

Ultrastar® Data60 and Data102 Hybrid Storage Platform SMR Drives in Combination with Swiss Vault Excels in Long-Term Active Archive Storage Strategy for the AI Age



Highlights

- Scalable high-density storage built on SMR HDDs
- Erasure coding flexibility, allowing mix-and-match parity, media volume, and hot-swapping of drives
- Resilient data using software driven performance from VaultFS unlocking 2x or greater read/write speeds
- Energy efficiency and longer hardware refresh cycles extending infrastructure life
- High Capacity & Performance from Ultrastar Data60 and Data102 Hybrid Storage Platforms

Challenges for the AI Age

The data landscape is evolving rapidly. Enterprise storage must keep up. With the surge in AI workloads, from large-scale model training to real-time inference and Retrieval-Augmented Generation (RAG), the nature of data access is shifting. These next-gen applications demand active data, with vast datasets that must remain always accessible, not just stored away in deep, inaccessible archives.

These data assets, often retained for decades, face slow but relentless threats: bit rot, file corruption, media degradation, and silent data loss. Over time, these issues erode the very foundation of data integrity, which is a must for archival data.

Traditional storage architectures like NAS and SAN, built for yesterday's workloads, struggle to scale economically or protect data over the long haul. To be resilient, enterprises need a new class of storage now, one that's flexible, scalable, and built for the future. A platform that not only grows with data volumes, but also safeguards data health, easily supports hardware refresh cycles, and preserves critical datasets over decades, without compromising data availability or performance.

Why Shingled Magnetic Recording Drives Matter

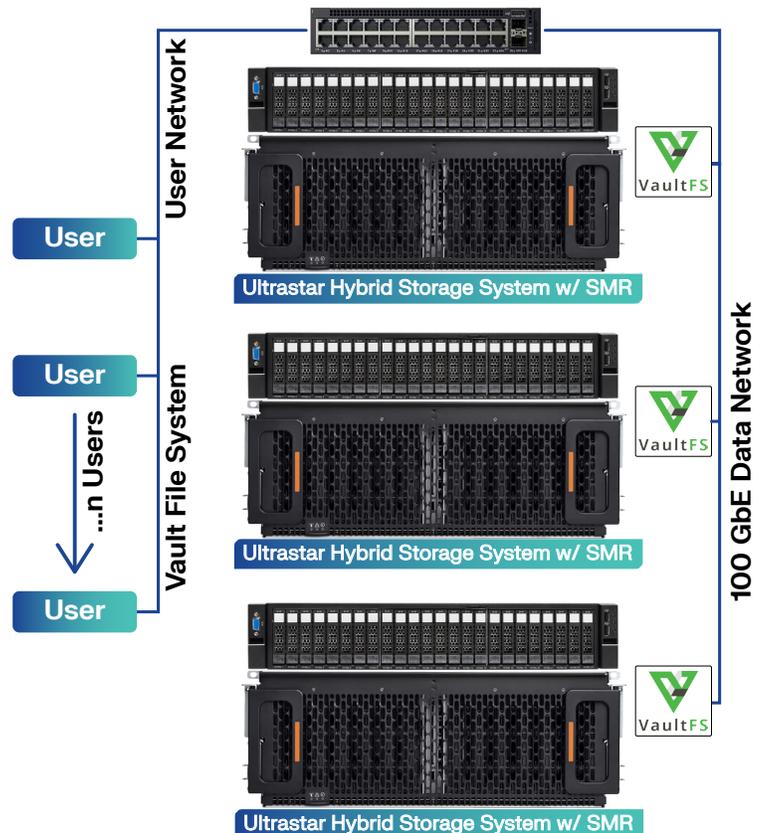
Shingled Magnetic Recording (SMR) technology is transforming enterprise storage economics by delivering up to 20% greater storage density in the same physical footprint as conventional drives. With SMR, organizations can pack more capacity into existing set-up, a critical advantage as data volumes continue to explode.

SMR is enabling sustainability and scalability. By enabling higher capacity per drive, SMR reduces rack space, power consumption, and cooling requirements, all while lowering the overall cost per Terabyte.

When paired with Western Digital's Ultrastar Data60 and Data102 hybrid storage platforms, SMR Ultrastar drives excel as *active archive storage*. This ensures large datasets remain online, accessible, and cost-effective. The true performance unlock comes with Vault File System (VaultFS), a software-defined storage layer that applies advanced erasure coding, parallel I/O, and distributed data processing. Together, these innovations turn SMR-based systems into high-performance, highly resilient storage platforms that meet the dual mandate of economy and performance.

Configuration for Economical Performant Storage

The example configuration includes servers, 3 Ultrastar data nodes, SMR Drives, Vault File System, and 1x100 GbE networking for intra-node cluster communication.



Ultrastar Data60 and Data102 and SMR with Swiss Vault Excels in Long-Term Active Archive Storage Strategy for the AI Age

VaultFS: Unlocking the Power of SMR drives on Ultrastar Data60

SMR drives deliver unmatched storage density. To fully realize their potential, enterprises need a software layer that can extract performance, resilience, and long-term value. That's where Vault File System (VaultFS) comes in.

1. Built for Flexibility and Scale: VaultFS removes the traditional constraints of file systems by unifying diverse storage resources into a single, limitless namespace.
2. Smarter Economics through Data Efficiency: With VaultFS, storage efficiency translates directly into cost savings. By reducing redundant copies and minimizing data volume, VaultFS cuts both infrastructure overhead and energy consumption, delivering a greener, more sustainable solution without compromising durability.
3. Erasure Coding, Reinvented: Organizations can define data-plus-parity (D+P) ratios per file, per directory, or per workload, tuning the balance of capacity, performance, and protection as business needs evolve.
4. High Availability: VaultFS distributes data across multiple nodes, ensuring uninterrupted operation even in the face of drive, node, or JBOD failures.
- T. Performance: Through parallel I/O and distributed data processing, VaultFS unlocks 2x or greater read/write speeds compared to replication-based traditional file systems, making SMR economical and performant enough for active archive and AI-driven workloads.

VaultFS – Capacity & Performance

Presenting here is VaultFS performance results on Ultrastar platforms+ SMR drives, showcasing capacity and throughput in a mid-tier configuration. The testbed consisted of three Ultrastar Data60 JBODs with 180 SMR drives¹, managed by three VaultFS servers over a 100 GbE network.

VaultFS Understands SMR Drives

VaultFS is a single file system that understands SMR drives and optimizes data transfers, just the way SMR drives prefer. That means no additional overhead, no fragmentation of performance, and no architectural bottlenecks. Figure 1 and Figure 2 show the results of testing various Erasure Coding configurations, e.g., 20 Data fragments + variable Parity.

VaultFS achieves 12,000 MiB/s write and 11,000 MiB/s read speeds for Data=20 + Parity=2.

VaultFS is built to optimize SMR characteristics, orchestrating writes and reads across clustered drives and clustered JBODs in parallel, with erasure coding tuned for speed and resilience.

VaultFS makes SMR drives simple to use, giving operators full control to dial in resilience (Parity (P) settings), performance (Data (D) setting), and cost optimization (D+P ratios) with ease.

Instead of one-size-fits-all limitations, VaultFS puts ultimate flexibility in your hands, letting you tune the system to match workload, budget, and business priorities. The result is a storage platform that adapts to you—not the other way around.

VaultFS : Achieving High Performance, High Capacity or Both

The choice is yours. An economical high-capacity storage configuration can be set up by using a minimum 8 CPU Cores, 32 GiB RAM, 1 HBA and 1x 100 GbE for each node. To further reduce costs, one can use 2 HBA + 2 JBODs per server. Higher data throughput performance can be achieved by adding more HBAs and SAS cables for connectivity to the Ultrastar JBODs as well as additional RAM and network connectivity to the server. The power of VaultFS's parallel and clustered processing means you can scale infinitely with Ultrastar Data60 or Data102 JBODs to 100's of exabytes and beyond, increasing the performance as well.

Figure 1: VaultFS SMR Write Throughput

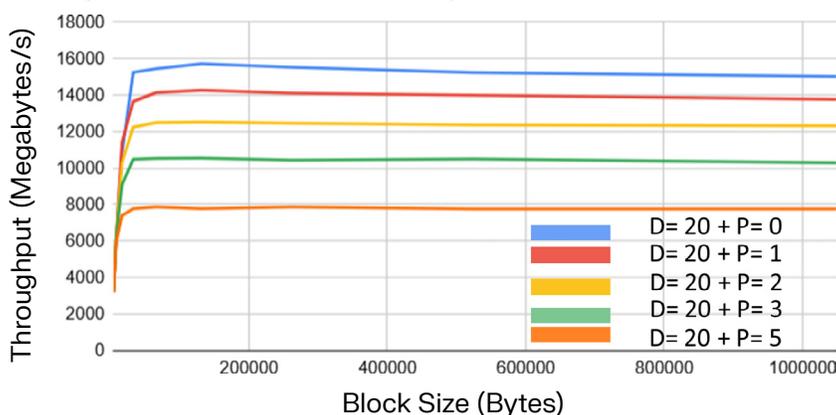
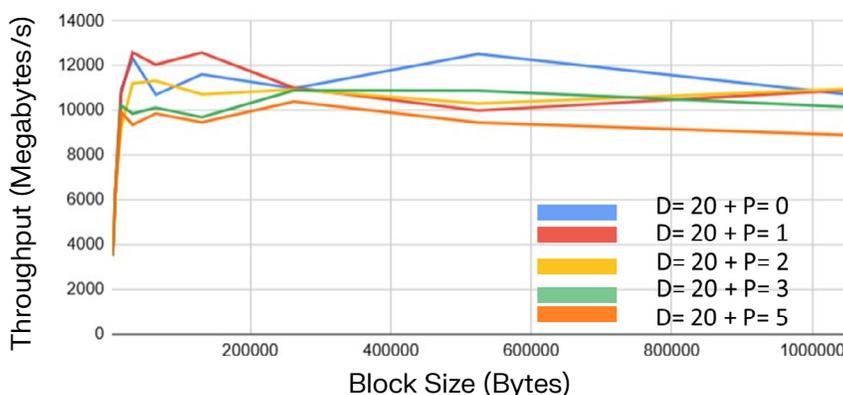


Figure 2: VaultFS SMR Read Throughput



Ultrastar Data60 and Data102 and SMR with Swiss Vault Excels in Long-Term Active Archive Storage Strategy for the AI Age

Hardware Flexibility and Automated Data Migration

Swiss Vault's customers can re-use existing standard networks, servers, and storage to extend the life of that infrastructure, and can also take advantage of recent hardware advances. By connecting existing and new systems in a single namespace, organizations can scale their data storage, and enable automated data migration to new disks and servers.

Ultrastar Data60³

- Up to 60 Ultrastar HDDs (SAS or SATA)
- Up to 1.8 PB² of raw storage up to 30TB SMR HDDs
- 4U form factor



Ultrastar Data102⁴

- Up to 102 Ultrastar HDDs (SAS or SATA)
- Up to 3.06² PB of raw storage up to 30TB SMR HDDs
- 4U form factor



Daisy-chaining for High Capacity

Up to 4 units may be daisy-chained for a total raw capacity of 7.2 PB (Data60) or 12.24 PB (Data102)

Innovations for Performance and Reliability

IsoVibe™: Patented technology improves isolation of vibration propagation both to and from each individual drive to help maximize performance even in heavy workloads.

ArcticFlow™: Improves cooling via discrete airflow channels that allow cool air to reach more components within the system, improving cooling effectiveness that can improve drive reliability.

Flexible

- Choose dual-port SAS for high availability or single-port SATA for low cost.
- Up to 12 x 12Gb/s SAS-3 host connections.

Designed for Serviceability

- Enterprise-grade redundant and hot swappable PSUs, IO Modules, and fans.
- Rack-mounted top cover for quick and easy service.

