

Boost AI performance and scale data cost-effectively

Optimize resources and performance with a powerful data server for midscale AI projects

PEAK:AIO

Created for the AI market, PEAK:AIO has been adopted by AI projects ranging from new-starts to world-leading healthcare solutions.

Simply plug-in at the beginning of a project and scale as and when needed.

Enables teams to focus on the project, the insights and innovation, not the storage.

Evolves the universally supported network filesystem (NFS) and modern NVMe-oF into a simple yet tremendously powerful ultra-low latency NVMe AI Data Server.

Leverages RDMA and NVIDIA's GPUDirect® to read and write directly to GPU memory, bypassing the CPU providing significant performance improvements over legacy or standard NAS and block storage.

Advanced, yet simple to use leading multipath technology provides over 40GB/s from a single Dell PowerEdge server.

Outperforms some of the largest enterprise and HPC class storage solutions, in a simple 2U cost-effective footprint.

Enough performance to drive several GPU servers at the cost of a single Dell PowerEdge server.

Proven, stable and relied upon by leading research institutions and worldwide organizations.

Dell Validated Designs for AI - built in collaboration with PEAK:AIO and NVIDIA, deliver an AI Data Server designed for mainstream AI projects, providing realistic capacity levels and ultra-fast performance, in a single 2U solution at a price that enables more investment in project and GPU resources.

The AI Data Server is a joint solution which includes PEAK:AIO and NVIDIA® GPUDirect® deployed on Dell PowerEdge servers to create a central shareable pool of ultra-low latency NVMe resource designed and tuned to speed AI performance while simplifying data shareability.

Streamline moving AI from concept to production

GPUs are shaping the future of artificial intelligence (AI). Whether you are creating medical breakthroughs, researching scientific problems, training neural networks or modelling financial markets, GPUs drive faster results, as long as the data throughput can keep up with them. However, today's AI systems consume and analyze data at much higher rates than many legacy storage solutions can deliver, resulting in low utilization of expensive GPU resources and dramatically extended training and project times.

As AI projects evolve, many find themselves in the gap between proof of concept and datacenter level scale, with many projects finding a viable solution with single or a limited number of GPU servers — however, the data performance still needs to keep pace, forcing a choice between feature-rich enterprise-class storage and parallel filesystems designed for HPC environments. Both options are complex and expensive, designed for much larger solutions and impact the budget for GPUs, draining resources away from innovation.

To break through the cost/performance gap, organizations need a cost-effective means to accelerate AI workloads while scaling and sharing data. Dell Technologies collaborated with PEAK:AIO to provide a scalable and powerful data server solution that pushes the boundaries of AI innovation and performance. The single 2U AI data server delivers the ultra-low latency and incredible bandwidth needed to optimize GPU utilization, while allowing organizations to start as small as 30TB and grow seamlessly up to 360TB per 2U as AI projects ramp. Along with performance and scalability, the solution reduces complexity so users and teams can focus on the project and the GPU functionality, not storage

“PEAK:AIO simply allowed KCL and the NHS trials to invest more funds in GPUs, where they provided for the best patient outcome.”

—Dr Jorge Cardoso is a reader in Artificial Medical Intelligence at King’s College London and CTO of the AICENTRE, the creators of FLIP/AIDE, the new generation of AI based healthcare.

[Learn more](#)

Dell Technologies and PEAK:AIO

Dell Technologies and PEAK:AIO are pioneers in AI-powered, scalable infrastructure solutions. Together, Dell Technologies and PEAK:AIO provide unique value, delivering integrated IT platforms and solutions that accelerate results and redefine what is possible. PEAK:AIO’s experience in deep learning, AI and tailored solutions, paired with the Dell Technologies best-in-class product portfolio and global scale, create a powerful and innovative technology partnership.

[AI InfoHub](#)
dell.com/ai

¹Dell Technologies case study, [AI-powered image detection for COVID improves patient outcomes](#), May 2022.

Key benefits

- **AI growth made easy:** No need to invest time evaluating or money into storage with costly features that are never used.
- **AI simplified:** Jointly engineered and validated solutions make it quick and easy to deploy a data server solution to accelerate AI workloads and mainstream data science environments.
- **Faster AI insights:** The AI Data Server architecture enables blisteringly fast performance for AI workflows while leaving more budget for GPUs.
- **Proven AI expertise:** Confidently deploy an engineering-tested data server solution backed by world-class Dell Technologies services and support. Select ProSupport Plus for a single point of contact for software and hardware support.
- **Multiple reference sites:** Buy with confidence, PEAK:AIO have been supplying the AI market for over two years and have a strong install base within a wide range of sectors with data scientists and professors testimonies.

Ideal for smaller or new-start projects with direct connect



Single 2U PEAK:AIO DELL PowerEdge AI Data Server (30TB – 360TB)



Drive up to 10 Dell GPU Servers via Dell switch for scaled projects

Technical specifications

Compute	Networking	Protocols	Software
PowerEdge R7525 2U with up to 24x NVMe drives	2/4 x 200Gbe/HDR	RDMA based	PEAK:AIO 3yrs License
Drive options: 7.69TB/15.3TB	Direct connect Ethernet	NVMe-of (Block) and NFS with RDMA	NVIDIA GPUDirect compatible
Dell details	Ethernet Switched InfiniBand switched	Multipath options	
		TCP also supported	

Validated Designs for AI with PEAK:AIO help you simplify storage to accelerate performance and improve AI results with confidence.