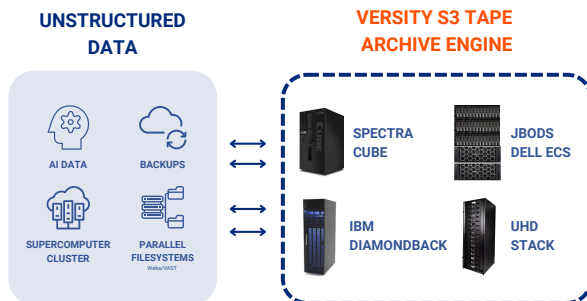# S3 Tape Archive Engine

Scalable Mass Storage For the AI-Era

Versity

Today, hyperscalers face the challenge of managing massive volumes of cold data while balancing cost efficiency, scalability, and long-term durability in storage solutions.

Introducing Versity's S3 Tape Archive Engine

The S3 Tape Archive Engine is an advanced solution tailored for hyperscalers and data-intensive enterprises, engineered to seamlessly integrate object storage with the unmatched reliability of tape. Optimized for efficient data transfers, it offers unparalleled scalability and cost-effectiveness for long-term archiving, enabling organizations to securely store vast volumes of cold data while significantly reducing total ownership costs. By merging the durability of tape with the flexibility of S3-compatible storage, it provides an ideal archival solution for hyperscale environments seeking both performance and efficiency.

A Seamless S3 to Tape Solution for AI & HPC

The S3 Tape Archive Engine is a highly scalable, modular solution designed for hyperscale environments, seamlessly integrating S3-compatible object storage with exabyte-level tape libraries such as Spectra Cube and IBM Diamondback systems.The engine simplifies archiving and retrieval by translating S3 commands into streamlined tape operations, allowing hyperscalers to effortlessly move cold data to tape while maintaining full S3 access and workflow continuity.

Tailored for AI and high-performance computing (HPC) workloads, it ensures secure, efficient long-term archival, offering cost-effective data management across storage tiers while supporting multiple ingestion protocols like NFS, Samba, FTP, and S3. Our deep integration with AI/HPC tools ensures minimal disruption, leveraging the full potential of tape storage for large-scale data archiving within complex AI and HPC workflows. By automating data movement and optimizing energy consumption, the platform provides flexible, rapid access to vast datasets without disrupting existing workflows.

Optimize Power and Cost

Leverage the unmatched cost efficiency of tape storage, offering a substantially lower cost-per-terabyte than disk-based solutions, especially for cold and infrequently accessed data. Our high-density, modular architecture reduces power consumption and operating costs, while seamlessly integrating with energy-efficient tape systems. With intelligent automation, the platform ensures easy access to secure, long-term storage for AI model archives, simulation data, and critical research outputs, preserving your data for decades without risk of degradation.

Support Warm and Cold S3 Data

Unlike solutions that are limited to Glacier, Versity enables seamless tiering between active and deep archival storage, allowing users to store frequently accessed data in cache and migrate colder data to tape as needed. This flexibility offers a distinct advantage for organizations needing dynamic storage options that optimize both performance and cost across AI, HPC, and other data-heavy environments.

Intuitive and Effortless

With a simple point-and-click interface, the S3 Tape Archive Engine requires no prior tape-specific knowledge, making it accessible to anyone, regardless of expertise. Installation is straightforward, allowing for quick setup and integration into existing systems without the need for extensive training. This ease of use empowers anyone in your organization to manage tape archiving effectively. Whether you're an IT professional or a non-technical user, managing data archiving has never been more accessible or efficient.

About Us

Versity is a leading independent technology company focused on delivering innovative mass storage solutions. With a commitment to performance, scalability, and open-source collaboration, Versity empowers organizations to efficiently manage their data in a rapidly evolving digital landscape.

Worldwide consumer base managing over **2 exabytes** of data across 95 sites

Accelerated innovation driven by customer-focused feature requests and a strong, next-generation engineering team

Unrestricted access to expert guidance and **personalized support** from our highly-skilled engineering team

Key Benefits

High Performance

- High throughput to archival storage
  - 40GB/s per 5 node cluster
- **Parallel transfer**: Our S3 Tape Archive Engine reads and writes data to mass storage devices through multiple channels per node simultaneously.
- **Fast metadata-driven queries:** Efficiently retrieve critical information without accessing tape, streamlining workflows for optimal performance.

Secure Long-Term Storage

- Supports compliance with **data governance and long-term archiving regulations**, including legal and industry-specific data retention requirements
- **High availability** even if servers are lost, with built-in failover and no need for external tools
- **Ultimate disaster recovery** directly from storage media that supports asynchronous remote replication for continuity

Advanced Management

- **Extensive API coverage** to ease integration with third party tools and site specific workflows
- Integrated software **automatically manages tape libraries**, optimizing storage density and access to meet hyperscale data retrieval requirements
- Exa-scale archival designed for long-term retention with minimal maintenance

Trusted by Industry Leaders

PAWSEY · UAF UNIVERSITY OF ALASKA FAIRBANKS · UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN · NIWA Taihoro Nukurangi · ThinkOn · NASA · wetaFX · DNEG · tgen · LIBRARY