

Primary Storage Optimization for OEMs

Permabit Extends Albireo to Storage OEMs

June 2010



Storage optimization is readily available in the secondary storage world. Common data reduction techniques like deduplication cost-effectively address the problem of hugely growing data stores from storage tiering, backup and archive operations. Storage optimization tools like deduplication are very important to this market, which enjoys compression rates of 20X or more.

However, deduplication is not as well mannered in the production environment. This environment's transactional traffic and random I/O requires high performance, high capacity and highly available storage. Garden variety storage optimization tools are not engineered to manage these environments and will not work well in the primary environment. Yet the production environment increasingly needs some level of storage optimization to manage fast-growing data stores and to control attendant performance issues.

Taneja Group believes that this emerging development is an important new approach to intelligently optimizing storage starting with the production environment. We call this technology class Primary Storage Optimization (PSO). PSO is not limited to a single optimization tool, but one very promising PSO technology is deduplication on primary storage. This is not your grandfather's secondary storage deduplication, as it requires a level of performance and integrity that are non-issues in the backup stream. This area is seeing some exciting growth in the marketplace, and the few vendors who already offer it are profiting from its growing popularity. Many other storage OEMs would love to climb on board the PSO bandwagon but primary deduplication is difficult and time-consuming to develop. How then can a storage OEM enter this promising new market of PSO with primary deduplication?

Primary Storage Optimization

With data growing at the rate of 50% per year and up, companies are buying ever more disk just to keep up with new data and multiple backup and replication

copies of existing data. More disk means more data center space, more power and cooling costs, more management headaches, more risky retention practices. Taneja Group recognized early

on that this state of affairs was unacceptable. We originally coined the term Capacity Optimized Storage (COS) as the technology class for safely shrinking stored data in the storage infrastructure. COS addresses uncontrolled data and storage growth, which lets companies control their storage-related CAPEX and OPEX.

As time went on and storage vendors developed COS technologies, it became clear that COS had two sides: one for the secondary storage environment and one for the primary. The two environments have different needs and requirements. Data compression and deduplication for secondary storage are frankly easier to develop, and vendors came out with optimization products for that environment first. The primary storage environment has lagged well behind the secondary in terms of these optimization tools. But production environments are groaning under the same load that bedevils secondary storage, and PSO has become more and more critical for managing primary storage environments.

This need creates real opportunities in the marketplace, and storage consumers are looking to their vendors to provide PSO capabilities to control runaway storage consumption in their production environments. But when their vendors cannot oblige, the consumers are going elsewhere to the vendors who can address their needs.

In an effort to reach this expanding market, many storage vendors are pouring time and money into R&D to develop products to meet the difficult optimization challenges of the primary storage environment. Meanwhile they are lagging behind the storage OEMs that already offer the capabilities.

PSO and Primary Deduplication

This is particularly true for primary storage deduplication, which holds great promise for PSO. Customers are already familiar with deduplication from their secondary environments, and may neither know nor care that deduplication is much more challenging to do in the primary environment. They just want it. And OEMs want to provide it – but they must first develop it or acquire it, and therein lies the rub.

Storage vendors risk losing market share when the R&D process is slow or unsuccessful. And even if they succeed in developing a deduplication product for their storage offerings, they may already be irrelevant in a marketplace that has moved from meat-and-potatoes deduplication into advanced high performance data optimization.

This is because the deduplication market is already moving beyond traditional tactical deduplication, where the deduplication process shrinks the size of backup volumes in order to extend storage capacity. This is a good start but it is only a start. Fast-growing corporate

data requires a strategic data optimization approach that starts in the primary storage environment and moves downstream in the secondary and even tertiary storage infrastructure.

This is why storage vendors who merely introduce simple deduplication will not be able to compete against next generation solutions, which leaves them vulnerable to vendor consolidation and margin pressures. The market is moving towards optimizing primary storage as part and parcel of optimizing the entire storage infrastructure.

Vendors like Oracle/Sun, NetApp, EMC and Dell have already moved beyond secondary deduplication in a tape/VTL-based backup stream, and have developed primary deduplication capabilities. This development not only preserves capacity on secondary storage tiers, but also optimizes primary storage capacity and performance. By operating deduplication at the primary storage tier, vendors offering primary deduplication can maintain storage performance across the entire storage infrastructure. They no longer relegate deduplication to the tape-based backup stream, but apply it to primary storage for dramatic performance gains at the transactional level.

Key Characteristics for Primary Deduplication

Storage vendors cannot merely insert traditional deduplication into the primary

environment and call it good. Instead they should observe the following requirements for developing or securing primary deduplication capabilities.

- *Key #1: Operate across all primary storage architectures.* Primary deduplication cannot be relegated to just block or just file. It must operate across both architectures and also across vendor-specific unified storage.
- *Key #2: Maximize deduplication choices.* Match primary deduplication to the vendor's architecture with flexible choices between parallel, in-line, or post processing. Deduplication in the primary environment must also operate at the sub-file level.
- *Key #3: Scale deduplication for large primary environments.* Terabyte scalability will not cut it in growing primary environments. Primary deduplication technology must include petabyte scalability. Storage vendors should also embed deduplication in storage management software so it becomes an integral part of a highly scalable offering.
- *Key #4: Support high performance and data integrity.* The deduplication tool should not compromise data integrity. It should operate outside of the storage data flow. This avoids performance penalties and does not risk data availability should the deduplication engine go down. The

technology should also eliminate data rehydration and must protect data against the risk of corrupted indexes. Ideally it should be content-aware by supporting both fixed and variable segmentation.

Permabit Albireo

Permabit has stepped up to offer sub-file level deduplication software that is purpose-built for storage OEMs. Permabit engineered their Albireo embedded solution to offer primary deduplication across all standard architectures.

Albireo delivers embedded next-generation OEM data optimization software that dramatically accelerates time-to-market for storage vendors and provides best-in-class performance. Storage vendors can leapfrog the expensive development process and quickly bring primary optimization products to market. These factors let storage vendors leverage their R&D investment, increase margin, and accelerate time-to-market for next-generation unified storage platforms including deduplication for production environments. And Albireo's native flexibility enables storage vendors to preserve competitive differentiators such as performance and advanced data protection schemes.

How Albireo Works

Hashing the Data

The OEM software manages incoming data and pushes data and metadata (i.e. file name, offset, block, LUN) to Albireo. Albireo's content-aware segmentation breaks larger objects into variable or fixed chunks optimally, and uses the SHA-256 algorithm to hash unique content fingerprints and compares them to existing hash keys.

Permabit's patented indexing technology determines if it has already seen the chunk. If an identical data chunk already exists then the Albireo API asynchronously (in the case of a parallel integration implementation) pushes previous placement information to the OEM software for unification. Albireo adds a pointer to its location. If the data chunk is a duplicate, then the OEM's storage application modifies its storage tables.


High Performance and Scalability

Albireo stores the keys in a key index table. The index sits in memory and enables Albireo to check for existing chunks extremely quickly (we are talking microseconds here). The index can be distributed across grid-based nodes and there is no additional latency.

This index, which can be a severe bottleneck on other deduplication systems, is highly available and resource-efficient using the Albireo High

Performance Index Engine. This technology maintains index entries at just 3.5 bytes in memory and around 40 bytes on disk space per entry. The compact indexes enable Albireo to reduce/eliminate disk access during deduplication processes, which enables high performance and operates efficiently in environments with constrained memory.

Albireo is also highly scalable well into petabyte levels. Existing primary deduplication vendors struggle to scale this far. Permabit has combined its extreme scalable architecture with high deduplication performance, making their storage OEM customers highly competitive in a marketplace that increasingly demands highly scalable storage.



The tiny size of the key entries enables the index engine to return exceptionally fast matches within microseconds. Data ingestion rates are at or better than 140 MB/sec/core. Cluster performance scales from a single quad core processor at 400 MB/sec and a dual quad core processor at 800 MB/sec.

Protecting Integrity

Albireo does not operate inside the storage data flow, which is a deliberate architectural decision on Permabit's part. Operating outside the flow ensures that Albireo will not degrade the storage application's performance. And since Albireo is not in the read path, there is no risk to data availability should Albireo not be operating.

This architecture also ensures that Albireo does not have to cause a performance hit from data reads. This eliminates rehydration processes, which is better done by the OEM's storage stack and core technology. And Albireo also comes complete with self-healing capabilities to protect the primary optimization process.

Extreme Flexibility

Albireo supports both fixed and variable block deduplication. Traditional deduplication frequently uses fixed block sizes to minimize processing load and

simplify their development process. Albireo does support this approach, which lets OEMs who are already developing fixed deduplication to leverage their R&D investment.

However, variable block deduplication is an important and valuable option in the Albireo SDK enabling more effective data reduction. In this scenario, Albireo analyzes data using content awareness technology. Albireo scans incoming data using content filters. These filters, or scanners, identify data from specific applications or data using file-based objects, such as MS Office documents, graphics, zip files, backup files, or VMware images. Knowing the content type enables Albireo to more efficiently deduplicate incoming data. Storage vendors may choose to use this feature.

Complete SDK for Easy Implementation

The Albireo SDK enables OEMs to quickly embed flexible and powerful primary optimization offerings into their core storage application products. Albireo enables OEMs to embed true PSO-based deduplication for block, file and unified storage platforms at a sub-file level. Even within these protocols OEMs may choose different approaches such as fixed or variable sizes for deduplication. OEMs may also integrate the tool at any point including parallel, inline, or post-process.

The SDK also comes with simple API integration for fast development and with

extensive documentation and developer support. Together these powerful and flexible development options enable storage OEMs to come to market quickly and confidently with competitive PSO offerings.

Key Takeaways

Albireo primary deduplication – or Dedupe 2.0 as Permabit calls it – offers serious storage cost savings across the storage infrastructure. With the Albireo SDK for storage OEMs, Permabit extends the high performance Albireo optimization engine into high stakes development environments.

- Key Takeaway #1: Albireo High Performance Index Engine enables low resource utilization, massive scaling, and incredible performance.
- Key Takeaway #2: Content segmentation is a valuable option for OEMs, who can choose variable-chunk segmentation with content aware scanning.
- Key Takeaway #3: The SDK enables easy implementation with the OEM's storage application. Flexible choices include inline, parallel or post-process integration; and block, file and unified storage. Simple APIs and extensive support round out the SDK environment.

Taneja Group Opinion

The storage market's needs around primary optimization are being driven by production environments. These environments must have high performance, high availability and high capacity in their storage systems. The same deduplication that works on backup streams is wholly inadequate to deduplicate active data in these demanding environments.

Primary Storage Optimization has exceptionally high value for managing the entire storage ecosystem. We consider it a game changer because it literally re-defines the competitive landscape for storage.

OEMs that are still sniffing around the edges of deduplication will suddenly have to provide the tool in secondary and primary environments in order to stay

competitive. The Albireo embedded development environment enables storage vendors to leverage Permabit's R&D investment in deduplication, increase margin and accelerate time to market for next-generation unified storage platforms.

With the release of Albireo embedded deduplication software for OEMs, Taneja Group identifies Permabit as a leading PSO vendor. Permabit has a long history of optimizing primary storage using deduplication, replication and compression. Now it has built on those strengths to offer the Albireo Data Optimization SDK to storage OEMs. OEM's need it today and Permabit is ready to provide it.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.