BLUESHIFT
DATA PROTECTION

Blueshift Data Protection White Paper

De-duplication: More Means Less

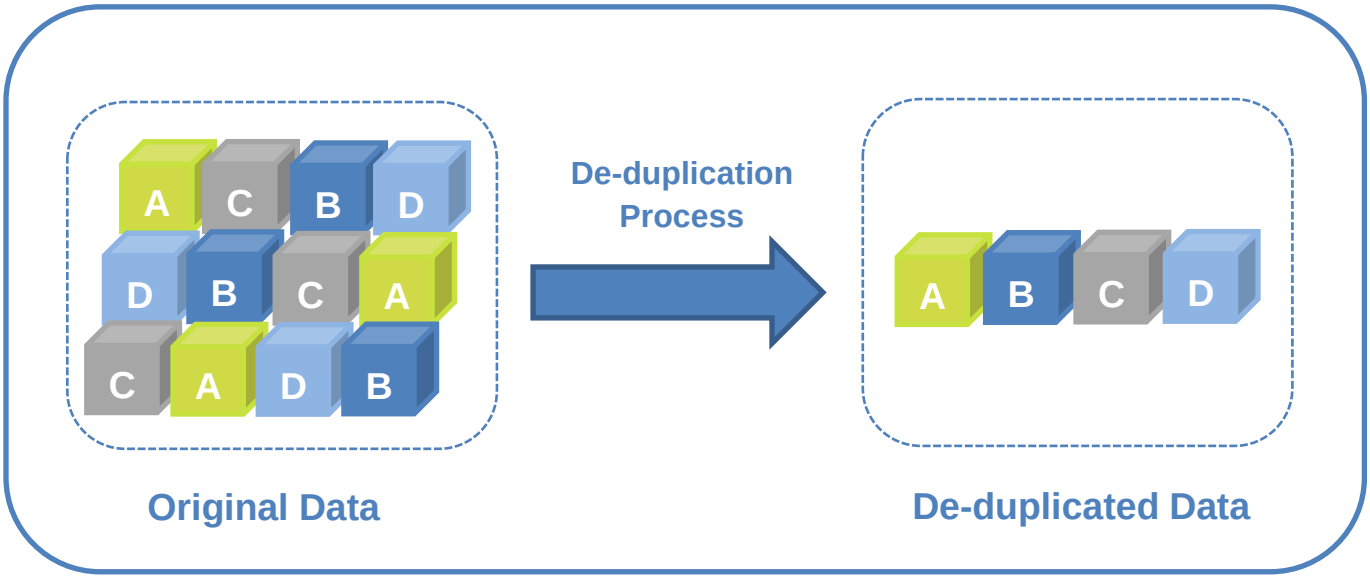

Table of Contents

Executive Summary……3
De-duplication Defined……3
Source vs. Target De-duplication……4
De-duplication Types……4
De-duplication Process Visual……6
Conclusion……7
About Blueshift Data Protection……8

Executive Summary

Data de-duplication technology crossed the chasm years ago. Initial adopters were large organizations needing to reduce the amount of data that was regularly consuming their high-end storage capacity. The first de-duplication solutions shipped were priced at thousands of dollars per terabyte as early developers and manufacturers required highly customized hardware to tackle the data de-duplication problem. The high per terabyte price and technical complexity made it nearly impossible for smaller companies to take advantage of this newly introduced data management technology.

As with most technology, time ushers in change. Many of the proprietary hardware de-duplication manufacturers from yesteryear have given way to less expensive more flexible software oriented solutions. New de-dup software typically uses standard off the shelf hardware with x86 processors to compete with expensive name brand hardware for enterprise IT budgets. The change to software driven data de-duplication solutions is saving enterprises hundreds of thousands of IT investment dollars annually. These new cost effective de-dup solutions have also found their way into the greater marketplace delivering savings to smaller businesses as well.

De-duplication technology continues to evolve and for good reason. Businesses see the benefit of eliminating same data stores across their digital infrastructure and the byproduct of this is overall lower bandwidth requirements as there is less data sent back and forth across networks. Less data also reduces backup volume which means daily “backup windows” are easier to maintain increasing the overall effectiveness of corporate data recovery objectives.

De-duplication Defined

Modern de-duplication is a method of efficiently and compactly storing data, by intelligently reducing the storage of recurring bytes within the data. This is similar to, but not analogous to traditional compressed archives (such as *.zip, etc.), which are usually better suited to single data store compress/decompress processes.

A de-duplicated data store consists of three main components - data store, metadata database, and user-space application. The data store is the processed, de-duplicated data, in its final resting form on disk. The metadata database is responsible for tracking all files that are stored. The user-space application or tools are used for adding, deleting, or retrieving files from the data store. It is also common to have a virtual file system layer, allowing mounting and accessing the de-duplicated archive as though it was a standard file system, with some overhead.

Depending on the type of data, de-duplication can deliver space savings from 2 to 200 times which is significant for organizations of all sizes. As an example, the same file may exist in multiple copies within a company’s storage network. Multiple versions of frequently edited files are also common. By saving only the unique data within these identical or similar files, de-duplication eliminates the need to store countless copies of the same data.

With most experts agreeing that the data being created worldwide each day is doubling in size every two years, we can expect that there will around 44 zettabytes or 44 trillion gigabytes by 2020. So where will this data be stored? The data deemed important by each and every person or business will need to be stored on personal devices, computers or servers. Without de-duplication’s space saving technology just think how much physical storage space would need to be used. It makes a lot of sense to only store the data that you need not additional copies. This also plays out in the sending or transferring of data across the networks that we use daily.

Source vs. Target De-duplication

Where the data de-duplication process actually happens is of importance. If de-duplication occurs near to where data is created, it is referred to as “source de-duplication”. Conversely, if de-duplication occurs closer to where the data is stored it is typically referred to as “target de-duplication”.

Source De-duplication ensures that data on the source server, desktop or mobile device is de-duplicated. This can take place directly within a file system, or specialized application storage mechanisms.

Advantages of source de-duplication include:

- Reduction of total data transfer volumes to data storage location

Disadvantages of source de-duplication include:

- The transfer process typically take longer as data de-duplication is being done before transfer
- Host CPU is being used by the backup application to perform data de-duplication, compression and encryption processes.

Target De-duplication is the process of removing duplicates when the data was not generated at that location. Example of this would be a server connected to a SAN/NAS, The SAN/NAS would be a target for the server. The server is not aware of any de-duplication, the server is also the point of data generation.

Advantages of target de-duplication include:

- Transfer times are much faster

Disadvantages of target de-duplication include:

- All data must be sent over the network to the storage area before de-duplication takes place

Source / Target Hybrid is a newer option offered by de-dup technology vendors and service providers where the de-duplication compute process is off loaded from the source device onto a purpose built appliance or VM. This reduces load on the source servers and when deployed near the source servers eliminates potential LAN/WAN network congestion challenges as faced with target de-duplication solutions.

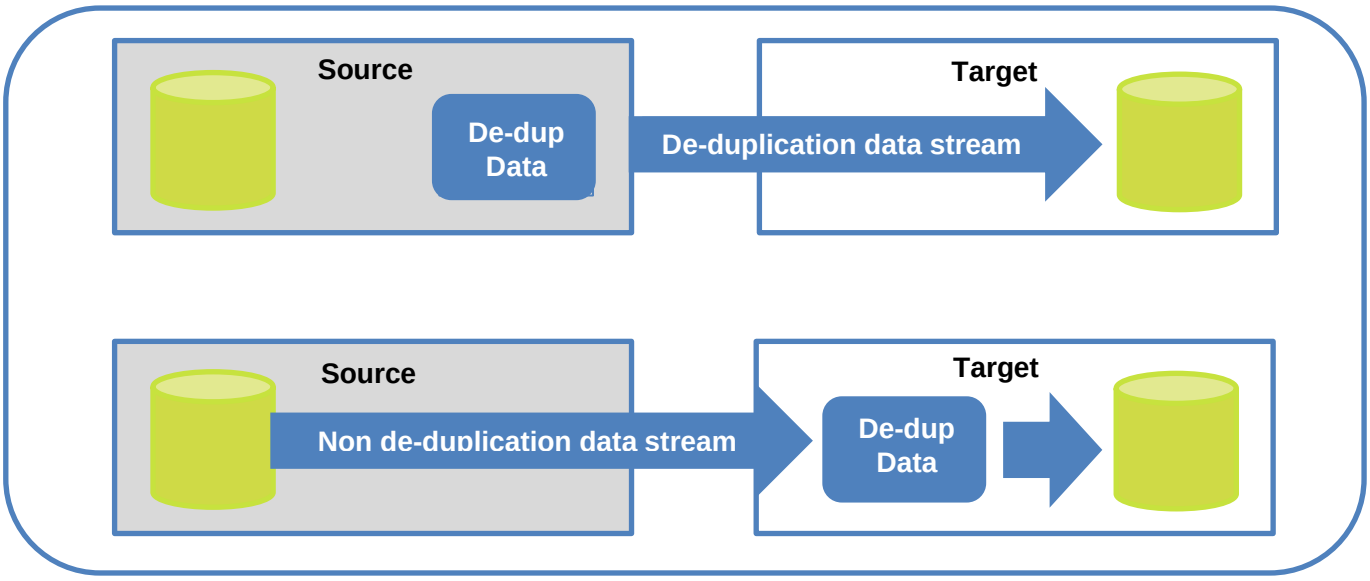

De-duplication Types

The term de-duplication refers to a few different generations of technologies. Within this group are three main evolutionary stages of de-duplication. The first is simple **file-level de-duplication**, and is considered legacy. The second is **differential calculation de-duplication**, and is also now considered legacy. The third and current generation is called **block-level de-duplication**, and is what this article is focused on.

File-level De-duplication works by simple file exclusion. For example, a batch of files may be added to the archive daily. Any pre-existing files are skipped when the next batch is added. Any new or changed files are added to the archive as complete files. This simple de-duplication can work well enough for certain needs, such as file-servers that have large amounts of small files. Because any changed files must be re-added in full, however, It is not well suited to large files that change a little from day to day (such as big databases, etc.). Even for file-servers, it's not very efficient in the long term. This generation is still common for source-based de-duplication, however, due to its simplicity.

Differential De-duplication was the first major advancement after file-level de-duplication. Files are initially added to the archive in full. When a file is changed and re-added to the archive, a differential calculation is performed, and only the changed blocks of the file are required to be stored. This is fairly efficient but it always requires a full "master" copy of the files in the data sore, in addition to the differential chunks.

There are two subcategories of differential de-duplication: master+delta, and master+incremental. Master+delta takes the initial "master" version of the file, then creates a "delta chain" from each subsequent version of the file added into the archive. Each delta in the chain contains the changed data since the last delta, all recursively linking back to the original master file. The deltas are relatively small in size, but the entire recursive "chain" of deltas is required order to retrieve any specific version of a file from the data store.

The second type, master+incremental, works a bit differently. Instead of a delta chain, each changed file that is added to the data store creates an "incremental" copy, which contains all changed data since the master copy. Each master plus incremental pair is self contained. This avoids the need to recursively walk the delta chain when retrieving a file from the data store, but the incrementals themselves are larger in size.

Block-level De-duplication is the current and most advanced level of de-duplication technology. There are a few important variations within this category, but the main components are the same. There must be an on-disk archive, or data store. There must be a database for keeping track of what data is currently stored within the data store. And finally, there must be a means to add, retrieve and remove data from this data store, which generally takes the form of proprietary software (though libre/free software exists).

To understand how block-level de-duplication works, it's necessary to understand the basic process of adding and retrieving files from the data store. When a file is first added to a new data store, it is broken down into smaller blocks of data, often called "chunks". The size of these blocks can be either fixed or variable, but the basic concept is the same. A file is broken into blocks, which are each checksummed and stored in the local database for tracking. At the same time the block is saved to disk, and the location is recorded in the database. This means that every file put into the data store becomes a number of virtual blocks of data, and the checksumming and block order is tracked in the database so that the file can be reconstructed later. If the same file is added to the data store again, after having been changed slightly, the same process occurs. The majority of the blocks will not have changed, however, so only the blocks which contain new unique data need be saved to disk.

The database will now contain the metadata necessary for reconstructing both the old and new files from their individual blocks. Since most of the blocks are the same in both versions of the file, however, most of the metadata for reconstructing the files will simply be pointing to the same blocks of data on disk. For common office files, and even large databases, there can be tens or hundreds of versions of a file that are sharing the majority of the same blocks, leading to space savings. When you retrieve a file, no matter which version, it will piece together the proper blocks necessary for reconstruction, regardless of when the blocks were added to the archive. This is the basic concept of how block level de-duplication works.

Fixed or Variable Block Lengths

Another element that requires consideration is the actual block size. This varies greatly per implementation, but is generally measured in kilobytes. Common sizes are standard byte multiples such as 32KB, 64KB, 128KB, 256KB, etc. When a changed file is de-duplicated, it is scanned for changes by the current block size. If any block is different, then that entire block is stored in the archive. Let's briefly look at the two options being offered today.

Fixed Block de-duplication involves a predetermined block size that is used for all files. If the block size is 64KB, and 8KB of the file is changed, the entire 64KB block is stored in the storage subsystem.

Advantages of fixed block length include:

- Less processor power required as less meta data is created and required to be tracked
- De-duplication process can be faster as less variable data chunks exists
- Effective with smaller data storage volumes

Disadvantages of fixed block length include:

- Wasted storage across block barriers, due to storing the entire block size when only a tiny portion of the file has changed
- Not efficient for tiny files, such as logs which may be smaller than the predetermined block size.
- Most implementations do not account for the less common case of data being appended to the front of the file

Variable Block de-duplication involves using algorithms to determine a variable block size. This could mean that each file has its own best block size, or each file contains multiple differently sized blocks. The data is split based on the algorithm's determination. Then, those blocks are stored in the subsystem.

Advantages of variable block length include:

- Reduction of data storage requirements especially with very large data stores
- Smaller variable blocks get better de-duplication than larger or fixed block ones
- Variable-length is a more advanced and anchors variable-length segments based on interior data patterns

Disadvantages of variable block include:

- More CPU intensive
- Can be slower, due to multiple passes being required to match multiple block sizes
- Space savings compared to fixed block de-duplication can be minimal for larger growing files
- Large amounts of metadata are required for tracking blocks, which can sometimes offset the more efficient de-duplication

De-duplication Process Visual

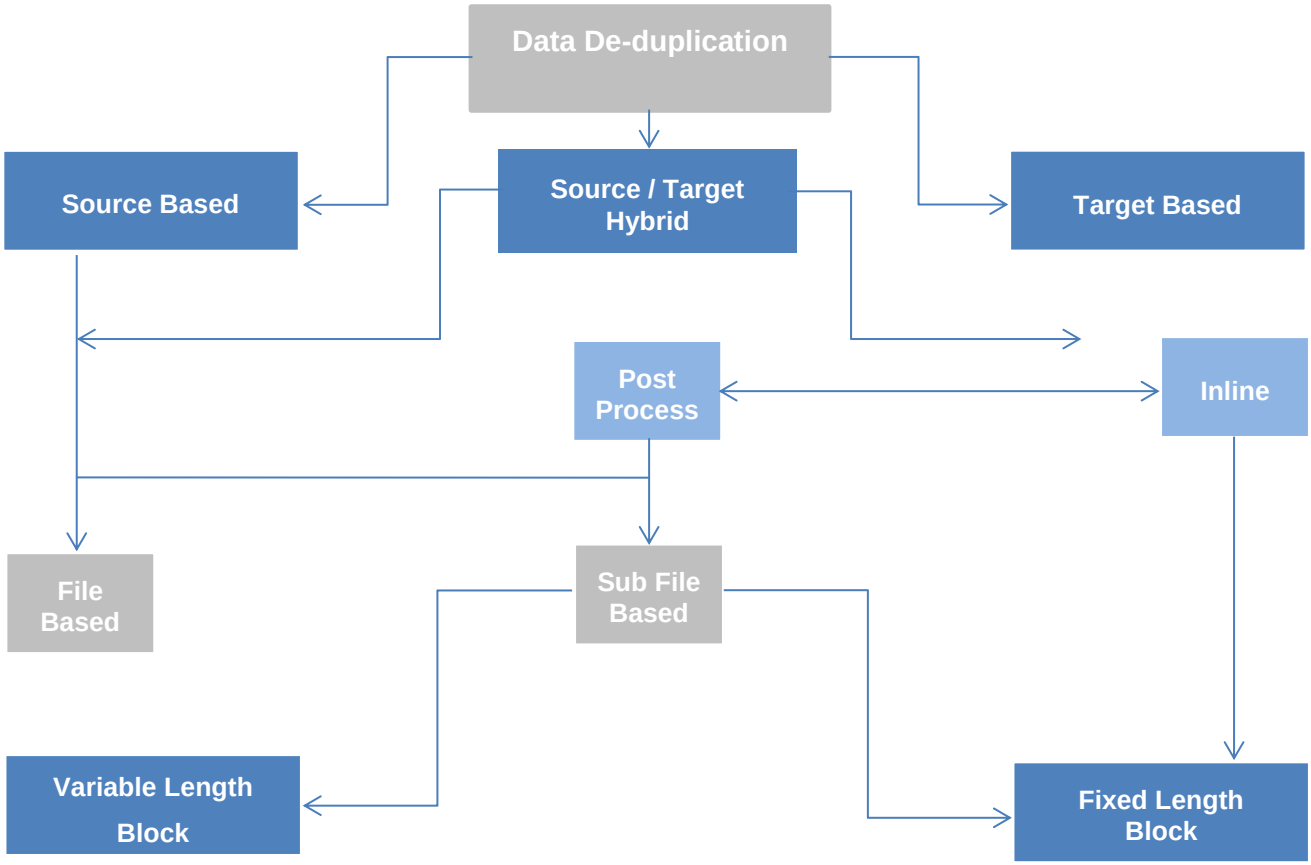

Conclusion

Implementing data de-duplication technology was the sole domain of large enterprises. However, over time the decoupling of software from proprietary hardware has made the solution affordable to businesses of all sizes. Typically used with secondary storage or data backup systems because these receive data from multiple sources making it the ideal place to identify and reduce duplicate data. We also see primary storage de-duplication solutions in the market. The use of de-duplication technology will continue to grow as will the following list of benefits:

- Less overall investment in storage hardware and software
- Increased storage space savings – industry average is 20x
- Decreased network bandwidth needs
- Faster data recovery time objectives
- Decreased daily, weekly and monthly backup window times
- De-duplication can be applied to a wide variety of applications and data

About Blueshift Data Protection

Founded in 2004, Blueshift is a pioneer in the private, public and hybrid cloud backup and recovery services market. Headquartered in Tokyo, Japan, Blueshift delivers secure, off-site, data backup and recovery solutions for clients ranging from small, medium-sized to large organizations. We develop and deliver unique agent-less cross platform solutions that effectively reduce client costs, decrease risks and improve service levels.

As a thought leader in the backup and disaster recovery market, Blueshift continues to invest in people, technology and markets throughout the Asia-Pacific region. If you should have any questions about our company or technology feel free to visit, follow or contact us using the information provided below.

Visit Us	**Email Us**	**Follow Us**
You can visit our website at www.dataprotection.co.jp and go to our financial industry page	Please send your emails to the following address info@dataprotection.co.jp	You can follow us on Twitter, Facebook, LinkedIn, Google+ or search for us on YouTube