

STORAGE

Buying Guide:



SOURCE DATA DEDUPLICATION PRODUCTS

Source dedupe products can efficiently protect remote site data by significantly reducing the amount of backup data that needs to be transmitted. The products are generally easy to implement and use, but you'll need to match their capabilities with your specific remote site needs.

inside

- 2 Key features of source data deduplication products
- 5 Special considerations

Key features of SOURCE DATA DEDUPLICATION PRODUCTS

Source dedupe products can protect remote site data by significantly reducing the amount of backup data that needs to be transmitted. The products are generally easy to implement and use, but you'll need to match their capabilities with your specific remote site needs.

BY W. CURTIS PRESTON

bACKUPS of remote sites and laptops have plagued IT departments for years. While remote sites typically have relatively small amounts of data in comparison to centralized data centers, the data that's "out in the field" is often just as important as home office data. Because producers of data "at the edge" are often closer to a company's customers, the data they create may be among the most important information a company has.

But as important as the data may be, safeguarding it is often left in the hands of non-IT personnel or staff members with little or no training. In many cases, the person entrusted with backing up the site's data lacks any significant computer experience and doesn't understand the importance of effective backups. But the traditional choices for backing up remote data require quite a bit of computer-related knowledge.

Remote backup options typically meant running some type of backup software at the site that was connected to a removable-media storage device, usually a tape drive. Managing those re-

KEY
FEATURES OF
SOURCE DATA
DEDUPLICATION
PRODUCTS

SPECIAL
CONSIDERATIONS

movable media devices is the biggest problem at many remote offices. Being certain that media is inserted, ejected, rotated and stored properly can be a challenge. Offsite storage of backup media is another issue, requiring pickup by a vaulting service or some other means of securing the media off-premises for each remote site. Some firms cut corners and skip shipping the media offsite altogether, while others simply put backup tapes in someone's trunk or living room.

Mobile users add logistics to the remote-office backup problem. Those users may attempt to back up their data by using rewritable DVDs, solid-state "thumb drives" or USB-connected external hard drives, but those

devices require a certain degree of expertise to use properly. And while they're made to be transported and are less prone to environmental issues than tape, they're often carried along with a laptop and can be lost or stolen just as easily.

None of these methods—whether they appear to be working or not—adequately addresses the issue of getting the data back to a company's central site.

Offsite storage of backup media is another issue, requiring pickup by a vaulting service or some other means of securing the media off-premises for each remote site.

HOW SOURCE DATA DEDUPLICATION CAN SOLVE THE PROBLEM

Backup and recovery of remote-office data is still a struggle, but in the past five years there have been a number of developments that can take some of the sting out of securing remote data. All of the options available today do a better job than traditional methods, and most are far easier. Source data deduplication has emerged as one of the most popular options as it not only takes care of backing up remote servers and workstations, but can efficiently ship the backup data to a central location using existing wide-area network (WAN) facilities.

Data deduplication—or dedupe—is a process that identifies and eliminates redundant data at a sub-file level. It doesn't just find redundant files; it identifies redundant segments of data within and among files. There are two primary types of dedupe systems: target deduplication and source deduplication.

Target deduplication systems are disk systems that are used as the destination for backup streams delivered by traditional backup applications. When native, non-deduplicated backup jobs are received by the target system, the data is then deduped. The deduped backups can then be replicated to another target dedupe system for offsite storage.

KEY
FEATURES OF
SOURCE DATA
DEDUPLICATION
PRODUCTS

SPECIAL
CONSIDERATIONS

A target dedupe system can be used to back up remote sites, but that would require some kind of target deduplication appliance at each site with the data then replicated to a central site. But purchasing a dedupe appliance for each remote site can be very costly. And the appliances can't be used to protect laptops when they're disconnected from the remote sites' network.

Source deduplication doesn't require special hardware, and remote-office sites and laptop users can be supported using a single product. Source dedupe can provide both onsite and offsite recovery mechanisms for larger remote sites, even with very demanding data recovery objectives.

WHAT YOU NEED TO KNOW ABOUT SOURCE DEDUPLICATION

As with typical backup applications, source deduplication products typically include a backup server and a piece of client software that's loaded on every server or PC that needs to be protected. But some products (most notably those from Asigra Inc.) operate in a clientless fashion where the source dedupe server logs in to each application server to perform the requisite operations. Therefore, there isn't an agent permanently residing on the server that needs to be backed up, which, especially for large installations, can make administration a little easier. (For our purposes, we'll use the term "client" to refer to the part of the system that enables a server to be backed up, even if it's not permanently installed on the server.)

The first backup a dedupe system does is often referred to as the "seed," and it's essentially a full backup, although it still works very differently compared to typical backup software. For file system backups, the client walks its way through each file system to be backed up, looking for files that are newer than the last backup. Because it's the first time that a particular system is being backed up, all of the files in the system will be "newer" and all of them will be backed up.

A typical backup product would just stop here and back up each new file to the server, but that might waste precious bandwidth. Although the backup server already knows that a particular file hasn't been backed up from that specific location before, it might have seen parts of the file somewhere else. It may even be the same exact file that has already been backed up from another server, or possibly a version of another file that's been backed up before.

To check for these possible redundancies, source dedupe products slice

The first backup a dedupe system does is often referred to as the "seed," and it's essentially a full backup, although it still works very differently compared to typical backup software.

KEY
FEATURES OF
SOURCE DATA
DEDUPLICATION
PRODUCTS

SPECIAL
CONSIDERATIONS

each file into subfile segments of various sizes. These segments are often referred to as chunks and their actual size can range from 8 KB all the way to 256 KB, depending on the product and the scenario. There's a direct relationship between segment size and performance, and an inverse relationship between segment size and the resulting deduplication ratio. The bigger the segment size, the better the performance, because there are fewer hashes to create and to look up in the hash table. Bigger segments also reduce the size of the hash table, which makes hash table lookups quicker. With smaller segment sizes, the system examines the data on a more granular level so it can yield a better dedupe ratio. Taken to the extreme, a source deduplication system could find a lot of commonality if the seg-

SPECIAL CONSIDERATIONS

SOURCE DEDUPLICATION was developed specifically for backing up remote-office site data. One of the ways this design intent manifests itself is in the aggregate restore capabilities of these systems. Let's assume, for example, that the fastest a given source dedupe system can restore is 200 MBps, which should be more than adequate to restore a single remote site.

And even if you have multiple remote sites, the odds that you'll need to restore more than one of those sites at the same time are improbable at best. It's not likely that you'll lose Wichita and Orlando on the same day.

But what happens if you're using one of these products in a virtualized server environment or even in a regional data center? The odds that you could lose more than one server if they're all collocated in the same data center are a lot higher. And in virtual server environments such as Citrix XenServer, Microsoft Hyper-V or VMware, the odds are pretty high that you might lose multiple virtual machines (VMs) at the same time with a disk failure in the array holding your VMFS file system. With such a failure, you would have to restore multiple virtual servers at the same time. A throughput of 200 MBps might be fine if you're restoring across a wide-area network (WAN) link, but when you have terabytes of data to restore and all kinds of available bandwidth, 200 MBps looks kind of puny: It's only about the speed of one fully streaming LTO-4 tape drive using 1.5:1 compression.

Bottom line: Be careful how you apply source deduplication technology in non-remote-office scenarios such as virtual server environments and data centers.

ment size was very small, but the number of hashes required would profoundly affect performance. Each source dedupe vendor has had to weigh these alternatives to arrive at what they consider the optimal balance of performance and deduplication efficiency in their products.

The hashing part of the process involves running each segment or chunk through a cryptographic hashing algorithm called SHA-1, which produces a 160-bit value referred to as a hash. The client sends this hash to the backup server, which checks its value against all of the hashes already stored in its hash table—a step referred to as a hash table lookup. Some products also store on each client a localized cache of the hashes that have been backed up by that client, making hash table lookups even quicker for hashes that have been seen before on that client.

If a hash has never been seen, the segment is backed up by transferring it across the network and storing it on disk. If it has been seen before, it's backed up by simply noting that another segment exists with the same hash.

All subsequent file system backups are very similar to the first backup. The file system is walked once again, with the system looking for files that have been created or modified since the last backup. Those files are sliced into segments and hashed; a hash table lookup is performed and the system takes the appropriate action based on the results of the lookup.

Some source deduplication products also handle special types of application data, such as that from Microsoft SQL Server, Microsoft Exchange, NDMP or Oracle. For those data types, the only difference is that file system walking is replaced by the appropriate command for a full or incremental backup. The dedupe product then slices and hashes the data that comes from that backup just as it does with file system data. Special software is required for each special data type a source deduplication application supports as it needs to know how the application's data streams are constructed.

Some source deduplication products also handle special types of application data, such as that from Microsoft SQL Server, Microsoft Exchange, NDMP or Oracle.

Slicing, hashing and looking up are a lot of work, so the client wants to do all of that on as little data as possible. In file system backups, the workload is reduced by only slicing and hashing files that are newer than the last backup. For special data types, it's usually done by asking the app for an incremental backup.

Restores work similarly to traditional restores, with the exception of how the data is assembled before being transferred across the network.

KEY
FEATURES OF
SOURCE DATA
DEDUPLICATION
PRODUCTS

SPECIAL
CONSIDERATIONS

Data that needs to be restored must be “reassembled” from the various distributed pieces spread out across the disks used by the backup server. This is essentially a very fragmented disk read, and can obviously take longer than reading the same data stored contiguously on disk. The time it takes to reassemble the requested data is often referred to as the “rehydration penalty,” and the penalty can vary widely between products. Source deduplication products tend to have slower restore rates than target dedupe systems, but the slower restore rates are typically still much faster than the WAN pipes the data is travelling across, so rehydration delays typically aren’t noticed by most users. Still, when evaluating source dedupe systems, it’s important to understand what their restore speed capabilities are.

Here are three questions you should ask source deduplication vendors:

- 1.** “Assuming 10 Gb Ethernet [10 GbE] of bandwidth to a given client, and assuming a file system that can write as fast as you can supply the data, what’s the fastest you can restore a single file system?” Source dedupe vendors are likely to be skeptical about conditions like unlimited bandwidth, and will probably say that file systems are always slower than their ability to write. But press them on the issue and make them give you an upper limit.
- 2.** Let’s assume the vendor told you their system could do a single restore at 200 MBps. That’s pretty good performance, which equates to restoring 700 GB in an hour. Your next question should be “What type of configuration of your product will I have to buy to get 200 MBps of restore throughput to an individual client?”
- 3.** Now let’s get to the worst-case scenario and the next question. “You can restore a single system at 200 MBps. How fast can you restore several systems simultaneously, assuming the same 10 GbE network and very fast clients?” The response will give you an idea of how long it would take to restore an entire remote office if you lost multiple servers.

Assuming the answers to questions 1 and 3 meet your recovery requirements, ask modified versions of question 2 based on those requirements. In other words, if your recovery requirements can be met with only 100 MBps performance, ask the vendor for a configuration that can support that level of throughput. If a product’s recovery capabilities don’t match your requirements, you may need to consider another product. If none of the available source deduplication products meets your recovery requirements, then you’re probably looking at the wrong technology.

Source dedupe is a relatively new technology, so some vendors may not be used to answering tough questions like those above. Be patient, but don’t let them off the hook without answering them.

TEST EVERYTHING AND BELIEVE NOTHING

Asking detailed, specific questions will get you a lot of useful information about a source deduplication product, but you won't truly know how it will perform for you until you get it into your environment. A bit of skepticism is required, so you should test each product in production to see if it can achieve the performance numbers provided by the vendor.

Build a test 10 GbE network and assemble the configuration the vendor said would meet the requirements outlined in the above questions—even if that level of performance is overkill for your environment. If nothing else, you'll find out if the product can meet future needs. Finally, create a configuration that more accurately meets your needs and test that as well. Assuming the tests go well, then the only thing left to do is haggle over the price. ☹

W. Curtis Preston is an independent backup expert.

KEY
FEATURES OF
SOURCE DATA
DEDUPLICATION
PRODUCTS

SPECIAL
CONSIDERATIONS