



Ultrastar® Data60 and Data102 Hybrid Storage Platform Powered by Swiss Vault



Delivering Fast, Resilient, & Sustainable Storage

“Every CIO should have a strategy for 100-year data retention.” Swiss Vault’s software (“VaultFS”) delivers data resiliency on top of Western Digital’s Data60 and Data102 JBODs. This solution provides an economically attractive archive for huge volumes of data that can last for decades. Retaining and actively accessing large data sets can be of great value to AI/ML workloads.

Highlights

- Swiss Vault’s VaultFS and Western Digital JBODs together deliver low-cost, high-density, sustainable storage.
- According to Swiss Vault’s chart below, VaultFS has been shown to Increase data Read and Write speeds by 2X or more over replication file systems such as ZFS.
- Provides software data resiliency with hardware features that together enhance the life of data and reduces the frequency of returned drives.
- Enables scheduling for drive technology updates to further capacity and energy cost savings.

Solution

Western Digital Ultrastar Data60 and Data102 with Swiss Vault’s VaultFS software, offer highly-scalable, highly-available storage solutions. Compatible with industry-standard X86 servers, Swiss Vault will scale on the server platform of your choice. Distributing the data across multi-nodes increases data resiliency with lower data overhead. Therefore, with VaultFS replacing failed hardware can be managed on a schedule rather than having to react immediately to hardware failures.

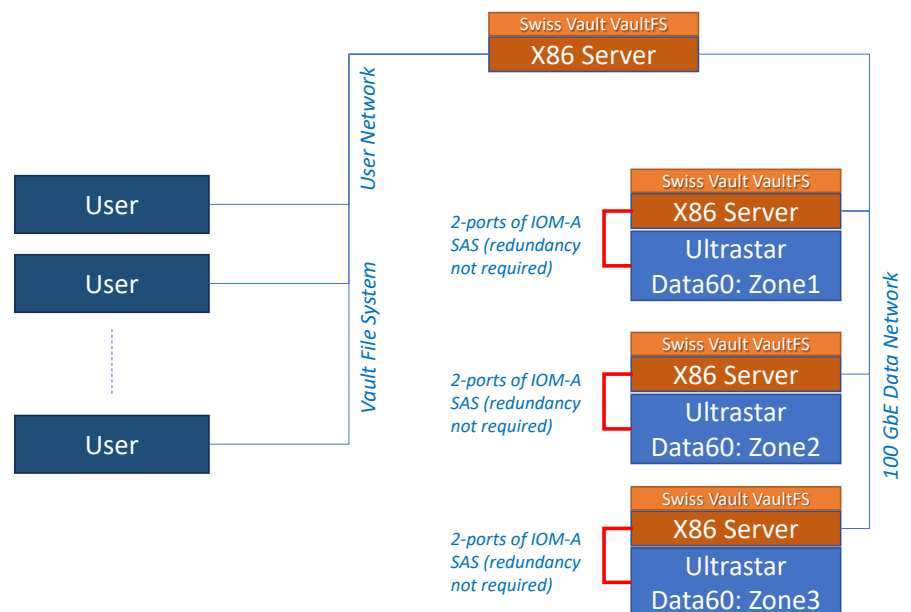
Configuration

This example configuration includes one (1) client node and three (3) data nodes. Each data node is assigned to one zone in the Ultrastar Data60 JBODs. Each server includes two, 100GbE ports for application and intra-node cluster communication. For fault tolerance, VaultFS employs erasure coding to provide uninterrupted operation with the failure of up to two servers or one JBOD improving reliability and robustness. The Peer-to-Peer (i.e. no master-node) setup also enhances security of the overall solution.

This configuration supports a POSIX presentation and access to capacity in all storage nodes.

Enterprise Storage Challenges

With the continued growth of data, resulting from the boom in AI, the need for flexibility and scalability of data storage solutions is also rising. Traditional enterprise solutions such as Network Attached Storage (NAS) and Storage Area Network (SAN) lack the capability to upscale for these next-gen applications. To ensure long-term economic viability, the data storage solution should be able to seamlessly scale capacity while maintaining excellent data durability and seamlessly upgrade to newer, higher density, more energy efficient technologies when they are available.



Ultrastar Data60 and Data102 Hybrid Storage Platform Powered by Swiss Vault

Swiss Vault VaultFS

Flexibility & Speed are Key

Swiss Vault's commitment to sustainability prioritized greater user flexibility. This flexibility enables data operators to: (i) fully utilize existing infrastructure; (ii) reduce costs by minimizing data volume, thus cutting down on infrastructure and energy expenses for data under management; (iii) incorporate diverse storage resources into a single namespace with no limits on the filesystem size; and (iv) adapt to the customer's needs to scale-up without forcing customers to fit within the technology provider's restrictions.

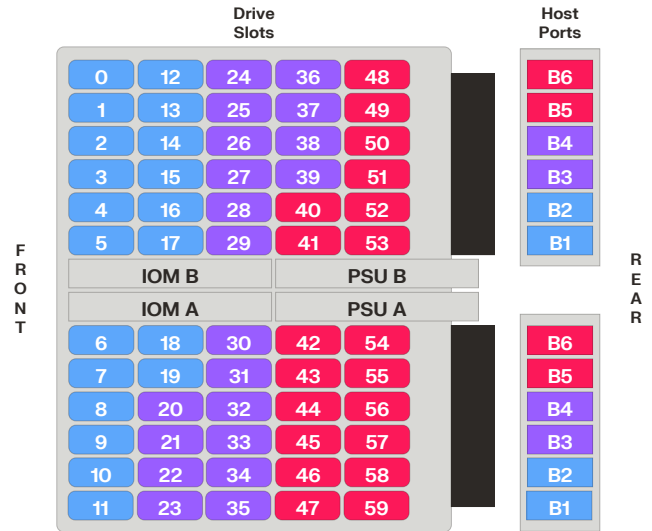
Erasure Coding Flexibility: Any Data, Any Parity, Any Time

Erasure coding enables higher resilience and more efficient capacity utilization than traditional RAID schemes. Further, VaultFS distributes the data over multiple servers, allowing for limited node failures while increasing reliability, security and robustness. The Vault File System's erasure coding is different in that it is easy to implement, manage and much more flexible than other solutions.

Swiss Vault's flexible erasure coding allows hot swapping of drives, permitting any level of resilience which means an organization can run existing HDDs to failure without risking downtime, maintenance and repair. Simply replace older low-capacity drives with more energy-efficient high-capacity drives to optimize for capacity and energy efficiency on a schedule that is optimal for the organization.

Customers can dynamically re-configure erasure coding using any combination of data plus parity (D+P) chunks. Most competitors only allow customers to choose from a few vendor-defined options. However, Swiss Vault customers, at any time can assign a different D+P per directory, file, file type, or file class to match the customer's current requirements. This flexibility applies across the life of the hardware infrastructure, regardless of media size, speed, or vintage.

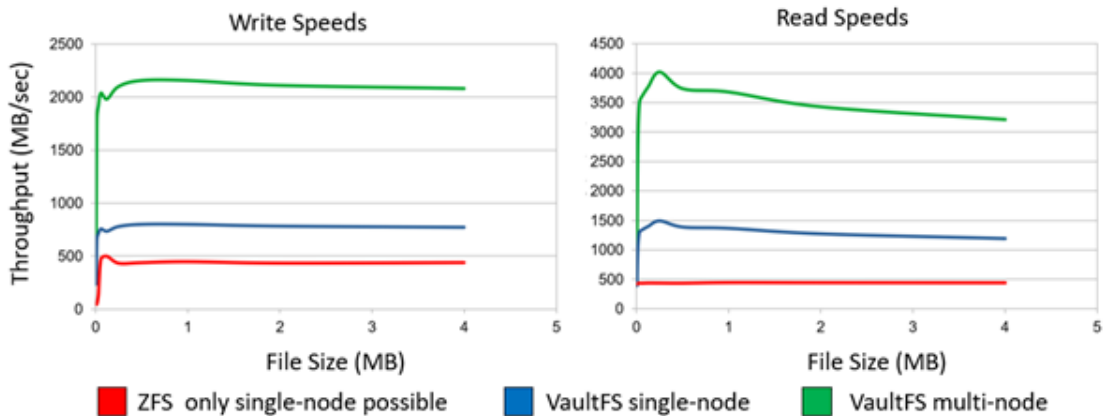
With VaultFS' erasure coding flexibility, the customer can set the desired r/w speeds.



Three Zone Configuration on a Data60

Disks	6 to 12 x 5TB ¹ HDD per node
CPU	56x Intel® Xeon® e5-2680 v4 @2.4 GHz
Memory	4 x 16 GiB 2400 MT/s DDR4
Node	1 to 5 x Dell® PowerEdge® R730xd
Ethernet	25 GbE

Single Node and Multi-Node Throughput



Source: Swiss Vault

Ultrastar Data60 and Data102 Hybrid Storage Platform Powered by Swiss Vault

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 Data60

- Up to 60 Ultrastar HDDs (SAS or SATA)
- Up to 1.44 PB² of raw CMR HDD storage
- 4U form factor



Ultrastar Data102

- Up to 102 Ultrastar HDDs (SAS or SATA)
- Up to 2.45² PB of raw CMR HDD storage
- 4U form factor



Daisy-chaining for High Capacity

Up to 4 units may be daisy-chained for a total raw capacity of 5.76 PB (Data60) or 9.8 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 24Gb/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.

¹ One terabyte (TB) is equal to one trillion bytes. Actual user capacity may be less due to operating environment.

² One terabyte (TB) is equal to one trillion bytes and one petabyte (PB) is equal to 1,000 TB. Actual user capacity may be less due to operating environment.

³ This solution was tested with 24Gb/s SAS host connections.

