

# OpenFlex™ Data24 4200 and PEAK:AIO Software-Defined Storage for GPU-Accelerated AI



PEAK  
AIO

## Challenges

- Keeping pace with exponential data growth with predictable performance and TCO.
- Avoiding head-of-line blocking and metadata bottlenecks that starve GPUs.
- Meeting data protection and recovery objectives while preserving throughput.
- Reducing operational overhead with automation, observability, and simple lifecycle management.

## Highlights

- Unified file (NFSv3/v4 over RDMA/TCP) and block (NVMe-oF over RDMA/TCP) from the same protected pools.
- Exceptional performance validated up to 149 GB/s sustained Read throughput
- Scales from 2.8PB to 23PB with industry-leading power efficiency and rack throughput
- Composable disaggregation: scale GPU servers, protocol bandwidth, and storage enclosures independently.
- PEAK:PROTECT (N+2) software-defined erasure encoding provides enterprise-class resilience with read performance comparable to RAID0.
- GPUDirect Storage enablement through PEAK:AIO enabling direct GPU-to-NVMe DMA data paths.
- Built-in Prometheus metrics and Redfish APIs for automation and SRE workflows.

## Solution

The SDS head composes NVMe namespaces from multiple OpenFlex Data24 4200 enclosures into virtualized, protected pools. Administrators can export shared file services for datasets and block volumes for device-level access—both concurrently. Policy-driven placement can isolate high-IO consumers, and optional tiering extends the plane across heterogeneous media for cost-efficient retention.

Modern AI pipelines demand both large, sequential bandwidth and sustained small-IO performance. This solution brief outlines a reference design that combines disaggregated NVMe-oF™ shelves (OpenFlex Data24 4200) with a GPU-aware software-defined storage (SDS) head (PEAK:AIO). The result is a protected, virtualized storage plane implemented entirely by the PEAK:AIO SDS layer, exposing both file and block services over RDMA/TCP, sustaining line-rate throughput to GPU servers while keeping administrative operations straightforward.

## OpenFlex Data24 4000 Series Storage Platforms

The OpenFlex Data24 4000 series NVMe-oF storage platform extends the high performance of NVMe flash to shared storage. The 4000 series provide low-latency sharing of NVMe™ SSDs over a high-performance Ethernet fabric to deliver similar performance to locally attached NVMe SSDs. Western Digital RapidFlex™ NVMe-oF controllers, allows up to six dual pathed hosts to be attached without a switch. The OpenFlex Data24 4000 series uses Western Digital's RapidFlex C2000 Fabric Bridge Adapters to provide up to 12 ports of 100GbE which can connect to RDMA and/or TCP configured host ports.

## PEAK:AIO Software-Defined Storage

Combining OpenFlex Data24 4200 with PEAK:AIO yields a software defined AI and HPC ready storage plane that sustains line-rate bandwidth and multi-million-IOPS behavior while remaining simple to deploy and operate. The design preserves local-device characteristics across an NVMe-oF fabric, delivers robust N+2 protection, and exposes automatable interfaces that align with modern operational needs.

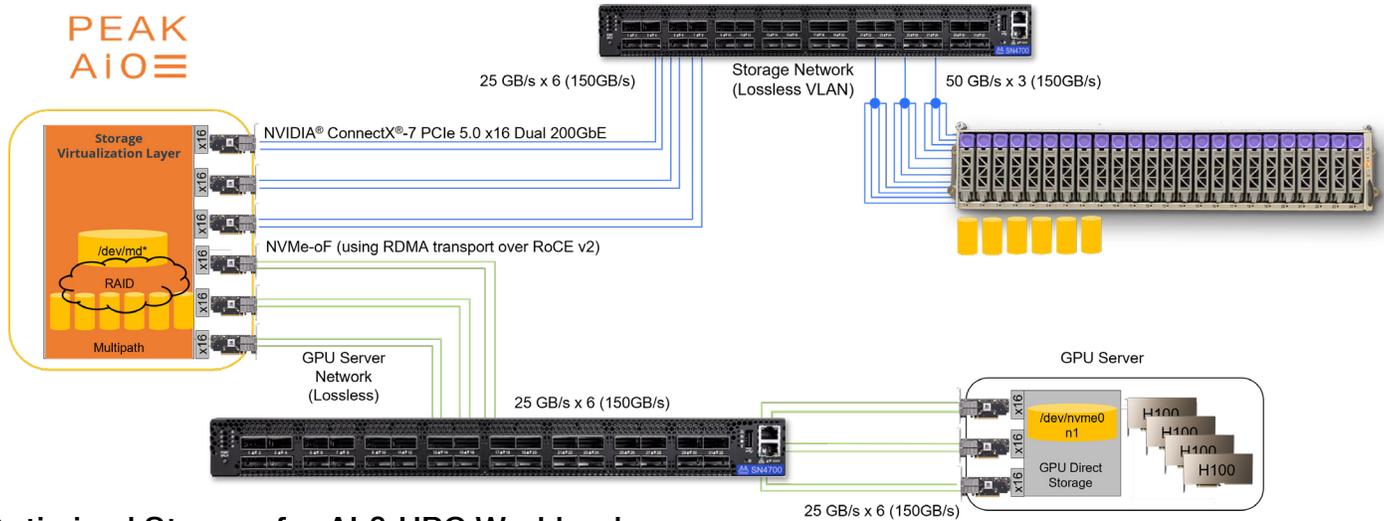
- Storage virtualization that composes one or more protected NVMe device groups into large or isolated pools.
- Front-end protocols optimized for GPU dataflows: NFS over RDMA and NVMe-oF with multipathing.
- PEAK:PROTECT erasure-coded (N+2) data protection implemented in software, tuned for NVMe media and parallel rebuilds.
- Optional Apex tiering for sub-file, block-level placement across flash/HDD/object tiers.

## Reference Architecture

A two-fabric design isolates the storage network from the client/GPU network. On the storage backend, the SDS head connects to a dedicated storage fabric and discovers NVMe-oF namespaces exported by two or more Data24 enclosures. On the client network, the SDS head exports file and block services to GPU servers over a separate RDMA-enabled fabric. VLAN or physical separation preserves deterministic behavior and allows distinct congestion-control and QoS policies per domain.

- Storage Network: NVMe-oF to OpenFlex Data24 4200 enclosures (multipath enabled).
- Client/GPU Network: NFS over RDMA for shared datasets; NVMe-oF for device-level access to scratch/checkpoints.
- The reference design can utilize a single-socket server architecture to avoid NUMA constraints.

# OpenFlex Data24 4200 and PEAK:AIO Software-Defined Storage for GPU-Accelerated AI



## Optimized Storage for AI & HPC Workloads

### AI Training & Fine-Tuning

- Sustained large-block sequential throughput for ingest and checkpoints.
- Parallel small-IO from data loaders served via RDMA file/blockaths with low jitter.
- GPU Direct Storage reduces CPU mediation and improves latency.

### Data Preparation & ETL

- High-throughput file access for decode/augment pipelines.
- Block volumes for temporary scratch and shuffle operations.

### Management & Automation

- Prometheus metrics and reference Grafana dashboards for throughput/IOPS/latency and health.
- Redfish API for capacity lifecycle; integrates with orchestrators (e.g., Kubernetes CSI).
- Decoupled telemetry and provisioning enable policy-driven placement and autoscaling workflows.

### Sizing & Scaling Guidance

- Start with one SDS head and two to eight Data24 enclosures; add shelves to scale capacity, add SDS heads for redundancy and failover.
- Increase RNIC count and link speed to scale protocol bandwidth and path diversity.
- Use isolated RAID sets or pools to segregate high-IO tenants and maintain QoS.

### Inference & Feature Stores

- Shared models and embeddings over RDMA NFS; high-IOPS block volumes for databases and vector stores.
- Consistent latency under concurrent read-heavy mixes.

### Performance Characteristics (Representative)<sup>1</sup>

- RDMA NFS sequential read: sustained ~149 GB/s from a single PEAK:AIO SDS head, backed by the OpenFlex Data24 4200 in controlled tests.
- NVMe-oF 4 KiB: multi-million IOPS across clients; suitable for parallel data loader patterns.
- Internal software data-protection layer: reads near raw NVMe maxima with N+2 parity; predictable parity-write behavior.

### Security & Data Protection

- N+2 parity protection with parallel rebuild to minimize degraded-mode impact.
- Export-level access controls aligned with enterprise segmentation patterns.
- Optional immutable/archive tiers for long-term retention and compliance workflows.

For technical details on this solution, see the reference architecture: [Western Digital OpenFlex 4200 and PEAK:AIO Software-Defined Storage for GPU-Accelerated AI Infrastructure](#).



<sup>1</sup>Actual results depend on CPU memory bandwidth, RNIC count/link speed, PCIe layout, and fabric configuration