DELL Technologies

Specification Sheet

PowerEdge XE9712

The Future of High-Performance Dense Compute

PowerEdge XE9712

The Dell PowerEdge XE9712 is a next-generation, rack-scale server platform engineered to power the most demanding AI workloads at unprecedented scale. Built on NVIDIA's Grace Blackwell Ultra GB300 NVL72 architecture, the XE9712 unifies up to 72 Blackwell Ultra GPUs per rack using high-bandwidth NVLink interconnects, delivering 1.8 TB/s of GPU-to-GPU communication.

Ideal for super-scale AI clusters, the XE9712 eliminates traditional bottlenecks with a scale-up design that integrates Grace CPUs, direct liquid cooling, and Dell's ORv3 rack infrastructure. By removing power distribution units (PDUs), optimizing east-west traffic, and simplifying deployment with modular, standards-based components, the XE9712 accelerates AI innovation, improves operational efficiency, and offers a turnkey foundation for enterprises building AI factories or hyperscale superpods.

Key highlights

- Accelerated AI Innovation: Train larger models faster with unified GPU memory and ultra-high bandwidth.
- Faster Time-to-Value: Rapid deployment using modular, standards-based rack designs.
- Operational Efficiency: Save energy and space through direct liquid cooling and integrated power delivery.
- Future-Ready Infrastructure: Seamless scaling with multi-rack topologies and NVLink interconnect.
- Enterprise Manageability: Simplify operations using Dell's OpenManage Enterprise and iDRAC.

IRSS and IR9048

Dell's Integrated Rack Scalable Systems (IRSS) program is designed for rack-scale AI and HPC deployments, maximizing space, energy efficiency, reducing costs and are delivered as fully built and tested sets of racks. As an option within the IRSS program, The IR9048 series features a 19-inch rack infrastructure with dense compute and liquid cooling for high TDP GPUs and CPUs.Benefits

Feature	Benefit	What you can expect
72x Blackwell Ultra GPUs via NVLink (1.8 TB/s)	Creates a massive unified GPU pool enabling ultra-large model training	Reduces training time for foundation models and GenAI workloads
Grace Blackwell Ultra architecture (GB300)	Combines CPU-GPU memory coherence for faster data access and compute efficiency	Maximizes throughput across complex AI pipelines
Direct Liquid Cooling (DLC)	Improves thermal management and allows higher rack power density	Increases performance while lowering cooling and energy costs
ORv3 Busbar & 33kW Power Shelves	Removes traditional PDUs and supports dense power delivery	Simplifies deployment and optimizes power distribution at scale
1:1 NIC-to-GPU & BlueField-3 DPU integration	High-speed east-west traffic for multi-rack GPU clusters with smart networking	Enables efficient scale-out AI systems with storage and networking acceleration

Cyber Resilient Architecture for Zero Trust IT environment & operations

Security is integrated into every phase of the PowerEdge lifecycle, including protected supply chain and factory-to-site integrity assurance. The Silicon-based root of trust anchors end-to-end boot resilience while Multi-Factor Authentication (MFA) and role-based access controls safeguard trusted operations.

Sustainability

From recycled materials in our products and packaging, to thoughtful, innovative options for energy efficiency, the PowerEdge portfolio is designed to make, deliver, and recycle products to help reduce the carbon footprint and lower your operation costs. We even make it easy to retire legacy systems responsibly with Dell Technologies.

Rest easier with Dell Technologies Services

Maximize your PowerEdge Servers with comprehensive services designed to meet you wherever you are. Accelerate time to value in achieving high AI use cases with Professional Services for AI, choose from tailored deployment options with the ProDeploy Suite, receive proactive and predictive support with our ProSupport Suite, and so much more with our services available across 170 locations and backed by our 60K+ employees and partners.

Increase efficiency and accelerate operations with autonomous collaboration

The Dell OpenManage systems management portfolio tames the complexity of managing and securing IT infrastructure. Using Dell Technologies' intuitive end-to-end tools, IT can deliver a secure, integrated experience by reducing process and information silos in order to focus on growing the business. The Dell OpenManage portfolio is the key to your innovation engine, unlocking the tools and automation that help you scale, manage, and protect your technology environment.

Feature	Technical Specifications for IR9048, XE9712	
Processor	Two NVIDIA Grace with 72 cores per processor	
Memory	• 480 GB of LP DDR5 memory with error-correction code (ECC) / CPU • 288 GB HBM3e / GPU	
GPU	• 4 Blackwell Ultra GPU's • 900GB/s of coherent memory through NVLINK CPU-GPU interconnect	
Front drive bays	Up to 8 x EDSFF E1.S hot-swappable NVMe drives	
Storage controller	Direct NVMe using CX8 mezz PCIe	
Integrated Rack solution	IR9048 ORv3 (OCP Open Rack version 3)	
Rack and Sled Form Factor	IR9048 is a 48 RU (Rack Units) rack Each XE9712 sled is 1 RU Note: 1 RU = 44.45 mm (1.75 inches) height, and 482.6 mm (19 inches) width.	
Power Supplies	• IR9048 rack consists of Power Shelf 33kW (PS33) that supports 6 x 5500 W AC PSUs • PS33 supplies up to 54 VDC to the XE9712 through Power bus bar located on IR9048 rack	
Cooling specifications	• Air cooling • Direct Liquid Cooling (DLC)	
Fans	Eight standard dual rotor cooling fans	
Dimensions and Weight	IR9048 • Height — 2294 mm (90.3 inches) • Width — 750 mm (29.52 inches) • Depth — 1200 mm (47.24 inches) — 1340 mm (57.75 inches) with stablizers • Cabinet wet weight — 1590 kg (3505.35 pounds)	XE9712 • Height — 43.6 mm (1.72 inches) • Width — 438 mm (17.24 inches) • Depth — 836.8 mm (32.9 inches) • Weight — 30 kg (66.14 pounds)
Embedded Management	For XE9712 • Aspeed AST2600 BMC DC-SCI compatible • OpenBMC • BMC-to-NVIDIA HMC for HPM management	For XE9712 with iDRAC • iDRAC 10 • BMC-to-NVIDIA HMC for HPM management
Network options	• 1G integrated LOM • 4 x16 CX8 OSFP ports	
Ports	For XE9712 • 1 x USB 3.0 Type - A port • 1 x RJ45 (BF3 management port - optional) • 1 x Mini-DisplayPort • 1 x BIOS Serial out • 4 x OSFP • 1x 1Gb LOM • 1 x RJ45 dedicated BMC port	For XE9712 with iDRAC • 1 x USB 3.0 Type - A port • 1 x RJ45 (BF3 management port - optional) • 1x Mini-DisplayPort • 1 x USB 2.0 Type-C • 1x 1Gb LOM • 1 x RJ45 dedicated iDRAC port
PCIe slots and Risers	• Up to 2 x16 FH HL Gen5, front risers only • 2 x Mezzanine - Gen5/Gen6	
PCIe Cards	• 1 x BF3 SuperNIC Full Height Half Length	
Boot Drives	• 1 x M.2 Boot	
Rail Support	Static rails for ORv3 IR9048 rack	
Operating System and Hypervisors	• Ubuntu	For specifications and interoperability details, see Dell.com/OSsupport.

Technical Abbreviations

Octal Small Form-factor Pluggable (OSFP)

Quad Small Form-factor Pluggable (QSFP)

BF3 is a device called BlueField-3, a networking card from NVIDIA.

APEX Flex on Demand

Acquire the technology you need to support your changing business with payments that scale to match actual usage. For more information, visit https://www.delltechnologies.com/en-us/payment-solutions/flexible-consumption/flex-on-demand.htm.

Discover more about PowerEdge servers

Learn more
Dell Networking solutions

Contact a Dell Technologies Expert

View more resources

Follow on Dell.com

Follow on X

Follow on Linkedin

Join the conversation with #DellTechnologies

April 2026