

Albireo Virtual Data Optimizer (VDO)

Universal Data Optimization for Linux-based Enterprise Storage Providers



Primary data is growing at over 75% per year and organizations are demanding new solutions that reduce primary storage expense. The market for Linux-based OEM storage has grown 300% over a 3-year period. But things are about to change. By leveraging data optimization techniques, "Big Storage" providers are beginning to deliver highly efficient storage at extremely low realized costs per GB, enabling them to be increasingly competitive. The challenge facing today's value-oriented, Linux-based storage OEM is: how can they continue to leverage open source to compete effectively with larger competitors who have data optimization capabilities?

Data optimization is a capability that Linux-based storage vendors must incorporate to deliver competitive, cost-efficient storage solutions. Albireo Virtual Data Optimizer (Albireo VDO) provides the fastest route to market for block, file, and unified storage solutions. 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 storage OEMs with a complete ready-to-run data optimization capability.

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

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 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.

"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

Highlights

- Provides out-of-the-box data optimization for Linux storage solutions
- Delivers data reduction benefits competitive with the Big Storage vendors
- Avoids requirements for proprietary hash generation hardware cards
- Provides "plug-and-play" integration

Value Proposition

- Provides immediate data optimization capabilities for standard Linux
- Delivers industry-leading throughput
- Allows the OEM to offer significant cost savings to end-user customers
- Extends and complements existing Linux storage features
- Requires fewer hardware resources than any other solution available today

Albireo VDO for Linux Architecture

The Albireo Virtual Data Optimizer 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.

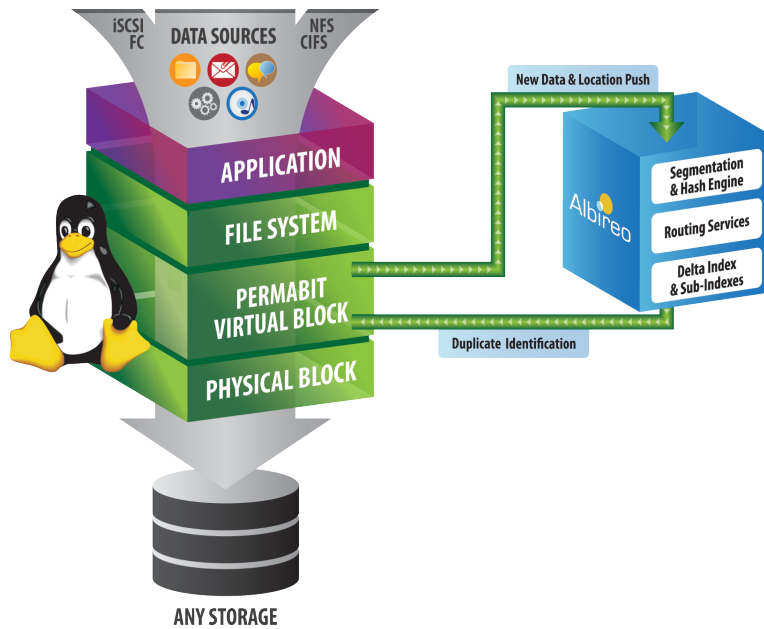


Figure 1: Albireo VDO Architecture

Key Benefits

- Delivers Advanced data optimization technology that substantially reduces effective storage costs.
- Incorporates a virtual block device module that works with existing storage features.
- Uses the Albireo block API to provide unsurpassed performance, scalability, and deduplication efficiency.
- Simplifies storage allocation by introducing thin-provisioning capabilities.

Fast Route to Market

- Delivers a virtual block device driver that “just works” out-of-the-box with existing file systems and data management features.
- Supports existing Linux-based NAS, SAN, and unified storage solutions without modification.
- Supports the most popular embedded Linux distributions for storage devices including CentOS, Debian, Red Hat, and SuSE.

Competitive Advantages

- VDO utilizes the industry-leading Albireo data optimization software to provide the fastest, most efficient, and most scalable data optimization solution on the market today.
- Industry-leading scalability grows with the OEMs’ storage product lines as their scale and capabilities expand to penetrate mid-market and enterprise accounts.
- Albireo VDO enables storage OEMs to implement leading data optimization technology today with an out-of-the-box solution for Linux.

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