June 2016

Three Recommendations for Evaluating Data Deduplication to Ensure Storage and Bandwidth Efficiency While Delivering Fast Backups, Restores, and Recoveries

By Jerome Wendt

DCIG

Three Recommendations for Evaluating Data Deduplication to Ensure Storage and Bandwidth Efficiency While Delivering Fast Backups, Restores, and Recoveries

Table of Contents

DCIG

Three Recommendations for Evaluating Data Deduplication to Ensure Storage and Bandwidth Efficiency While Delivering Fast Backups, Restores, and Recoveries

Executive Summary

Data deduplication's entrance into the backup landscape captured the fancy of organizations like few other technologies. It became the trigger that effectively introduced disk backup into the backup process by lowering disk's price, versus tape backup, as a backup target. This freed up organizations to use disk to shorten their backup windows, improve their backup success rates and, in essence, "solve" backup for them. These successes led to deduplication's rapid adoption and attaining the status of a silver bullet that solved all backup problems.

Over time, deduplication found its way into almost every layer of the backup stack. Whether it was on the target disk appliance, the backup software media servers, or in the backup software installed on clients, deduplication became almost ubiquitous in terms of how and where organizations could deploy it.

While this ready and easy availability of deduplication was great on one level, the deduplication technology itself created misconceptions among many organizations about its true capabilities. This resulted in unexpected issues that they were sometimes ill equipped to handle. For instance:

- Many came to conclude that all deduplication algorithms worked the same regardless of what deduplication algorithm was used and where it was implemented in the backup stack. They only found out later this was not the case.
- Others became convinced that deduplicating all backup data as it was stored to disk would have no impact on recovery times. This belief was found to be incorrect.
- Still others looked at the capacity of some appliances and placed more weight on that than on the performance metrics associated with the ingestion of backup data. These individuals now realize that ingest performance should be examined more carefully as it impacts the backup window length up front and over time as data grows.

These and other misconceptions about the implementation of data deduplication largely arose because organizations failed to recognize the three new challenges for implementing data deduplication in the backup process. Although deduplication solves disk storage and offsite replication challenges, deduplication by its very nature is highly compute intensive and creates three new challenges around backup and restore. By not understanding these three new challenges, many organizations find themselves in a position where they have implemented deduplication in a manner that is neither optimal for their environment nor do they have a clear path forward to success.

The good news is that whether or not an organization has already implemented deduplication, they can reevaluate their position at any time. Regardless of where they are currently, they can put themselves on a path to implementing it correctly. The ExaGrid family of deduplicating backup appliances provides aggressive deduplication for low storage and WAN bandwidth use and also solves deduplication's three new compute challenges. ExaGrid provides organizations with a straightforward path to implementing deduplication while successfully backing up and recovering data in the time and manner required without the unexpected costs and obstacles that may result from other implementations.

The Silver Bullet of Deduplication Has Left Holes in Restore and Recovery

No one disputes that data deduplication has revolutionized the backup process. Backups require retention *(keeping multiple copies of backups over time: weeks, months and years)* and as a result can use 40 to 100 times the storage of the primary data copy. For example: if an organization keeps 13 weekly backups as well as each monthly backup for three years, the total is 49 backup copies. This is a tremendous amount of data redundancy from backup to backup as only about 2% of the data changes from week to week.

Instead of keeping each full backup, data deduplication identifies and then only stores the unique blocks or bytes of data across multiple backup jobs. Data deduplication effectively brings the price of disk down to or below the cost of tape when measured on a per-gigabyte (GB) basis. A standard retention period of 18 weeks will use approximately 1/20th of the disk with data deduplication versus storage that does not use data deduplication (a 20:1 reduction).

Using disk as a backup target has contributed heavily to organizations realizing:

- Improved backup success rates
- Increased ingest performance
- Shorter backup windows *(versus using tape as a target)*

Together disk and deduplication form the silver bullet that organizations have sought to solve their backup problems. Using disk and deduplication as a single solution, organizations get the faster, more reliable backups and restores that they need while largely staying within their existing backup budgets.

However, this silver bullet of data deduplication has left holes in restores and recoveries for many organizations. Once organizations back up and deduplicate their data, the following three issues surface in many environments:

1. ***Data restores are not as fast as expected.*** People typically expect to restore their data and/or applications within minutes of—and certainly no more than an hour or two after—requesting a restore. Yet deduplicated data takes time to restore because the data must be rehydrated before it is usable. Depending on the amount of data to be restored, this may take hours or even days to complete.
2. ***Virtual machine (VM) boots are not instantaneous.*** Organizations of all sizes have heavily virtualized their environments and want and/or need to recover VMs directly from the disk backup target on which they reside. Here again the same problem resurfaces; rehydrating data to boot a specific VM may take hours or even days to complete.

Today's Data Deduplication Conundrum

Organizations may find themselves caught between a rock and a hard place when it comes to trying to determine how to best implement deduplication. The data reduction ratios that data deduplication provides certainly get their attention. Whether it is 7:1, 10:1, 20:1 or even greater, the potential reductions in storage capacity and associated cost savings that deduplication offers make it easy for even the most cynical IT professional or business owner to get excited about implementing it.

The approach of deduplicating all data works extremely well in cases where there is time to complete the restore or recovery of applications and/or data. Whether it is restoring a file that was deleted weeks ago but its deletion was just noticed; restoring data that was corrupted months ago but was just noticed; recovering from a natural disaster such as an earthquake, hurricane or a tornado; or recovering from a man-made disaster caused by human error or negligence; there is generally an expectation that it will take time to recover from these types of events. Under these conditions, most end-users understand and tolerate the time required to first rehydrate the deduplicated data before it is recovered.

However, focusing solely on data reduction numbers holds little merit when the deduplication solution fails to provide speedy restores and recoveries. Backing up data to disk prompts organizations to expect faster restores and recoveries, even if the solution stores data in a deduplicated state. These reasons include:

- Up to 95% of application and/or data restores come from a recent backup
- Up to 98% of VM boots initiate from the most recent backup
- Up to 100% of offsite tape copies originate from a recent backup

This puts many of today's backup software, backup appliances and deduplicating appliances at a disadvantage. In striving to deliver ever higher data reduction ratios, these solutions forgot, neglected and/or overlooked this basic tenet of backup: the need for fast restores, VM boots and/or tape copies. Successfully meeting these expectations can only be done at scale by using non-deduplicated backup storage that is implemented as a disk staging zone.

This is the conundrum many organizations often encounter when they look to implement deduplication in their environment. Doing so increases their backup success rates, shortens their backup windows, and controls their storage costs but they unknowingly sacrifice faster restores and recoveries and face unpleasant surprises when they go to perform these tasks.

3. ***Tape copies remain a reality in many organizations.*** Many organizations still use tape for three purposes: an offsite copy for disaster recovery, a tertiary copy in addition to the onsite and offsite backup storage appliances, and for extremely long periods or retention. The challenge that organizations encounter here is the same: copying deduplicated data from disk to tape requires that it be rehydrated first. Due to the growing amount of data that can be stored on a single tape cartridge *(the latest LTO-7 tape cartridge holds up to six terabytes (TBs) of data),* the process of first rehydrating data and then copying it all to tape could again take hours or days to complete.

These conflicting requirements for effectively utilizing backup data put organizations in a quandary. On one hand, they see the value of using disk as a backup target and then deduplicating the data as absolutely critical to reducing their backup storage costs.

However, as they encounter these various restore and recovery challenges, they rightfully question how they should implement deduplication in their environment. While they see the value of deduplication, its implementation needs to be re-examined in light of their other needs to do restores, VM boots and making offsite tape copies in a timeframe that better aligns with their business requirements.

Five Common Misconceptions about Deduplication

To understand what one needs to fix when it comes to implementing deduplication, it is prudent that one first understands how so many have come to view deduplication as a silver bullet to solving most, if not all, of their data protection challenges.

While many rightfully view deduplication as an excellent, even optimal, technology to cost-effectively introduce disk into the backup process, not every approach to deduplication works equally well. To implement it appropriately, one first needs to understand some of the misconceptions that have emerged over the years and how these viewpoints prevent organizations from maximizing the benefits they may achieve from their implementation of deduplication.

Misconception #1: All Deduplication Algorithms Are Created the Same

Deduplication can be found at almost every layer in the backup stack. Whether it is in the backup software, included with a backup appliance or on a deduplicating appliance, organizations can find a solution that offers deduplication. The ready availability of deduplication in these solutions could lead one to assume that each solution uses a similar, or even the same, underlying deduplication algorithm. That assumption would be incorrect.

Data deduplication is not compression. Compression works on a single copy of data and can reduce the data to about half its size, a 2:1 reduction.

In backup, many copies of data are kept over weeks, months and years. It is not uncommon to keep 40 to 100 copies of backup data at different points in time *(historical versions).* If an organization has 40 copies of a 100TB backup and simply used compression, it would only reduce its total data by 2:1. Where data deduplication differs from compression is that it compares one backup to the next and only stores the unique blocks or bytes, i.e. changed data. This approach achieves a reduction of up to 10:1 to as high as 50:1, depending on the data type and mix.

The first backup of data would be compressed or deduplicated about the same, 2:1. Thereafter compression would continue to compress each copy at 2:1. The overall result with compression will never exceed a total of about 2:1.

Data deduplication also reduces the first copy of data at about a 2:1 ratio. Thereafter it only stores about 2% of the backup data as, on average, only about 2% of the blocks or bytes change from week to week. The more weeks of retention, the better the deduplication ratio. At three weeks of retention the deduplication ratio will be 3:1 or 4:1. At 8 to 10 weeks, closer to 10:1. At 18 weeks of retention, the data deduplication ratio would be approximately 20:1.

This is why the first backup of 100TB of data will not result in a 5TB backup or achieve a 20:1 data reduction ratio. A more realistic result for an initial backup of 100TB of data is a 50 – 60TB backup with a 2.0:1 or even a 1.8:1 data reduction ratio.

Deduplication ≠ Compression

It is only after deduplication completes that compression kicks in. Compression removes the spaces and leading zeros in the deduplicated chunks of backup data and complements deduplication by further reducing the total amount of data stored. However, compression may be the only underlying technology that the various deduplication algorithms share in common.

Purpose-built, Target-based Appliances Offer the Best Deduplication Algorithms

The best algorithms *(those that deliver the highest deduplication ratios possible)* examine data at a very granular level. They rely upon adjustable, variable length, content splitting deduplicating algorithms or zone stamps that either do block-level compares using relatively small block sizes (~8KB block sizes) or they may even do comparisons at the byte level.

This technique is very effective. By examining strings of data for matches using different intervals of data, the likelihood of finding matches increases significantly and maximizes the possibility of identifying like segments. This approach helps organizations achieve the average 20:1 or even greater deduplication ratios that they ideally want as they will consume less disk and use less network bandwidth to replicate deduplicated data offsite for disaster recovery.

But this best approach to deduplication comes with a price. The better the deduplication algorithm, the more CPU and memory it consumes. Data deduplication is a compute- and memory-intensive process, especially when called upon to parse through the tens or hundreds of terabytes that an enterprise shop may have to protect on a daily or weekly basis. This is why the *"best"* deduplicating algorithms are typically found on purpose-built, target-based deduplicating backup appliances as they have dedicated processors and memory available to process this volume of data.

Media Server-based Deduplication is Less Aggressive

An alternative to running the deduplication algorithm on a purpose-built, target-based deduplicating backup appliance is to run it on a backup software media server. The appeal of running deduplication here is that the server's sole job is to manage and run backups. Further, organizations may use low-cost disk as opposed to acquiring a purpose-built, target-based deduplicating backup appliance.

The trade-off is that the server hardware on which the media server runs is not optimized for running deduplication. Since the best deduplication algorithms demand large amounts of CPU and memory to run quickly and efficiently, the backup software media server must scale back the deduplication algorithm it uses to align with the resources it has available.

To do so, the backup data is chunked up into large fixed-length blocks. By way of example, Commvault's default is a 128KB fixed block length.[1] Veritas NetBackup uses 32KB to 1MB fixed block lengths *(it refers to them as "segments"*) depending on the data type. By using these fixed block lengths, they reduce the amount of CPU and memory required to deduplicate the data. However, this results in an increase in the amount of disks needed to store the data and bandwidth needed to move the data by up to four times more than purpose-built, target-based, deduplicating backup appliances.

To deliver deduplication on some level, backup software providers purposely scale back the robustness of their deduplication algorithms to preserve the CPU and memory of the application and file servers on which they run. To accomplish this, they use larger fixed-length block sizes that range from 32KB to 1MB in size. By eliminating the need to split up backup data into chunks of various sizes and then examine them, they consume fewer CPU cycles and use less memory on the server.

Table 1
Ranking the Deduplication Algorithm Implementations

BEST	Purpose-built, Target-based Deduplicating Backup Appliances
Average	Media Server

Misconception #2: All Data Deduplicates Equally Well

Achieving the highest deduplication ratios goes well beyond choosing the best deduplication algorithm. Organizations first need to understand what types of data that they intend to deduplicate and what they hope to accomplish before they implement deduplication in their environment.

By way of example, deduplication may now be found on multiple devices and in multiple products within today's data centers. Whether it is in archival solutions, backup software, deduplicating backup devices, primary storage arrays or replication software, all of these may use deduplication on some level. Yet how effectively they deduplicate the data stored will vary widely.

Deduplication is most often associated with backup data in large part because backup is so conducive to achieving high deduplication ratios. By way of comparison, deduplicating data on primary storage may, in a best case scenario, reach a 3:1 deduplication ratio. Conversely, the long retention times associated with backup data may result in deduplication ratios that can reach 20:1 or even higher.

Achieving these high deduplication ratios largely hinges on the deduplication algorithm being used, the types of data, the mix of data within the backups and the retention

1. Commvault Systems, Inc. Deduplication Building Block Guide. N.p., n.d. Web. 08 Mar. 2016. http://documentation.commvault.com/commvault/v10/article?p=features/deduplication/deduplication_building_block.htm

period *(how many historical copies are being kept)* of these backups. Consider:

- If much or all of the primary data being backed up is compressed or encrypted, organizations will see little to no data reduction as these data types are not readable by deduplication solutions
- If the backup jobs largely contain unstructured data such as files, deduplication ratios will reach approximately 7:1 over an 18-week period
- In a best case scenario where the backup data consists entirely of structured data such as databases or email, the deduplication ratios could be 100:1 or even greater

Setting Realistic Expectations for Deduplication Ratios

Setting realistic expectations for deduplication ratios requires that organizations first take the following three steps:

1. Determine what percentage of their data is classified as:
   - Compressed and/or encrypted
   - Structured data
   - Unstructured data
2. Determine how long they will retain data in their backup repository. The longer the retention period, the higher the deduplication ratio generally becomes.
3. Determine where to implement deduplication: on a purpose-built, target-based deduplicating backup appliance, on the media server, or at the client.

Table 2
Backup Data Retention Periods

Retention Period	Deduplication Ratio
1 Week	1.8 – 2:1
3 Weeks	3:1
18 Weeks	20:1

Table 3
Approximate Deduplication Ratios by Data Type

Data Type	Deduplication Ratio
Compressed/Encrypted	1:1
Unstructured Data (Files)	7:1*
Structured Data (Databases/Email)	100:1*

*Assumes minimum 18 week retention period

The type of data, the mix of data types, the retention period for the backup data, and the algorithm used to deduplicate will combine to determine the final deduplication ratio that any organization achieves. Assuming an organizational data mix of little to no compressed or encrypted data, 30 percent consisting of email and/or databases, and the remaining 70 percent comprised of files all retained for a minimum of 18 weeks, the potential exists for organizations to achieve a deduplication ratio of 20:1. Achieving this 20:1 ratio will depend upon the type of deduplication solution used. As a reference point:

- Current backup applications achieve about 2:1, 4:1, 6:1 and 8:1 deduplication ratios, depending on the backup application
- Current target-side deduplication appliances achieve 10:1, 14:1 or 20:1 depending on the appliance
- A low deduplication ratio results in more storage and WAN bandwidth needed for replication which, in turn, increases storage and bandwidth costs

Misconception #3: Deduplicated Data Has No Impact on LAN/WAN Bandwidth

A primary reason to deduplicate data before moving it offsite to a cloud storage provider or to another site owned and/or leased by the organization is to minimize the amount of data sent over the LAN/WAN link. Yet organizations may not grasp that failing to use the best deduplication algorithm will, in turn, lower the deduplication ratio that they will achieve. This will result in the need to replicate more data and use more LAN/WAN bandwidth. If, for example, a solution achieves a 6:1 deduplication ratio, it will require three times the WAN bandwidth (and potentially 3X the cost) for replication versus a solution that achieves a 20:1 deduplication ratio.

This need to replicate data makes it imperative for organizations to identify the best deduplication algorithm for their environment. The higher that they can drive their deduplication ratios, the less data that they will need to transmit. The less data they transmit, the faster the data will replicate offsite as it uses lower amounts of bandwidth.

Failure to select the best approach to deduplicating data has the following effect on their disaster recovery strategy. Initially or over time, organizations may require larger network pipes to move their data in a timely manner. These larger pipes translate into increased monthly costs. If more bandwidth is not an option, then they may find themselves in a predicament where they either cannot move data offsite in a timely manner or move it offsite in its entirety. Either way, not having all data offsite negatively impacts recoverability.

Misconception #4: Capacity Matters as Much as or More than Performance

All providers of deduplicating backup appliances promote the capacity levels to which their systems scale. These high capacity numbers, whether based on the appliance's usable capacity or its effective capacity after data deduplication is applied, may leave the impression that the appliance's published performance numbers align with their published capacity numbers and that they can scale to fully utilize this capacity. This is rarely the case.

Almost all deduplicating backup appliances use inline deduplication to process incoming backup data. While the amount of capacity is certainly relevant in determining how much data the appliance can actually store, organizations must first calculate if the rated ingest for the appliance aligns with its rated capacity. In other words, an appliance's ability to scale to 250TB, 500TB, 1PB, etc., primarily matters if it can actually back up that amount of data during the backup window. The real question then becomes, *"Do they have enough ingest performance to keep up with the amount of backup storage so that backups finish in the allotted backup window time?"*

Without a means to add more compute resources to the appliance, inline deduplication caps ingest throughput rates. The appliance controllers that contain the CPU and memory only contain fixed amounts of these resources. Once the appliance reaches these limits, it cannot ingest data any faster no matter how much more capacity is added to it.

To illustrate, a popular line of deduplicating backup appliances, the EMC Data Domain, uses inline deduplication in its models. Its high-end model, the DD9500, has a maximum usable capacity of up to 864TB and natively performs at an ingest rate of up to 27.7TB/hour. One of its mid-tier models, the DD2500, has a maximum usable capacity of up to 133TB and natively performs at an ingest rate up to 5.6TB/hour.[2]

It is when one does the math that despite the ability of both of these appliances to scale to 864TB and 133TB respectively, neither one supports an ingest rate that can fully use its capacity within a 12-hour backup window. If either of these models had to actually back up that amount of data, the DD9500 running at an optimal 27.7TB/hour would take over 31 hours to complete its backups while the DD2500 would take about 24 hours running at its optimal throughput of 5.6TB/hour.

To improve performance, the EMC Data Domain models use software known as "DD Boost" that is installed on backup media servers *(or on database servers if a database backup utility such as SQL dumps or Oracle RMAN is used)* to accelerate back up throughput. DD Boost starts the process of data deduplication on the media or database server by leveraging some of the compute resources on the server. Using

Inline Deduplication's Drawbacks Persist

In recent years inline deduplication has emerged as the predominant way in which most deduplicating backup appliances implement deduplication. Driving its adoption is the reality that inline deduplication is often technically easier for providers to implement *(relatively speaking)* than competing approaches.

Yet the drawbacks associated with inline deduplication persist. Inline deduplication requires a tremendous amount of compute power to analyze and deduplicate incoming data. While initially *(in the first one to four weeks)* its processing requirements may not be noticed, as more data is backed up, analyzed, and stored over time, the compute resources needed to deduplicate the data increases which, in turn, slows down backups and lengthens backup windows.

The lack of available resources is exacerbated if and when data restores occur. To restore data, the appliance has to rehydrate or re-assemble the various deduplicated chunks of data. This again consumes resources and takes time which slows down application restores, VM boots and the creation of offsite tape copies.

Further, as the amount of data stored increases, background processes that rebase or defragment the deduplicated data in the most optimal manner on the system also tap into available resources. This defragmentation helps to reduce the time it takes to recover but also competes for and consumes resources that are needed to deduplicate incoming data.

These appliances with their scale-up architecture only accommodate the introduction of more capacity, not more compute resources. This approach does not help and masks the real issue: the need for more compute resources to effectively deduplicate more data. The scale-up architecture of these appliances *(fixed resources with a front end-controller and disk shelves)* offers no means to deliver more performance other than to upgrade to a new front-end controller with more performance. However, that still relies upon the same built-in architectural limitations. These restrictions will again begin to surface as the new front-end controller continues to store more data.

2. EMC DATA DOMAIN DEDUPLICATION STORAGE SYSTEMS. (2015): 1-6. 2015. Web. 9 Mar. 2016. http://www.emc.com/collateral/specification-sheet/h11340-datadomain-ss.pdf

this technique of offloading some of the compute from the appliance to the backup media server increases the maximum throughput rate to 58.7TB/hour for the DD9500 and to 13.4TB/hour for the DD2500.

Even with these improved ingest rates, only one of these two models can fully utilize available capacity within a 12 hour backup window. The DD9500 still needs over 14 hours to back up 864TB of data while the DD2500 can back up 133TB of data in about 10 hours. Further, these numbers are only met if the optimal throughput conditions are satisfied. As most organizations only realistically expect to achieve 50 to maybe 80 percent of any vendor's published numbers, the actual ingest performance numbers they realize in their environment will likely be lower than those examined here.

This example illustrates how an appliance's ability to scale up capacity has little to no bearing on its ingest rates. The capacity of an appliance primarily comes into play if the amount of data that an appliance ingests within a 12-hour backup window exceeds its overall capacity. Otherwise, an appliance's maximum usable or logical capacity is window dressing on its data sheet that should only be used as a reference point and rarely as a determining factor when making a buying decision.

Misconception #5: Deduplicated Data is Recovered as Fast as Data Stored in its Native Format

When organizations think about backup data stored on disk, they may naturally assume that they will recover their data quickly,

The 30% YoY Data Growth Problem

As organizations size deduplicating backup appliances for their environments, they have a tendency to primarily look at their current data stores and maybe a year or so into the future. Yet to ensure that the selected appliance can match their real backup requirements two, three, or more years into the future, they need a handle on what their real data growth rates are.

This is where deduplicating backup appliances and data growth rates fail to align. In the last couple of years, it almost seems as if analyst firms have thrown in the towel in terms of trying to forecast how fast data is going to grow, in part because actual data growth continually outstrips their predictions.

If an organization assumes a 30 percent annual growth rate for its data *(which for many organizations may prove to be conservative),* it will double every 32 months *(or about every 2.5 years).* Assuming data growth continues at this rate, 100TB of data will grow to 200TB of data in under three years, 400TB of data in about five years and about 1PB of data in nine years.

To put an exclamation point on how problematic this data growth is when using deduplicating backup appliances that use inline deduplication, they may be up to 70 percent full after just their first day of backups. This leaves little to no room for future data growth and puts organizations in a position where they may be forced to either upgrade to a new appliance in as little as 18 months or introduce new appliances into their environment that create silos of data and management overhead.

Chart 1 shows the outcome of following a path of upgrading to new appliances with higher capacities as the existing one fills up. Taking this route ensures that organizations only have one appliance to manage, but they could end up buying up to five or more appliances over a ten-year period. *(Each hash mark indicates the purchase of a new, higher capacity appliance to keep up with data growth.)*

Chart 1: Appliance Upgrades

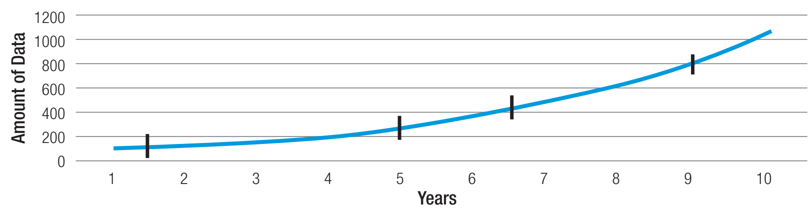

Each hash mark represents the need for an appliance upgrade or replacement.

Chart 2 illustrates the outcome of taking a different path that adds new appliances with the same amount of capacity as the original one to the environment. In this scenario, an organization could potentially be looking at the purchase of up to eight or more appliances over a ten-year period. *(Each hash mark indicates the purchase of another appliance of the same capacity as the original one to keep up with data growth.)*

Chart 2: Adding Equivalent Capacity

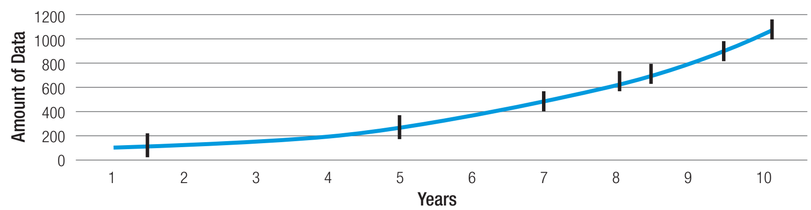

Each hash mark represents the need for an additional, new appliance.

Granted, these two charts do not tell the entire story in that the storage capacities of the appliances in the years to come will likely be higher than the ones currently available. However, they do accurately illustrate the challenge that organizations face in scaling the capacity of their deduplicating appliances to protect and house new backup data in the years to come.

Note: These two illustrations were created using EMC Data Domain's available models and their maximum usable capacities as of March 9, 2016. For purposes of this illustration, it was assumed that every 1TB of production data would consume about 1TB of capacity on the deduplicating backup appliance. Actual results in each organization's backup environment will vary.

regardless of whether or not the data is deduplicated. This would be a false assumption. If anything, the introduction of deduplication into the backup process has taken some organizations backwards in this respect: it may actually take them ***longer*** to recover.

Unlike data stored in a non-deduplicated state, deduplicated data must be rehydrated before it can be recovered and used for application or file restores, VM boots, or tape copies. This process of rehydrating data for restoration by reassembling deduplication blocks of data scattered throughout the system takes time. Many of today's deduplication solutions attempt to mitigate the time associated with data rehydration by rebasing or defragmenting the data so that related blocks of deduplicated data are positioned close to one another for faster retrieval. However, this only works relatively well and still does not match the speed or performance of data restoration when data is stored in its native state.

This is why tape libraries use disk staging or a disk cache as it serves three purposes:

1. It accelerates backups by temporarily storing data on disk until it moves data off to tape.
2. It allows for more reliable backups since writing to disk is more reliable than writing to tape.
3. It ensures that backups are ready for fast restores and helps to minimize the need to go to tape to restore from tape.

Deduplicating incoming data inline slows down backups over time as well as slows down restores versus storing backup data in a non-deduplicated state on disk to eliminate time-consuming data rehydration.

The Three Primary Purposes for Deduplicating Backup Data

These five aforementioned misconceptions about the purpose of data deduplication in the backup process have served to obfuscate deduplication's three primary purposes:

1. Simplify the introduction, use, and management of disk as a backup target in the backup process.
2. Drive down the cost of disk to make it a viable, cost-effective alternative to tape for long-term data storage and data retention.
3. Effectively move data across LANs/WANs to keep copies of data offsite for DR.

While data deduplication is sometimes billed, promoted, and/or perceived as delivering on more than these three basic premises *(and sometimes it does),* if organizations implement deduplication in their environment expecting more than what it was originally intended to deliver, they set themselves up for disappointment and even failure. To avoid these potential pitfalls and realize the real benefits of deduplication, organizations are best served by following these three recommendations when implementing it.

Three Recommendations for Evaluating Data Deduplication

Recommendation #1: Back Up and Store Data in a Non-Deduplicated Form for a Period of Time

Inline deduplicating backup appliances have made impressive advancements in the last few years to improve their backup and restore performance. By adding more powerful CPUs, more memory, and introducing solid state drives (SSDs) into them to host the deduplication metadata, they have significantly improved the rates by which they can ingest and deduplicate backup data.

Despite these numerous upgrades and improvements, backing up non-deduplicated data to disk remains as fast or faster than using inline deduplication. Further, restoring non-deduplicated data is faster in almost every circumstance as it incurs minimal performance overhead. Further, backing up or restoring non-deduplicated data directly to disk eliminates the need to install agents on either backup clients or media servers to start the process of deduplicating or rehydrating data when it needs to be restored.

Granted, storing non-deduplicated data on disk is only recommended for a short period of time *(typically no more than a week).* Storing it longer than a week starts to unnecessarily incur additional disk storage costs and serves few practical purposes. However, storing non-deduplicated data in this form for a short period of time addresses the immediate needs that organizations have to quickly recover their data for application or file restores, VM boots, and creating tape copies for disaster recovery, long-term retention and offsite archival.

Recommendation #2: Choose a Target-based Deduplication Appliance with Scale-out Architecture

Storing backup data in a non-deduplicated form for a period of time before it is deduplicated necessitates that organizations select a target-based deduplication appliance with a scale-out architecture. While other backup architectures and configurations exist, a target-based deduplication appliance with a scale-out architecture represents the simplest and easiest solution to implement, manage, and maintain initially and over time for the following reasons.

First and foremost, a scale-out configuration gives organizations the flexibility to size a target-based deduplicating backup appliance according to their specific backup window requirements. Then as they grow, the more backup jobs that they need to complete, the more appliances they may deploy since each appliance contributes to linearly increasing the solution's overall backup ingest rates.

Each appliance added to the configuration introduces more compute and network ports to improve ingest rates while the capacity contained in each appliance can then store more backup data. In this way, organizations may add as few or as many appliances at any time to their existing configuration to match their actual or anticipated backup ingest or capacity requirements.

Second, the scale-out architecture must create a single, logical configuration that globally deduplicates data across all of the appliances in the scale-out configuration. This avoids the creation of silos of data and eliminates the need for administrators to individually manage which appliance each backup job is sent to initially and over time. This approach also ensures that the time associated with configuring and managing the system—whether it is two appliances or twenty-two—remains roughly the same.

Third, a scale-out architecture mitigates product obsolescence as it better accommodates the situation where a deduplicating backup appliance can no longer handle the backup workloads or capacity demands of the backup environment. In contrast, scale-up architectures require either a product replacement or the introduction of a new appliance into the environment that must then be managed separately.

A scale-out architecture ensures that organizations can continue to use the existing appliance so they continue to realize its value without increasing the management complexity of their environment. Further, once it is time to decommission an existing appliance, data on an individual appliance can be easily migrated to any of the other appliances that are part of the scale-out system.

Finally, a scale-out architecture natively offers what almost no scale-up deduplicating backup appliance currently offers—a highly available configuration. Every scale-out architecture with two or more appliances provides additional resiliency. Configured this way, any appliance may go off-line due to a mechanical failure or be taken off-line for maintenance even as backups continue to the remaining appliances. This creates a high probability that the majority of the solution and its data remain online for both backups and recoveries. Some backup applications are becoming aware of storage failures and will automatically re-route the backups, or portions of backups, before they fail to a different on-line appliance.

Here again, the single controller configurations commonly found with scale-up architectures do not provide these same high levels of availability. Should they need to go offline for maintenance, all backup and recovery functions need to stop to include access to the data until the appliance is back online. Scale-out architectures do not have this limitation.

Recommendation #3: Deduplicate Backup Data as Quickly as Possible after it is Stored to Disk

Initially storing data in a non-deduplicated form on disk for a period of time does not mitigate the need for organizations to get a copy of it into a deduplicated form as quickly. If anything, because organizations want to minimize the size of their data stores long term as well as replicate a copy of data offsite as quickly as possible, they need to deduplicate this data once it is stored on disk.

To achieve this objective, deduplication of the backup data must occur during the backup window and in parallel with the backup after the data is committed to disk. Furthermore, the deduplication of the data residing on disk should only be done in a manner where it does not compete for the compute resources being used by the backup process. Successfully completing the backup should always take priority over the deduplication of the data since the greatest probability of a data recovery occurring is in the days immediately following the backup—not weeks or months later.

The ExaGrid System: A Deduplication Solution Built for Backup

The current landscape of providers primarily introduces deduplication from a design perspective into their solutions as a feature. Whether they implement deduplication as part of their backup software in a media server or include it as part of their appliance that does inline deduplication using a scale-up architecture, they do not fully address the compute-intensive nature of deduplication and how it impacts the rate of ingest and restore.

The ExaGrid system differs from all of the deduplication solutions currently available on the market. ExaGrid adheres to these recommendations for leveraging deduplication in the backup process as it most closely follows the three primary purposes that data deduplication serves in this process.

ExaGrid's mindset was to develop a solution specifically tuned to provide fast backups and restores of applications and files as well as facilitate rapid VM boots. Its implementation of deduplication complements and enhances existing backup and recovery processes without putting organizations in a positon where they must make trade-offs in their backup and recovery procedures.

Landing Zones and Deduplication Repositories

The ExaGrid system accomplishes this by adhering to all of these recommendations for implementing deduplication. Backup software may initially back up data in its native format and store it on the landing zone found on all ExaGrid appliances.

This technique helps ExaGrid systems achieve the fastest ingest rates and shortest backup windows among deduplicating backup appliances using published benchmarks. For example, a fully configured ExaGrid EX40000E can scale-out to 25 appliances and achieve 200TB/hour in ingest performance. This is 3X more throughput per hour that the EMC DD9500, resulting in much shorter backup window times.

Creating a landing zone on each of its deduplicating backup appliances also solves another design problem. ExaGrid eliminates the need for organizations to introduce a separate disk array to temporarily store backup data before it is deduplicated.

By default, the capacity on each ExaGrid system is split equally *(50-50)* between being configured as a landing zone and as a repository for deduplicated data. However, these percentages may be altered in either direction to make them smaller or larger. If organizations expect to have large backup jobs, they may increase the size of the landing zones to facilitate larger backups. If organizations need to retain their data for lengthier periods of time, they can increase the size of the repository used to store the deduplicated data.

Yet what makes this design so powerful is that organizations may deploy a single solution that increases their backup and recovery times even as they automate the deduplication of the data after the backup is complete. Any other approach that tries to introduce this flexibility of backing up data in its native form to disk requires much more upfront design work, is a more complicated configuration to implement and manage, and costs more. The ExaGrid system delivers the implementation of these recommendations in a single solution that provides an intact landing zone integrated with a deduplication repository configured as a scale-out GRID so it can be expanded to meet increased volumes in backup data.

Scale-out Architecture that Globally Deduplicates Data

ExaGrid's scale-out architecture is worth examining. Scaling up to 25 appliances in a single GRID, an organization can mix and match ExaGrid appliances of any size or age in a single GRID. In this way, an organization can expand its system in a manner that best aligns with its specific data growth needs.

If it experiences only a small amount of data growth that exceeds the capacity of its existing appliance, the organization can opt to add an appliance as small as the ExaGrid EX2000 with up to 4TB of usable capacity to its GRID. Conversely, if an organization suddenly needs to manage a large increase in backup data volumes, it can add an EX40000E that provides up to 78TB of usable capacity and 8TB/hour of backup throughput. Regardless of which appliance model or models an organization needs, they all may be added to the same GRID and managed as one logical instance.

This logical management extends beyond a web-based GUI that displays all of the appliances on a single screen while leaving each appliance to be independently managed as its own entity. These appliances function together as one. In this way as each appliance deduplicates data, it works with the other appliances to globally deduplicate the data in that GRID. Then, as the data is deduplicated, the deduplication metadata moves with the deduplicated data to eliminate a scaling limitation often found in scale-up architectures.

To accomplish this feat, ExaGrid used a deduplication algorithm tuned for its scale-out architecture. Alternative deduplication algorithms typically use block level compares to deduplicate data. This technique excels at deduplicating data but it cannot scale because its hash tables grow too large.

Conversely, some deduplication algorithms use byte level compares to deduplicate data which are also very good at deduplication. The drawback with these algorithms is that in order for them to work effectively, they need to know the format of how every application formats the data inside of its backup data streams. When they have access to the data in the backup streams, they work great—arguably better than block level compares. If they do not, then their ability to deduplicate data is severely limited.

DCIG

Three Recommendations for Evaluating Data Deduplication to Ensure Storage and Bandwidth Efficiency While Delivering Fast Backups, Restores, and Recoveries

ExaGrid essentially takes a hybrid approach to data deduplication. ExaGrid stamps logical zones *(large blocks of data),* does a byte level compare within those zones and stores only the bytes that change from backup to backup. ExaGrid delivers the scalable approach that organizations need in a scale-out deduplicating backup appliance. It offers the granularity in deduplication to achieve the higher deduplication ratios that organizations expect.

"No-wait" Data Deduplication

Deduplicating data after the backup is complete is known as "post-process." A potential hazard of this approach is that data cannot be swiftly replicated to a disaster recovery site in another location. This results in a recovery point that is not as up-to-date as the recovery point of inline deduplication.

ExaGrid mitigates this by using a form of "adaptive" deduplication which starts deduplicating data as soon as backup data is committed to disk. However, ExaGrid does not wait until all backup jobs are done before it starts to deduplicate data. While backup jobs are occurring, the system monitors for lulls in the backup process or when its system's resources are not needed in part or in full.

During those periods of time, data deduplication and replication commence on the ExaGrid system. As soon as the data is deduplicated, the system may then concurrently replicate the data offsite. The net effect of this approach is that backups and backup data deduplication complete in less time than using inline deduplication while simultaneously creating an offsite copy of data that has an up-to-date recovery point objective (RPO).

ExaGrid Positions Organizations to Implement Deduplication without Compromising on Backups, Restores, or Recoveries

Deduplication is now an essential part of every organization's backup strategy. It ensures cost-effective, long-term retention for backup data on disk and facilitates the efficient offsite replication of data to conserve and optimize available network bandwidth. Yet deduplication's ability to perform these tasks has resulted in organizations treating it as a silver bullet that can solve all of their backup problems. This has led to it appearing anywhere and everywhere in the backup process—and not always for the better.

The ExaGrid system implements deduplication as part the backup process in a more thoughtful way that best satisfies the largest number of backup use cases that most organizations encounter.

The ExaGrid system:

- Creates a landing zone that facilitates the initial backup of data in a non-deduplicated form
- Stores data in this state for a period of time
- Immediately begins to deduplicate data once it is committed to disk
- Concurrently replicates deduplicated data offsite
- Non-disruptively grows its capacity and performance with its scale-out architecture

Organizations that select the ExaGrid system get a solution that solves all three of the compute challenges of implementing deduplication. In so doing, they get all of the benefits of deduplication without some of the unpleasant drawbacks. Maybe most importantly, ExaGrid positions organizations to implement deduplication without compromising on the speeds of their backups, restores, or recoveries. This is a scenario every organization wants to implement and which the ExaGrid system is uniquely configured to deliver. ■

DCIG

DCIG, LLC // 7511 MADISON STREET // OMAHA NE 68127 // 844.324.4552
dcig.com