Demystifying Deduplication for Backup with the Dell DR4000

This Dell Technical White Paper explains how deduplication with the DR4000 can help your organization save time, space, and money.

John Bassett

SMB Marketing

January 2012| Rev 1.0

Contents

Tables

Figures

Executive summary

In this paper, we examine the new technology of data deduplication. Data deduplication, or dedupe for short, can yield enormous savings, making disk-to-disk backup more feasible. However, not all deduplication technologies are equal. We will look at the different approaches to data deduplication, as well as the costs and benefits of each. In this context, we will present Dell's Ocarina Deduplication technology and demonstrate why it is a best-of-breed choice to handle your data deduplication.

Introduction

According to Gartner, data growth presents the biggest challenge for a data center's hardware infrastructure.[1] As data continues to grow at a rate of 50 percent per year,[2] IT organizations are struggling to find a backup solution to keep up with their growing storage environments.

In addition to the sheer volume of data, there are other pressures on the storage infrastructure. Access to current data is essential to the running of any business. A disruption in the availability of data can break business continuity. Any good business impact analysis must consider how quickly access to data must be restored (the recovery time objective or RTO) and how recent the data must be (the recovery point objective or RPO).

Tape backup is no longer an effective solution for all data protection scenarios. Many organizations are turning to disk-based backup solutions to assist in augmenting or in some cases, replacing their existing backup solutions. However, this backup method presents challenges of its own. The need for greater retention of data creates an enormous and expensive appetite for storage.

Deduplication is so powerful because nearly 75 percent of all data that a typical organization stores is a copy.[3] Sending an e-mail attachment to multiple recipients creates a copy of the attachment for every recipient. Cloning virtual machines creates multiple copies of identical operating system images. Many other essential processes create identical copies of data. This is why deduplication technology can be so powerful—simply put, deduplication eliminates the copies. For obvious reasons, writing a single copy to the disk, rather than 100 copies, can save an enormous amount of space. Deduplication makes it appear that the user or application has its own copy. This is essential, because it makes the space savings totally transparent. All existing processes continue to work without modification, while using considerably less storage.

Removing redundant data brings many benefits. For example, with deduplication, backing up data from remote locations requires far less data to be transmitted across the WAN. By reducing bandwidth requirements with deduplication, organizations gain flexibility in architecting their backup and recovery plans. Greater automation and more frequent replication become possible, backing up data from remote offices to a central location is more efficient and data recovery operations become more reliable and require less supervision than ever before.

Much like virtualization, data deduplication has the potential to be a transformative technology in the data center. Before virtualization, setting up a system with a special purpose was expensive and labor

---

[1] *User Survey Analysis: Key Trends Shaping the Future of Data Center Infrastructure Through 2011* Gartner.
[2] *A Digital Universe Decade – Are You Ready?* IDC.
[3] *Ibid.*

intensive. Adding memory or processors to a server used to require physical access to the box, but now requires only a few clicks of a mouse. In many cases, users can even reconfigure their VMs without involving IT at all. Because deduplication radically reduces the storage needed for backup, it can reduce hardware costs considerably. The benefits of deduplication don't end with hardware costs. Deduplication can also mean more frequent, targeted backups and disaster recovery designed for an organization's specific needs, without restriction by the data center infrastructure. In the same way that virtualization gets more work from the same hardware, deduplication lets you get better recovery from the same storage.

What is deduplication?

Dedupe reduces the amount of stored data by removing each redundant copy of data, leaving a pointer to the original data, now stored once, in its place. This concept is a familiar one, because a URL is a type of pointer; when someone shares a video on YouTube, they send the URL for the video instead of the video itself. There's only one copy of the video, but it's available to everyone.

Deduplication uses this principle in a more sophisticated, automated way. Data is deduplicated as it is written, and the reduction in space is much greater than in the example above.

Deduplication can be applied to data in primary storage, backup storage, cloud storage or data in flight for replication, such as LAN and WAN transfers. Deduplication has the potential to decrease disk capacity and bandwidth requirements by as much as 98 percent.

Figure 1. Deduplication reduces the amount of stored data.

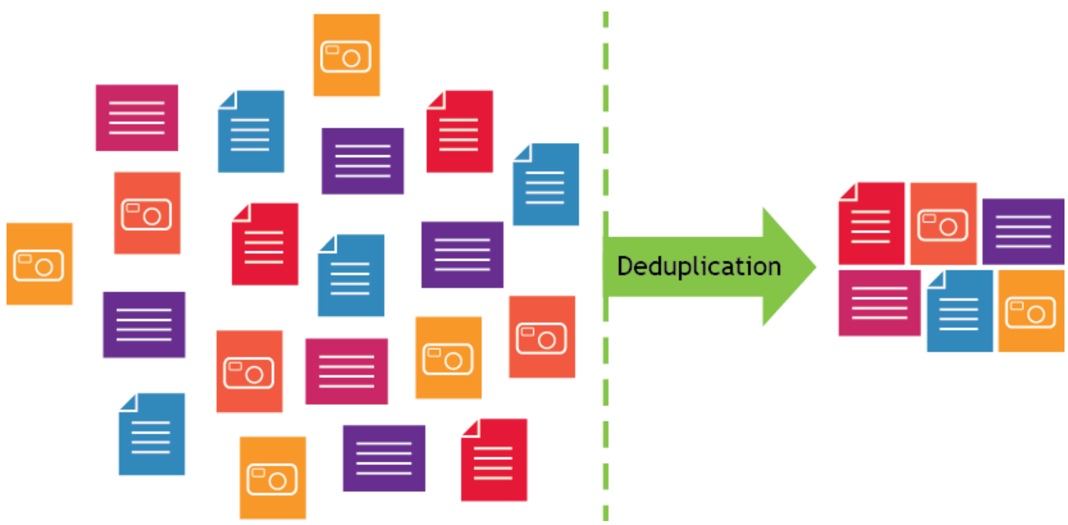

When considering deduplication solutions, a key variable to consider is the effect on the ingest rate, the rate at which the server can accept and store data without falling behind. Any process added to the storage pipeline has the potential to slow the ingest rate, and thus reduce server performance.

Chunking: A fundamental concept

Simple data deduplication works on fixed-size blocks. More advanced data deduplication methods work by organizing data into chunks. A chunk can be as large as is appropriate. It might be 4KB, 32KB, or whatever size works best for the data at hand. The use of the word "chunk" is similar to that in psychology; people learn to process information more efficiently by grouping it in ever larger "chunks."

The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

Figure 2. Deduplication uses chunking to sort through data being stored.

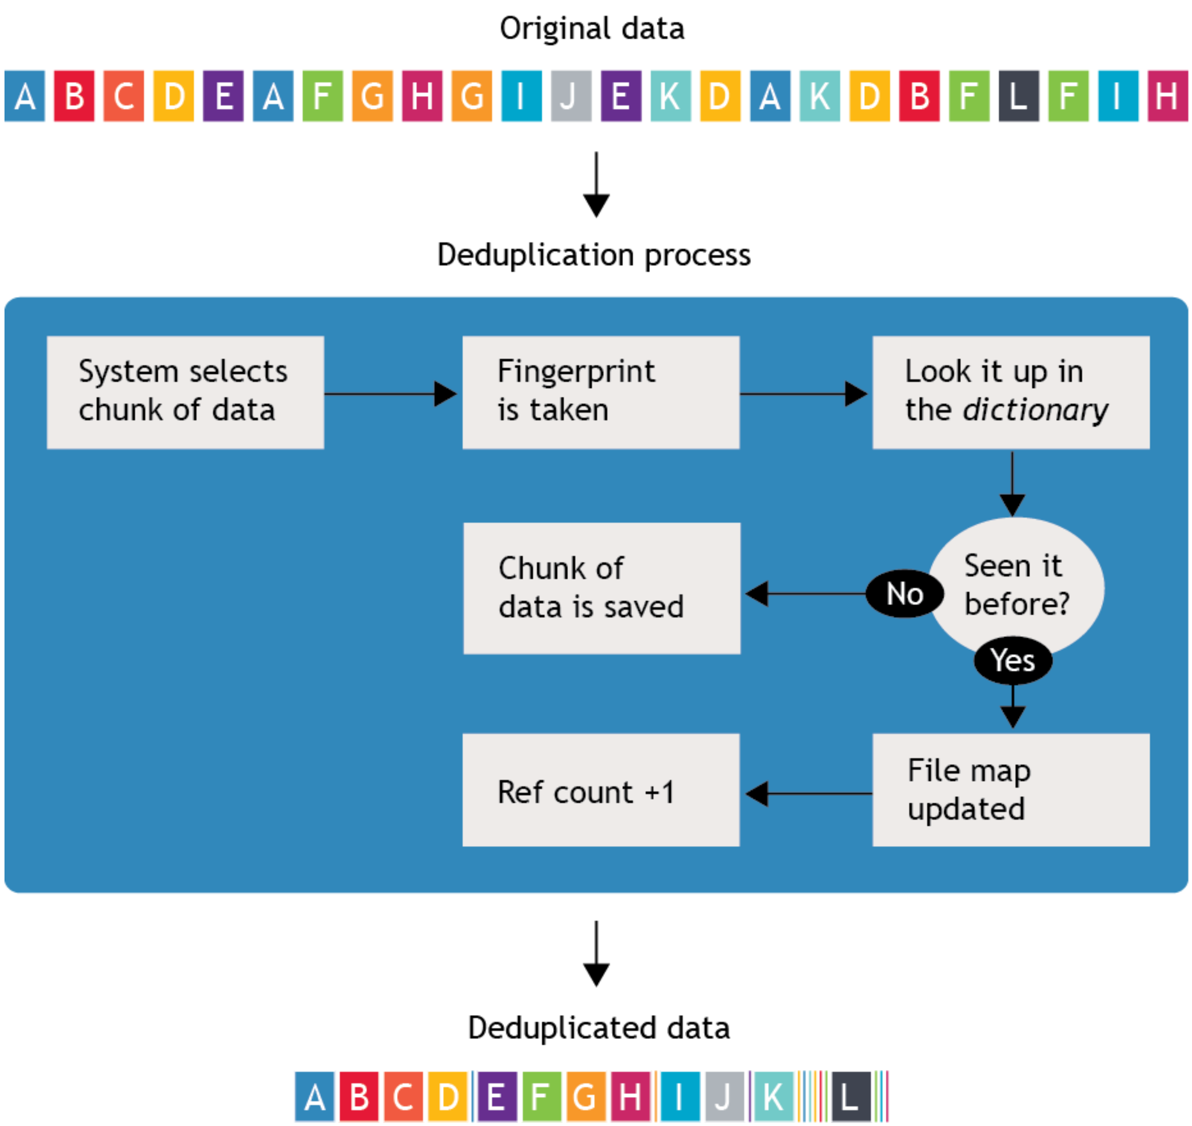

Chunking methodologies for deduplication

Most data deduplication engines work by breaking the data into chunks. However, the ways that the different products break the data into chunks differ considerably. As in all things, there are tradeoffs. More advanced methods can yield significantly better reduction ratios, but may also require more CPU, disk or memory resources. There are two fundamental methods of block-level chunking exist: fixed window/fixed block and sliding window/variable block. Each has its advantages and disadvantages.

Fixed window/fixed block

The simplest method of block-level chunking is fixed window/fixed block. In this method, the block size never changes. While different devices may use different block sizes, the block size for a given device using this method remains constant. The device always calculates a fingerprint on a fixed block and sees if there is a match. After a block is processed, the window advances by exactly the size of a block and the process repeats. While this method requires the minimum CPU overhead, because it does not account for shifts in data, it is less effective than more advanced methods.

Figure 3. How the fixed window/fixed block method chunks data.

Sliding window/variable block

The variable block method has all the advantages of the fixed block method, because of its ability to increase the block size allowed in the dedupe process to find sections of duplication much larger than a block. As explained in the section Chunking: A fundamental concept, the Dell DR4000 automatically determines the optimum chunk size. The Dell DR4000 uses semantic information and the structure of the content in the backup stream to find the natural boundaries of the data. This lets the Dell DR4000 compensate for shifts in the data, and still find the optimal chunks in a repeatable manner. This method yields better reductions than the other two methods, but requires greater CPU overhead.

Figure 4. How the sliding window/variable block method chunks data.

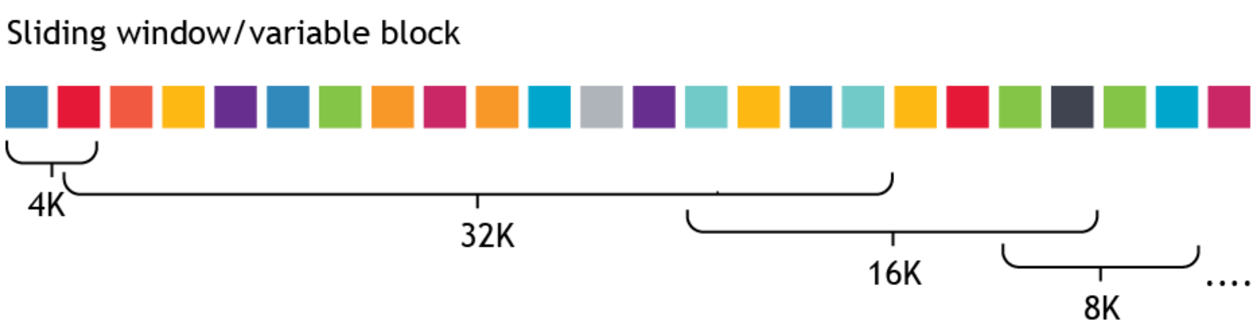

Dell has designed the DR4000 to utilize the sliding window/variable block method because it offers the best deduplication ratios. Consider a PowerPoint presentation that changes slightly from one backup to the next, for example by adding one slide. This shifts the data in the backup stream and can defeat fixed block methods. However, as explained above, the Dell DR 4000 is able to use the content and structure of the data to find the natural boundaries, and so get the best deduplication across backups.

Because the DR4000 is a hardware appliance, it can offer these reductions without loading the CPU resources of the servers in the data center. Dell specifically optimized the underlying hardware in the DR4000 for this chunking method, and designed it for maximum data ingest rates.

Inline vs. post-process deduplication

By definition, deduplication requires processing data to determine and reduce data copies. Data can be processed at three points: before it is written to disk (inline deduplication), after it is written to disk (post-process deduplication) or both before and after it is written to disk. Each has its advantages and disadvantages, which are important to understand because your chosen method affects your ingest speed and data center footprint.

Inline deduplication processes the data before it is written to disk. This means that only the deduplicated data goes to the disk, yielding more predictable performance and making it easier to properly size a solution. However, inline processing governs ingest speed. Because of the demands of processing data in real time, keeping performance at acceptable levels can be very computationally intense.

Post-processing deduplication processes the data after it is written to disk. After the data is written, the deduplication process starts and reclaims the disk space. Because the data does not have to be processed in real time, the computational demands are lighter. However, post-processing has the following disadvantages:

- Greater space requirements: An interim period exists during which you need sufficient available storage to hold both the original data and the deduplicated data. At a minimum, this leads to more complicated sizing calculations, greater storage requirements, and more configuration. At the worst, it is possible to run out of space using this method, especially if the dedupe process from the previous backup did not complete.
- Greater bandwidth requirements: In order to post process the data, you must first write the duplicate data and read it back to be deduplicated, and then write the deduplicated data to storage again. The greatly increases the I/O requirements for the backup.
- Greater time requirements: Because the backup consists of two separate, sequential processes, the amount of time required for the backup increases. This can limit the opportunities to perform a backup without affecting operational efficiency.

A third method is to deduplicate the data inline, and then apply a post-processing step. This allows for predictable sizing and even greater space savings. With inline deduplication, replication can be performed as data is ingested, saving considerable time over the two-phase post-process method. It also has the advantage of faster restores of the most recent backup.

The Dell DR4000 has made further optimizations, and can actually dedupe and compress as part of the same inline process. This provides the benefits of compression without requiring that space be dedicated to staging uncompressed data.

Figure 5. The four approaches to deduplication.

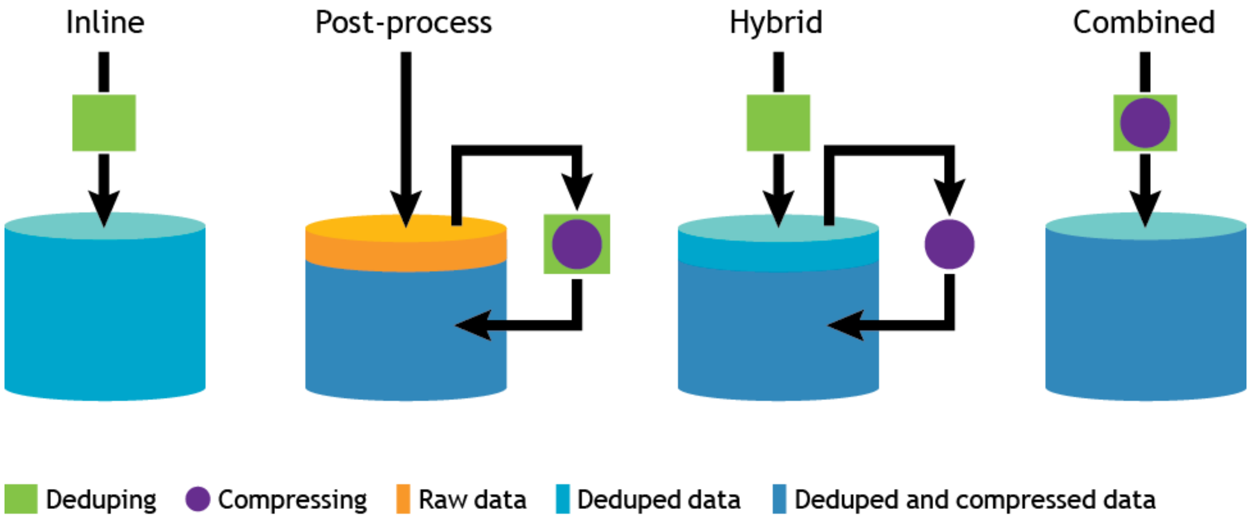

Target vs. source-based deduplication

Now that you've decided how to deduplicate your organization's data, the next question is *where* to do it. You can deduplicate the data at the source, before it goes over the network, or you can dedupe it when it arrives at the target storage. Selecting the correct method for your organization is critical to the success of the deployment.

Target-based deduplication requires that the backup server and/or target dedupe appliance handles all of the deduplication. This means no overhead on the client or server being backed up. This solution is transparent to existing workflows, so it creates minimal disruption. However, it requires more network resources because the original data, with all its redundancy, must go over the network. Target-based solutions can be provided via software or a dedicated hardware appliance such as the Dell DR4000.

Source-based deduplication processes the data before it goes over the network. This reduces the amount of data that must be transmitted, which can be important in environments with constrained bandwidth. This method is often used in situations where remote offices must back up data to the primary data center. In this case, deduplication becomes the responsibility of the client server.

Dell believes that replication using target-based dedupe solutions at both the primary facility and the remote office is the best option, offering the greatest flexibility, the most predictable performance and the shortest RTO. By deploying systems in both locations, backup data is deduplicated at the target backup device at the remote office. Once the remote backup completes, the deduplicated data can be replicated to the primary data center. With this architecture, you can restore the data locally at either the primary datacenter or the remote office, significantly decreasing the recovery time objective for the remote office recovery situations, because the data is available locally.

Software vs. hardware appliance solutions

The deduplication process can be implemented in hardware or software. Because deduplication makes demands on CPU resources, RAM and storage, and can significantly affect ingest performance, it is important to understand the tradeoffs and make an informed decision.

Some vendors have integrated dedupe functionality into their backup software offerings. The obvious advantage with this scenario is that you get deduplication with very little effort because it is already integrated into the backup software. However, it does not come without cost; these costs can be direct or in the form of added hardware requirements.

The direct costs for software deduplication come in the form of licenses and fees. The software ISVs offer dedupe capability in many different pricing models. For example, some charge a flat fee to license the capability as an add-on, while others offer a usage-based model where your cost is based on the number of terabytes of data you deduplicate. Some even provide the functionality for free. An existing data center is unlikely to change its existing backup software based solely on deduplication, but knowing how your ISV charges for deduplication is one of the cost factors to consider.

As for added hardware requirements, it is important to remember that in the software solution model, deduplication is the responsibility of the media/backup server. This means the server not only has to run the traditional functions of the backup software, it also has the added work of deduping the data. In order to reduce the strain on the backup server, the ISV dedupe algorithms typically use a basic fixed window/fixed block algorithm. While this does reduce the demands on the server, it does not produce the best compression, as noted above. So, the resulting deduplicated data may be larger than it would be with a more sophisticated solution.

Even when software dedupe uses simpler algorithms, having all the overhead on a single server limits ingest performance. When the ingest performance is too severely constrained, multiple media servers must be deployed to provide the performance needed to maintain backup windows. Because each media server has its own dedicated storage, this adds to hardware sprawl. Running multiple copies of the deduplication application can lead to deduplication silos. In this situation, each dedupe stream is deduplicated, but duplications across the silos are not considered. This makes the suboptimal deduplication from the fixed window/fixed block algorithms even worse. With multiple instances of deduplicating products, impacts of 20 percent to 50 percent, or even more, can be expected.[4] As explained above, the Dell DR4000 builds on this by adding inline compression for even greater storage savings.

Figure 6. How data siloing uses more storage.

[4] *Maximizing Data Efficiency: Benefits of Global Deduplication* Advanced Storage Products Group

In contrast, a hardware solution such as the DR4000 is dedicated to only one job, deduplicating data. It offloads the CPU-intensive dedupe process, reducing the requirements for the backup server, to the point where it may not require a dedicated server and could be virtualized.

Centralizing deduplication to the DR4000 eliminates unnecessary hardware. You no longer require the extra media servers, or the extra pools of storage that they would require. Beyond the hardware savings, a hardware-based deduplication solution like the DR4000 yields higher deduplication ratios than individual storage pools because more data is compared against the same dedupe dictionary.

Figure 7. Software deduplication requires more hardware.

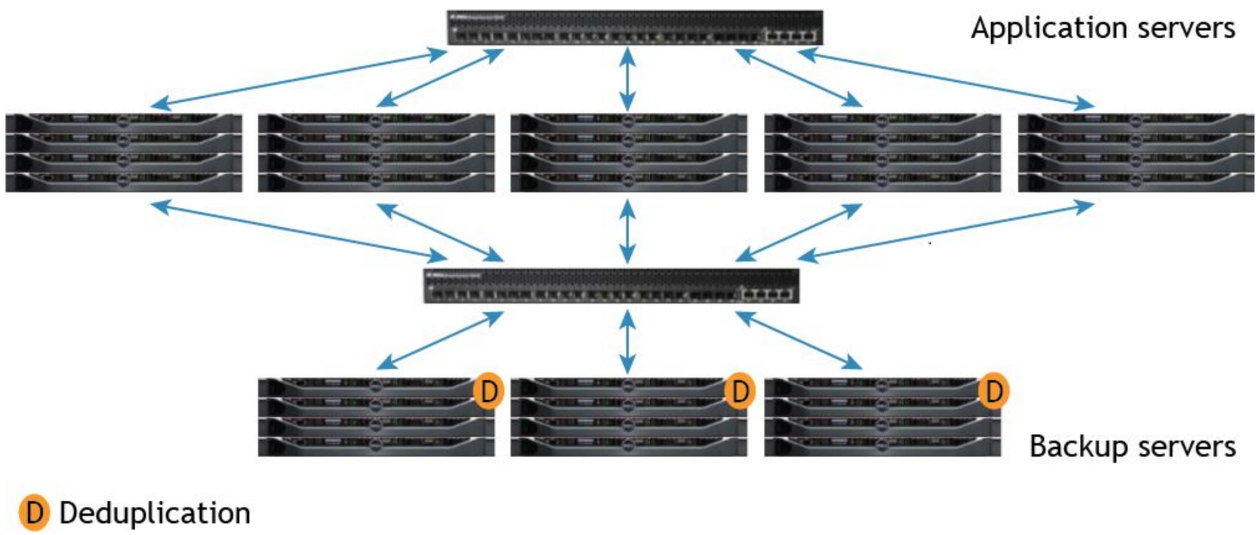

Figure 8. The Dell DR4000 reduces your need for multiple backup servers.

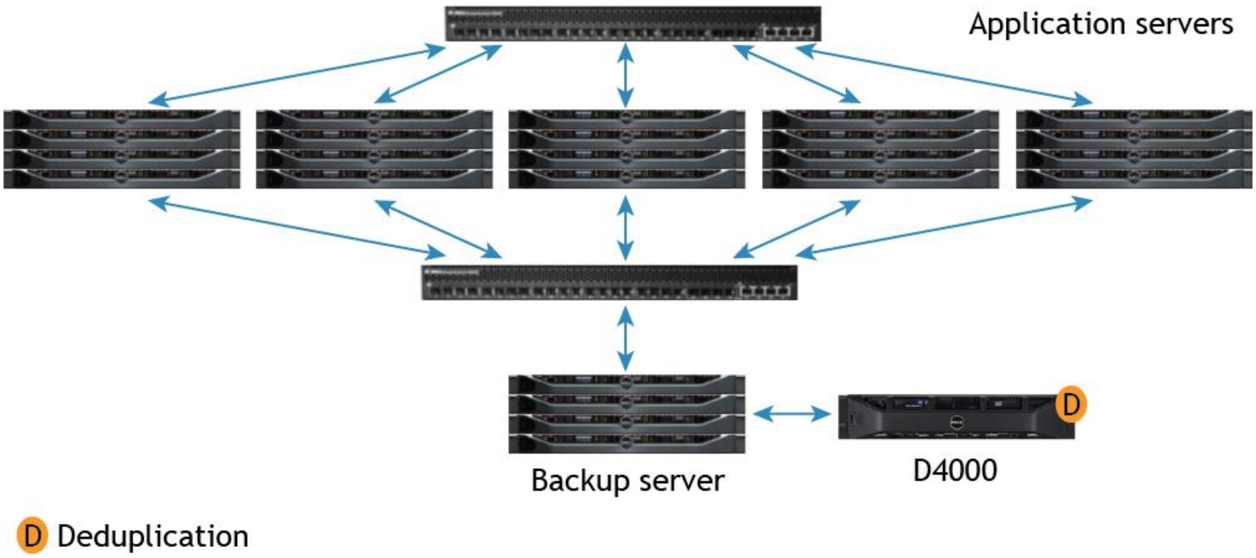

The Dell DR4000 is a target-based deduplication appliance specifically engineered to handle streaming backup workloads. Dell tests and certifies the ingest performance with third-party validation.

Of course, superior performance is no good if the data is not protected. The DR4000 includes a number of software and hardware features designed to protect data integrity. For example, the DR4000 utilizes RAID 6 as the foundation of the hardware protection. RAID 6 improves on RAID 5 by using two parity blocks with block-level striping across all member disks. It also uses NVRAM to protect data in the event of power loss. These data protection features make this a hardened appliance suited for enterprise data protection.

Calculating deduplication ratios

In the backup-to-disk arena, deduplication ratios are hotly debated and very controversial. The deduplication ratio is the measure of how effective the deduplication process is. It is calculated by dividing the amount of Cumulative Logical storage by the Cumulative Physical Storage. For example, if the dedupe process allows 400 GB of data to be stored in 200 GB of physical storage, the dedupe ratio is 2x. If the dedupe process stores 400 GB of data in 100 GB of physical storage, the dedupe ratio is 4x.

It seems each new vendor of deduplication hardware claims to offer higher dedupe ratios than the last vendor did. In reality, because many variables determine the dedupe ratio for a given data set, there is no guarantee that you will see the same ratio another data center saw. When comparing deduplication ratio claims from vendor to vendor it's important to understand what these ratios mean. Dell estimates a 15x ratio which is about 93% reduction. Dell chose a 15x average based on a retention period of 12 weeks. Most vendors claim between 10x-30x reduction which represents 90-97%.

The variables that influence the dedupe ratio for a given workload include the following:

- The type dedupe algorithm. As noted above, sliding window/variable block algorithms typically provide better ratios than fixed window/fixed block.
- The type of data being backed up. Not all data sets have the same amount of duplicate data or compressibility.
- The frequency of backups. More frequent backups will build the dedupe dictionary more quickly.
- The retention period for backup jobs. Longer retention yields higher ratios.

Consider the effect of backing up data over time. On the first full backup, the deduplication ratio is limited to the amount of redundancy in that single backup. Most of the time, a small amount of commonality will be found in the data set, and deduplication rations of 2x to 4x are typical.

As additional backup jobs run, the deduplication is based on the data set at hand and a dedupe dictionary built over all previous backup jobs. Thus, the more full backups that are stored, the higher the deduplication ratio will be.

Most organizations retain backup data on disk for 30 to 90 days. When a deduplication device is used with this type of retention policy, a deduplication ratio of 15x is typical. Organizations that retain data longer should see higher ratios than organizations who only store data for a short time. If you were to continue to store the backup data for a year, it is very possible to see much higher deduplication ratios.

Table 1. How deduplication ratios can improve over time.

Time	Backup data	Cumulative logical storage	Estimated reduction[1]	Cumulative physical storage
First full	5 TB	5 TB	2x	2.5 TB
Week 1	5.2 TB	10.2 TB	4x	2.7 TB
Week 2	5.4 TB	15.6 TB	6x	2.9 TB
Week 3	5.4 TB	21 TB	7x	3.1 TB
Month 1	5.4 TB	26.4 TB	8x	3.3 TB
Month 2	22.6 TB	49 TB	12x	4.1 TB
Month 3	24 TB	73 TB	15x	4.8 TB
Month 4	25 TB	98 TB	18x	5.4 TB
Month 5	26 TB	124 TB	21x	5.9 TB
Month 6	27 TB	151 TB	23x	6.5 TB
	TOTAL	**151 TB**	**23x**	**6.5 TB**

[1] Savings will vary based on data.

Summary/Why Dell DR4000 for backup deduplication?

In this paper, we have explained deduplication and how it has the potential to revolutionize the data center. We discussed basic concepts such as chunking, as well as different strategies for implementing deduplication. We also looked at the impacts different deduplication strategies have on storage savings and ingest rate, and what strategies Dell chose to implement when designing the DR4000 appliance. By utilizing inline, target-based deduplication with the sliding window/variable block chunking method, the Dell DR4000 is a best-of-breed solution for your deduplication needs.