x 13.82)
Operating Temperature	<ul style="list-style-type: none"> • 10°C – 35°C
Expansion Slots	<ul style="list-style-type: none"> • 8 half-heights + 4 full-height PCIe 5.0 x16 slots
Software	<ul style="list-style-type: none"> • ROCm: 6.2.4 • RCCL: 2.20.5 • PyTorch: 2.3, 2.2, 2.1, 2.0, 1.13 • TensorFlow: 2.16.1, 2.15.1, 2.14.1 • JAX: 0.4.26 • ONNX Runtime: 1.17.3