OP-M

The OP-M is cooled by water or by air, in the same chassis.

The OP-M runs Threadripper on TRX50. The water-cooled build takes up to four RTX PRO 6000 Max-Q GPUs on one loop, and an air-cooled build is also offered. It is for teams that have outgrown one or two GPUs and still want the computer in the room with them.

SPECIFICATIONS

GPUs	Water-cooled: Up to 4 × NVIDIA RTX PRO 6000 Blackwell Max-Q · Air-cooled: Up to 3 × NVIDIA RTX PRO 6000 Blackwell Max-Q (or RTX PRO 5000, 4500); or up to 2 × NVIDIA GeForce RTX 5090 (Founders Edition only); 5080, 5070 (Founders Edition or PNY, whichever is in stock)
GPU maker	Water-cooled: PNY, by default · Air-cooled: PNY or NVIDIA Founders Edition, by default
Processor	AMD Threadripper 9960X, Threadripper 9970X or Threadripper 9980X
Platform	AMD Threadripper (TRX50)
Other processors	Intel Xeon builds use a different board and are on request. Ask through https://optimistpc.com/quote?system=op-m.
Motherboard	ASUS Pro WS TRX50-SAGE WIFI A
Memory	64 GB to 512 GB of DDR5 ECC registered (4 × 16 GB to 4 × 128 GB); on the board: four channels, four DDR5 ECC registered DIMM slots (ASUS’s specifications)
Storage	Up to three NVMe drives, 500 GB to 8 TB each; on the board: four M.2 (PCIe 5.0 x4 or 4.0 x4 by slot); two SlimSAS; four SATA at 6 Gbps; three M.2 free with every GPU slot in use
Cooling	Water-cooled (the processor and every GPU) or air-cooled
Handles	Two, removable
Slots	Four full-length: three PCIe 5.0 x16 and one PCIe 4.0 x16 (ASUS’s specifications)
Network	10 Gb Ethernet (Marvell AQtion); 2.5 Gb Ethernet (Intel I226-LM)

Wireless	Wi-Fi 7, 2×2, on 2.4, 5 and 6 GHz; Bluetooth 5.4
USB	Two USB4 at 40 Gbps (USB-C); six USB-A at 10 Gbps; one USB 2.0
Display inputs	Two DisplayPort in (for USB4 display out)
Audio	Two jacks
Buttons	BIOS FlashBack, Clear CMOS
Remote management	No management chip on the board; a header for an IPMI expansion card (ASUS's specifications)
Outlet	Water-cooled: Up to two GPUs run from a standard 15 A outlet, three need a 20 A circuit on 100–120 V or any 200–240 V outlet, and four need a 200–240 V circuit. By rated power, not measured; method: https://optimistpc.com/how-we-measure#outlet · Air-cooled: With RTX PRO 6000 Max-Q or RTX PRO 5000, up to two GPUs run from a standard 15 A outlet and three need a 20 A circuit on 100–120 V or any 200–240 V outlet; with RTX PRO 4500, up to three GPUs run from a standard 15 A outlet; with GeForce RTX 5090, one GPU runs from a standard 15 A outlet and two need a 200–240 V circuit; with GeForce RTX 5080 or GeForce RTX 5070, up to two GPUs run from a standard 15 A outlet. By rated power, not measured; method: https://optimistpc.com/how-we-measure#outlet
Power inlet	C14, with a C13 cord in the box
Cooling (water-cooled)	One loop for the processor and every GPU, two 280 mm radiators, four 140 mm fans
Cooling (air-cooled)	A tower cooler on the processor, four 140 mm fans and one 120 mm fan

ORDERING

Status	Available to order
Warranty	3-Year Limited Warranty; terms: https://optimistpc.com/terms#warranty
Ships	Water-cooled: 20 days from order · Air-cooled: 6 days from order
Ships to	The United States

BY INDUSTRY

The industry pages that recommend this system (optimistpc.com/industries), as each page ranks it. A figure comes from a datasheet and names its source.

Training simulation	Pick 2 of 3, for a simulator with more screens than one GPU drives, with an RTX PRO 4500. Screens it drives: Four outputs on each RTX PRO 4500; add GPUs as the simulator grows (from NVIDIA’s datasheets)
Flight simulation	Pick 2 of 3, for several channels from one system, with an RTX PRO 4500
Racing simulation	Pick 2 of 2, for a team simulator that will add projector channels, with an RTX PRO 4500
Live visuals	Pick 2 of 2, for a fixed LED wall or dome, with an RTX PRO 4500
Robotics and autonomy	Pick 2 of 3, for a team training several policies at once, with an RTX PRO 4500
Private AI	Pick 2 of 3, for a department’s larger model
AI development	Pick 2 of 3, for a team running several jobs at once, with an RTX PRO 6000 Max-Q
Research	Pick 2 of 3, for a lab sharing one system, with an RTX PRO 6000 Max-Q
Universities and teaching	Pick 2 of 3, for a whole class sharing one system, with an RTX PRO 6000 Max-Q
Field research	Pick 2 of 2, for the lab back home, with an RTX 5080
Drug discovery	Pick 2 of 3, for a group running several simulations, with an RTX PRO 6000 Max-Q
Engineering simulation	Pick 2 of 3, for meshes larger than one GPU holds, with an RTX PRO 6000 Max-Q
Life sciences	Pick 2 of 3, for a lab’s genomics and cryo-EM, with an RTX PRO 6000 Max-Q
Mapping and reality capture	Pick 2 of 2, for projects over 30,000 images (from Esri’s memory table), with an RTX 5090

Rendering and VFX	Pick 2 of 2, for a studio that renders all day, with an RTX 5090. Largest scene in GPU memory: 32 GB on each RTX 5090; 96 GB on each RTX PRO 6000 Max-Q, up to four. Each GPU holds the whole scene, and more GPUs render it faster (from NVIDIA’s datasheets)
Office workstations	Pick 2 of 2, for the few engineers past 128 GB, air-cooled, with an RTX PRO 4500
Video and post-production	Pick 2 of 2, for a second GPU, more drives and faster Ethernet, air-cooled, with an RTX 5090

How we measure: optimistpc.com/how-we-measure