

EXASCALE STORAGE FOR HIGH-PERFORMANCE WORKLOADS

ExaStor Scale-out Storage

GLUESYS



CHALLENGES FOR MODERN HIGH-PERFORMANCE WORKLOADS

THE STORAGE BOTTLENECK ISSUE

Storage can easily become a main cause of bottleneck when it is not prepared for fast growing AI/ML workloads

ESCALATING STORAGE COST

AI/ML workflow requires tremendous amount of storage space and network traffic which leads to more operational costs

INCREASING OPERATIONAL OVERHEAD

To fully support the dynamic needs of AI applications, storage must provide transparent monitoring and insights throughout the pipeline

EXASCALE STORAGE FOR HIGH-PERFORMANCE WORKLOADS

High-performance Parallel File System

Provide high-speed parallel I/O to every data in the storage cluster

NVIDIA® GPUDirect® Storage Support

Support NVIDIA GPUDirect storage for direct memory access between GPU and storage

Storage Type in All Scenarios

Leverage NVMe all-flash, SSD and HDD hybrid configuration

Flexible Deployment and Scalability

Choose how to deploy or scale your storage that suits your environment

Up to 40GB/s per Node

Delivers 40GB/s sequential read throughput with all-NVMe storage system

RDMA-powered File Service

Supports NFS over RDMA protocol with InfiniBand EDR and HDR network

HIGH-PERFORMANCE DISTRIBUTED FILE STORAGE OPTIMIZED FOR AI/ML WORKLOADS

HIGH PERFORMANCE

High-performance Parallel File System

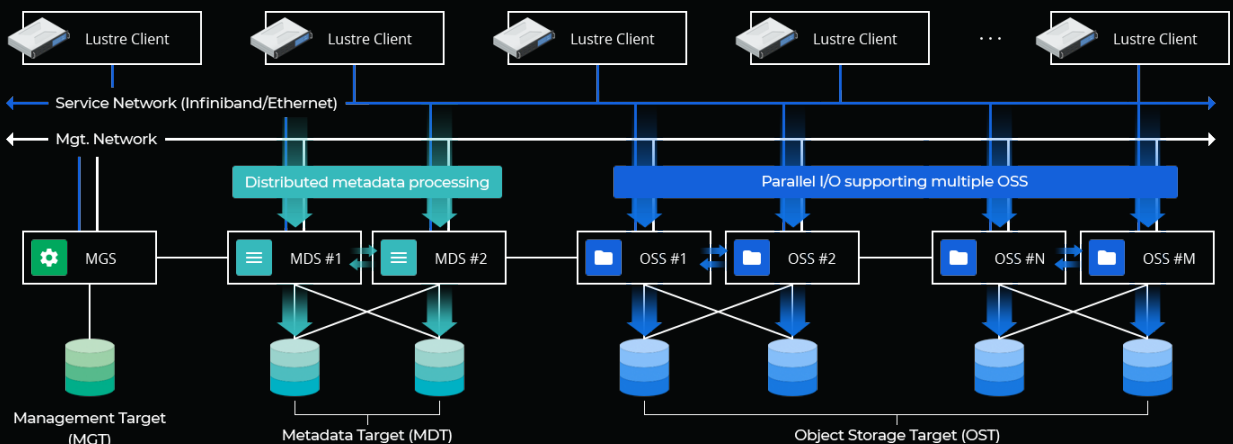
Allows parallel data access through Luster parallel file system, with high-speed file sharing and petabyte-level scalability through its parallel I/O architecture. Files are distributed and stored in objects across the cluster through global namespace.

NVIDIA® GPUDirect® Storage Support

Provides a storage interface to accelerate GPU applications, minimizing CPU and memory load. Maximizes GPU performance by providing a distributed file system IO driver module for applications running on GPGPU.

Flash-based Level 2 Cache

Can leverage SSD drives as secondary read cache, extending the main memory cache. When the main memory is full, the system can perform read request on the allocated cache drive, instead of using slower hard drives.



Brochure

HIGH AVAILABILITY

Hardware Redundancy

Redundancy based on dual controllers ensures service continuity through automatic failover in the event of a controller failure. This architecture effectively eliminates the risk of split-brain scenarios during data synchronization, ensuring data integrity and reliability. (* DC model only)

Multi-Rail Network

Supports parallel network I/O with multi-rail configuration utilizing both TCP/IP and RDMA network. Network traffics are automatically distributed through load balancing to prevent bottleneck while optimizing the overall network performance.

High Availability Architecture

Ensure uninterrupted services by active-active storage redundancy architecture. Minimize interconnectivity between resources as well as the complexity of active-active configuration to provide instant recovery through partial failover.

Flexa Snapshot Manager

Leverage the power of the Global Write Barrier of Flexa Snapshot Manager to take instant snapshots across all storage clusters without any system downtime. Snapshots are swiftly generated through copy-on-write by capturing only the changes, operating deep within the file system layer to ensure minimal impact on metadata performance.

SCALABILITY AND EFFICIENCY

Scale-out Architecture

Effortlessly scale your storage capacity by dynamically adding nodes online. You can achieve massive storage space in petabyte-level in a single system, while having linear performance boosts through parallel scaling. This approach not only minimizes the scale of initial adoption but also ensures cost-effective storage management tailored to your business needs and data growth.

Flexible Volume Management

Enhance your storage efficiency with thin provisioning, to allocate logical space beyond the current physical capacity by anticipating future growth trends. This maximizes capacity utilization while providing flexible and elastic capacity management, ensuring your storage resources adapt seamlessly to evolving demands.

Dynamic Storage Scaling

Optimize your high-capacity, high-density storage with a cutting-edge storage enclosure that supports up to 108 SAS drives in a compact 4U form factor. When scaling up with the E484 model, you can install up to 408 disks, configuring a massive 8.16PB of physical capacity. This scalable architecture ensures your storage infrastructure can grow seamlessly with your data demands.

Storage Space Optimization

Achieve optimal storage efficiency with real-time data compression, which reduces I/O bandwidth and maximizes storage space utilization. This advanced technology also delivers integrity checks on compressed data, ensuring the detection of data corruption and maintaining data integrity.

STREAMLINED MANAGEMENT

Multi-Protocol Support

ExaStor supports a wide range of protocols, including NFS, CIFS, GPUDirect Storage, POSIX-compliant interfaces, and RESTful APIs, ensuring that heterogeneous users and applications can seamlessly access its global namespace.

Comprehensive System Monitoring

Monitor system usage and availability across clusters and local volumes, including NFS, CIFS, and file system clients. Our dedicated management software enables real-time tracking of disk health and location within the system, ensuring efficient and effective management.

Advanced Performance Monitoring

With our dedicated management software, you can effortlessly track the IOPS and bandwidth for clusters and local volumes. It provides real-time monitoring of network I/O performance for every controller, file system usage, power consumption, and CPU/memory utilization per controller, ensuring visibility of your system's performance metrics.

USE CASES



Educational Institutions

- AI-related data
- Online courses



Research Institutions

- Research data
- Unstructured data



Bio and Life Sciences

- Clinical data, Genetic big data
- Medical image analysis, DNA sequencing



4th Industrial Revolution Fields

- Autonomous driving
- Smart cities
- Robotics, Drones, Weather data



Internet Companies

- LLM-based ChatGPT services
- AI services



Media/Entertainment

- VR games, VOD services
- 3D rendering, video editing

GLUESYS CO.,LTD.

5F, 11-31, Simin-daero327beon-gil, Dongan-gu,
Anyang-si, Gyeonggi-do, Korea

T. +82.70.8787.5370

W. www.gluesys.com



GLUESYS