


Deduplication
Storage

Dedupe-Centric Storage

Hugo Patterson, Chief Architect, Data Domain

September 2007



INTRODUCTION 3

DEDUPLICATION: THE POST-SNAPSHOT REVOLUTION . 3

BUILDING A BETTER SNAPSHOT 4

DEDUPLICATION AS A POINT SOLUTION. 4

DEDUPLICATION AS A STORAGE FUNDAMENTAL 5

ABOUT DATA DOMAIN 8

People, by their nature, tend to build on what is already there to improve it, repurpose it, or just to keep it up to date. They also tend to share the product of their work with other people. Such activity, when applied to the digital world, generates multiple successive versions of files, multiple copies of those file versions and, in general, proliferates essentially similar data and so fills corporate data centers. And people like to keep all those versions and copies around because they serve different purposes, or protect against human or computer errors. Proliferating and preserving all these versions and copies drives much of the tremendous data growth most companies are experiencing. IT administrators are left to deal with the consequences.

Deduplication is the process of recognizing common elements in the many versions and copies of data and eliminating the redundant copies of those common elements.

Deduplication reduces total data size and so simplifies the data management problem. There is less data to store, less to protect, less to replicate, less to index and search, and less to preserve for compliance. For IT administrators, this means there are fewer and/or smaller storage systems to manage, smaller data centers to run, fewer tapes to handle, fewer tape pickups, smaller network pipes, and cheaper WAN links. People tend to proliferate similar data, and deduplication makes it easier to manage the data when they do.

With effective deduplication, people can make and preserve all the versions and copies they'd like with

less concern for storage space and cost. Because it directly addresses one of the key engines of data growth, deduplication should be at the heart of any data management strategy; it should be baked into the fundamental design of the system. When deduplication is done well, people can create, share, access, protect and manage their data in new and easier ways, and IT administrators will not be struggling to keep up.

Deduplication: The Post-Snapshot Revolution in Storage

The revolution to come with deduplication will eclipse the snapshot revolution of the early 1990s. Since at least the 1980s, systems had been saving

versions of individual files or snapshots of whole disk volumes or file systems. Technologies existed for creating snapshots reasonably efficiently. But, in most production systems, snapshots were expensive, slow, or limited in number and so had limited use and deployment.

About 15 years ago, Network Appliance (NetApp) leveraged the existing technologies and added some of their own innovations to build a file system that could create and present multiple snapshots of the entire system with little performance impact and without making a complete new copy of all the data for every snapshot. It did not use deduplication to find and eliminate redundant data, but at least it didn't duplicate unchanged data just to create a new logical view of the same data in a new snapshot. Even this was a significant advance. By not requiring a complete

When deduplication is done well, people can create, share, access, protect, and manage their data in new and easier ways and IT administrators aren't struggling to keep up.

copy for each snapshot, it became much cheaper to store snapshots. And, by not moving the old version of data preserved in a snapshot just to store a new version of data, there was no performance impact to creating snapshots. The efficiency of their approach meant there was no reason not to create snapshots and keep a bunch around. Users responded by routinely creating several snapshots a day and often preserving some snapshots for several days. The relative abundance of such snapshots revolutionized data protection. Users could browse snapshots and restore earlier versions of files to recover from mistakes immediately all by themselves. Instead of being shut down to run backups, databases could be quiesced briefly to create a consistent snapshot and then freed to run again while backups proceeded in the background. Such consistent snapshots could also serve as starting points for database recovery in the event of a database corruption, thereby avoiding the need to go to tape for recovery.

Building a Better Snapshot

Despite the many benefits and commercial success of space and performance efficient snapshots, it was over a decade before the first competitors built comparable technology. Most competitors' snapshots still do not compare. Many added a snapshot feature that delivers space efficiency but few matched the simplicity, performance, scalability and ultimately the utility of NetApp's design. Why?

The answer is that creating and managing the multiple virtual views of a file system captured in snapshots is challenging. By far the easiest and

most widely adopted approach is to read the old version of data and copy it to a safe place before writing new data. This is known as copy-on-write and works especially well for block storage systems. Unfortunately, the copy operation imposes a severe performance penalty because of the extra read and write operations it requires. But, doing something more sophisticated and writing new data to a new location without the extra read and write requires all new kinds of data structures and completely changes

the model for how a storage system works. Further, block-level snapshots by themselves do not provide access to the file version captured in the snapshots; you need a file system that integrates those versions into the name space.

Creating and managing the multiple virtual views of a file system captured in snapshots is challenging.

Faced with the choice of completely rebuilding their storage system, or simply bolting on a copy-on-write snapshot, most vendors choose the easy path that will get a check-box snapshot feature to market soonest. Very few are willing to invest the years and dollars to start over for what they view as just a single feature. By not investing, their snapshots were doomed to be second rate. They don't cross the hurdle to that makes snapshots easy, cheap and most useful. They don't deliver on the snapshot revolution.

Deduplication as a Point Solution

On its surface, deduplication is a simple concept: find and eliminate redundant copies of data. But computing systems are complex, supporting many different kinds of applications. And data sets are vast, putting a premium on scalability. Ideally, a deduplication technology would be able to find and eliminate redundant data no matter



which application writes it, even if different applications have written the same data. Deduplication should be effective and scalable so that it can find a small segment of matching data no matter where it is stored

in a large system. Deduplication also needs to be efficient or the memory and computing overhead of finding duplicates in a large system could negate any benefit. Finally, deduplication needs to be simple and automatic or the management burden of scheduling, tuning, or generally managing the deduplication process could again negate any benefits of the data reduction. All of these requirements make it very challenging to build a storage system that actually delivers the full potential of deduplication.

Early efforts at deduplication did not attempt a comprehensive solution. One example is email systems that store only one copy of an attachment or a message even when it is sent to many recipients. Such email blasts are a common way for people to share their work and early email implementations created a separate copy of the email and attachment for every recipient. The first innovation is to store just a single copy on the server and maintain multiple references to that. More sophisticated systems might detect when the same attachment is forwarded to further recipients and create additional references to the same data instead of storing a new copy. Nevertheless, when users download attachments to their desktops, the corporate environment still ends up with

Vast requirements make it very challenging to build a storage system that actually delivers the full potential of deduplication.

many copies to store, manage and protect.

At the end of the day, an email system that eliminates duplicate email attachments is a point solution. It helps that one

application, but it doesn't keep many additional duplicates, even of the very same attachments from proliferating through the environment.

The power of an efficient snapshot mechanism in a storage system is its generality. Databases could all build their own snapshot mechanism, but when the storage provides them they don't have to. They and any other application can benefit from snapshots for efficient, and independent, data protection. NetApp was dedicated to building efficient snapshots so all the applications didn't have to.

Dedupe as a Storage Fundamental

Data Domain is dedicated to building deduplication storage systems with a deduplication engine powerful enough to become a platform that general applications can leverage. That engine can find

duplicates in data written by many different applications at every stage of the lifecycle of the file, can scale to store many hundreds of terabytes of data and find duplicates wherever they exist, can reduce data by a factor of 20 or more, can do so at high speed without huge memory or disk resources, and does so

Such an engine is only possible with an unconventional storage system architecture that breaks with traditional thinking. Can vendors who try to bolt on deduplication as merely a feature of their existing system, ever deliver?

automatically while taking snapshots and so requires a minimum of administrator attention.

Such an engine is only possible with an unconventional system architecture that breaks with traditional thinking. Here are some important features needed to build such an engine.

Variable Length Duplicates

Not all data changes are exactly 4KB in size and 4KB aligned. Sometimes people just replace one word with a longer word. Such a simple small replacement shifts all the rest of the data by some small amount. To a deduplication engine built on fixed size blocks, the entire rest of the file would seem to be new unique data even though it really contains lots of duplicates. Further, that file may exist elsewhere in the system, say in an email folder, at a random offset in a larger file. Again, to an engine organized around fixed blocks, the data would all seem to be new. Conventional storage systems, whether NAS or SAN, store fixed sized blocks. Such a system which attempts to bolt on deduplication as an afterthought will only be able to look for identical fixed size blocks. Clearly, such an approach will never be as effective, comprehensive, and general purpose as a system that can handle small replacements and recognize and eliminate duplicate data no matter where in a file it appears.

Data Domain is dedicated to building deduplication storage systems with a deduplication engine powerful enough to become a platform that general applications can leverage.

Local Compression

Deduplication across all the data stored in a system is necessary, but it should be complemented with local compression which typically reduces data size by another factor of 2. This may not be as significant a factor as deduplication itself, but no system that takes data reduction seriously can afford to ignore it.

Format Agnostic

Data comes in many formats generated by many different applications. But, embedded in those different formats is often the same duplicate data. A document may appear as an individual file generated by a word processor, or in an email saved in a folder by an email reader, or in the database of the email server, or embedded in a backup image, or squirreled away by an email archive application. A deduplication system that relies on parsing the formats generated by all these different applications can never be a general purpose storage platform. There are too many such formats and they change too quickly for a storage vendor to support them all. Even if they could, such an approach would end up handcuffing application writers trying to innovate on top of the platform. Until their new format is supported, they'd gain no benefit from the platform. They would be better off sticking with the same old formats and not creating anything new. Storage platforms should unleash creativity, not squelch it. Thus, the deduplication engine must be data agnostic and find and eliminate duplicates in data no matter how it is packaged and stored to the system.

Multi-Protocol

There are many standard protocols in use in storage systems today from NFS and CIFS to blocks and VTL. For maximum flexibility, storage systems should support all these protocols since different protocols are needed for different applications at the same time. User home directories may be in NAS. The exchange server may need to run on blocks. And backups may prefer VTL. Over the course of its lifecycle, the same data may be stored with all of these protocols. A presentation that starts in a home directory may be emailed to a colleague and stored in blocks by the email server and then archived in NAS by an email archive application and backed up from the home directory, the email server, and the archive application to VTL. Deduplication should be able to find and eliminate redundant data no matter how it is stored.

CPU-Centric vs. Disk-Intensive Algorithm

Over the last two decades, CPU performance has increased 2,000,000x*. In that time, disk performance has only increased 11x* (*Seagate Technology Paper, Economies of Capacity and Speed, May 2004). Today, CPU performance is taking another leap with every doubling of the number of cores in a chip. Clearly, algorithms developed today for deduplication should leverage the growth in CPU performance instead of being tied to disk performance. Some systems rely on a disk access to find every piece of duplicate data. In some systems, the disk access is to lookup a segment

Deduplication has leverage across the storage infrastructure for reducing data, improving data protection and, in general, simplifying data management.

fingerprint. Other systems can't find duplicate data except by reading the old data to compare it to the new. In either case, the rate at which duplicate data can be found and eliminated is bounded by the speed of these disk accesses. To go faster, such systems need to add more disks. But, the whole idea of deduplication is to reduce storage, not grow it. The Data Domain SISL™ (Stream-Informed Segment Layout) technology does not have to rely on reading lots of data from disk to find duplicates and organizes the segment fingerprints on disk in such a way that only a small number of accesses are needed to find thousands of duplicates. With SISL, backend disks only need to deliver a few megabytes of data to deduplicate hundreds of megabytes of incoming data.

Deduplicated Replication

Data is protected from disaster only when a copy of it is safely at a remote location. Replication has long been used for high-value data, but without deduplication, replication is too expensive for the other 90% of the data. Deduplication should happen immediately inline and its benefits applied to replication in real time, so the lag till data is safely off site is as small as possible. Only systems designed for deduplication can run fast enough to deduplicate and replicate data right away. Systems which have bolted on deduplication as merely a feature at best impose unnecessary delays in replication and at worst don't deliver the benefits of deduplicated replication at all.

Deduplicated Snapshots

Snapshots are very helpful for capturing different versions of data and deduplication can store all those versions much more compactly. Both are fundamental to simplifying data management. Yet many systems that are cobbled together as a set of features either can't create snapshots efficiently, or they can't deduplicate snapshots, so users need to be careful when they create snapshots so as not to lock duplicates in. Such restrictions and limitations complicate data management not simplify it.

Conclusion

Deduplication has leverage across the storage infrastructure for reducing data, improving data protection and, in general, simplifying data management. But, deduplication is hard to implement in a way that runs fast with low overhead across a full range of protocols and application environments. Storage system vendors who treat deduplication merely as a feature will check off a box on a feature list, but are likely to fall short of delivering the benefits deduplication promises.

About Data Domain

Data Domain® is the leading provider of Deduplication Storage systems for disk backup and network-based disaster recovery. Over 1,000 companies worldwide have deployed Data Domain's market-leading protection storage systems to significantly reduce backup data volume, lower their backup costs and simplify data recovery. Data Domain delivers the performance, reliability and scalability to address the data protection needs of enterprises of all sizes. Data Domain's products integrate into existing customer infrastructures and are compatible with leading enterprise backup software products. To find out more about Data Domain, visit www.datadomain.com. Data Domain is headquartered at 2300 Central Expressway, Santa Clara, CA 95050 and can be contacted by phone at 1-866-933-3873 or by e-mail at sales@datadomain.com.

Copyright © 2007 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-DCS-0907

Data Domain – Dedupe-Centric Storage

Data Domain
2300 Central Expressway
Santa Clara, CA 95050
866-WE-DDUPE
sales@datadomain.com

datadomain