



Other cool reads to check out!

Deduplication, Incremental Forever,
and the Olsen Twins

7 Shortcuts to Losing Your Data
(and Probably Your Job)

Six Fairy Tales of VMware and
Hyper-V Backup

Oh, and look at the blog, too!

Don't Get Duped By Dedupe

Unitrends™

7 Technology Circle
Suite 100
Columbia, SC 29203

Phone: 866.359.5411
E-Mail: sales@unitrends.com
URL: www.unitrends.com

Unitrends™

Don't Get Duped By Dedupe: Introducing Adaptive Deduplication

The purpose of deduplication is to provide more storage, particularly backup storage, for less money, right? Then wouldn't it be ridiculous if deduplication vendors were demanding that their customers pay more per terabyte of storage? Or, if they were simply pushing the task of integrating, monitoring, and managing deduplication back onto their users?

The purpose of this white paper is to help you understand the various approaches to deduplication, the strengths and weaknesses of each, and to introduce a different approach to deduplication, which we term Adaptive Deduplication. Adaptive Deduplication delivers the advantages of deduplication without the capital expense of hardware-based deduplication devices but with better performance and manageability than software solutions.

Data Reduction Concepts

There's a lot of confusion concerning deduplication - so much so that we've entitled this section "Data Reduction" in order to be able to discuss various techniques for reducing the amount of data necessary to represent a larger data set. We'll discuss two primary data reduction concepts: compression and deduplication. We'll also discuss two other data concepts that have a major influence on data reduction: encryption and physical storage density and pricing trends.

Compression

Compression is just the encoding of data using less data than that which made up the original. There are two fundamental types of compression: lossless and lossy compression. Lossless data compression means that you can recover every bit of your original data; lossy data compression means that your original data is lost in the compression process. For the purposes of our discussion, we're only going to consider the lossless form of data compression.

Lossless data compression typically exploits the statistical redundancy of underlying data to represent the original data more concisely and yet with the ability to fully and accurately reconstitute that data at a later date when that data is uncompressed. Statistical redundancy exists because almost all real-world data isn't random but instead have specific underlying patterns.

Here's a trivial example. In standard character sets, all letters of an alphabet are represented by the same number of bits. In English, the letter "e" has a frequency of use of 12.7% while the letter "t" has a frequency of 9.06%. If it takes eight bits to represent each, then if you use fewer bits to represent these characters, you're going to be able to show a significant reduction just by encoding these two letters.

Compression is typically a trade-off between the utilization of the microprocessor and primary memory versus the utilization of secondary memory and a transmission line, such as a WAN.

Deduplication

Deduplication is a specific form of lossless compression in which redundant data is eliminated. I realize that this sounds strange - deduplication is compression? The reason it sounds so strange is that there have been so many hundreds of millions of dollars spent representing deduplication as magical and revolutionary.

In deduplication, duplicate data is deleted leaving only one copy of the data to be stored. The deleted duplicate data is said to be “stubbed.” Stubbing is a process by which an indication that some data has been deleted and a pointer to the “index” of that data (the one copy that is not deleted) is put in place of the deleted data.

Here’s a trivial example of deduplication. If you have ten servers, and each server has the same 1GB file, data deduplication across all of those servers should allow nine of those ten copies to be stubbed. That means your data has been deduplicated by a factor of 10:1.

While deduplication is a version of compression, it also is typically at odds with the more traditional version of compression. Deduplication tends to achieve better data reduction efficacy against smaller backup sets (the amount of data being backed up each time), while compression tends to achieve better results against larger data sets.

Encryption

Encryption is the process of transforming data to make it unreadable to anyone except those who possess a key that will decrypt that data. In order to make the data unreadable, encryption algorithms attempt to eliminate any discernible patterns in the underlying data. Since compression and deduplication both tend to work on patterns of data, both are negatively impacted by encryption.

Physical Storage Density and Pricing

The rapid increase in available physical storage density and the rapid trend toward decreasing price per terabyte negatively impacts deduplication. In this sense, the increasing physical storage density and decreasing price per terabyte is analogous to encryption - both make conventional deduplication more difficult. The reason is that conventional deduplication requires a great deal of processors, processor cores, and physical memory in order to efficiently map and process each segment of information (which is typically called a “block”.) As the size of the physical storage grows, so does the amount of processors and memory required for deduplication.

The technique that most deduplication device vendors use is to keep the amount of disk space low on a per device basis and to recommend multiple devices in order to scale their aggregate storage capacity higher. This technique works effectively for avoiding forklift upgrades; however, it is costly on a per-storage device because of the additional capital expenditure costs as well as the operational expenditure associated with managing the devices. In addition, if you go this route, you want to make sure you understand federated storage and local versus global deduplication (discussed in a section by that name in the next chapter.)

Data Deduplication Background

Now that we’ve discussed the various concepts associated with data reduction, it’s time in this chapter to discuss on a more detailed basis data deduplication specifically. In this chapter we’re going to examine the following deduplication concepts:

- Source versus target deduplication
- Inline versus post-processing deduplication
- File- versus block- versus byte-level deduplication
- Hardware versus software deduplication
- Federated storage and local versus global deduplication

Each will be explored on a general basis. The purpose of this is to give you a general background in the “lingo” before presenting in more detail specific advice regarding avoiding getting “duped by deduplication.”

Source Versus Target Deduplication

Source deduplication means that deduplication occurs at the client (i.e., the server, PC, workstation - but not at the SAN or NAS) being protected while target deduplication means that deduplication occurs on the backup medium. Note that “backup medium” can mean the server upon which the backup software resides, on a dedicated deduplication device attached to that backup server, or on a backup appliance.

Source deduplication has the advantage of using less LAN (on-premise) bandwidth but has the disadvantage of using more client resources; target deduplication has the advantage of using less client resources but has the disadvantage of using more LAN bandwidth.

Inline Versus Post-Processing Deduplication

Inline deduplication means that the deduplication occurs while the backup data is being ingested by the backup device. Post-processing deduplication means that the deduplication occurs after the data has been ingested. In short, inline deduplication means that the deduplication occurs before the data is written to disk while post-processing deduplication means that deduplication occurs afterward.

Advocates of both forms claim that their techniques are superior. However, inline deduplication tends to decrease ingest rates against real-world data (be very careful about quoted ingest rates - they are often done with tuned data that allows faster deduplication ingest rates) while post-processing deduplication only increases ingest rates if the “landing site” for the data is large enough to accommodate all backups and deduplicate them before the next backup arrives.

File- Versus Block- Versus Byte-Level Deduplication

File-level deduplication works at the file level by eliminating duplicate files; block-level deduplication works at a block level (which may be a fixed-size block or a variably-sized block) by eliminating duplicate blocks.

The advantage of file-level deduplication is that it takes less resources and thus may be deployed over larger amounts of physical storage; the disadvantage of file-level deduplication is that it can't eliminate smaller redundant “chunks” of data than a file. The advantage of block-level deduplication is that it can eliminate chunks of data smaller than a file; the disadvantage is that it can't be deployed over larger amounts of physical storage.

Byte-level deduplication combines the advantages of file-level deduplication and block-level deduplication. Byte-level deduplication is a form of block-level deduplication that understands the content, or “semantics”, of the data. These systems are sometimes called CAS - Content Aware Systems. Typically, deduplication devices perform block-level deduplication that is content-agnostic - blocks are blocks. The problem of course is that certain blocks of data are much more likely to change than other blocks of data. For backup systems, the “metadata” (data about data) that contains information about the actual backup tends to change continuously while the backup data statistically changes much less often. The advantage to byte-level deduplication is that by understanding the content of the data the system can more efficiently deduplicate the bytes within the data stream that is being deduplicated.

Ironically, file-level deduplication is a form of byte-level deduplication since there must be some degree of content-awareness in order to detect a file versus some other form of data. But of course the problem as described above is that file-level deduplication can't handle unstructured data and can't handle changes at the sub-file level. Byte-level deduplication can handle changes at the sub-file level.

Hardware Versus Software Deduplication

(Preface: This section refers to hardware deduplication at the storage device - which is the most common form of hardware deduplication. For more information concerning hardware deduplication in a vertically integrated appliance, see the "Introducing Adaptive Deduplication" chapter later in this paper.)

The term "hardware" deduplication refers to a storage device such as a NAS that is dedicated to deduplicating that storage. "Software" deduplication refers to non-integrated commodity software packages that support deduplication of some form in software for a server and some attached storage.

The advantage of software deduplication is that it doesn't cost that much money to purchase. The disadvantage is that it costs a whole lot of money when you add in not only all of the components of the solution but the expense of trying to make it work. What the software vendor is doing is basically pushing back onto the buyer all of the integration, management, and monitoring work.

The advantage of hardware deduplication is that the vendor has done all of the integration, management, and monitoring work for the deduplication for the storage device. The disadvantage is that the capital expenditure is higher and that the integration, management, and monitoring work must be done between the rest of the backup solution (i.e., it's only storage that is integrated.)

Federated Storage and Local Versus Global Deduplication

Earlier it was noted that one way that deduplication device vendors handled some of the shortcomings of their approach was to sell multiple deduplication devices. These multiple deduplication devices are a pretty expensive way to solve the ingest rate issues of both inline and post-processing deduplication as well as the processor, memory, and general resource requirements of block-level deduplication when done at the target.

If you require multiple deduplication devices, either now or in the future, you want to ensure that these devices appear as federated storage - which simply means that when you add a second physical 1TB device to the first physical 1TB device that the device appears as the aggregate (2TBs) of those devices.

In this situation, you want to make sure you understand local versus global deduplication. Local deduplication in this environment means that each separate device performs deduplication for only that device; global deduplication means that deduplication applies across all of the devices. Global deduplication tends to have a superior data reduction rate because only a single deduplication index (or cache of actual blocks) is needed; the disadvantage is typically one of performance.

Shortcuts on the Road to Getting Duped By Deduplication

Rather than warn you about all the things that you should be concerned about regarding backup and deduplication, this section dictates to you precisely what you should do to ensure that you will get duped by deduplication.

Be Oblivious to Physical Storage Costs

One of the best ways to ensure that you will get duped is to buy based on the existence of the term “deduplication” - after all, it’s a pretty sexy buzzword - rather than focusing on the end result of deduplication - which should be more backup storage for less cost.

The best way to do this is to ignore the price per terabyte of commodity storage (like the 1TB 3.5” 7200RPM SATA drives available at the time this document was published for less than \$90) and more mundane technologies such as compression and focus on whether or not for \$12,500 you’ll really get a 20:1 data deduplication ratio from 750GB of physical space (note: this is a real product from a real deduplication vendor.) If you assume 2:1 compression ratios (very conservative) you find that you’re going to be able to achieve 20TB for \$900 with compression and 10 1TB drives - and have a lot left over not only for the NAS chassis, motherboard, memory, and even for backup software as well! Of course, that leads us to the next shortcut to getting duped...

Blindly Believe Data Reduction Ratios

In the previous section, in order to portray some deduplication solutions in the best possible light, we made the assumption of a 20:1 data reduction ratio (also called a data deduplication ratio.) In order to be duped by deduplication, you should blithely and blindly make that assumption as well.

In truth, data reduction ratios vary widely based on the following primary factors:

- The type of data being deduplicated (unstructured data is better)
- The type of backup being deduplicated (multiple master and differential backups)
- The frequency and degree of data that is changed (the lower, the better)
- The retention of the data (the longer, the better)
- The specifics of the data reduction algorithms being used

Anyone who tells you that they can predict the specific data reduction rate that you’ll achieve for your data is misleading you. The best that can be done is to ask questions and make assumptions about “normal” data. Thus, you’re always better off assuming a lower data reduction ratio.

Focus on Technology Instead of Price Per Effective Terabyte

This is a combination of ignoring physical storage costs and blindly believing data reduction ratios and is one of the most common ways to bamboozle yourself. There’s nothing “magic” about deduplication - so if you can achieve your goals with respect to backup size, retention, and price per terabyte then it doesn’t matter whether you use a dedicated deduplication device or a commodity disk drive.

Blissfully Accept Published Ingest Rates from Vendors

There are two tricks that deduplication device vendors play with respect to ingest rates (the rate at which data can be stored on the deduplication device):

- Post-processing deduplication vendors quote ingest rates that assume that the data being transferred will fit within a “landing site” and can complete before deduplication begins.

- Inline deduplication vendors quote ingest rates for the “best-case” data. By “best-case”, we mean data that causes the least amount of delay in processing. Selecting the right data means that an inline deduplication algorithm can show basically no delay - while data that isn’t crafted for ingest rate performance (for example, that doesn’t consist of all zeros) will always show ingest rate penalties because deduplication is a relatively expensive operation.

So if you’re going to get duped, treat the vendor ingest rates as the holy writ and suspend all skepticism - after all, they wouldn’t talk about ingest rates if it were really a problem, right?

By the way, one other way to get duped - make sure that you don’t take into account the overhead of the backup process itself (which in the case of a backup device, operates on another server) on the ingest rate. In other words, ignore the fact that the backup server has to first get the backup data and then send it over to the deduplication device.

Neglect Backup Size, Backup Type, and Retention

Another lesson on getting the wool pulled over your eyes - stick to the credo that “data is data” and don’t pay attention to the overall size of your backup or to the retention desired of your backup (retention just means the number of backups from which you want to be able to recover.)

In practice, backup size not only dramatically impacts ingest rates (particularly in post-processing deduplication devices - which as noted earlier start fast but quickly get bogged down if and when you exceed the landing site size) but also radically changes your expected data reduction ratios. More retention typically means more redundant data which means higher data reduction ratios.

Of course, the degree to which retention will impact the data reduction ratio depends upon the backup type in use. Doing full masters every day means maximum data reduction; doing block-level incrementals every hour means that your data reduction ratio will probably be small.

Disregard Recovery

Deduplication is the process of removing duplicate data. In order to reconstitute deduplicated data so that it may be recovered, the inverse of the deduplication process must occur. This process, which is commonly termed “hydration” or “re-duplication”, tends to have two negative consequences:

- It takes time to transmogify the data from the deduplicated state to the original state.
- There is a risk that something could go wrong and data could be lost.

In order to decrease your risk, you have to pay attention to what your recovery time objectives are with respect to deduplication. You also have to focus on the underlying reliability of the physical storage on which the deduplicated data is stored.



Overlook Integration, Management, and Maintenance Costs

Deduplication is far from simple. You have several choices regarding dealing with the complexity of deduplication. The first is to purchase some backup software with deduplication capabilities built into it. You will then be responsible for integrating, managing, and maintaining your backup server, its operating system, and your backup software. This is viable - but it means that you are going to need people who not only perform the initial installation but who handle making this work on an on-going basis. And if something goes wrong, you're pretty much on your own.

Another choice is to buy a separate deduplication device. At least with this approach you are not attempting to integrate, manage, and maintain your deduplication software. Of course, you're still going to be handling those tasks with respect to your backup server and your deduplication device.

A third choice is to buy an integrated backup appliance with integrated deduplication. This approach works best if you don't have the staffing or inclination to have more of a dedicated focus on backup.

Pay No Attention to Server, Workstation, PC, and Notebook Overhead

Data reduction, whether it's compression or deduplication, is simply the substitution of processor, memory, and disk I/O for disk storage space. The processor, memory, and disk I/O have to take place somewhere. As we discussed earlier, deduplication can take place at either the source (the element being protected) or the target (the backup server/deduplication device.)

If you're going to use source deduplication, then you need to make sure that you are protecting computers and other information technology devices that have pretty low utilization rates. In other words, be careful using source level deduplication on a busy computer.

Ignore the Impact of Encryption

As noted earlier, encryption is the process of transforming data to make it unreadable to anyone except those who possess a key that will decrypt that data. In order to make the data unreadable, encryption algorithms attempt to eliminate any discernible patterns in the underlying data. Since compression, and deduplication, both tend to work on patterns of data, both are negatively impacted by encryption. Ignoring the impact of encryption is a wonderful shortcut to receiving an unpleasant surprise regarding the impact of deduplication.

Introducing Adaptive Deduplication

It was due to the issues discussed previously in this paper that Unitrends set out to implement data deduplication. We of course understood quite well source-level block-based deduplication - we've been using the technology for years between our on-premise appliances and our off-premise appliances (or vaults) in order to reduce the amount of bandwidth necessary for disaster recovery protection over WANs. However, we knew from our shipping products how much overhead source-level block-based deduplication incurred on our on-premise appliances - and we knew that most of our customers wouldn't stand for our requiring that much resource to be dedicated on the computers that they wished to be protected.

We then evaluated simple block-level deduplication performed at the target (on the backup appliance itself.) However, what we found was that we needed to increase the pricing associated with our backup appliance in order to have enough processor and memory for this approach to work well - and even if we did that we ended up using much smaller amounts of secondary disk. This simply didn't make sense to us when we examined what was occurring in terms of disk drive storage density and pricing.

That's the reason that we created Adaptive Deduplication.

What Do You Mean, Adaptive?

Adaptive deduplication is a content-aware deduplication technique that adapts dynamically based on the following:

- The content of the data being protected
- The data reduction possible for that data using compression
- The data reduction possible for that data using deduplication
- The amount of secondary (disk) storage that may be used by Adaptive Deduplication
- The performance of the secondary storage that may be used by Adaptive Deduplication
- The amount of primary (RAM) storage that may be used by Adaptive Deduplication
- The number of processor cores that may be used by Adaptive Deduplication
- The number of processor threads that may be used by Adaptive Deduplication

What Are the Specifications of Adaptive Deduplication?

The specifications and features of data reduction techniques that were discussed conceptually earlier are discussed in this chapter specifically with respect to their implementation in Adaptive Deduplication.

Source Versus Inline Data Reduction

Adaptive Deduplication is a target-based data reduction technique. This means that all data reduction occurs on the backup appliance itself rather than taking resources from the computers and other devices that are being protected.

Inline Versus Post-Processing Data Reduction

Adaptive Deduplication uses inline compression combined with post-processing data deduplication.

File- Versus Block- Versus Byte-Level Data Reduction

Adaptive Deduplication uses an inline variable byte-by-byte compression technique that has a storage overhead of 0.03% plus 6-bytes per stream. This practically works out to 5 bytes per 16KB block. The theoretical compression limit is 1032:1. The actual compression is dependent upon the data stream and the data type and more commonly ranges from 2:1 to 4:1 for unstructured data (e.g., files) and from 4:1 to 9:1 for structured data (e.g., databases.)

Adaptive Deduplication then selectively uses a post-processing byte-level-based deduplication scheme to perform deduplication if the system determines that deduplication would increase the overall data reduction ratio. The amount of data reduction that occurs due to byte-level deduplication depends upon both the data and the underlying attributes of the hardware.

Hardware Versus Software Data Reduction

Adaptive Deduplication is a hardware-based data reduction scheme that is integrated; it is not a commodity-based software data reduction scheme. Like other hardware-based data deduplication devices, Adaptive Deduplication is implemented in concert with the storage itself. Unlike storage-oriented data deduplication devices, however, Adaptive Deduplication is also integrated with the backup server itself - in other words, it is integrated into the Unitrends backup appliance.

Federated Storage and Local Versus Global Data Reduction

Adaptive Deduplication is a local data reduction technique; it does not support federated storage and it does not perform global data reduction. In other words, the post-processing deduplication is performed on a per-backup appliance basis only.

Am I Going to Get Duped By Adaptive Deduplication?

Adaptive Deduplication was created in order to keep you from being duped by deduplication. For each of the "shortcuts" to getting duped we discussed previously, below we discuss what we've done to address each.

Physical Storage Costs

Adaptive Deduplication works with our current generation of backup appliances: from desktops with 3TB of raw storage to rack mounts with 15TB of raw storage.

Data Reduction Ratios

As we noted previously, anyone who tells you that they can predict the specific data reduction rate that you'll achieve for your data is misleading you. The best that can be done is to ask questions and make assumptions about "normal" data. Thus, you're always better off assuming a lower data reduction ratio. Adaptive Deduplication typically provides 6:1 to 10:1 real-world data reduction. We've seen more in the labs; but no one cares what we see in our labs.

Price Per Effective Terabyte

Adaptive Deduplication is available on all current generation Unitrends backup appliances at no additional cost and for new purchases is included in the backup appliance purchase price.

Ingest Rates

Due to the way it has been implemented, Adaptive Deduplication does not have any negative impact on Unitrends line-speed backup ingest rates.

Backup Size, Backup Type, and Retention

As with all other data reduction schemes, Adaptive Deduplication works best when there is more redundant data. Redundant data typically increases with more computers and devices being protected, more retention, and backup types with more redundant data. The more redundant data, the higher the data reduction ratio.

Recovery

Adaptive Deduplication is designed to have minimal impact on the most often used recovery - the last backup period. This is achieved because the "backup cache" (also called a backup index) of fully hydrated data is always that of the last backup period. This also applies to D2D2x vaulting and archiving.

However, recovery from prior backups that have been deduplicated will be slower due to hydration.

Integration, Management, and Maintenance Costs

Adaptive Deduplication is completely integrated with the already integrated all-in-one Unitrends backup appliances.

Server, Workstation, PC, and Notebook Overhead

Adaptive Deduplication is a target-based deduplication technique and puts no data reduction load on protected computers or devices.

Encryption

As with any data reduction scheme, encryption has a negative impact on data reduction rates in Adaptive Deduplication.

What's the Bottom Line?

Adaptive Deduplication is offered on Unitrends' appliances ranging from desktops with 3TB raw storage to 3U rack units with 15TB raw storage. Adaptive Deduplication typically provides 6:1 to 10:1 real-world data reduction which translates on our 3U rack unit backup appliances to 90TB to 150TB of raw storage. Actual data reduction ratios are dependent upon the data being protected, the size of the backup, the retention, and all of the other factors that were discussed earlier in this paper.

About Unitrends

Unitrends offers a family of affordable, all-in-one on-premise backup appliances that support virtual and physical system backup and disaster recovery via disk-based archiving as well as electronic vaulting to private- and public-clouds. Unitrends is customer-obsessed, not technology-obsessed, and is focused on enabling its customers to focus on their business rather than on backup.

For more information, please visit www.unitrends.com or email us at sales@unitrends.com.



Unitrends™

7 Technology Circle, Suite 100
Columbia, SC 29203

Phone: 866.359.5411
E-Mail: sales@unitrends.com
URL: www.unitrends.com

Copyright © 2011 Unitrends. All Rights Reserved.

Unitrends™

Don't Get Duped By Dedupe: Introducing Adaptive Deduplication
DGDBDP-20110921-01