



/ Technology Brief /

FlexiRemap Technology

Software that unleashes the true performance of flash storage

Enterprise Data-Access Performance Bottleneck

Enterprise-level IT departments require high-performance systems to handle the data-access traffic of today's massive volumes of data and growing diversity of applications. It's not surprising that in such environments, conventional storage systems become the bottleneck for data access; they simply can't handle so much data instantaneously. Simply adding more storage without making fundamental changes is likely to defeat performance scaling and quickly render the IT budget unaffordable.

Modern All-Flash Arrays vs. Conventional Disk Arrays

Modern all-flash arrays are designed specifically to address the demand for scaling storage performance. All-flash arrays, sometimes referred to as "solid-state appliances" (SSAs), belong to a new category of enterprise-storage systems comprised of flash memory, or, more popularly, solid-state drives (SSDs). In contrast to traditional hard disk drives (HDDs), which typically handle hundreds of random I/O requests per second, SSDs are generally able to handle tens of thousands of random I/O requests per second. Plus, SSDs are much more energy efficient, consuming only half to one-third of the electrical power, and the latency can be as low as one fiftieth as compared to HDDs. These

advantages make all-flash arrays a better choice for storage performance.

However, simply replacing HDDs with SSDs in a disk array does not create an all-flash array. This kind of "fix" does not deliver the best storage performance, nor can such performance be sustained over time. The RAID algorithms originally designed for HDDs and commonly used in conventional disk arrays actually limit the overall performance and often shorten the lifespan of SSDs.

What Is FlexiRemap Technology?

FlexiRemap technology is flash-oriented software developed by AccelStor. This software features an architectural re-design tailored for the fundamentals and internal workings of flash memory. Unlike RAID algorithms, FlexiRemap remaps the data whenever beneficial before passing it to the underlying flash memory, avoiding unnecessary overhead and extending the lifespan of SSDs. Such remapping is especially crucial for random-write access patterns, which comprise most enterprise-data traffic and are a major cause of storage-access performance bottlenecks. Plus, with data protection built in, FlexiRemap provides the most desirable benefits of RAID algorithms, without the drawbacks.

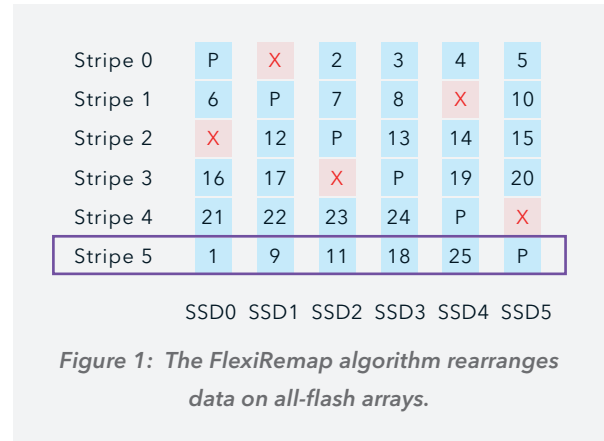
How Does FlexiRemap Technology Work?

Before understanding the methodology of FlexiRemap, it's essential to know how flash memory works. Flash memory is an electronic non-volatile storage medium that can be electrically erased and reprogrammed. In a SSD or all-flash array, NAND flash is usually adopted. NAND flash is made up of blocks, and each block is formed by pages on which data will be written. Unlike HDDs, NAND flash has pages with data that cannot be directly overwritten with new data. Existing data must first be copied to a new location, while data in the original location is erased; this is called "garbage collection." The original location is then available for new data. Reading and programming are performed on a "page" basis, but erasure can only be done in "blocks." This difference leads to a storage-performance challenge, especially when new data is saved on used SSDs. The random write operation presents an even bigger barrier to storage performance.

FlexiRemap Technology

FlexiRemap is a unique technology that can handle these storage-process challenges. To accelerate random writes, FlexiRemap adopts an efficient strategy to speed up data processing. FlexiRemap rearranges all new data to sequential block addresses before passing it to SSDs, which then receive and process the sequential data. In sum, FlexiRemap reduces the workload of SSDs by enabling them to directly write the sequential data into consecutive pages.

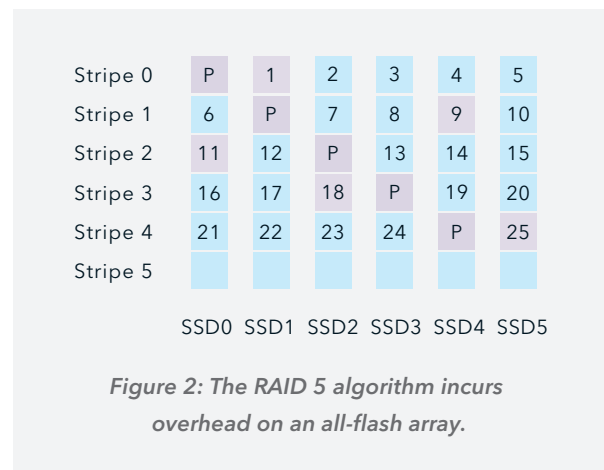
In the example below, new data does not need to be stored at the original location, and there is no calculation needed to produce new parity. All new data is placed at the new location. Thus, data-access performance is not impacted. Those blocks marked with an X (Figure 1) won't be used and are cleaned by FlexiRemap's garbage collection function; the new space is then released for upcoming data. This eliminates all the penalties of random writes.



FlexiRemap vs. RAID

In storage systems, drive failure usually means simultaneous loss of read and write operations. RAID was developed to recover from a failed drive by segmenting across multiple physical drives. In most RAID algorithms, calculating and storing redundant data can be used to reassemble the information lost on one or more failed HDDs. However, the RAID algorithm is not the best solution for all-flash arrays.

In the following example, six SSDs are grouped as a RAID 5 drive. When the hosts update new data at locations 1, 9, 11, 18 and 25 (Figure 2), the current data and parity codes need to be retrieved and recalculated with new data to produce new parity codes, which are then put back into their original locations. The processing time for the additional reads and writes is longer, and data access is slowed.



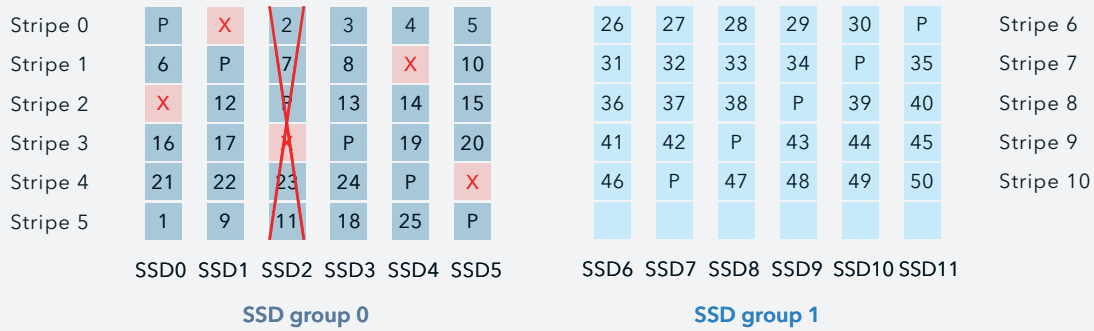


Figure 3: FlexiRemap features built-in data protection.

With RAID 5, random write operations are the heaviest, because they usually involve multiple reads and writes to fulfill a single request. To accelerate system performance, a more efficient method is required.

What Are the Advantages of FlexiRemap Technology?

Enterprise IT is always concerned about data protection, storage performance, and system durability. AcceleStor FlexiRemap technology answers these needs by providing the following high-value features:

▶ Data Protection

Unlike the data redundancy of traditional RAID, FlexiRemap provides a proprietary algorithm for data protection. It divides SSD drives into two groups of independent drives (Figure 3). Data is evenly distributed across SSDs of the two groups, with fault tolerance of single-drive failure.

For example, as shown below, if one SSD fails in SSD group 0, this group will then enter read-only mode to avoid the heavy loading of further writes. This decreases the chance of failure on the second drive in the same group. That is, FlexiRemap prevents data corruption by avoiding I/O writes to the group with the failed SSD, in this case group 0.

▶ High Performance

When the system attempts to update data in flash memory, the RAID algorithm reads the current data and parity in a stripe first. After the calculation, the new data and parity are written on the original striping location. This process generates two read operations for the previous data and parity and two write operations for the new ones. RAID is unable to efficiently handle a large load of I/O requests, particularly in random write mode, resulting in poor storage performance.

FlexiRemap features an exclusive data remapping algorithm based on the physical nature of NAND flash memory. Instead of writing the new data to the original location, FlexiRemap organizes the inbound writes to sequential logical addresses. 4KB random write performance is significantly improved, and total performance is dramatically enhanced for 1,020K IOPS as compared with RAID algorithm, which only generates around 5K IOPS to 78K IOPS (Figure 4).




Figure 4: FlexiRemap vs. RAID - Performance in IOPS

Maximization of SSD Lifespan

Unlike HDDs, SSDs have a finite number of program-erase (P/E) cycles. Frequent programming to a particular SSD will often lead to bad blocks and shorten the lifespan of the device. FlexiRemap's global wear leveling algorithm guarantees that data are evenly distributed across each SSD in every group. When new data is rewritten on the storage system, FlexiRemap checks the mapping table record in advance to ensure that each SSD stores a similar amount of data.

Benefits of evenly distributed data across SSDs in groups:

- Extended SSD endurance, due to the limited P/E cycles of flash memory.
- Better performance. Since every SSD stores a similar amount of data, there is less chance of traffic bottlenecks.

RAID does not address repeated programming and wear on flash memory. The nature of RAID is to generate repeated operations on the same locations; the frequent erase and write operations cause some SSDs to endure more loading than others. As noted below, test results reveal a difference of about 1.5 times the number of requests received between the most-used and least-used SSD in a RAID system. Under the management of FlexiRemap, however, each SSD is afforded a similar number of requests (Figure 5).

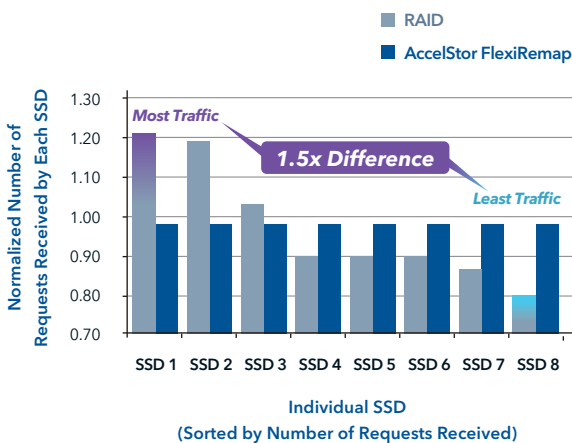


Figure 5: FlexiRemap vs. RAID - Number of requests received on SSDs

Conclusions

AccelStor FlexiRemap is flash-oriented software. Its unique algorithm prevents unnecessary I/O overhead, while providing excellent storage performance, enhanced data protection, and maximized lifespan. FlexiRemap ensures that AccelStor products meet the stringent performance demands of all kinds of mission-critical applications. Conventional HDD-based disk arrays or RAID systems are good for manageable working sets and sequential writes. However, when it comes to consolidated missions and high-performance analytics, all-flash arrays with FlexiRemap technology are the best solutions for handling intensive random writes and additional I/O workloads (Figure 6).

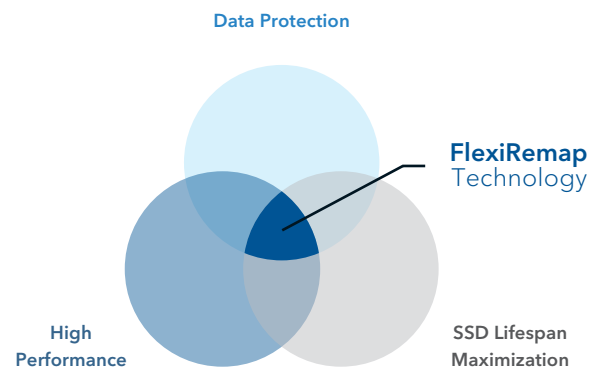


Figure 6: FlexiRemap advantages

Note: Test results mentioned in this document are generated by AccelStor and may vary in different test environments.

www.accelstor.com
inquiry@accelstor.com

