



Hyper-converged Secondary Storage for Backup with Deduplication | **Q & A**

The impact of data deduplication on
the backup process



Table of Contents

Introduction.....	3
What is “data deduplication”?	3
Is all data deduplication created equal?	4
What factors most impact the deduplication ratio?	5
Does data deduplication impact backup windows, restores, VM boots, and offsite tape copies?	5
Should I deploy “inline” deduplication?	6
Should I deploy “post-process” deduplication?	8
Should I deploy “adaptive” deduplication with a landing zone?	8
Should I use a scale-up or scale-out storage architecture?	9
Can I use the cloud for my disaster recovery (DR) data?	9
What is the “right requirements” list?	10
To achieve the right requirements, what should a solution offer?	11
Summary.....	11

Introduction

There are many misconceptions about data deduplication, and making decisions based on those misconceptions can produce undesirable (and unplanned) results. For instance, deployment of the wrong type of deduplication typically results in:

- excessively high disk usage and using as much as three times the bandwidth for offsite replication,
- slower backup storage ingest due to inline compute-intensive data deduplication that greatly slows backups down and expands the backup window,
- slower restores, VM boots, and tape copies that can take hours or even days due to the time-consuming rehydration of deduplicated data, and
- backup windows that continue to expand with data growth.

Choosing a hyper-converged secondary storage for backup with data deduplication solution will have a major impact on the cost and performance of your backup environment for the next three to five years.

What is “data deduplication”?

Data deduplication looks at incoming data, breaks it into smaller block or zone sizes, and then utilizes different techniques to compare the blocks or bytes within. Only unique blocks and bytes are stored so that redundant data doesn't take up valuable disk space. For primary storage or archive storage, the deduplication ratio ranges from 1.2:1 to as high as 1.8:1. This is actually the same as, or worse than, standard data compression, so data deduplication doesn't bring much value to primary or archive storage – but backup is different.

With backup, multiple copies are kept for version history. It isn't uncommon to keep 12 weeks of backups and then monthlies for 3 years. In this case, you'd have over 40 copies, and data retention is where data deduplication has the most value. If you have a 100TB full backup and keep 40 copies, you'd need 4PB of disk. Even if you compress the data at 2:1 using standard compression, you'd still need 2PB of disk. Using data deduplication, the first 100TB copy can be stored in about 50TB of disk, and each subsequent copy will require about 2TB since approximately 2% of the data changes each week (i.e., 2TB on a 100TB full backup). In this example, you'd have 50TB plus 78TB (2TB per week x 39 copies) equals 128TB of disk.

The deduplication ratio is calculated by taking the amount of disk required without data deduplication divided by the amount of disk required with data deduplication. In the example, 4PB divided by 128TB equals a deduplication ratio of 31:1. The longer the retention period, the greater the deduplication ratio. If you keep one copy for retention, you may achieve 1.8:1, and if you keep four copies you might achieve 3:1; however, at 18 copies, the industry average is about 20:1. The industry average deduplication ratio of 20:1 is calculated at about 18 weeks of retention.

Is all data deduplication created equal?

No, and the differences are significant. Unlike standard data compression where the most you can achieve is 1.8:1 or 2:1, data deduplication is all over the map based on the deduplication method used. Vendors employ vastly different approaches, depending on how much resource (processor and memory) can be applied.

Backup applications that perform data deduplication on the media server use less aggressive algorithms to minimize processor and memory usage and, as a result, they achieve a much lower deduplication ratio. Target-side appliances that have built-in, dedicated processor and memory use much more aggressive algorithms and therefore achieve higher deduplication ratios. The lower the deduplication ratio, the more disk will be used over time (especially with longer-term retention) and the more bandwidth will be required for replication. Using deduplication in the backup application may save money up front, but over time the cost will be three to four times that in additional disk and bandwidth.

Typical backup application deduplication:

- 1MB blocks = 2:1
- 128KB blocks = 6:1
- 64KB blocks = 8:1

Typical target-side appliance deduplication:

- 8KB blocks = 12:1
- 8KB blocks with variable-length content splitting = 20:1
- Byte level = 20:1
- Zone stamps with byte level compare = 20:1

It's important to know what approach is being used by the backup application or target-side appliance since the related costs for disk and bandwidth can vary greatly.

What factors most impact the deduplication ratio?

There are three variables that have the biggest impact on deduplication ratio:

1. **Retention period** - A longer retention period provides more backup copies in which to find repetitive data, so the deduplication ratio will be higher.
2. **Algorithm** - A more aggressive algorithm will produce a higher deduplication ratio.
3. **Data type and mix** - Different data types deduplicate at different ratios. For instance, unstructured files achieve a ratio of 6:1 or 7:1, but databases can achieve 100s:1 assuming, of course, an aggressive algorithm. Compressed and encrypted data does not deduplicate and achieves a 1:1 ratio. Many vendors achieve a ratio of 10:1 to as much as 50:1 with an average of 20:1 (at an average 18 weeks of retention), depending on the mix of data types.

Does data deduplication impact backup windows, restores, VM boots, and offsite tape copies?

Yes. Although data deduplication reduces disk storage and bandwidth requirements as well as the related costs of each, it also creates three additional compute problems. Data duplication by its very nature is extremely compute intensive. It takes a great deal of processor and memory to break data into smaller segments to compare, and it takes even more compute to put the deduplicated data back together again when it's needed, a process called "rehydration." In addition, as data volumes increase, the amount of data to be deduplicated also increases as does the time required to do so, resulting in an ever-expanding backup window.

To summarize, the three new problems that arise as a result of compute-intensive deduplication are:

1. slow ingest resulting in slow backups and long backup windows,
2. slow restores, VM boots, and offsite tape copies due to the compute-intensive nature of data rehydration, and
3. a backup window that continues to expand to accommodate growing data with eventual spilling over into time that is outside of the allotted backup window.

Should I deploy “inline” deduplication?

Inline deduplication means that the data is being deduplicated on the way to disk. There are three approaches to inline deduplication:

1. **Adding inline deduplication to the media server software of the backup application.**

In this case, the media server platform is shared with core media server tasks and now also compute-intensive deduplication, slowing backups down considerably. To offset slower backups, backup applications use less aggressive algorithms in order to use less compute resources. As a result, however, they use more disk over time (longer retention periods) and more bandwidth to replicate. Most backup applications achieve a 2:1, 4:1, 6:1, or 8:1 deduplication ratio, depending on the size of the fixed-length blocks they use. They further require an expensive media server with flash storage, numerous dual core processors, and a lot of memory.

Some backup application vendors recommend which servers to buy and allow you to use your own preferred vendor for disk. Others package the media server software, the physical server, and the disk in a single solution. In all cases, the backups will be slower than target-side appliances that deploy dedicated hardware, and the amount of disk and bandwidth will be three to four times greater. In addition, all data is stored in deduplicated form so for each restore, VM boot, or offsite tape copy request, the data has to be rehydrated, resulting in VM boots that can take hours and offsite tape copies that can take days.

2. **Adding inline deduplication to a dedicated appliance with a scale-up storage architecture (i.e., front-end controller and disk shelves).**

This approach is faster than performing deduplication on the backup media server because all the system resources are dedicated to data deduplication. These appliances employ a more granular and aggressive algorithmic approach, which achieves a much higher deduplication ratio and saves additional storage and bandwidth. However, the ingest speed is still slow since the more aggressive approach uses more compute.

This approach is faster than inline deduplication in the backup software but is not fast enough to stay within allotted backup window times. Backup speed is further compromised when replication is turned on, as replication competes for processor

and memory along with deduplication, and if encryption is turned on, performance drops yet again. In addition, all data is stored in deduplicated form, so for each restore, VM boot, or offsite tape copy request, data has to be rehydrated, a process that can take hours for a VM boot and days for an offsite tape copy.

3. **Adding inline deduplication to a dedicated appliance, with a scale-up storage architecture (front-end controller and disk shelves) with a software option that is installed on the media and application servers.**

There is software that can be installed on the media server or database server if utilities such as SQL dumps, Oracle RMAN, etc. are used. This approach increases the ingest rate by doing some of the deduplication work on servers on the networks, thus borrowing compute from somewhere else. The drawback is that this takes compute from the media or database server, pushing the bottleneck further up the chain.

This approach does improve the ingest rate of dedicated inline deduplication appliances but has two shortcomings. The first shortcoming is that the media server or backup server will be slower since you have to install and run software on the media servers and production database servers. These appliances employ a more granular and aggressive algorithmic approach, which achieves a much higher deduplication ratio to save additional storage and bandwidth. When replication is turned on, the backups further slow down as replication competes for processor and memory along with deduplication, and if encryption is also turned on, performance drops yet again. The second shortcoming is that even though ingest will increase with the software add-ons, all data is still stored in deduplicated form, so for each restore, VM boot, or offsite tape copy request, data has to be rehydrated, a process that can take hours for a VM boot and days for an offsite tape copy. In other words, running software on media and database servers to increase ingest does not change the fact that all of the data is deduplicated and will need to go through the same time-consuming rehydration process for restore requests.

Should I deploy “post-process” deduplication?

Post-process deduplication allows backups to write direct to disk, avoiding the compute-intensive overhead of deduplication, resulting in fast ingest. In addition, the most recent backups are stored in their complete non-deduplicated form for fast restores, VM boots, and tape copies; there’s no need for time-consuming data rehydration since over 95% of restores, VM boots, and tape copies come from the most recent backup.

The downside to this approach is that backups need to complete before deduplication and replication begin, which creates a poor RPO (recovery point) at the disaster recovery site.

Should I deploy “adaptive” deduplication with a landing zone?

Adaptive deduplication is the best of all worlds. Backups are as fast as post process and three times faster than inline because they’re sent direct to a disk landing zone, avoiding the compute overhead of deduplication while backups are running. The most recent backups are stored in the landing zone in their complete non-deduplicated form for fast restores, VM boots, and tape copies; there’s no need for time-consuming data rehydration since over 95% of restores, VM boots, and tape copies come from the most recent backup. However, in contrast to inline deduplication, which slows backups down and only stores deduplicated data, or post-process deduplication, which doesn’t occur until all the backups are complete, adaptive deduplication with a landing zone begins deduplicating and replicating as data commits to disk in parallel with backups coming in, providing full system resources to the backups for the shortest backup window.

Adaptive deduplication provides:

1. fast direct-to-disk, high-speed backup performance for the fastest ingest and shortest backup window,
2. fast restores, VM boots, and tape copies from the most recent backups that are stored in a front-end landing zone,
3. deduplication and replication that occur in parallel with backups for a strong RPO (recovery point) at the disaster recovery site, and
4. a repository of all deduplicated data that sits behind the non-deduplicated data in the landing zone for cost-efficient storage of long-term retention.

Should I use a scale-up or scale-out storage architecture?

A scale-up architecture (front-end controller with disk shelves) presents numerous challenges to disk backup with data deduplication. The first is continued expansion of the backup window. Data deduplication is compute intensive; however, in a scale-up system, only storage capacity is added as data grows so the backup window grows as the data does. When the backup window becomes too long, a bigger, faster controller is required in order to apply more compute, called a “forklift upgrade.” This is expensive and disruptive.

Secondly, if the front-end controller fails, no backups can occur. In contrast, a scale-out approach adds full appliances in a scalable system so that processor, memory, and network ports are added along with storage capacity. As the data doubles, triples, quadruples, etc., all required resources are also doubled, tripled, and quadrupled, called “adding compute with capacity.” This scale-out model allows you to add appliances of various sizes as data grows while maintaining a fixed-length backup window. Adding appliances of varying sizes allows you to pay as you grow, fixes the length of the backup window even as data grows, and eliminates forklift upgrades. In addition, older and newer appliances can be mixed and matched in a single scale-out system, eliminating product obsolescence and resulting suspension of support. Lastly, if a single appliance fails in a scale-out system, all of the other appliances remain in production and can continue to receive backups. The majority of the backups continue to run, which is not the case when the front-end controller fails in a scale-up system.

Can I use the cloud for my disaster recovery (DR) data?

Smaller companies that have a few terabytes and don’t have a second site to house an offsite appliance use solutions where the DR data is stored at a cloud provider. The true challenge for cloud-based DR is that backup data changes daily, and therefore you have less than 18 hours to get the data into the cloud, requiring a lot of bandwidth. From a security standpoint, your data is co-mingled on storage that also stores the data of others. Furthermore, when it comes time to recover data, it is virtually impossible to do so in any reasonable amount of time. Companies with tens of terabytes to petabytes of data have a second data center to house a second-site backup storage system. Running their own offsite DR appliances costs less than replicating DR data to a cloud when compared over a three-year period and provides much faster recovery times. It’s also far more secure because data is kept behind the

organization's physical and network security, and the data isn't intermingled with that of other organizations.

What is the "right requirements" list?

- Fast ingest for the shortest backup window
- Most recent backup kept in a landing zone in non-deduplicated form for fast restores, VM boots, and offsite tape copies
- As data grows, the backup window remains fixed in length
- Efficient long-term retention of deduplicated data with a 20:1 deduplication ratio
- Data encryption at rest performed at the drive level, which only takes nanoseconds, versus software encryption that takes CPU cycles away from data deduplication
- Efficient use of bandwidth by moving the least amount of data offsite with a 20:1 deduplication ratio
- A strong RPO for offsite DR data

To achieve the right requirements, what should a solution offer?

1. Backups that write direct to a disk landing zone, avoiding compute-intensive inline data deduplication.
2. Storage of the most recent backups in a landing zone in a non-deduplicated form for fast restores, VM boots, and offsite tape copies.
3. A scale-out storage architecture that adds landing zone capacity and repository disk along with network ports, processor, and memory, i.e., adds compute with capacity versus a traditional scale-up solution that has a fixed front-end controller and only adds capacity as data grows.
4. Adaptive deduplication of data after it commits to disk but in parallel with the backups coming in so that a strong offsite recovery point (RPO) is maintained.
5. Encryption at rest performed at the drive level versus in software.
6. Deduplication algorithm that deploys one of the following:
 - 8KB blocks with variable length content splitting = 20:1
 - Byte level = 20:1
 - Zone stamps with byte-level compare = 20:1
7. Deduplication algorithm that only moves changed data and deploys one of the following:
 - 8KB blocks with variable length content splitting = 20:1
 - Byte level = 20:1
 - Zone stamps with byte-level compare = 20:1

Summary

In summary, the important things to look for are:

- A disk landing zone
- A storage architecture that adds full appliances into a scale-out system
- A 20:1 deduplication ratio
- Adaptive deduplication

Choosing a disk storage with data deduplication solution will have a major impact on the cost and performance of your backup environment for the next three to five years. Take the time to ask the questions above to ensure that you make the right choice.



United States: 2000 West Park Drive | Westborough, MA 01581 | (800) 868-6985

United Kingdom: 200 Brook Drive | Green Park, Reading, Berkshire RG2 6UB | +44 (0) 1189 497 051

Singapore: 1 Raffles Place, #20-61 | One Raffles Place Tower 2 | 048616 | +65 6808 5574

ExaGrid reserves the right to change specifications or other product information without notice. ExaGrid and the ExaGrid logo are trademarks of ExaGrid Systems, Inc. All other trademarks are the property of their respective holders.

©2018 ExaGrid Systems, Inc. All rights reserved.