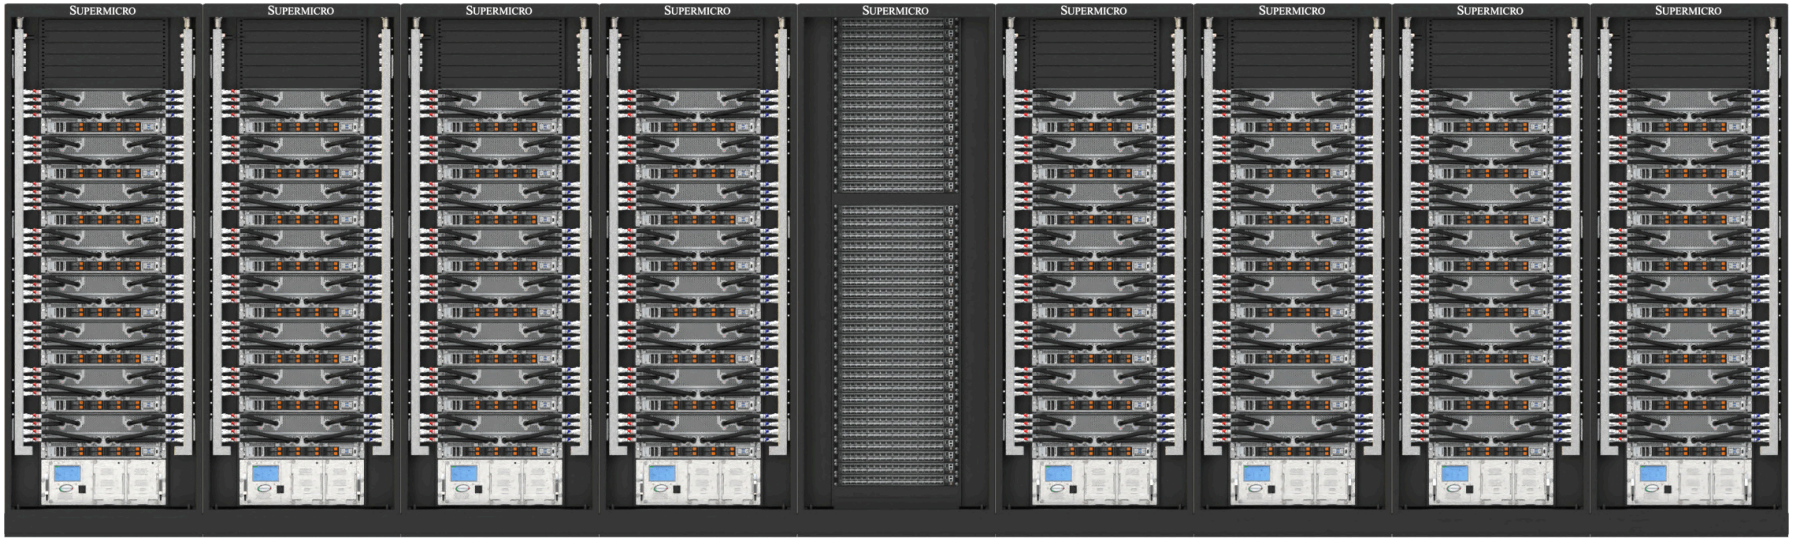
NVIDIA GPU systems · October 2026

Available to quote.

Fifteen servers, racks and desk-side stations from Supermicro, Gigabyte, HPE, Dell, MSI and ASUS. New, with manufacturer warranty, and priced per project.



A scalable unit of HGX Blackwell servers around a central network rack. Image: Supermicro.

At a glance

Fifteen systems. One process.

Reference	System	GPUs	GPU memory	Cooling	Availability	Lead time	Page
HGX B300 server nodes							
GPU-01	**SYS-822GS-NB3RT** Supermicro	8 × B300	288 GB HBM3e	Air	High volume	2 to 4 weeks after KYC and escrow	03
GPU-02	**SYS-422GS-NB3RT-LCC** Supermicro	8 × B300	288 GB HBM3e	Liquid	Strong	Confirmed at order	04
GPU-03	**G894-SD3-AAX7** Gigabyte	8 × B300	288 GB HBM3e	Air	High volume	2 to 4 weeks after KYC and escrow	04
GPU-04	**Compute XD690 (S6G97A)** HPE	8 × B300	288 GB HBM3e	Air	Available	Confirmed at order	05
GPU-05	**PowerEdge XE9780 or equivalent** Dell or Tier-1	8 × B300	HBM3e, per OEM	Air or liquid	On request	Confirmed at order	05
HGX B200 and H200 nodes							
GPU-06	**HGX B200 node** Tier-1 OEM	8 × B200	180 GB HBM3e	Air	Available	2 to 4 weeks after KYC and escrow	06
GPU-07	**HGX H200 node** Tier-1 OEM	8 × H200	141 GB HBM3e	Air	Very limited	4 to 6 weeks after KYC and escrow	06
Rack-scale NVL72							
GPU-08	**GB300 NVL72 rack** Tier-1 OEM	72 × B300	288 GB HBM3e	Liquid	High volume	Confirmed at allocation	07
GPU-09	**SRS-GB200-NVL72-M1** Supermicro	72 × B200	186 GB HBM3e	Liquid	Limited	Confirmed at allocation	07
Vera Rubin. The next generation							
GPU-10	**Vera Rubin NVL72 rack** Tier-1 OEM	72 × Rubin	288 GB HBM4	Liquid	Limited	From December 2026	08
GPU-11	**HGX Rubin NVL8 node** To be confirmed	8 × Rubin	HBM4	Air or liquid	Limited	From December 2026	08
Inference and enterprise servers							
GPU-12	**AS-5126GS-TNRT** Supermicro	8 × RTX PRO 6000	96 GB GDDR7	Air	Available	2 to 4 weeks after KYC and escrow	09
A GB300 on your desk							
GPU-13	**XpertStation WS300** MSI	1 × GB300	748 GB coherent	Desk-side	On request	To be confirmed	10
GPU-14	**Super AI Station** Supermicro	1 × GB300	748 GB coherent	Liquid, closed loop	On request	To be confirmed	10
GPU-15	**ExpertCenter Pro ET900N G3** ASUS	1 × GB300	748 GB coherent	Desk-side	On request	To be confirmed	10

Lead times run from the date KYC clears and funds are in escrow, unless shown otherwise. Nothing is reserved by this list.

Blackwell Ultra

HGX B300 server nodes.

Eight Blackwell Ultra GPUs per node, joined by NVLink. Built by Supermicro, Gigabyte, HPE and Dell, air-cooled or liquid-cooled, with high volume available now.



Image: Supermicro.

High volume available

Supermicro NVIDIA HGX B300

SYS-822GS-NB3RT

8U air-cooled server with eight Blackwell Ultra GPUs and eight 800 Gb/s SuperNICs.

GPUs	8 × NVIDIA B300 SXM
GPU memory	288 GB HBM3e per GPU
CPU	2 × Intel Xeon 6 (e.g. 6768P, 64 cores)
System memory	DDR5-6400, 2 to 3 TB
Networking	8 × 800 Gb/s ConnectX-8 SuperNIC, BlueField-3 DPU
Cooling	Air
Form factor	8U

Lead time
2 to 4 weeks after KYC clears and funds are in escrow

Quote reference
GPU-01

Blackwell Ultra

HGX B300 server nodes.



Image: Supermicro.

● Strong; liquid-cooled allocation confirmed at order

Supermicro NVIDIA HGX B300

SYS-422GS-NB3RT-LCC

4U direct liquid-cooled server: the same eight GPUs in half the height.

GPUs	8 × NVIDIA B300 SXM6
GPU memory	288 GB HBM3e per GPU
CPU	2 × Intel Xeon 6788P (86 cores)
System memory	DDR5-6400, 2 TB
Networking	8 × 800 Gb/s ConnectX-8 SuperNIC, BlueField-3 DPU
Cooling	Direct liquid cooling
Form factor	4U

Lead time
Confirmed at order

Quote reference
GPU-02



Image: Gigabyte.

● High volume available

Gigabyte NVIDIA HGX B300

G894-SD3-AAX7

8U air-cooled server with the eight 800 Gb/s SuperNICs built onto the board.

GPUs	8 × NVIDIA B300 SXM
GPU memory	288 GB HBM3e per GPU
CPU	2 × Intel Xeon 6776P (64 cores)
System memory	DDR5-6400, 2 to 4 TB
Networking	Onboard ConnectX-8, 8 × 800 Gb/s OSFP; BlueField-3 DPU
Cooling	Air
Form factor	8U

Lead time
2 to 4 weeks after KYC clears and funds are in escrow

Quote reference
GPU-03

Blackwell Ultra

HGX B300 server nodes.



Shown: Dell PowerEdge XE9780. Image: Dell Technologies.

HPE, Dell and other Tier-1 manufacturers NVIDIA HGX B300

Tier-1 HGX B300 nodes.

The same eight-GPU platform from HPE, Dell or another Tier-1 manufacturer, built to your specification.

● Available

HPE NVIDIA HGX B300

HPE Compute XD690

S6G97A

GPUs	8 × NVIDIA B300
GPU memory	288 GB HBM3e per GPU
CPU	2 × Intel Xeon 6776P (64 cores)
System memory	2 TB DDR5 (32 × 64 GB), up to 3 TB
Networking	8 × ConnectX-8 on the HGX board; BlueField-3 DPU; fabric quoted per design
Cooling	Air
Form factor	10U

Lead time
Confirmed at order

Quote reference
GPU-04

● On request

Dell or another Tier-1 NVIDIA HGX B300

Dell PowerEdge XE9780 or equivalent

GPUs	8 × NVIDIA B300
GPU memory	HBM3e, per manufacturer
CPU	Intel Xeon 6 or AMD EPYC
System memory	Per buyer specification
Networking	ConnectX-8, BlueField-3 DPU
Cooling	Air or liquid
Form factor	Per manufacturer

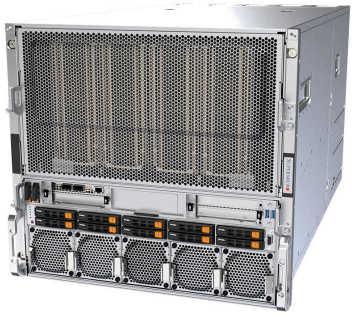
Lead time
Confirmed at order

Quote reference
GPU-05

Blackwell and Hopper

HGX B200 and H200 nodes.

Eight-GPU nodes on the previous two NVIDIA generations, built by the Tier-1 manufacturer you prefer.



Shown: Supermicro SYS-A22GA-NBRT. Image: Supermicro.

● Available; lower volume than B300

Tier-1 OEM NVIDIA HGX B200

HGX B200 node

Eight Blackwell GPUs in the Tier-1 system of your choice: Dell, HPE, Supermicro or Gigabyte.

GPUs	8 × NVIDIA B200 SXM, NVLink 5
GPU memory	180 GB HBM3e per GPU
CPU	2 × Intel Xeon 6 or AMD EPYC
System memory	DDR5, 2 TB typical
Networking	8 × 400 Gb/s NICs, BlueField-3 DPU
Cooling	Air
Form factor	Per manufacturer

Lead time
2 to 4 weeks after KYC clears and funds are in escrow

Quote reference
GPU-06



Shown: Supermicro SYS-821GE-TNHR. Image: Supermicro.

● Very limited; Hopper stock is scarce

Tier-1 OEM NVIDIA HGX H200

HGX H200 node

Eight Hopper GPUs for teams standardised on H200. Stock is scarce.

GPUs	8 × NVIDIA H200 SXM, NVLink
GPU memory	141 GB HBM3e per GPU
CPU	2 × Intel Xeon or AMD EPYC
System memory	Per buyer specification
Networking	8 × 400 Gb/s NICs
Cooling	Air
Form factor	Per manufacturer

Lead time
4 to 6 weeks after KYC clears and funds are in escrow

Quote reference
GPU-07

Grace Blackwell

Rack-scale NVL72.

Seventy-two GPUs and thirty-six Grace CPUs in one liquid-cooled rack that works as a single NVLink domain. Delivered factory-integrated.



Shown: Supermicro SRS-GB300-NVL72. Image: Supermicro.

● High volume available Recommended

Tier-1 OEM, factory-integrated NVIDIA GB300 NVL72

GB300 NVL72

Seventy-two Blackwell Ultra GPUs and thirty-six Grace CPUs, liquid-cooled and joined in one NVLink domain.

GPUs	72 × NVIDIA Blackwell Ultra
GPU memory	288 GB HBM3e per GPU
CPU	36 × NVIDIA Grace (Arm)
System memory	LPDDR5X (Grace)
Networking	NVLink 5 rack domain; ConnectX-8 per compute tray
Cooling	Direct liquid, about 130 to 140 kW per rack
Form factor	Full rack

Lead time
Confirmed at allocation

Quote reference
GPU-08



Image: Supermicro.

● Limited; we recommend GB300

Supermicro NVIDIA GB200 NVL72

SRS-GB200-NVL72-M1

Seventy-two Blackwell GPUs in eighteen 1U compute trays, with its own 250 kW coolant distribution unit.

GPUs	72 × NVIDIA B200
GPU memory	186 GB HBM3e per GPU (13.4 TB per rack)
CPU	36 × NVIDIA Grace (72 cores)
System memory	LPDDR5X (Grace)
Networking	NVLink 5; ConnectX-7 or BlueField-3 per compute tray
Cooling	Direct liquid; in-rack 250 kW CDU
Form factor	Full rack, 18 × 1U ARS-121GL-NBO trays

Lead time
Confirmed at allocation

Quote reference
GPU-09

Vera Rubin

Vera Rubin. The next generation.

NVIDIA's next platform, with HBM4 memory, sixth-generation NVLink and 1.6 Tb/s ConnectX-9 networking. Allocation is limited, and first shipments are expected from December 2026 at the earliest.



A Vera Rubin NVL72 scalable unit of 1,152 GPUs. Image: Supermicro.



Shown: Supermicro SRS-VR-NVL72. Image: Supermicro.

● Limited allocation

Tier-1 OEM, factory-integrated NVIDIA Vera Rubin NVL72

Vera Rubin NVL72

Seventy-two Rubin GPUs and thirty-six Vera CPUs in one liquid-cooled rack.

GPUs	72 × NVIDIA Rubin
GPU memory	288 GB HBM4 per GPU
CPU	36 × NVIDIA Vera
System memory	Per manufacturer
Networking	NVLink 6; ConnectX-9 SuperNIC, 1.6 Tb/s
Cooling	Direct liquid
Form factor	Full rack

Lead time
December 2026 at the earliest; January 2027 is more realistic

Quote reference
GPU-10



Shown: Supermicro 2U liquid-cooled HGX Rubin NVL8 system. Image: Supermicro.

● Limited allocation

Manufacturer to be confirmed NVIDIA HGX Rubin NVL8

HGX Rubin NVL8 node

Eight Rubin GPUs on one NVLink 6 board, for halls built around eight-GPU nodes.

GPUs	8 × NVIDIA Rubin, NVLink 6
GPU memory	HBM4
CPU	Per manufacturer
System memory	Per manufacturer
Networking	Per manufacturer
Cooling	Air or liquid
Form factor	Per manufacturer

Lead time
December 2026 at the earliest; January 2027 is more realistic

Quote reference
GPU-11

RTX PRO Blackwell

Inference and enterprise servers.

PCIe GPU servers for inference, visual computing and simulation in a standard air-cooled data hall.

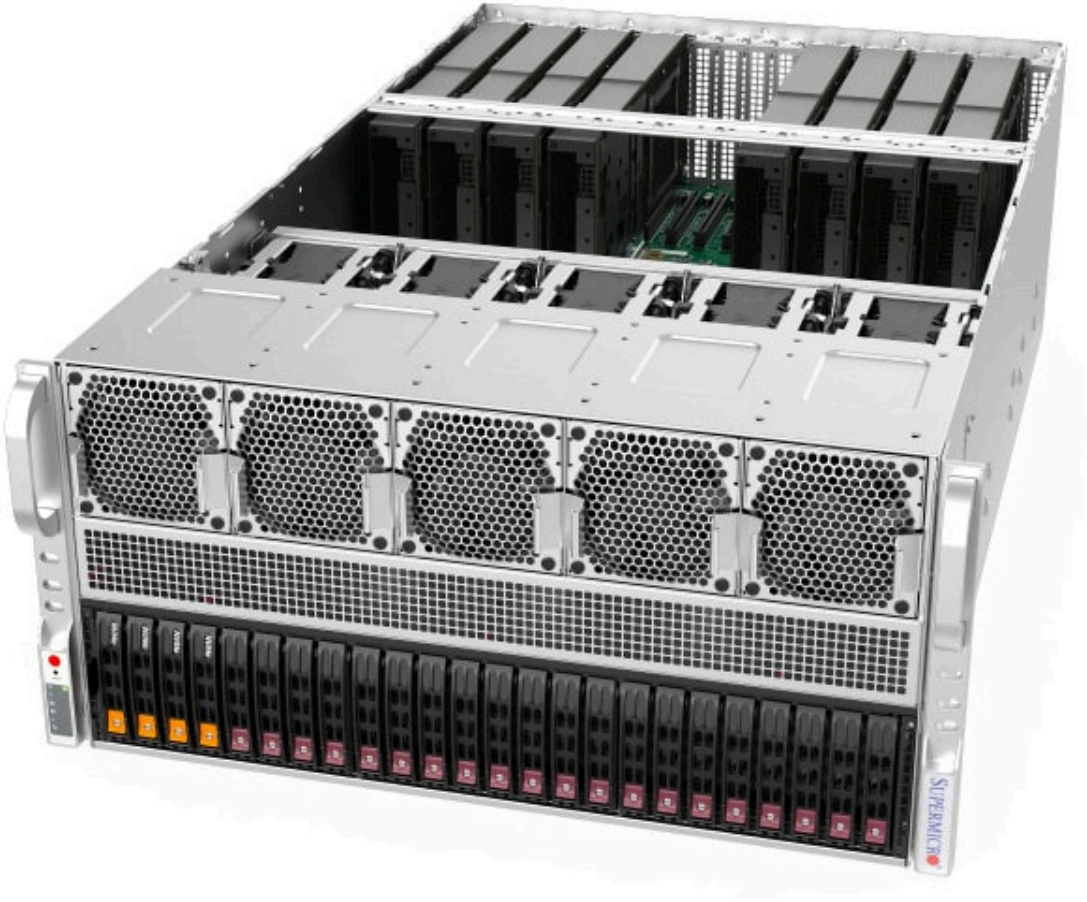


Image: Supermicro.

Available

Supermicro, or an equivalent NVIDIA RTX PRO 6000 Blackwell Server Edition

AS-5126GS-TNRT

5U air-cooled server with eight RTX PRO 6000 Blackwell Server Edition GPUs.

GPUs	8 × RTX PRO 6000 Blackwell Server Edition (PCIe 5.0)
GPU memory	96 GB GDDR7 per GPU
CPU	2 × AMD EPYC 9555 (64 cores)
System memory	768 GB DDR5-6400
Networking	ConnectX-6 Dx 100 GbE; more NICs optional
Cooling	Air
Form factor	5U

Lead time
2 to 4 weeks after KYC clears and funds are in escrow

Quote reference
GPU-12

DGX Station architecture

A GB300 on your desk.

Desk-side AI systems built on NVIDIA DGX Station architecture: one GB300 Grace Blackwell Ultra Desktop Superchip with 748 GB of coherent memory. Offered on request while availability is confirmed.



Image: MSI.

● On request; availability being confirmed

MSI NVIDIA GB300 desk-side

XpertStation WS300

A desk-side AI supercomputer on NVIDIA DGX Station architecture.

Superchip	1 × NVIDIA GB300 Grace Blackwell Ultra
Memory	748 GB coherent: 252 GB HBM3e and 496 GB LPDDR5X
CPU	NVIDIA Grace
Networking	2 × 400 GbE ConnectX-8 SuperNIC
Power	1,600 W
Form factor	Desk-side tower

Lead time
To be confirmed

Quote reference
GPU-13



Image: Supermicro.

● On request; availability being confirmed

Supermicro NVIDIA GB300 desk-side

Super AI Station

ARS-511GD-NB-LCC

Liquid-cooled inside, in a tower that can also be rack-mounted as a 5U server.

Superchip	1 × NVIDIA GB300 Grace Blackwell Ultra
Memory	748 GB coherent: 252 GB HBM3e and 496 GB LPDDR5X
CPU	NVIDIA Grace, 72 cores
Networking	2 × 400 GbE ConnectX-8 SuperNIC
Cooling	Closed-loop direct liquid
Form factor	Tower or 5U rack

Lead time
To be confirmed

Quote reference
GPU-14



Image: ASUS.

● On request; availability being confirmed

ASUS NVIDIA GB300 desk-side

ExpertCenter Pro ET900N G3

ASUS's DGX Station-class workstation, with room for an RTX PRO graphics card.

Superchip	1 × NVIDIA GB300 Grace Blackwell Ultra
Memory	748 GB coherent: 252 GB HBM3e and 496 GB LPDDR5X
CPU	NVIDIA Grace, 72 cores
Networking	2 × 400 Gb/s ConnectX-8 (QSFP112)
Power	1,600 W
Form factor	Desk-side tower

Lead time
To be confirmed

Quote reference
GPU-15

Cluster networking and data-centre kit

Everything around the GPUs.

Switches, optics, cabling and cooling, supplied with the servers once the cluster design and site are confirmed. Quoted per cluster design, not as standalone items.

Item	Model	Description	Used for
InfiniBand switch	**NVIDIA Quantum-X800 Q3400**	144 ports of 800 Gb/s XDR	GPU compute fabric
Ethernet switch	**NVIDIA Spectrum-X SN5610**	64 ports of 800 GbE	Front-end and storage network
Ethernet switch	**NVIDIA Spectrum-6 SN6600-LD**	128 ports of 800 Gb/s, liquid-cooled	Next-generation fabric and spine
Management switch	**NVIDIA SN2201**	1 GbE management	Out-of-band management
Optics	**800 Gb/s and 1.6 Tb/s OSFP transceivers**	Compatible modules; NVIDIA LinkX on request	GPU and switch links
Cabling	**MPO-12 APC single-mode fibre**	Cut to the lengths in the design	Fabric cabling
Cooling	**In-row CDU, about 1.35 MW, N+1**	Serves liquid-cooled racks	NVL72 and direct liquid-cooled halls

How to buy

Four details. Five steps.

Manufacturers release pricing and allocation only against a named project. These four details start a quotation.

1

End customer

The legal entity name and its company registration number.

2

Installation site

The data centre, delivery or colocation address and country.

3

Indicative quantity

The number of nodes or racks, and the models from this list.

4

Target delivery date

When the systems need to be on site.

Then, from agreement to order.

1

Signed NCNDA

A non-circumvention and non-disclosure agreement opens the file.

2

Indicative quote

Price and lead time for your models and quantity, for planning.

3

KYC and end-use certificate

The manufacturer's checks on the end customer, site and use.

4

Firm price and allocation

Confirmed against the manufacturer's allocation on the day KYC clears.

5

Order and escrow

You order direct with the manufacturer or supplier. Funds are held in escrow.

Our role

World One acts as introducing and facilitating broker. The supply contract is between the buyer and the manufacturer or supplier.

Terms of this list

Clear terms, from the start.

Configuration

The builds shown are the standard configurations we have priced. CPU, memory, storage and network options can be changed to your specification.

Availability

Availability and lead times are confirmed against the manufacturer's allocation on the date KYC clears. Nothing is reserved by this list.

Delivery

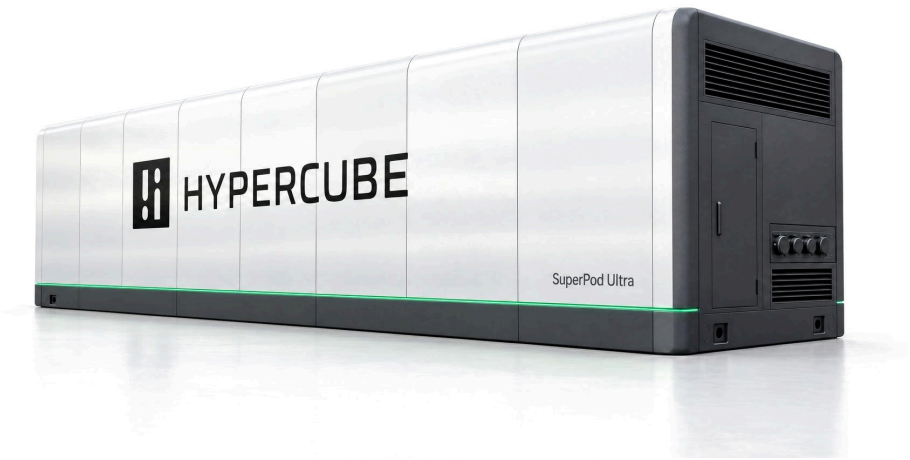
Ex works, United States, unless agreed otherwise; delivered at place (DAP) to your site on request. Freight, duties, taxes, racks, power and installation are quoted separately.

Export control

Systems ship to the named end user and installation site only. No re-export or diversion.

Warranty and pricing

New equipment with the manufacturer's warranty. No prices are published; pricing is issued per project.



HYPERCUBE SuperPod Ultra. Styling concept. Configuration not finalised.

No data hall? Take them in a module.

If you have power but no suitable building, the same GPUs can arrive inside a HYPERCUBE module: racks installed, liquid cooling connected and the system tested in the factory before it ships.

Request a quotation

Send the four details and the quote references of the systems you need to **info@worldoneindia.net**.

World One Enterprises Pvt. Ltd. · NESCO IT Park, NH 8, Goregaon East, Mumbai 400063, Maharashtra, India · worldoneindia.net

In partnership with HYPERCUBE QUANTUM AI