Managing the information that drives the enterprise

STORAGE

Buying Guide:

TARGET DATA DEDUPLICATION BACKUP SYSTEMS

One key difference among target data deduplication systems is whether they dedupe inline or post-process; but there are a lot of other issues to consider that may be more important.

inside

What you need to know about target data deduplication

Special factors to consider

TechTarget

Buying Guide:
TARGET DATA DEDUPLICATION BACKUP SYSTEMS

One key difference among target data deduplication systems is whether they dedupe inline or post-process; but there are a lot of other issues to consider that may be more important than when your data gets deduped.

BY W. CURTIS PRESTON

KEY FEATURES: WHAT YOU NEED TO KNOW ABOUT TARGET DATA DEDUPE BACKUP SYSTEMS

SPECIAL FACTORS TO CONSIDER

FOR MOST DATA CENTERS, and especially larger ones, target data deduplication (or "target dedupe") is the easiest way to store all of your backups on disk without significant changes to the way you do backups. It could be as easy as deploying a target dedupe product and just pointing your backups at it—assuming it was properly architected and sized. Having said that, choosing the right target data deduplication system for your environment is perhaps the most crucial decision you must make. The following describes the key issues you'll have to consider to help you make the most appropriate buying decision.

TARGET DEDUPE DEFINED

Target dedupe (as opposed to source data deduplication) deduplicates data once it has arrived at your backup system. Data is sent unmodified from the backup client to the backup server and is then deduplicated. With typical target dedupe systems, no changes to backup software or processes are required. For

source data deduplication, you must use source dedupe-knowledgeable backup software on both the client and server, which may require some modification to your backup operation. With source dedupe, the client talks to the server during a backup and identifies redundant segments of data before sending them to the server; if a segment of data is determined to be redundant, it's not sent across the network.

There are two significant differences between target dedupe and source dedupe. Source dedupe requires a change in backup software (assuming you're not already using a backup application that has source dedupe capabilities), but can significantly reduce the amount of data sent across your LAN during backups. Target dedupe doesn't require any changes to backup software, but it also doesn't reduce the amount of traffic sent across your LAN from the backup client to the backup server.

Target dedupe doesn't require any changes to backup software, but it also doesn't reduce the amount of traffic sent across your LAN from the backup client to the backup server.

KEY FEATURES: WHAT YOU NEED TO KNOW ABOUT TARGET DATA DEDUPE BACKUP SYSTEMS

SPECIAL FACTORS TO CONSIDER

VARIETIES OF TARGET DATA DEDUPLICATION

There are different "flavors" of target data deduplication products. The most common type is an appliance that consists of a dedicated server running specialized software with one or more disk arrays connected to it. The server is often referred to as the dedupe "head" or "controller." Some target dedupe appliance companies may also sell a "gateway" version of their product that's only the dedupe controller without any disk attached to it; gateways allow customers to use the disk they feel is most appropriate for their environment. EMC Corp., FalconStor Software Inc. and IBM all offer gateway models of target dedupe products.

Target data deduplication is also starting to appear in backup software. Because some knowledge of the backup format can assist in the deduplication process, some backup software vendors think it's most appropriate for their applications to handle the deduplication process. This typically consists of some additional code running on the backup server that dedupes the data once it arrives at the backup server.

TARGET DEDUPE VS. SOURCE DEDUPE

While we've already discussed the differences between target dedupe and source dedupe, people are still often confused as to which technology is most appropriate for their environment. To get a better understanding, you should consider the problem that each technology was designed to solve.

Source deduplication was designed to overcome the difficulties of backing up remote data. By eliminating redundant data at the source, source dedupe lets you back up large servers over relatively small pipes, allowing you to have a centralized backup of several remote data centers. Target dedupe was designed to replace high-speed tape systems, so bandwidth usage wasn't a concern. What target deduplication systems have to be able to do is accept terabytes of backups in a relatively short amount of time—backups that often run at hundreds or even thousands of megabytes per second. So target dedupe is very good at going fast, but it doesn't save bandwidth.

Generally speaking, people use target dedupe systems for large data centers where bandwidth isn't a primary consideration, and source dedupe for backing up remote sites and mobile data. There may be some overlap, but that's the general rule of thumb.

APPROPRIATE USE OF TARGET DEDUPE

When purchasing a target dedupe system, it's important to remember why we're using disk in the first place—and it's not that disk is fast and tape is slow. The real reason is that it has become increasingly difficult to satisfy the throughput needs of today's high-speed tape systems by backing up directly across the LAN to those tape drives. Originally, disk was used in the backup process as a high-speed cache to help stream high-speed tape drives when the cache was copied to tape. Because most people who will deploy target dedupe systems will still copy their backups to tape, the ability of a target deduplication system to supply data fast enough to stream a modern tape drive should be one of the primary evaluation criteria for target deduplication systems.

When purchasing a target dedupe system, it's important to remember why we're using disk in the first place—and it's not that disk is fast and tape is slow.

IS TARGET DATA DEDUPLICATION OLD NEWS?

Some pundits are already declaring that target data deduplication is passé. They build their case on the premise that deduplication is now included everywhere (e.g., primary file-system deduplication) and that the source deduplication functionality available in backup software makes target dedupe systems obsolete. There are several flaws in that argument. First, while deduplication and high-speed compression is likely to be included everywhere eventually, it has barely begun to be deployed in most data centers, and many (although not all) of the products that offer this functionality are either new or still in beta. Second, with very few exceptions,

KEY FEATURES: WHAT YOU NEED TO KNOW ABOUT TARGET DATA DEDUPE BACKUP SYSTEMS

SPECIAL FACTORS TO CONSIDER

primary data deduplication does nothing to reduce the amount of duplicate data stored in backups, so deduplication is still needed in the backup system. Finally, while many smaller data centers can indeed deploy source deduplication systems in lieu of buying a target deduplication product, even the market leader in source deduplication only scales up to 36 TB in a single grid. In addition, the restore speeds of most source deduplication systems look great when you're restoring over a WAN connection, but they don't look as good when you're trying to restore a 20 TB database in your data center. While source deduplication is absolutely the way to go for backing up remote data centers and mobile data, it's not yet ready for today's mega-data centers.

For these reasons, target data deduplication is still very much relevant and hardly old news. That may change in five or 10 years, but you're dealing with backup problems today, and target deduplication can solve them for you.

TARGET DEDUPE vs. IMMUTABILITY

Another claim against target dedupe is that because it technically changes the data, it can't be used in environments with a requirement to store data in an immutable form (meaning it can't be changed after it's created). Does target data deduplication modify the data? Absolutely; so does power, sector-based storage, compression and even IP. The question is: Can you prove the email you're presenting in a court of law is exactly the same email (metadata and all) that was sent or received seven years ago? It's the job of your data archive application to prove that, and the fact that it was deduplicated, compressed, put into tar format, or chopped up into packets and transmitted across a network is totally irrelevant.

CHUNKING AND HASHING vs. DELTA DIFFERENTIALS

Target dedupe systems must examine large streams of backup data and identify which segments are redundant and can be discarded. There are two primary methods of identifying redundant segments. The first, and most common, is the chunking and hashing method. Backup streams are first sliced into large segments (chunks) that are typically anywhere from 8 KB to 256 KB. These chunks are then run through a cryptographic hashing algorithm to produce a "hash." That hash is then checked against all other hashes that the deduplication system has ever seen (this is referred to as the hash table lookup), and if the hash has been seen before, the data is considered redundant. If the hash is unique, then the segment/chunk is unique and is written to storage.

Delta differential systems work very differently. They must first identify which backups are similar to other backups. For example, tonight's full backup of the Exchange database from a server named Elvis will obviously be very similar to the last full backup of the Exchange database from Elvis. The

two backups are compared byte by byte looking for segments that are new; new segments are stored, while redundant segments are discarded.

Hashing vendors like to point out that they compare every chunk to every other chunk that has been seen, so they'll find redundant segments of data the delta differentials won't find. If a large spreadsheet is stored in the file system and is then sent as an attachment via Exchange, the same data might be in three or four different places; a hash-based deduplication system will detect these redundant segments and eliminate them. Because a delta differential system will only compare like backups to each other, it wouldn't find data that's redundant between different types of backups. However, they do tend to get better deduplication ratios on the data they do compare.

Delta differential vendors like to raise the issue of hash collisions. Much has been written about this topic, so there's plenty of available research. However, the chances of a hash collision with a hash-based deduplication system using SHA-1 (the most common algorithm) are so remote that it shouldn't be a concern when planning a purchase.

MULTINODE DEDUPE (AKA GLOBAL DEDUPE)

One very important feature to consider with target data deduplication systems is what happens when you buy more than one appliance so that you can load balance your backups across multiple appliances. If your target dedupe system supports multinode—or global—dedupe, then load balancing can be done easily. Redundant backup data will be identified, regardless of which node it was sent to. However, if your target deduplication system doesn't support global dedupe, your alternative is to scale in one of two (and perhaps difficult) ways. The first method of scaling is to replace your dedupe controller with a faster controller, keeping your existing disks with the deduped data stored on them. This may not seem like a terribly difficult process, but it's the definition of a "forklift upgrade" and means you'll be discarding the slower controller in favor of the quicker one at what may be considerable cost. The second way to scale is to create multiple dedupe "islands," split your backups into multiple subsets and then distribute them across these dedupe islands. While this might not be terribly difficult to do,

One very important feature to consider with target data deduplication systems is what happens when you buy more than one appliance so that you can load balance your backups across multiple appliances.

it's an arrangement that's extremely difficult to maintain over time. Because backup sets are constantly changing in size, users will have to deal with constantly moving data among the islands. Multinode dedupe is therefore an important feature to consider.

There's one more issue related to multinode dedupe. As with clustering, there's significantly more complexity between a two-node system and a system that supports three or more nodes. If you're considering a vendor that offers only a two-node global dedupe system, don't assume that a three or more node system will come anytime soon. It may be years before the vendor moves beyond two-node global dedupe, so keep that in mind if you anticipate significant growth in your backup operation.

The most hotly debated difference among target dedupe products is whether they do their deduplication inline or post-process.

INLINE vs. POST-PROCESS DEDUPE

Perhaps the most hotly debated difference among target dedupe products is whether they do their deduplication inline or post-process. To be considered an inline system, a target dedupe product must dedupe the data before it's written to disk; all deduplication is handled by the product's CPU and RAM, and the original, native data is never written to disk. A post-process system writes the backup data to a staging area before it begins deduplicating it. The data may be deduped seconds, minutes or hours later, although it's most common for deduplication to begin while backups are still going on. This is why many people refer to inline dedupe as synchronous deduplication, and to post-process dedupe as asynchronous deduplication.

The primary argument against the post-process architecture is the staging area, which is disk capacity that an inline system doesn't require. A post-process system typically needs at least enough disk for a single night's backup. But despite needing that additional capacity, the staging area is one of the primary benefits of using a post-process architecture because it allows for extremely fast restores from the latest backup, which can be a significant advantage.

One caveat about post-processing systems is that these types of vendors tend to advertise only their products' ingest speeds, or the rates at which you can send backup data to them, rather than their dedupe speeds. If a product can only dedupe at 100 MBps, it doesn't matter if it can ingest data at 10,000 MBps; at 100 MBps, it can only dedupe 8.6 TB per day. Be sure to check out the speed at which the product dedupes data.

The primary argument against the inline architecture is that the compute-intensive nature of deduplication could cause the inline system to throttle incoming backups and slow them down. This is true for some inline deduplication systems, but not all of them, so testing is critical.

Whether a product does inline or post-process deduplication isn't necessarily an important issue. What matters is how the product performs in your environment with your data. Don't believe the hype; believe the numbers you see with your own eyes.

Whether a product does inline or post-process deduplication isn't necessarily an important issue. What matters is how the product performs in your environment with your data. Don't believe the hype; believe the numbers you see with your own eyes.

WHICH INTERFACE: NAS, VTL OR OST?

Another feature to consider is what interface you want to use when sending data to the target data dedupe system. The most popular is certainly network-attached storage (NAS) because it's the easiest to set up and doesn't require the creation of virtual tape drives and tapes. The one downside to NAS is that it assumes you're going to use IP as the transport, and many people have moved to LAN-free backups where backup data is sent via Fibre Channel (FC). However, more and more multiprotocol products are emerging that support FC and IP.

If your shop uses Symantec's NetBackup or Backup Exec, you have a third interface to consider: Symantec OpenStorage (OST). Target vendors that want to integrate with the OST API must write a plug-in that's installed on the NetBackup or Backup Exec server. The backup software then sends backups to the OST API and the plug-in decides how they're transmitted to the target device. Most target data deduplication systems using OST over NAS or a VTL have shown improved performance, so that's the first reason to consider it.

The other, and perhaps more important, reason to consider OST as an interface is that it allows your backup software to control replication of your backups from one target data deduplication system to another. Each backup sent via OST has a unique ID, and your backup software can tell one dedupe appliance to copy that unique ID to a second dedupe appliance, and the appliance can then use replication to perform the copy. That way, NetBackup or Backup Exec can monitor and report on the replication process while staying aware of both copies. Without OST, NetBackup or Backup Exec

SPECIAL FACTORS TO CONSIDER

Aggregate vs. single-stream throughput
Most data deduplication system vendors advertise their product's aggregate throughput. What they don't advertise is the product's ability to accept a single stream of data, otherwise known as the single-stream throughput rate. If you're going to back up very large servers, the target data deduplication systems single-stream rate may significantly impact your ability to back them up quickly. If you're going to restore large servers, or copy from the deduplication system to tape on a regular basis, then the ability to create a large single stream during restore is very important. Make sure to investigate and test the target deduplication system's ability to accept and create a large single stream of data.

The "dedupe tax"
You may have heard about the "dedupe tax" that makes restores from a deduplication system slower than if the data hadn't been deduplicated. But just as the taxes you pay depend on the area where you live, the dedupe tax you'll pay depends on the product you choose. Some dedupe systems have almost no dedupe tax, while others may suffer significant performance challenges during large restores. They may also suffer from a greater dedupe tax depending on the backup used to create the restore; restores may be faster depending on whether or not you're restoring from a recent backup or an older backup. Make sure to learn about a product's dedupe tax before you consider purchasing it.

wouldn't know if the replication worked or not, nor would they know that a second copy of the data exists.

TEST EVERYTHING

All of the factors described here are likely to figure into your final decision, but the most important criteria should be the results of your product testing. Believe nothing, and test everything. What matters is how big it is, how fast it is and how much it costs—what it does to get those results is irrelevant. ⊙

W. Curtis Preston is an independent backup expert.