

Kumulus Distributed File System

A high-performance, scalable parallel file system designed for demanding applications like HPC, AI and large-scale data processing.

Empowering High-Performance Data Storage Solutions

Kumulus Distributed File System is a versatile parallel filesystem used globally in performance-driven sectors like HPC, AI, life sciences, media, and energy. Its userspace architecture offers flexibility and scalability, supporting demanding applications and maximizing hardware performance. It enables scalable infrastructure investments, accelerating results and supporting advanced data analysis without disrupting workflows.

Use Cases

- ✓ High Performance Computing (HPC)
- ✓ AI/ML Workloads
- ✓ Big Data Analytics
- ✓ Media Production and Rendering
- ✓ Financial Services Data Archiving
- ✓ Government and Defence
- ✓ Research institutions
- ✓ Cloud Infrastructure Providers
- ✓ IoT Data Storage

Key Benefits of Kumulus Distributed File System

- 1 **Efficient Distribution of Files and Metadata**
UniStor FS eliminates bottlenecks by distributing file contents and metadata across multiple servers, optimizing resource utilization and enhancing performance for large-scale and metadata-intensive applications.
- 2 **Advanced HPC Technologies**
Kumulus Distributed File System is designed for high-performance computing, leverages scalable, multithreaded components with native RDMA support, ensuring reliable operation and automatic failover for seamless connectivity.
- 3 **High Performance**
Kumulus Distributed File System provides parallel I/O capabilities, allowing multiple nodes to access and write data simultaneously, significantly boosting performance.
- 4 **Scalability**
Kumulus Distributed File System is highly scalable, supporting growth in both storage capacity and performance by simply adding more servers or storage devices.
- 5 **Optimised for Small Files**
Kumulus Distributed File System is optimised to handle workloads involving numerous small files efficiently.

