WHITE PAPER

HIFN

PROTECT · SECURE · ASSURE

Hardware-Based Data Deduplication

AN INTRODUCTION

Bob Farkaly

1.1 Introduction

1.1.1 The Economic Drivers of Data Deduplication

The adoption and acceptance of data deduplication in storage was a natural byproduct of the acceptance of disk-based backup and recovery. Storage managers, hungry for more reliable backup and tired of the arcane processes surrounding tape, made backing up to disk a datacenter imperative almost overnight.

Storage managers quickly learned that recovering data from disk – the real point of backup – was far easier, faster and more convenient than recovering data from tape. And the longer backup data was retained on disk, the greater the odds that data could be recovered at disk speed. But while disk reduced or eliminated the need for tape, many of backup's operational processes remained in place.

The strategy of full, incremental and differential backup translated nicely from tape to disk media. The downside, however, was that more and more disk storage was required to hold all that data. And since disk storage is not inexpensive, consumes power, requires cooling and occupies valuable datacenter real estate, something had to be done.

The first attempts to reduce the physical capacity requirements of disk-based backup storage revolved around compression. But software-based compression was very slow and consumed great quantities of processor resource. So backup to disk vendors took a cue from tape drive manufacturers and equipped their products with hardware-based data compression from Hifn. Hifn's hardware-accelerated compression enjoyed rapid adoption by the manufacturers of most backup to disk appliances and systems.

But the industry quickly recognized that nominal 2:1 compression ratios were only enough to slow the growth in the number of disk drives and their expensive ecosystems. What was needed was a technology that was capable of dealing with multiple full and incremental backups without a corresponding increase in the size of the underlying storage hardware infrastructure. In other words, a paradigm change was required. That paradigm changing technology proved to be data deduplication, giving rise to a new industry segment.

1.1.2 How Data Deduplication Works

Data deduplication simply means the elimination of redundant data. In the deduplication process, data is scanned and unique patterns are identified, assigned corresponding unique data "fingerprints", indexed and retained. Duplicate copies of data already "fingerprinted" are deleted, leaving only one stored copy of each unique data pattern along with its corresponding fingerprint.

The fingerprints for all unique data patterns along with their corresponding full data representations are retained to enable reconstituting the deduplicated data for subsequent use. Deduplication is able to reduce the required storage capacity since only unique data is stored.

At the heart of the process most frequently employed to identify unique data patterns and assign fingerprints is a computational process called a "checksum". If a calculated fingerprint already exists, then the deduplication software knows that a data pattern has been previously processed and stored. Therefore, only a pointer to the previously stored data needs to be saved. If the fingerprint is new, then the block is unique, the new fingerprint is added to the index and the unique block of data is stored.

Data deduplication product manufacturers employ cryptographic hash functions to perform checksum processing because of their ability to produce unique fingerprints from unique data patterns.

It is the computationally intensive task of checksum processing (performing the cryptographic hash calculation) that is at the heart of Hifn's approach to hardware-based data deduplication. Hifn's data deduplication process employs Application Specific Integrated Circuits (ASICs) to offload checksum/fingerprint processing from the general purpose processors and software used by deduplication vendors who do not use Hifn technology.

Data deduplication is being commercially deployed in both storage and communication products. In storage products, deduplication is employed to reduce the amount of physical capacity required to store data. In communication products, deduplication is employed to reduce the amount of data that has to be transferred over a WAN connection to produce and maintain synchronized copies of data stored at other physical locations.

The value propositions of data deduplication are therefore:

- In storage applications, data deduplication reduces the storage capacity required for a given amount of data, thus reducing acquisition costs, rack space requirements, power requirements and heat dissipation (reducing cooling needs). Data deduplication enables more data to be stored on a given amount of disk, lowering storage management costs by reducing the amount of physical capacity requiring management.
- In communication applications, data deduplication is the optimal (and is currently the only) technology for transporting large quantities of data over low-bandwidth communications facilities, enabling remote data applications that include remote office backup consolidation and the creation and maintenance of disaster recovery archives.

Hifn's hardware accelerated data deduplication can be employed in a wide range of storage and networking applications by manufacturers currently shipping software based deduplication as well as by manufacturers seeking to incorporate deduplication into their products.

1.1.3 Why Specialized Data Deduplication Hardware is an Imperative

Today, the computational heavy lifting of data deduplication is handled by software in general purpose computers, just as compression once was. The "lifting" is indeed heavy and this makes the process slow. Hifn changed the storage landscape by replacing the computational "heavy lifting" of software data compression with specialized hardware.

Current products featuring software-based data deduplication in the backup data path ("in-line") are capable of accepting data streams measured in the 10's of megabytes per second as opposed to the 100's of megabytes per second easily achieved by non-deduplicating products.

This performance shortfall has limited the size of the market for data deduplication systems because the #1 concern of storage managers continues to be the backup window – or the length of time it takes to complete the backup. Indeed, a reduction in backup performance frequently outweighs the benefits of adding deduplication technology.

The IT industry has long accepted the use of specialized hardware to accelerate computationally intensive processes. For instance, RAID controller parity calculations and data compression—once done in software—are now routinely performed in customized silicon.

It is this technology—the acceleration of the data deduplication process in customized silicon—that Hifn is now delivering.

1.2 Data Deduplication Methods and Terms

1.2.1 File vs. Block

Data deduplication can be performed at the file or block (sub-file) level. File-level deduplication—also referred to as "Single Instance Stores", or "SIS"—eliminates identical files either within or across systems. For example, Windows servers all have the same executables, .dlls and so forth. Storing them only once can shrink the amount of disk required for their storage and backup.

File level deduplication is not considered to be the most efficient method of performing data deduplication. This is because files must be identical in their entirety to be eligible for deduplication. Two 100 megabyte PowerPoint presentations having only a title page change would not

be considered duplicates and each file would be stored in its entirety even though the files were identical down to the title page difference.

Block-level deduplication works below the file level. As its name implies, files are broken down into segments -- blocks -- using either fixed-length block sizes or logical, variable-length blocks and searched for patterns of ones and zeroes that are examined for redundancy against previously stored information. In the PowerPoint example described in the preceding paragraph, block-based deduplication would identify the blocks within the PowerPoint files that had been previously indexed and would store a pointer to the fingerprint of those blocks. Finally, the deduplication process would store the difference(s) between the blocks constituting the PowerPoint files.

Hifn's technology performs fixed length, block-level deduplication.

1.2.2 In-Line vs. Post Processing

In-line data deduplication takes place in real time and "on the fly" by computational intelligence located in the data path.

Today, in-line data deduplication requires a great deal of processing power to keep from becoming a bottleneck in the data path. Large quantities of CPU resources are required to perform the checksum algorithmic processing used to calculate the data fingerprints. An index of all data is retained to enable reconstituting the deduplicated data for subsequent use.

"Post-processing" data deduplication is a technique that performs deduplication between the time data is written to a non-deduplicating disk and the time the deduplicated data is finally committed to permanent media. Post-processing data reduction appliances do not reside in the data path.

In-line data deduplication approaches are considered superior to post-processing approaches because data only needs to be "handled" once. In the backup application, in addition to the second data "touch", a post-processing approach requires the full backup to be stored on disk in a non-deduplicated form in order to achieve reasonable backup performance. This approach defeats the economies of disk storage promised by deduplication.

Hifn's technology enables in-line data deduplication to be performed directly in the data path without compromising performance.

1.2.3 Data Deduplication and Data Integrity

Hifn technology employs a hardware-based Secure Hash Algorithm (SHA-1) that produces a 160-bit hash "fingerprint" to enable in-line data deduplication directly in the data path without com-

promising performance or data integrity. SHA-1 is specified by the U.S. Department of Commerce and the National Institute of Standards and Technology for use "wherever a secure hash algorithm is required" (FIPS PUB 180-1).

1.4 Applications for Data Deduplication Technology

1.4.1 Backup Deduplication

Of all the applications for data deduplication identified to date, the application with the most stunning claims is backup to disk. DataDomain's (and others) claims arise from the characteristics of the backup application itself.

The key to understanding how DataDomain, Quantum and other vendor's products deliver on their promise of 10X – 100X data reduction is the critical phrase "over time". Deduplication products like those from DataDomain and Quantum DO NOT compress the first full backup by anything approaching 10X (or 50X).

From the ESG Lab Report on DataDomain:

"One of the key features of the DD200 is its global compression functionality. It works as follows:

1. The system administrator performs a backup of a 2 TB data set.
2. The DD200 analyzes the data to see what areas it can compress through a proprietary methodology.
3. For the first backup, the DD200 compression typically provides a 3 or 4 to 1 compression ratio. The result is that the 2TB data set is compressed to approximately 570 GB of actual storage on the DD200.
4. DataDomain claims that over time most customers will experience a 20 to 1 compression ratio. This means that a customer can perform additional backups of that 2 TB data set and only store 100 GB on the DD200 at a time.

The DD200 supports just under 1.3 TB of capacity but with compression can keep backing up a data set while using minimal capacity on the disks. If a customer needed to back up 2 TB of data and had a 3 to 1 compression rate on the first backup, they would use a total of 0.7 TB of the 1.3 TB on the DD200. If the customer experienced the 20 to 1 compression rate that Data Domain claims their customers should get over time, then the system administrator can back up that 2 TB data set for several months without using up the 1.3 TB of capacity on the DD200. Other systems would actually take up 2 TB of storage every time a full backup was performed. Typical compression rates from the field as claimed by Data Domain are presented in the following table:

Description	Compressio Rate
First full backup	3 – 4 to 1
Incremental backups	6 – 7 to 1
Following full backups	50 – 60 to 1
Aggregate over time – weekly full, daily incremental	20 to 1

A graph produced by ESG that highlights the way in which these compression ratios reduce the amount of backup data that has to be stored is shown below.

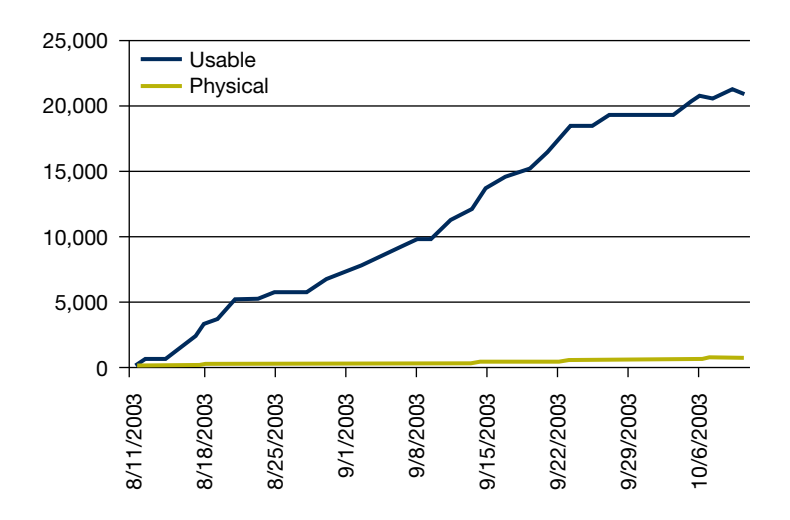

Source: ESG Lab Validation Report on DataDomain

1.4.2 Primary Storage Data Deduplication

In a world of unstructured data that is rarely accessed, the deduplication of primary storage is an obvious next step. But the objection to data deduplication for primary storage is the performance penalty. That is precisely the problem that Hifn has solved.

1.4.3 Wide Area File Services (WAFS) and Wide Area Data Services (WDS)

WAFS allows for the centralization of remote location backup data by allowing applications at a central location to appear as if they are local to the remote location. WAFS enables the “serverless remote office” in which all of a company’s servers and data are concentrated at a central location.

The goal of WAFS-based backups is to keep the backup application and data at a central location in order to cut down on administrative costs at the remote offices. This also ensures that the data used by the remote locations is protected and recoverable.

One central office appliance can support many field appliances, which is presented as an opportunity to consolidate all previously remote data into the central location. The appliances use a variety of techniques, which vary by vendor, to streamline their communication across the WAN. These techniques include deduplicating, compressing and encrypting the data that is sent. The net result is to provide secure, LAN-like performance via a WAN.

1.4.4 Backup Software Vendors (Backup ISVs)

Backup software vendors are increasingly focusing on the protection of remote location data, thus increasing their interest in technologies that bring the performance of the WAN closer to that of the LAN.

Appliances that employ Hifn's hardware-based data deduplication technology are highly responsive to both to disk storage savings and WAN transfer efficiency.

1.4.5 Storage as a Service (SAAS)

Storage as a Service (SaaS) is a business model in which third-party providers rent space on their storage to end users that lack the capital budget and/or technical personnel to implement and maintain their own storage infrastructure. There is nothing new about third-party storage services, but the complexity of today's backup, replication, and disaster recovery needs has renewed their popularity, particularly among small and medium-sized businesses.

1.4.6 Virtual Tape Libraries

A virtual tape library (VTL) is a disk-based storage system with processors and software that enable it to emulate a traditional tape library – except that it operates at disk speed and with disk reliability. Users of VTLs reap the benefits of backing up to disk, such as easier management and improved performance, over tape. Because VTL systems emulate tape libraries, they can be seamlessly integrated into existing enterprise backup environments. The backup application simply points to the VTL and writes to it just as it would a traditional physical tape library. Tape virtualization allows for multiple backup streams to run in parallel, thereby improving backup performance. Restore from a VTL is fast because it is direct from disk and does not require physical tape to be found, loaded, and positioned appropriately, as VTLs are fundamentally random access devices.

Because a VTL appears to the backup application as a physical library, the backup application can copy or clone data from the VTL to a physical tape library during off peak periods. However, in an environment with remote offices, transferring data across a WAN to a central location is extremely inefficient if data deduplication is not employed in the replication strategy.

1.4.7 Continuous Data Protection (CDP) Vendors

CDP is a data protection technology that automatically saves a copy of every change, essentially capturing every version of the data that the user saves. It allows the user or administrator to restore data to any point in time. Writes are saved to a journal file along with corresponding file changes. If data loss from human error, virus corruption or disk failure occurs, a CDP system can restore any number of protected files to any moment in time thus offering unparalleled granularity for data recovery. Restores are quick because the content is recovered directly from disk.

1.4.8 Fibre Channel and iSCSI SAN switches

A recurring theme in storage networking is the placement of functionality within the SAN topology. The industry has been debating this topic within the context of storage virtualization for years with the server vendors claiming the optimal location for virtualization is (not surprisingly) in the server, the storage vendors (end point/target manufacturers) claiming the storage sub-system is the optimal location and the switch manufacturers claiming the optimal location is in the storage network itself.

All their positions are supportable.

Industry speculation suggests that a similar technology skirmish will take place with the placement of data deduplication. Experts generally concur that a single topological placement should be selected to the exclusion of any other.

1.4.9 FC and iSCSI Host Bus Adapters

Under the server covers, every FC or iSCSI HBA has the potential to offer accelerated data deduplication services through a component on their card itself or via a “companion” card.

1.4.10 Deduplicating Gateways

Riverbed has just provided a look at their newest product which represents a sea-change in their strategic positioning. Riverbed calls its new product Atlas, and says it can eliminate up to 95% of redundant data in corporate storage systems. In a recent press release, the company asserted that the Atlas appliance “will do for its customers’ data at rest” what the company’s Steelhead

products “have done for their data in motion,” eliminating redundancy “which typically burdens IT infrastructure by slowing down access to data and applications and increasing costs and operational overhead of data management.”

Products described as “storage routers” and “storage gateways” can be built and sold as adding deduplication capabilities to existing primary and secondary storage devices.

Deduplicating gateways enable the buyer to take advantage of existing storage capacity rather than forcing the buyer to add new capacity. This capability is most significant in primary storage applications where the cost of enterprise-class storage capacity is extremely high.

1.4.11 RAID Controllers

Rounding out the potential points of placement for accelerated data deduplication technology in storage networks is at the end points. While a storage networking switch might have “visibility” of all data traversing the storage network, the server and RAID end points only have visibility of the data they exchange. This fact can lead to performance and deduplication limitations that would not be observed in a deduplicating storage switch.

1.4.12 VMware ESX, Microsoft Hyper-V and Other Virtualized Server Deduplication Conduits

A deduplicating, compressing, encrypting (DCE) conduit can provide services to VMware ESX, Microsoft Hyper-V and other virtualized servers and applications.

1.5 Summary

Data deduplication is rapidly becoming a required feature in backup to disk customer requirements and is starting to gain prominence in NAS and other primary storage applications. The value proposition of using data deduplication to accelerate network connections is equally compelling.

Widespread adoption will only occur when the performance and economics of the solution can be implemented with minimal operational disturbance. Today, the industry performs data deduplication using software in general purpose computers. This limits the technology’s performance, hence limiting its market potential.

The industry has long accepted the use of specialized hardware to accelerate computationally intensive processes. RAID controller parity calculations and data compression – once done in software – are now routinely performed in customized silicon.

It is this technology—the acceleration of the data deduplication process in customized silicon—that Hifn is delivering and has delivered to solve these critical business problems.

HIFN
PROTECT · SECURE · ASSURE

750 University Ave
Los Gatos, CA 95032
408.399.3500 tel
408.399.3501 fax
info@hifn.com
www.hifn.com