

Data Domain SISL™ Scaling Architecture

Technical Brief

Abstract

While tape has been the dominant storage medium for data protection for decades because of its low cost, it is steadily losing ground to disk-based deduplication storage systems.

The CPU-centric design of Data Domain systems takes the pressure off of disk accesses as a bottleneck. Over the last 20 years, CPUs have improved in speed by a factor of millions, while disks have improved by about 10x. CPU vendors appear poised to continue these benefits well into the future. It is reasonable to imagine that each doubling of cores could mean Data Domain systems can improve speed by about a factor of 50%.

In SISL (Stream Informed Segment Layout), Data Domain has developed a proven architecture to deliver high throughput deduplication storage systems with economical storage hardware. Over time, this will allow the continued scaling of CPUs to add direct benefit to system scalability.

Data Domain SISL Scaling Architecture

Table of Contents

ABSTRACT	1
CPU-CENTRIC, FAST DEDUPLICATION WITH SMALL SYSTEM HARDWARE	3
THE CHALLENGE: FINGERPRINT RECOGNITION AND SPEED ..	3
SPEED PER DISK	3
APPROPRIATE CAPACITY	3
TYPICAL RESULTS	4
DATA DOMAIN SISL	4
FUTURE SCALABILITY	6
CONCLUSION	6

CPU-centric, Fast Deduplication with Small System Hardware

While tape has been the dominant storage medium for data protection for decades because of its low cost, it is losing ground to deduplication in disk systems. Deduplication is an approach that can deliver an order of magnitude greater data reduction than traditional compression over time. This should mean that a deduplication system needs fewer disks. It should also mean that configured costs of a disk storage system are comparable to tape automation.

Customers are sometimes surprised that most emerging deduplication products use a lot more disks than expected. At its heart, deduplication is a disk-intensive activity. The conventional way to increase disk systems' performance is to use more disks and to use faster, more expensive disks. This can spread the load across the relatively low per-spindle access and transfer speeds. Unfortunately, using this approach in a dedupe array can quickly make it more expensive than the tape library against which it will be compared. When using low-cost, high-capacity SATA drives, it would also mean most of the capacity would be wasted, because each disk comes with a lot of space. By adding disks, the customer could pay for a lot of unnecessary capacity.

Data Domain solved this problem early on with the SISL (Stream-Informed Segment Layout) scaling architecture within DD OS. It optimizes deduplication throughput scalability and minimizes disk footprint by minimizing disk accesses. Doing so allows the system throughput to be CPU-centric. Speed increases directly as CPUs improve in performance. This paper explores the Data Domain SISL approach and compares it to emerging alternatives.

The Challenge: Fingerprint Recognition and Speed

Speed Per Disk

The basic algorithm for deduplication is to break the incoming data stream into segments in a repeatable way and compute a unique fingerprint for the segment. This fingerprint is then compared to all others in the system to determine whether it's unique or redundant. Only unique data is stored to disk. To its clients, the system appears to store the data in the usual way, but internally it does not use disk space to store the same segment repeatedly. Instead, it creates additional references to the previously stored unique segment.

For good data reduction the segments should be small to maximize the data reduction effect. Smaller segments are more likely to be found in more places. But smaller segments mean there are more segments and therefore more fingerprints to compute and compare. Data Domain deduplication technology uses a relatively small segment size (8KB average, varying in size). This provides the best deduplication results and provides a flexible, application-independent store. After identifying unique segments, local compression (e.g. LZ, gzip) is applied and only that data is stored to disk.

The fingerprint index in this kind of approach can be more than an order of magnitude bigger than RAM. As a result, it is typically stored on disk. So in doing an index lookup, the system will typically perform a disk read for each incoming segment. That's where things can get bad.

This would mean that for 100 MB/second throughput, a typical hash-based system would need about 100 disks. Here's why. A 500GB SATA disk can sustain approximately 120 disk accesses for index lookups per second¹. For a segment size of 8 KB, that means

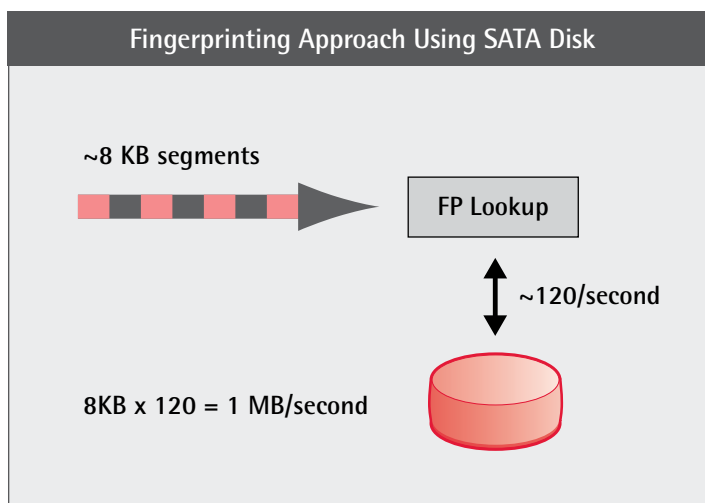


Figure 1. In a fingerprinting approach using high-capacity, low cost SATA disk, random lookups of fingerprints for segments that average 8KB, limit throughput to about 1 MB/second per disk.

a single disk could support an incoming data rate of about 120 x 8 KB, or about 1 MB/sec. To go faster, more disks would be required to spread the access load. Such a system would be too expensive to compete with tape in most situations.

Simple alternatives are sub-optimal. One resolution could be to use a much bigger segment size on average, but that would provide significantly worse deduplication effects and again make the system uncompetitive with tape automation on configured price. Or faster Fibre Channel disks could be used – but for twice the speed, they often cost 3x-5x more per gigabyte.

Appropriate Capacity

If 100 disks is too much, how much is enough? A dedupe process with traditional compression and a conventional onsite retention period generally achieves > 10x aggregate data reduction. A common onsite retention period stores 10x the amount of data in the starting set (for example, a weekly full backup with daily

¹ E.g. a Seagate Barracuda-ES 500GB drive at 7.2K RPM (model ST3500630NS) supports an average read seek time of 8.5 msec, so would support 117 read seeks / second).

incrementals, over two months). So it is reasonable to assume the dedupe store should be about as big as the starting set of primary data being backed up to it.

If performance is the limiting factor in matching a data set to a backup window, the weekly full backup and backup window often determines how fast a system is needed. The most challenging throughput configuration is when all full backups are on one weekend day. If using a 16 hour weekend window to allow for a restart on finding a problem, at 100 MB/sec., a starting data set would have to be less than 5.75 TB (16 hours times 100 MB/sec.). A dedupe storage system using 500GB drives should only require 12 drives for storage, apart from RAID, spares etc. Even with RAID 6, adding 2 parity disks for a total of 14, this would mean 100 MB/sec / 14, or about 7 MB/second/disk. Projecting forward, when each disk can store 1 TB, half as many disks would be needed, so they'd have to go twice as fast to stay on the curve dictated by capacity. Disks themselves won't do this.

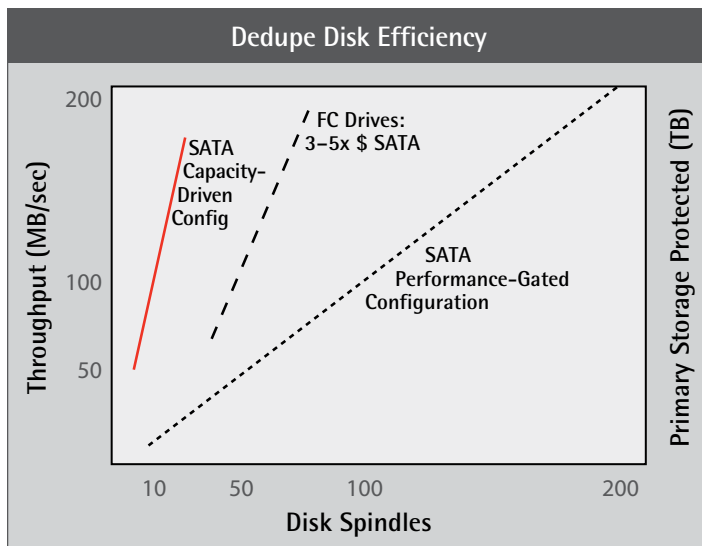


Figure 2. In a scalable deduplication system, fingerprints need to be indexed in an on-disk structure. To achieve speed, the system needs to seek to disk to determine whether a fingerprint is new and unique, or redundant. With current average seek speeds and a small segment size for good data reduction, more disks are required to get to speed than the desired number for capacity (above assuming 500 GB, 7.2k RPM SATA disks and an average 8KB segment size).

Typical Results

There's obviously a big difference in price and manageability between a configuration that needs 12 drives and one that needs 100 drives. As disks get more capacity and more data is addressed, this difference only becomes more profound. A system that scales speed based on disk drive count will waste enormous amounts of capacity. It would also create a need to manage many more drives than necessary. It will clearly cost more.

In addition, if data is spread across drives, accesses would be scattered across all disks and not optimized for read throughput. The store would be fragmented, so the read performance required

for recovery or streaming a copy to tape would be compromised. Modern LTO-4 drives require speeds approaching 60 MB/sec/stream or they will shoe-shine and slow down dramatically. At 1 MB/sec/disk, the system would need at least 60 disks (30 TB with 500GB disks, or 45 TB with 750 GB disks) to perform well enough on reads, not counting additional disks for spares, etc.

To succeed, a dedupe system would need to overcome the challenges of comparison across such a large logical space:

- ▶ minimizing RAM;
- ▶ using only enough inexpensive SATA disks to support the capacity required by the bandwidth and retention of the system, enabling CPU improvements to directly improve throughput;
- ▶ avoiding fragmentation in the disk store to allow fast read/recovery/copy speed.

Data Domain SISL

Data Domain's SISL technology includes a combination of approaches that solve these problems. First, it identifies 99% of the duplicate segments in RAM, inline, before storing to disk. Second, it stores related segments and fingerprints together, so large groups can be read at once. With these patented² techniques, Data Domain can utilize the full capacity of large SATA disks for data protection and, without increasing RAM, minimize the number of disks needed to deliver high throughput. In the long term, SISL allows DD OS-based system performance to track dramatic CPU speed improvements.

Uniqueness Identification

SISL includes a series of techniques performed inline in RAM, prior to data storage to disk, for quickly filtering both new unique segments and redundant duplicate segments: the *Summary Vector* and *Segment Localities*.

The Summary Vector is an in-memory data structure used by DD OS to quickly identify new, unique segments. Identifying new segments saves the system from doing a lookup in the on-disk index only to find the segment isn't there.

Based on a Bloom filter, the Summary Vector is a bit array in RAM initially set to all zeros. When a new segment is stored, a few bit locations in the array are set to 1. The locations are chosen based on the fingerprint of the segment. When a subsequent segment arrives, its chosen locations are checked. If any locations are 0, the system knows conclusively that the segment has not previously been stored, and it can stop looking.

The Summary Vector is not, by itself, sufficient for declaring a segment redundant. A small fraction of the time, typically less than 1%, all of the chosen locations have been set to 1 by different segments even though the new segment is unique. When this happens, the system needs to rely on other mechanisms to conclude recognition.

² These techniques are covered by patents 6,928,526; 7,065,619; and other patents pending.

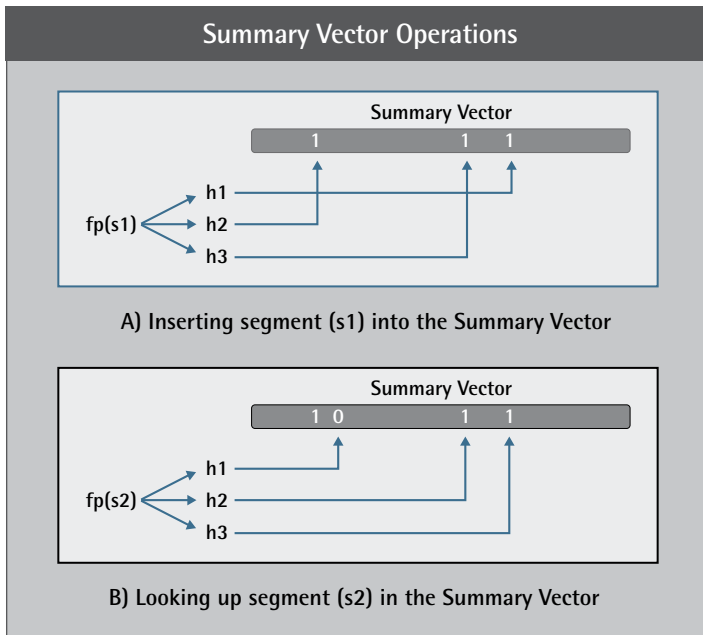


Figure 3. The Summary Vector can identify most new segments without looking up the segment in the on-disk fingerprint index. Initially all bits in the array are 0. On insertion, shown in (a), bits specified by several hashes, h1, h2, and h3 of the fingerprint of the segment are set to 1. On lookup, shown in (b), the bits specified by the same hashes are checked. If any are 0, as shown in this case, the segment cannot be in the system.

Redundancy Identification and Read Speed

The problem with finding duplicates exclusively with index lookups is that every disk access only retrieves one segment. One key to disk efficiency is to retrieve many segments with each access. Generally, a given small segment of data in most backups will tend to be stored sequentially with the same neighboring segments before it and after it most of the time. The Data Domain system stores these neighbors together as sequences of segments in units called **Segment Localities**, which are packed into containers. The Data Domain file system is a log structured system, at the heart of which is the log of containers storing Localities.

A Locality keeps segments close together on disk when they are neighbors. The system can access all the fingerprints or a whole Locality with a single disk access. This means many related segments or segment fingerprints can be accessed very efficiently.

The SISL Process

When considering a newly arrived segment, the system first checks the Summary Vector. If the Summary Vector indicates the segment is new and needs to be stored, the system, informed by the stream itself, adds the segment to the current Segment Locality in the order it appears in the stream for later storage to disk. Otherwise, the segment is probably a duplicate and the system looks in a fingerprint cache held in RAM.

In backup/restore, most segments are accessed very infrequently. A full backup passes an entire file system serially through the backup process and references huge numbers of segments that won't be referenced again until the next full backup. Therefore,

a conventional caching strategy based on data recently accessed would not be effective. Instead, with SISL, when a segment is not found in the cache, the system looks it up in the on-disk index and then prefetches the fingerprints of an entire stream-informed Locality into the cache. The vast majority of the following segments in the incoming backup stream are then typically found in the cache without further disk accesses.

Together, these techniques and others make it possible to find duplicate segments at high speed in an application-independent way while minimizing array hardware. It requires neither huge amounts of RAM nor large numbers of disk drives. The Summary Vector avoids pointless index lookups for new segments. Localities organize segments and segment fingerprints on disk so each disk access fetches data that is relevant for a sequence of segments. Prefetching brings these localities into the cache so that most duplicate segments are found at high speed in the cache. On long-running experiments with real backup data, these techniques together eliminate up to 98% of the disk reads and deliver balanced performance using the full capacity of low-cost SATA disk drives.

SISL Performance in Data Domain Systems

Data Domain systems have always incorporated revisions of SISL technology. How are they lining up against the desired speed and scalability chart above?

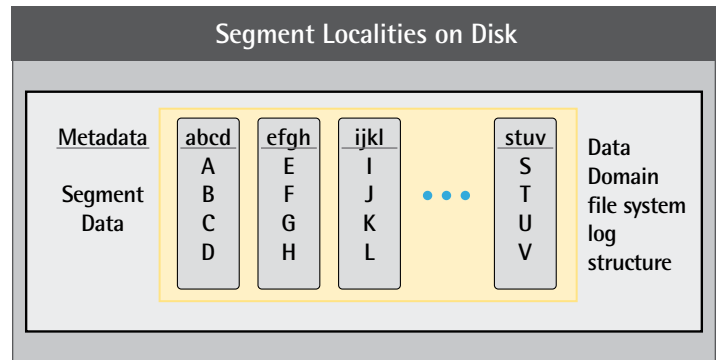


Figure 4. New data segments for a backup stream are stored together in units called Localities which, along with their fingerprints and other metadata, are packed into a container and appended to the log of containers. The fingerprints for the segments in the Localities are kept together in a metadata section of the container, along with other file system structural elements. This keeps fingerprints and data that were written together close together on disk for efficient access during writes when looking for duplicates and during reads when reconstructing the deduplicated stream.

As of this writing, a DD565 includes 15 500GB SATA drives in its base appliance chassis. It supports 100 MB/sec. read and write throughput in that configuration. Of the 15 drives, 1 is a spare, and 2 are for RAID 6 parity, so there are 12 useable drives for user data and metadata. This delivers the optimal match of capacity and performance for the 100 MB/sec. case discussed above.

The higher-end DD580 can deliver over 200 MB/sec throughput. Based on the assumptions above, it should call for 2x the capacity of the 560 base model, or 11.5 TB useable. The most popular base

configuration includes 12 TB (24 500 GB drives after RAID 6 and spares) useable for user data and metadata. This delivers balanced backup and restore speeds with a minimum of low-cost, high capacity drives.

Future Scalability

SISL in DD OS takes the pressure off of disk accesses as a bottleneck, so the remaining system design is CPU-centric. That's a good thing. Over the last 20 years, CPUs have improved in speed by a factor of millions, while disks have improved by 10x or so³. CPU vendors appear poised to continue these benefits well into the future. It is reasonable to imagine that each doubling of cores could mean Data Domain systems can improve speed by about a factor of 50%.

For example, the recent introduction of the DD690 appliance confirms the 50% performance improvement expected by doubling the CPU cores when using the SISL architecture. The new 8 core DD690, comprised of two quad core processors, delivers up to 170 MB/s (or ~ 600 GB/hour) for inline deduplication of a single large backup job as well as up to 1.4 TB/hour aggregate throughput over many backup policies. The entry 12 TB (24 500 GB drives after RAID 6 and spares) configuration delivers over 300MB/s of balanced backup and restore speeds with a minimum of low-cost, high capacity drives.

It is also straightforward to imagine methods to integrate controllers to create multi-node aggregates for larger datasets. A fingerprint-based dialog between controllers can be lightweight. Some approaches exist today that prove the merit in this scalability approach, but they do not solve the price/performance challenges addressed by SISL. Instead, they just make it possible to connect all those spindles needed for good performance. By solving those problems at the outset, the Data Domain architecture provides the foundation for cost-effective aggregation in the future.

Conclusion

Deduplication can help gain an order of magnitude more data reduction than traditional local compression such as LZ or gzip for backup users. But to be cost-neutral when compared to a tape automation system, the deduplication system needs to be CPU-centric and require very few disk accesses, so that it can be built with the minimum number of low-cost, high-capacity disks.

In SISL, Data Domain has developed a proven architecture to deliver high throughput deduplication storage systems with economical storage hardware. Over time, this will allow the continued scaling of CPUs to add direct benefit to system scalability.

³ http://seagate.com/docs/pdf/whitepaper/economies_capacity_spd_tp.pdf

Data Domain
2421 Mission College Blvd.
Santa Clara, CA 95054
866-WE-DDUPE; 408-980-4800
sales@datadomain.com
22 international offices: datadomain.com/company/contacts

Copyright © 2008 Data Domain, Inc. All rights reserved.

Data Domain, Inc. believes information in this publication is accurate as of its publication date. This publication could include technical inaccuracies or typographical errors. The information is subject to change without notice. Changes are periodically added to the information herein; these changes will be incorporated in new additions of the publication. Data Domain, Inc. may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time. Reproduction of this publication without prior written permission is forbidden.

The information in this publication is provided "as is". Data Domain, Inc. makes no representations or warranties of any kind, with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Data Domain and Global Compression are trademarks of Data Domain, Inc. All other brands, products, service names, trademarks, or registered service marks are used to identify the products or services of their respective owners.
WP-SISL-0508