

Only one copy!

By Glyn Bowden, Data Management Forum

Picture the scene, a large family gathered around a photograph album enjoying the many memories the pictures inspire. Now one of the children leans over and points to a picture that they like, and three of the gathering all agree. With this news, the owner of the album promptly takes out the photo, makes four copies and places them all back in the album so each person has their own copy. Sound ridiculous? Well every day, businesses are doing exactly this. The album is the shared resource of a storage pool and the photos are the files we send to each other every day, in e-mail, over instant messaging and in our home directories. The difference with the real life scenario is that our album is a far more expensive resource and it's usually a lot bigger than a family's collection of pictures. Imagine a family of thousands instead of tens and you can see the way in which the problem grows.

There has been a real explosion in the volume of stored data in recent years and this is largely due to the duplication of this data. So wouldn't it be remarkable if all of those duplicates could be identified and consolidated? Well, the technology is available to do this today; all that is missing, in most cases, is the understanding of how best to employ it. The first hurdle to overcome is the fear factor so let's deal with that first. In order to do that it is important to know how de-duplication works at the fundamental level.

What is data de-duplication?

One definition of "what is a duplicate" is based upon the method used to identify duplication which may include meta-data, content, hashing, or other comparative techniques – the process of examining a data-set or I/O stream and removing duplicate data while maintaining data integrity and authenticity.

So, based on this definition, the end result is that when performed, patterns of data are only stored once and referenced by instances that would otherwise require a copy of that data. The granularity of the de-duplication depends on the method and technology used to achieve this. The process can in fact happen in one of two ways. Either before the data is written to the de-duplication system disks otherwise referred to as "inline" or "in band", or after the data is already stored on a media, also known as "out of band" or "post-process". Regardless of where the process occurs it needs to be transparent to the client accessing that data. Each of these methods has pros and cons, but both support the final goals of making more efficient use of storage media resources, enabling better use of network bandwidth for transmitting the processed

data set for disaster recovery or even migration and finally, to help build a greener data centre thanks to reducing the requirement for additional storage capacity.

The challenges with data de-duplication

Alongside these benefits lie some challenges. To start with, the system is essentially virtualising data sets and with any virtualisation comes another layer of abstraction. However, simplicity and keeping the de-duplication process transparent to the user should be design goals of any viable solution. Another difficulty faced by the early adopters was the complication of gathering metrics about the process. In order to deem the de-duplication exercise worthwhile, it must be possible to measure how much efficiency is gained. Solutions exist that report on what has taken place, and this feature should be considered when choosing a technology. This leads to the chargeback question. If multiple customers are affected by a single de-duplication process, who pays for the used blocks and how is that measured? Finally, the process, by its very nature, can impact the performance of the storage device or network and could have a different use or impact than an existing service. It is important to understand the impact of the de-duplication process on the ability to deliver the storage service customers expect.

Some of the key questions that should be addressed by any investigation into data de-duplication are as follows.

The first is where the processing should occur. When inline de-duplication is used, only the de-duplicated data set is stored, making this method the most efficient in storage capacity, as the media does not need to be sized to cope with the full data set. This will also make any kind of snapshot technology more effective on the storage device, as it is impossible to capture the duplicate blocks in a snapshot. Inline de-duplication also happens in real time as the de-duplicated data is transmitted from client to storage or between storage tiers. However, this requires a considerable amount of processing power from the de-duplication device which will undoubtedly add latency to the transmission of the data. In applications that are latency sensitive, such as transactional or primary storage, this can cause issues. There is also the question of redundancy and how efficient the process can be if data is passing over more than one physical path. In order to be as efficient as possible, each redundant path would need to be aware of the data being processed on the other. It is unlikely that the database used for identifying duplicates in this case would be as efficient as that used on a device once the data is at rest. In this instance the metadata used to identify duplicates can be reproduced from simply scanning the data set. This is only possible when the data is already at rest on the device. Finally, inline de-duplication will only affect new data being transmitted, so for existing data-sets, unless a migration or retransmission occurs, this will not be effective.

Post-processing on the other hand offers the advantage of working with the full data set and the ability to re-examine any part of it at any time. This approach is also a batch process, meaning that it happens at regular intervals or on-demand, but not in real time. This allows the administrator to choose when the best time to perform the process is, for example prior to taking a snapshot or backup, or when the data is in less demand and the processing power of the device can be used for de-duplication. The disadvantage of this approach is that the planning of capacity to handle the

unprocessed data set can lead to wasted storage capacity. Also this post-processing must be carefully timed so that snapshots and backups only occur once the de-duplication process has completed. This can be quite tricky to estimate until some meaningful metrics have been gathered.

What to consider?

The questions that span both implementations are those of positioning, interoperability and risk. If a tiered storage model is in place, there is somewhat of a dilemma. One might expect that the top tier of storage that uses expensive media would benefit the most from de-duplication. However, performing this process might degrade performance significantly enough to impact delivery of service levels expected for this tier. The disaster recovery or near-line locations might be a better choice from a performance perspective, but then the disks or tape might be more inexpensive here and so the savings perhaps not as great. However, the near-line or archive devices usually are at a location where data is being consolidated. In this instance, data de-duplication can reap huge benefits in extending retention capabilities yet reducing the capacity requirement and creating efficiencies in the amount of disk or tape and even storage enclosures that are required. And in addition, less capacity means less power and cooling consumption which helps the data centre costs and the environment!

Vendor interoperability is also important. In order to protect the data from a vendor lock-in scenario it is important that all related vendor equipment can read the de-duplicated data or that the data is exported or migrated in a fashion that is open and able to be read by third party devices.

Some of these questions are still being addressed by the industry, but exciting de-duplication products already exist that offer an easily deployable solution to a very pertinent problem. Data de-duplication is good for the bottom line, good for administrators and good for the environment!

About SNIA and data de-duplication

The SNIA's Data Management Forum (DMF) has formed a new special interest group to focus specifically on data de-duplication. This new group is called Data De-Duplication & Space Reduction (DDSR) Special Interest Group (SIG).

The SNIA DDSR SIG is dedicated to advancing space reduction in all networked storage technologies. This mission addresses the continued exponential growth of data and the increasing need for storage technologies utilising data de-duplication and other space reduction techniques. By defining and promoting efficient networked storage solutions and common implementations, the DDSR SIG is enabling sustainable data storage operations that reduce both storage costs and the environmental impact of data centre infrastructures.

For more information about the DDSR SIG – please visit
http://www.snia.org/forums/dmf/programs/data_protect_init/ddrsig