

Vertiv™ Open Rack for NVIDIA MGX v1.1

Powering AI: High Density IT Infrastructures



Highlights


Purposely built for NVIDIA GB200 and GB300 NVL72 SuperPods, Vertiv™ Open Rack for NVIDIA MGX accelerate the deployment of AI factories and high-performance computing (HPC) workloads.

Benefits

- NVIDIA-approved design: Enables compatibility, standardization, and future-ready integration with NVIDIA AI deployments.
- High-density power support: Can incorporate scalable and efficient 19" DC Power Shelves to reliably deliver power for GPU-intensive clusters.
- Mid-density power support: Can incorporate vertical or horizontal rPDUs to deliver power for storage or networking assets.
- Standards-aware engineering: Designed for standard rack unit (44.45 mm) unit height, to enable seamless integration within 19" native AI servers and network devices such as the NVIDIA Quantum InfiniBand.
- Advanced thermal management: Supports cutting-edge cooling methods including in-rack manifolds to efficiently manage GPU heat loads.
- Scalability and reliability: Provides a strong infrastructure foundation that enables enterprises to scale AI deployments with confidence.




Vertiv™ Open Rack for NVIDIA MGX v1.1



Flexible configuration:

DC Power Shelves mounted to optimize space usage while maintaining balance and structural integrity.



AI-optimized design:

Purposedly built to handle the extreme power density of modern AI and machine learning workloads.

Product description

Vertiv™ Open Rack for NVIDIA MGX is designed to meet the unprecedented demands of modern artificial intelligence by delivering a comprehensive solution purpose-built for high-performance GPU clusters and large-scale AI deployments.

Engineered to support the extreme power densities required by AI and machine learning platforms, Vertiv™ Open Rack for NVIDIA MGX can incorporate a 1400A DC power busbar, up to eight (4+4) scalable, highly efficient DC Power Shelves. Vertiv Open Rack for NVIDIA MGX are also available as prebuilt racks, suitable for different applications such as computing and supporting racks for NVIDIA SuperPods. This enables seamless integration while maintaining structural integrity and maximizing usable space.

As power density per rack rises, Vertiv Open Rack for NVIDIA MGX can also feature advanced thermal management technologies including direct-to-chip liquid cooling with Vertiv™ CoolChip Fluid Network manifolds with blind-mate quick disconnect, to dissipate the significant heat generated by GPUs.

Vertiv Open Rack for NVIDIA MGX is also available as Vertiv™ SmartIT Solutions, a configurable, high-capacity rack system developed to simplify integration for system builders and support a variety of server and equipment types, ready to be further integrated with IT components.



Specifications

Product Description

OCP Specs	MGX v1.1
Compliance Level	OCP-Inspired
Description	Open Rack V3 NVIDIA Vertiv Open Rack for NVIDIA MGX v1.1 for GB200 NVL72

Mechanical Specifications

Height (mm), on castors	2286 mm
Height (mm), without castors	2236 mm
Width (mm)	600 mm
Depth (mm)	1068 mm
IT Gear Width	19" (482 mm)
OCP Rack Units (OU)	44OU / 48RU
Net Weight (kg)	180 kg (397 lbs.)
Transport Weight (kg)	200 kg (440 lbs.)
Static Load Rating (kg/lbs.)	1,600 kg (3,527 lbs.)
Dynamic Load Rating (kg/lbs.)	1,400 kg (3,086 lbs.)
Frame Material	Steel
Finish	Powder Coat
Color	Black (RAL9005)
Castors & levelling feet	Included

Applications for Vertiv™ Open Rack for NVIDIA MGX

- AI Factories: Modular setup for training LLMs and Generative AI at scale.
- HPC & Research: Running complex scientific simulations and data analytics.
- Digital Twins: Real-time industrial modeling using NVIDIA Omniverse.
- Edge & Telco: Deploying 5G/6G networks and AI inference in edge data centers.

Vertiv™ Open Rack for NVIDIA MGX v1.2

Powering AI: High Density IT Infrastructures



Highlights


Designed for deployment of the Rubin GPUs, the latest generation of Vertiv Open Rack for NVIDIA MGX natively integrates of DC power and liquid cooling to support compute, network and storage configurations, while providing a scalable foundation for AI factories as well as enterprise AI and high- performance computing environments.

Benefits

- **Rubin-ready architecture:** supports the mechanical and liquid-cooling interfaces used by next-generation NVIDIA rack-scale systems.
- **High-density power delivery:** integrates a liquid-cooled DC busbar and scalable DC power shelves for GPU- intensive configurations.
- **Advanced liquid cooling:** natively supports dual manifolds, for direct-to-chip liquid heat removal.
- **Structural resilience:** reinforced frame, cross-stiffeners and rack stiffener trays support heavy payloads and insertion forces.
- **Faster deployment option:** pre-integrated solutions reduce on-site assembly, interfaces and vendor coordination.




Vertiv™ Open Rack for NVIDIA MGX v1.2



Flexible configuration:

DC Power Shelves mounted to optimize space usage while maintaining balance and structural integrity.



AI-optimized design:

Purposedly built to handle the extreme power density of modern AI and machine learning workloads.

Product description

This design is engineered for the mechanical, electrical and thermal demands of next-generation accelerated computing for agentic AI.

The new design maintains 19" form-factor while expanding structural support, rear service volume and fixed interfaces required for dense liquid-cooled payloads. The rack architecture can accommodate power shelves selection, up to four 110 kW to support a rear, liquid-cooled DC busbar rated up to 5000 A, while the rack remains AC-fed at the shelf input.

For thermal management, dual full-height manifolds can be arranged for top or bottom-feed service. Blind-mate universal quick-disconnect interfaces simplify tray insertion and service, while the rear extension creates controlled space for manifolds, cable cartridges, power distribution and leak-detection components.

Available as components or as pre-integrated SmartIT solutions, the platform can be configured for Vera Rubin NVL72 compute racks, Rubin NVL8 systems and high-density Spectrum-X or Quantum InfiniBand network support. This gives system builders a repeatable rack boundary for power, cooling, cabling and controls.



Specifications

Product Description

OCP Specs	MGX v1.2
Compliance Level	OCP-Inspired
Description	Open Rack V3 NVIDIA Vertiv Open Rack for NVIDIA MGX v1.2 for Vera Rubin designs

Mechanical Specifications

Height (mm), on castors	2285 mm
Height (mm), without castors	2236 mm
Width (mm)	600 mm
Depth (mm)	1320.6 mm
IT Gear Width	19" (482 mm)
OCP Rack Units (OU)	44OU / 48RU
Net Weight (kg)	270 kg
Transport Weight (kg)	1943 kg
Static Load Rating (kg/lbs.)	1843 kg
Dynamic Load Rating (kg/lbs.)	
Frame Material	Steel
Finish	Powder-coat
Color	Black (RAL9005)
Castors & levelling feet	Included

Applications for Vertiv™ Open Rack for NVIDIA MGX

- **Agentic AI development:**
Training and inference for large language, reasoning, multimodal and mixture-of-experts models.
- **Enterprise AI:**
Dense accelerated infrastructure for enterprise AI applications.
- **High-performance computing:**
Tightly coupled GPU compute for scientific simulation and research workloads.
- **Digital twins and physical AI:**
Industrial simulation, robotics development and real-time digital-twin workloads using accelerated computing platforms.