

Albireo Virtual Data Optimizer (VDO) for SOHO Storage Solutions



Data storage for the small office and home office (SOHO) is now reaching tens of terabytes in scale. All of the data being created by imaging and audio applications, coupled with the existing data growth, makes SOHO a perfect opportunity for data deduplication. Much of the data created is copies of copies and the duplicate rates can be very high. Eliminating these duplicates with Albireo deduplication technology dramatically reduces overall storage costs by lowering data storage requirements by 5-35X. In a SOHO implementation, that cost reduction is especially significant relative to the overall budget.

Albireo Virtual Data Optimizer (Albireo VDO) provides the fastest route to market for SOHO storage deduplication on systems powered by the popular Linux operating system. Albireo VDO is a complete data optimization solution powered by Permabit's Albireo data optimization software, a powerful sub-file deduplication solution, which is integrated with Linux to provide Linux-based storage OEMs with a complete ready-to-run data optimization capability.

The Albireo VDO Linux kernel module provides deduplication and compression and is exported as a normal block device that can be used directly as block storage or presented through one of the many available Linux file systems (such as ext3 and XFS). Before writing new blocks to the underlying disk device (or at scheduled intervals if post-processing is in effect), duplicates are found with the Albireo deduplication technology. Any unique blocks are then (optionally) compressed and written to disk. Data optimization occurs both through block-level deduplication (not storing duplicate blocks) and through compression, increasing the overall capacity of the underlying device.

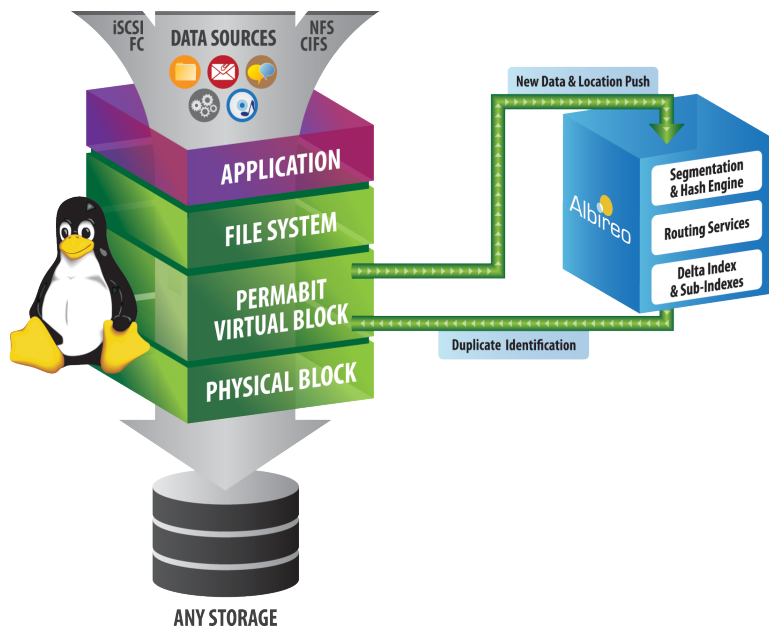


Figure 1: Albireo VDO Architecture

Permabit Albireo — Ready-to-Run Data Optimization for Linux-based SOHO Storage

Albireo VDO provides data deduplication and (optional) compression capabilities for Linux, enabling OEMs to continue leveraging all of their storage solutions' existing features, including existing Linux file systems, virtualization features, and data protection capabilities. Because Albireo VDO uses Permabit's patented Albireo deduplication technology it is able to avoid costs associated with today's Big Storage deduplication solutions that typically require large amounts of system memory, expensive SSD-based index repositories, and proprietary PCI Express cards to achieve even a fraction of Albireo's scalability and performance.

High performance data deduplication provides a truly competitive feature set for mixed applications and use cases. Albireo VDO's straightforward block-level, content-agnostic approach to data optimization

"There are dozens of Linux based NAS vendors that are looking for ways to differentiate themselves from their competitors. If it is available, most of these vendors are frustrated with the sub-par capabilities of the deduplication that comes with the OS they integrate into their NAS. Building the Albireo VDO technology into their devices is an excellent way to differentiate themselves and deliver a truly usable deduplication solution. By offering a turn-key data optimization solution for the Linux platform, Permabit is taking primary storage deduplication to the masses."

— George Crump
Founder, Storage Switzerland

Business Benefits

- Reduces storage growth and associated costs
- Works seamlessly at the server level without need for user intervention
- Improves quality of service and increases customer satisfaction
- Reduces the space required for storage devices Lowers cost of SOHO storage by reducing the amount of data stored

Technical Benefits

- Background server operation delivers storage reduction without user intervention
- Increases efficiency by lowering the effective cost per GB stored
- Enables data growth without expanding the physical space required

provides an effortless solution that is both transparent and non-disruptive to end-user customers. With Albireo's record-breaking performance, Linux-based storage OEMs can extend their deduplication capabilities and out-compete even the Big Storage players by providing data optimization capabilities for mission-critical application storage while effectively leveraging Linux open source to maximize value.

Albireo Index Engine

The Albireo Index Engine can identify duplicates across large storage pools in memory more than 99.95% of the time, eliminating the largest deduplication bottleneck, disk-based fetches. Index lookup averages just 10 *microseconds* – orders of magnitude faster than other deduplication solutions. This enables sustainable ingestion rates of over 1 GB/sec with a single 6-core processor.

Complete Solution for Easy Implementation

Albireo VDO provides a complete, ready-to-run solution for Linux-based storage OEMs, with support for all major Linux versions used by the storage industry. The solution delivers everything needed for deploying enterprise data optimization capabilities, including binaries, open source code modifications, management scripts, and full documentation. The Permabit development team provides full technical assistance to ensure rapid and seamless integration. Typically, Linux-based Storage OEM customers will have an operational solution on the first day of receiving the kit and can immediately begin performance tuning and testing for their specific configurations.

Features and Benefits

Albireo Feature	Albireo Benefit
High Performance Index Engine	Albireo handles heavy workloads and can identify duplicate data in <i>microseconds</i>
Mechanical Disk Optimization	Provides optimal performance for hard disk-based environments
Solid State Disk Optimization	Provides optimal performance for flash-based storage environments
Thin Provisioning	Allows space savings to be realized from the deduplication process, effectively making more virtual space accessible than is physically available
Virtual Block Mapping	Supports deduplication out-of-the box by allowing multiple identical 'virtual' blocks to point to a single physical data location

Find out More

To learn more about Permabit Albireo technology, or to license our products, visit our website at www.permabit.com or call us directly at 617.252.9600.

About Permabit

Permabit is a recognized leader in data efficiency technology. We enable OEMs to leverage their R&D investment, increase margin, accelerate time to market and achieve competitive advantage. Permabit Albireo software massively improves performance and efficiency of data creation, transmission and storage. Solutions built with Albireo are being delivered by leading hardware, software, and service providers.

Minimum System Requirements

- **CPU:** Single dedicated, Intel-compatible CPU core
- **RAM:** 512 MB of dedicated memory
- **Disk Space:** 52 GB
- **Linux Distribution:** CentOS, Debian, Red Hat, SuSE



Ten Canal Park
Cambridge, MA 02141
Phone: 617.252.9600
FAX: 617.252.9977

info@permabit.com
www.permabit.com