ExaGrid Product Description

Cost-Effective Disk-Based Backup with Data Deduplication

Contents

Introduction

The movement from tape-based to disk-based backup is well underway. Disk eliminates all the problems of tape backup. Backing up to disk is faster than backing up to tape, for and disk-based backups are more reliable and more secure than tape-based backups. Unlike tape, disk resides in a hermetically sealed case and in a data center rack so it is not exposed to heat and humidity and is secure as it sits behind data center physical security and network security.

In the past, the movement from tape- to disk-based backup has been less compelling due to the expense of storing backup data on disk instead of tape. Despite the disadvantages of tape, tape's cost advantage has allowed it to maintain its presence in the data center. With the advent of data deduplication, however, tape's cost advantage has been largely eroded.

When disk-based backup is used in conjunction with data deduplication, only the unique data is stored and depending upon the type of data deduplication and the specific implementation, reduction rates of 10 to 1 to as much as 50 to 1 can be realized. This allows only a fraction of the disk to be required versus straight disk without data deduplication.

Let's look at a simple example below to see how this works. Assume the backup environment is as follows:

- 5 TB of primary data
- Full backups of all data every Friday night (weekly); incremental backups on files and full backups on e-mail and databases on a nightly basis (Monday-Thursday, or four nights per week)
- 12 weeks of onsite retention for the full weekly backups; 4 weeks of onsite retention for the nightly backups

To back up this amount of data with straight disk would require the following:

- Each nightly backup is about 25% of the full backup (so 25% of 5TB) = 1.25 TB per night; four nights per week over 4 weeks = (1.25 TB) x (4 nights) x (4 weeks) = **20 TB**
- Each weekly full is 5 TB, so over 12 weeks, this would require (5 TB) x (12 weeks) = **60 TB**
- Total disk needed is 20 TB + 60 TB = 80 TB of usable disk; adding in disk required for RAID brings up the total amount of disk needed to about **100 TB**.

Using 100 TB of disk for this type of backup environment would of course be prohibitively expensive for most organizations. However, with data deduplication, one can expect to reduce the amount of disk needed to about 5% to 10% of the amount of disk needed for straight disk. This means that with data deduplication you can perform disk-based backups with only **5 TB – 10 TB** of disk, in contrast to the 100 TB needed in the example above. And it is this drastic reduction in the amount of disk needed that has put disk-based backup on a comparable footing with tape in terms of cost, and that has enabled the wave from tape- to disk-based backup that is now underway.

Considerations When Examining Disk-Based Backup Approaches

Now that it is economically feasible to move from a tape-based to disk-based backup approach, a large number of vendors with varying approaches have developed disk-based backup systems employing data deduplication. This has caused a great amount of confusion for IT managers looking to adopt a disk-

based backup system for their organization. One cannot assume that all disk-based backup approaches are created equally – and the following areas should be examined closely when evaluating various disk-based backup systems.

1. Backup Performance
2. Restore Performance
3. Deduplication Approach
4. Scalability
5. Support for Heterogeneous Environments
6. Support for Backup App Features Such as Granular Level Restore (GRT) and OST
7. Offsite Data Protection
8. Total Cost of Ownership

ExaGrid – A Disk-Based Backup Approach

ExaGrid started with the premise of building the best disk-based backup solution taking into account all of the considerations discussed above. While ExaGrid achieves some of the best deduplication rates in the industry, as you can see from the discussion above, a disk-based backup solution is not about deduplication alone. It is about addressing all of the key components related to backup and applying them to a disk-based backup solution. In the remainder of this paper, we will discuss how ExaGrid addresses these various aspects of disk-based backup.

1. Backup Performance: The Fastest Backup Performance with Post-Process Deduplication

ExaGrid employs *post-process deduplication* that allows the backups to write directly to disk at disk speeds. This produces a faster backup and shorter backup window. The rationale here is to defer the compute-intensive process until after the backup has landed, so as not to impact the time it takes to perform the backup. Another approach in the market is *inline deduplication*, which deduplicates data on the fly, before it lands on the disk. Because inline deduplication can potentially cause a bottleneck at the point where data is streaming into the backup appliance, inline deduplication can result in slower performance and a longer backup window (see figure at right).

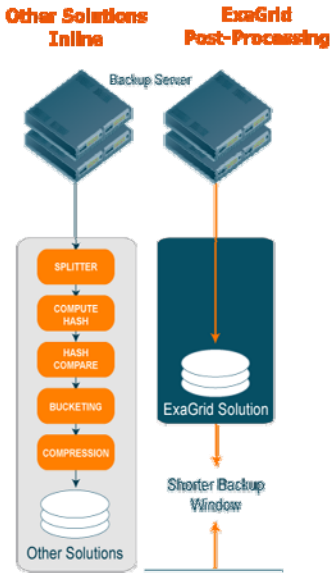

Proponents of inline deduplication often argue that their approach requires less disk and is therefore less expensive. However, because inline deduplication must rely on faster and more expensive processors--and more memory--in order to avoid being prohibitively slow, any cost differences in the amount of disk used are overcome by the need for more expensive processors and memory. Post-process deduplication also provides additional advantages with respect to restores and tape copies, which will be discussed further below.

2. Restore Performance: Quick Restores and Offsite Tape Copy

Full system restores are the most important restores as hundreds to thousands of employees at a time can be down when a full system is down. The longer it takes to recover the more lost hours of productivity. Nearly all disk-based backup appliances and backup software based deduplication implementations use inline deduplication (discussed above), however. Unfortunately, in addition to the issues mentioned in the previous section, the inline deduplication method requires a given backup to be rehydrated, or put back together from its deduplicated state, in order to be restored. This approach takes time – and time is typically not a luxury when a full system is down!

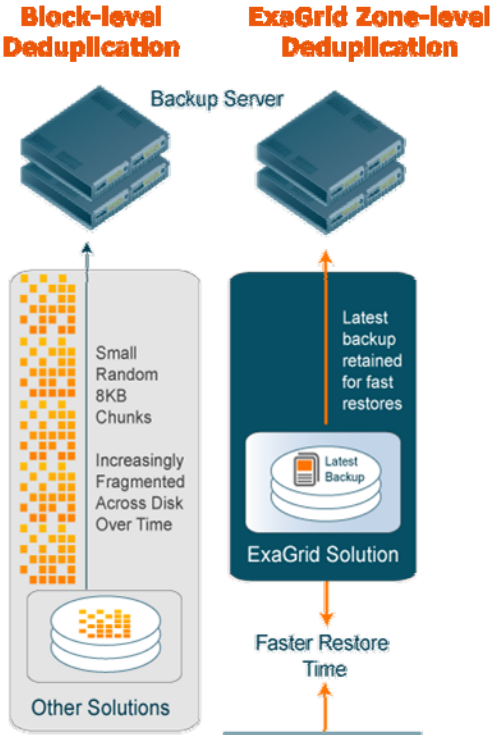

ExaGrid's post-process approach, because it allows backups to land to disk prior to deduplication, is able to keep that backup on disk in its complete form. And as you proceed with nightly and weekly backups over time, the ExaGrid appliance maintains a complete copy of your most recent backup on disk in its complete form, so that it can be rapidly restored when or if it is needed. This approach saves valuable time and productivity in the event of a full system restore. It is also quite useful with virtual server backups using server virtualization such as VMware. In this case, because a complete backup typically consists of one or more virtual servers in their entirety, ExaGrid's ability to enable rapid restores of the most recent backup effectively gives you the ability to restore multiple virtual servers very quickly.

Finally, ExaGrid's post-process approach is very useful for making fast tape copies. Because the ExaGrid appliance keeps your most recent backup in its complete form, this same backup can very easily be used to quickly generate a tape copy. With an inline deduplication approach, a tape copy would require the backup to be put back together again (rehydrated) prior to being sent to tape, even if the tape copy were scheduled mere moments after the backup itself took place. The result using inline deduplication, then, is a much slower tape copy – and a longer period of time until the data is fully copied to tape and protected offsite.

3. Deduplication Approach: Zone-Level Deduplication

There are several key areas to look at when evaluating a vendor's deduplication approach. The first, most basic aspect of deduplication is how well it reduces data. After all, the whole point of deduplication with respect to disk-based backup is to reduce the amount of disk needed such that the cost of backing up to disk can remain low. But there are other key aspects of deduplication that can have a profound impact on the ability of the solution to support a variety of backup applications, and on how well the solution is able to scale as a customer's data grows.

One common method of deduplication is known as "block-level" deduplication. This method takes a block of bytes and then looks for other blocks that match, storing only the unique blocks. The key to block-level

deduplication is the size of the block. Smaller block sizes, say around 8 KB, can be more easily matched and hence will result in a higher deduplication rate than larger block sizes (e.g., 64 KB). Block-level deduplication when used with smaller block sizes achieves excellent data reduction. This method, because it is generic in nature, also lends itself well to supporting a variety of different applications.

The problem with block-level deduplication, however, is its lack of scalability. Because block-level deduplication stores and matches unique blocks, a tracking table (known as a hash table) is required to manage all of the backup data that is stored. And the smaller the block size, the larger the hash table that is needed – such that with 8 KB size blocks, one billion entries are needed to deal with just 10 TB of data! This forces an appliance architecture consisting of a controller unit with multiple disk shelves – an inferior configuration that will be discussed further in the section on scalability, below.

ExaGrid utilizes a type of deduplication called "zone-level" deduplication. With zone-level deduplication, the backup jobs are broken into large 50 MB to 100 MB segments (instead of blocks). These segments are then broken down into zones, or areas, and the deduplication algorithm looks for the bytes that have changed from one backup to the next. Like block-level deduplication, zone-level deduplication achieves excellent data reduction, and lends itself well to supporting a variety of different applications.

Unlike block-level deduplication, however, the tracking tables required for zone-level deduplication are much smaller. The tracking tables can therefore be easily copied across appliances, allowing for a much more scalable grid-based architecture, as discussed below.

4. Scalability

Scalability is another important aspect in evaluating a disk-based backup appliance, and there are a couple of important things to consider when examining how well a given solution scales. First is the ability of the solution to scale as the amount of backup data grows, and second is how easily the system can be upgraded when additional horsepower is needed.

When looking at the various disk-based backup configurations available on the market, there are two basic alternatives – a controller / disk shelf model, and a grid-based system. With the controller-shelf model, all of the processing power, memory, and bandwidth are contained in the controller. Some disk may be contained in the controller as well, but when there is more data and a need for expansion, additional disk shelves are added to the controller. This implies a static amount of processing power, memory, and bandwidth for a given system even as the amount of data is growing, which in turn results in one or both of the following negative effects: (i) as the amount of backup data grows with a constant level of processing power, memory, and bandwidth, the backup starts to take longer and longer; (ii) the amount of processing power, memory, and bandwidth must be over provisioned when the system is first acquired, to allow for future data growth, but resulting in a more expensive system at the time of purchase. In addition, each controller can only handle a certain amount of disk, and when the customer's data increases beyond that level, the entire system must be swapped out for a new one in a costly "fork lift" upgrade.

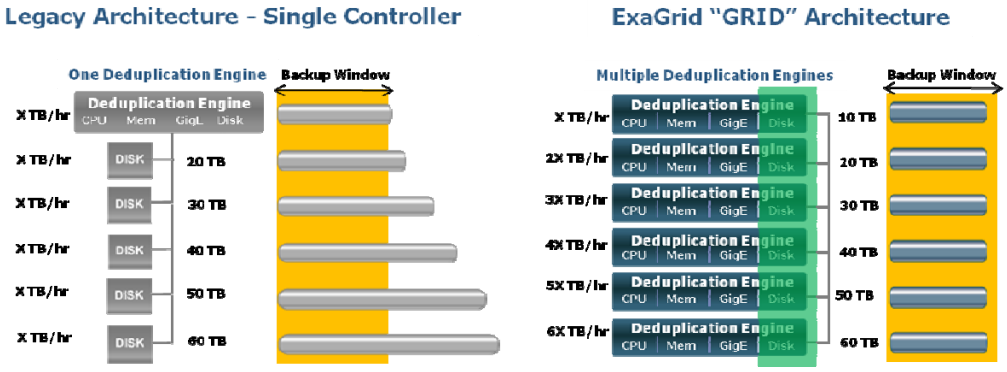

ExaGrid instead uses a grid-based configuration, where each appliance contains processing power, memory, bandwidth, and disk. When the system needs to expand, additional appliance nodes are attached to the grid, bringing with them additional processing power, memory, and bandwidth, as well as disk. This type of configuration allows the system to maintain all the aspects of performance as the amount of data grows – you are no longer simply adding disk to a static amount of processing power, memory, and bandwidth – and you are only paying for the amount of processing power, memory, and bandwidth as you need it, rather than up front. A grid-based approach also avoids the costly fork lift upgrades that come with controller / disk shelf configurations.

In addition to maintaining backup performance as your data grows and allowing you to seamlessly upgrade to larger and larger systems, ExaGrid's grid-based configuration automatically load-balances available capacity in the grid, maintaining a virtual pool of storage that is shared across the grid. All of the systems can also be managed using a single user interface that is able to access all of the systems on the grid. This provides a simple, single-dashboard view to give you a quick view of deduplication and replication status for any system in the grid.

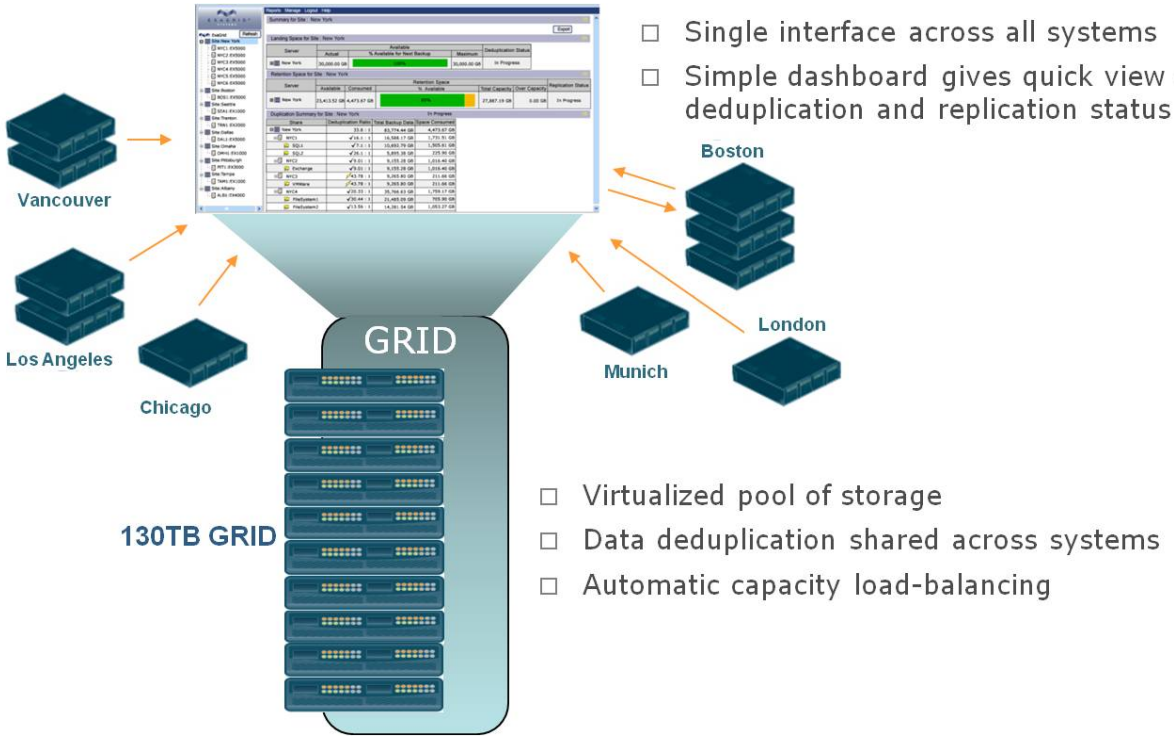

5. Heterogeneity

Customer environments are made of many backup approaches, backup applications and utilities, and different disk-based backup approaches support these in different ways. Customers may have any number backups occurring in their environment, including traditional backup applications, specialized VMware backup utilities, direct-to-disk SQL dumps or Oracle RMAN backups, and specific UNIX utilities such as UNIX TAR.

Backup application software solutions that have incorporated deduplication by definition only support their own backup application, with its own backup server software and its own backup client agents. These solutions are not able to support backup data from other backup applications or utilities.

Disk-based backup appliances with data deduplication such as ExaGrid's, however, are able to support backup data from multiple sources, including a variety of backup applications and utilities. Performing deduplication in the backup software limits the ability to have all data from all sources stored and deduplicated in a single target device. Unless 100% of your backup data passes through that particular backup application, a purpose built disk-based backup appliance such as ExaGrid's is the best choice to meet the requirements of your entire environment.

6. Support for Backup Application Features such as GRT and OST

Another area to consider when looking at disk-based backup solutions is how well a particular solution supports advanced backup application features such as GRT (Granular Level Restore) and OST (Symantec Open Storage). Some solutions do not integrate well with these features – poorly-implemented GRT solutions, for example, may take hours to restore an individual e-mail, or may not work at all.

Symantec's Open Storage is another popular feature that allows for more integrated offsite data protection, and it is important to check whether these features are supported if you are using this with Symantec NetBackup or Backup Exec.

7. Offsite Data Protection

There are many reasons to keep a complete set of backups offsite at a second location. This can be accomplished by making offsite tape sets or by replicating data from the primary site disk-based backup system and the second site disk-based backup system. There are many questions to ask when considering offsite data protection in a disk-based backup system:

First, what is the deduplication rate? As discussed earlier, deduplication rate in the determining factor in the amount of data that is reduced and the amount of disk that is required. But deduplication rate also impacts the amount of bandwidth that is needed to maintain a second site, for a given backup window, since it is only the data that has changed that is sent over a WAN to an offsite system. The poorer the deduplication rate, the more data must be sent to maintain the offsite backup, the more bandwidth is required, for a given backup window. Deduplication rates can vary greatly, particularly when looking at backup application software deduplication. ExaGrid achieves the highest deduplication rates and requires the lowest bandwidth.

Second, does the offsite system keep only deduplicated data or some form of already rehydrated data in order to offer quick disaster recovery restores? Any system that does inline deduplication only stores the deduplicated data and therefore results in slower disaster recovery (DR) times. As mentioned earlier, ExaGrid performs post-process deduplication, which makes the most recent backup available in its complete form. And as an ExaGrid appliance performs replication to an offsite ExaGrid appliance, that most recent backup is also maintained in its complete form on the offsite system as well. The result is that the data is ready to quickly restore from either the local or the remote system.

Third, can data from the offsite system be used to restore any lost or corrupted data on the primary site? ExaGrid owns the patent for this technology such that if anything happened to any of the backup data at the primary site the offsite system can be used to restore / replace the lost or corrupted data. This creates an added level of safety.

Fourth, does the system allow you to demonstrate DR restores for an auditor? ExaGrid has a dedicated function to be able to demonstrate DR restores for an internal or external audit.

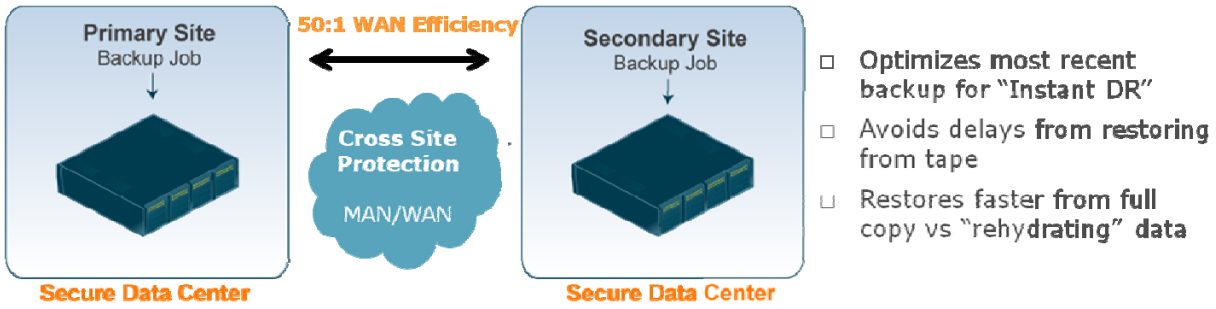

8. Total Cost of Ownership

Cost needs to be examined both up front and over time. You want a lower price up front but also over time you don't want to have to repurchase any part of the system. Because ExaGrid performs post process deduplication, an ExaGrid configuration does not require the same level of processor and memory that the controller / disk shelf approach uses. This allows ExaGrid to be more cost effective up front. Because ExaGrid's zone-level deduplication requires smaller tracking tables and allows the system to scale via a grid-based configuration, you can add servers into a grid and keep adding as you grow. There are no points where you must perform a fork lift upgrade, so your investment is protected as a result.

It is also important to look at cost effectiveness when comparing ExaGrid to non-appliance based deduplication systems, such as deduplication via the backup application software. On the surface, it would appear that backup application software deduplication would be fairly low-cost. After all, you can just turn on deduplication from within the backup server and you're good to go, right? Not exactly. It is important to keep in mind that using backup application software deduplication typically requires greater resources on the backup server – that is, more processing power, more memory, and more disk. In addition, how do you determine exactly how much of these components are needed to optimally work in your particular backup environment? And what do you do as the amount of data you're backing up grows? To answer these questions likely means additional costs in the form of professional services. When you get to the bottom line, then, the additional costs of all these items will typically exceed what you will be paying for an ExaGrid appliance.

So whether you're comparing ExaGrid to other appliance solutions, or to non-appliance solutions such as backup application software deduplication, ExaGrid is the most cost effective disk-based backup solution, up front and over time.

Conclusion

The movement from tape-based to disk-based backup is indeed underway. There are a great variety of approaches in the market today. While deduplication is an important factor, disk-based backup does not begin and end just with deduplication. It is important to take into account all of the various factors discussed above when evaluating these systems.