Storage Switzerland, LLC

Future-Proofed Backup For A Virtualized World

Prepared by: Colm Keegan, Senior Analyst
Prepared: January 2014

Like death and taxes, growing backup windows seems to be an inevitable fact of life. To stave off the backup window, IT architects have resorted to a myriad of tactics to stay one step ahead. The most popular tactic in recent years, has been utilizing a disk based backup solution that leverages deduplication for storage efficiency.

Data deduplication has become a standard offering in backup product portfolios. As a technology that has been mainstream for over five years, some IT decision makers may tend to treat deduplication as an afterthought. In actuality, how data deduplication is implemented is a critically important factor towards determining its ultimate success with solving backup ills in the data center or creating them.

While these solutions can help improve cost per GB of disk backup, they can also introduce latency and extend backup processes. This is especially true with respect to how deduplication can potentially impact data recovery times.

Selecting a disk-based backup solution should be more than going with the “default” vendor. Both the architecture of the disk backup system and how it integrates data deduplication, can directly impact backup and recovery windows. Understanding and selecting the right disk backup architecture therefore, is critical to designing a solution that meets business needs today but also, more importantly, in the future.

Disk and Tape Conundrum

Many organizations have utilized disk to disk to tape backup solutions (D2D2T) to decrease backup windows and enable fast copies to tape media. Rotational disk drives can handle multiple backup steams from backup applications for rapid backups and can efficiently stream tape drives to allow for fast copies to tape.

The first challenge with D2D2T infrastructure, however, is that as backup data grows, the front-end disk storage repository also has to increase to accommodate this growth. For example, a 10TB environment needs 20TB’s of disk storage to manage one week’s worth of backup data (one full and four daily incremental copies). To store an additional week’s worth of data, this figure must double to 40TB’s, and so on. As a result, most organizations can typically only afford to store between 1-2 weeks worth of backups on disk.

The second challenge is that D2D2T infrastructure is still very tape centric. In order to satisfy offsite data requirements for DR purposes, data has to be exported to tape nightly and then transported via courier to an alternate location. Replicating data that is not first deduplicated over WAN links is simply not an option due to all the bandwidth that would be required to move this information.

Dedupe Driven Delays

The challenges of D2D2T have sent many IT Planners in search of an alternative. Disk based data deduplication provides a solution to these issues since multiple week’s worth of backup data can now be efficiently stored on disk as data is stored in a highly compressed or deduplicated format.

From week to week there is about a 2% data change rate. Deduplication only stores or replicates the changed data and therefore is both disk storage and WAN bandwidth efficient. Data can be efficiently stored to reduce the amount of disk required and backup data can be efficiently replicated offsite over WAN links, to electronically vault data. This enables many organizations to dramatically reduce their reliance on tape.

It is important to note, however, that most organizations don’t replace tape in its entirety. They still will make tape based copies on a monthly or quarterly basis for an extra layer of protection and for additional data retention. This is also true in the mid-tier to small enterprise data centers. Many of these organizations can’t justify the costs associated with implementing electronic data vaulting - like bandwidth, secondary disk backup appliances and secondary data center infrastructure. For these organizations, tape is still the primary method of moving data off-site and the ability to create a rapid tape copy is critical to the protection process.

Dedupe Induced Latency

While data deduplication provides significant benefits, it can also have some unintended consequences. Data duplication is a computational and memory intensive process and as a result, it is possible for backup and recovery windows to actually increase when data deduplication is performed. In other words, the added workload of deduplicating data during the backup process before it is safely landed on disk storage can elongate backup windows.

Moreover, deduplication can also add latency to the recovery window and impact application recovery time objectives (RTOs). The fact is, reconstituting or “rehydrating” deduplicated backup data back to its native format, introduces latency into the recovery process; especially if that recovery is large or has many files associated with it.

If a recovery operation is taking place from a backup image, presumably the information is not available elsewhere. Which means that the affected business application will remain down until the data can be fully recovered. As a result, the data deduplication dehydration/rehydration process can potentially have a detrimental impact on RTOs.

Scale-Up Latency

To compensate for the added latency of deduplication, some disk backup appliances over-provision CPU and memory to increase the speed of deduplication processes. But this only serves to inflate the cost of these offerings and reduce their return on investment (ROI) when compared to comparable disk and tape backup systems. This is especially true in “scale-up” backup appliance disk architectures that require an up-front investment in the maximum CPU and memory resources supported by the array.

Hybrid Powered Recoveries

One way to circumnavigate the latency issues that data deduplication introduces into the recovery process, is to implement a hybrid disk backup solution. In other words, a backup platform which incorporates both non-deduplicated disk storage to enable rapid backup and recoveries, and a separate deduplicated storage partition which efficiently stores backup images for longer-term data retention and offsite replication. Think of it as a recovery zone and a backup archive zone.

The non-deduplicated disk storage partition would only need to accommodate a week’s worth of backup images (one full backup and 4 nightly incrementals) since the vast majority of restore requests typically take place within the first 7 days of data creation.

While this does require sacrificing some disk storage efficiency, the benefit is that there would be no deduplication induced performance penalty during the restore operation. Furthermore, since only one week’s worth of data would need to be maintained in the native area, this portion of the disk footprint could be reasonably small. The remaining portion of the disk backup system could then be reserved for efficient deduplicated backup images.

Making this small architectural change can have a big impact on how quickly organizations can respond to critical data restore requests; especially large ones like a full virtual machine or server recovery. In fact, the potential time difference between performing full application restores from native disk and deduplicated storage can be quite dramatic.

In some cases, it can be merely a matter of minutes to restore a large data set from non-deduplicated storage versus the number of hours it could take to restore the same information from deduplicated disk. Clearly this can be the difference between successfully meeting RTO’s versus needlessly extending application downtime on critical business systems.

Scale-Up vs. Scale-Out Backup Systems

In scale-up storage architectures, all the CPU, memory and disk are contained within a single chassis or design frame. In scale-up designs, storage I/O performance is shaped like a bell curve. Performance gradually increases as disk drives are added to the array but then eventually performance plateaus and then decreases as the array approaches maximum disk capacity.

More critically, performance does not increase as the data and resulting deduplication load increases and therefore the backup window continues to expand. The only way to then reduce the backup window is to replace the front-end controller with a larger and faster system.

This need to periodically refresh scale-up backup appliances with newer models places an increased burden on IT operational staff to plan for and manage upgrades. In addition, there is often a higher total cost of ownership (TCO) with investing in scale-up architecture designs, especially when compared to alternatives, like scale-out systems.

Scaled-Out Performance

Scale-out disk systems are a completely different architectural paradigm from traditional scale-up disk platforms. Scale-out disk systems consist of independent appliance nodes that contain their own discrete containers of storage, CPU, bandwidth and memory. The appliance nodes are then automatically integrated together via software to form a scalable GRID system.

With scale-out systems there is no need to over-provision CPU and memory resources on the initial deployment. Instead, the system can start out small and be gradually scaled out over time to meet additional backup capacity requirements as the environment grows.

One of the major benefits to this approach is that it dramatically simplifies the upgrade process. In fact, older appliances typically can be intermixed with newer appliances so there is never an end of life requirement. This helps organizations maintain investment protection in their backup storage assets.

A second benefit to scale-out backup systems is that as the data and resulting deduplication load grows, the appropriate resources (processor, memory and bandwidth) are also added. This keeps the backup window at a fixed length and eliminates costly fork lift upgrades. This is a key advantage for organizations that struggle with containing their backup windows.

With scale-out systems like those from ExaGrid, full server appliances can be added into a scalable GRID. All appliances onsite and offsite are viewed through a single user interface. In addition, the data across the appliances automatically load balances across all appliance nodes. There are various sized appliances allowing you to add the right amount of compute and capacity, as you need it.

Hybrid Backup Flexibility

Additionally, to further enhance backup and recovery speeds, some scale-out backup architectures incorporate a separate disk “landing zone” with native disk that is unencumbered by the workload of data deduplication processes. Backups are sent directly to disk avoiding the compute intensive process of deduplication during the backup process. This helps to speed up backup times. These systems maintain 7 days of backup data on native disk storage and then migrate data older than 7 days to a deduplication partition for longer-term data retention.

Since backup jobs do not have to go through the data dehydration process, this allows for backup copies to land very quickly on disk. And since the most recent backups are always in their complete non-deduplicated native format in the landing zone area, offsite tape copies can be quickly made without having to go through a lengthy re-hydration process. This is particularly important for organizations that plan to continue to use tape as their primary means for offsite disaster recovery. For those organizations that desire to eliminate tape at the disaster recovery site, deduplicated data can be replicated to a secondary offsite system. Since only deduplicated data traverses the WAN, it is extremely bandwidth efficient.

Virtual Environment Image Booting

There are additional practical considerations for implementing a disk backup solution which comprises a separate non-deduplicated disk partition. Virtualization backup applications like Veeam and Dell's VRanger, now provide support for booting virtual machines (VMs) directly off the backup images on disk. The concept, often called instant recovery, allows administrators to point their users directly at VM backup images on disk, to achieve rapid application recovery. Needless to say this is an extremely popular feature and many times a new backup application is selected solely for this capability.

This may be a valuable option for those environments that don't have high availability clustering solutions deployed, however, if VM operating system backup images and application data are only stored in a deduplicated format, instant recovery could be a non-starter. The time it would take to "rehydrate" deduplicated data before it could be presented back to the VM on the backup platform in its native format could simply take too long.

In short, performing "in-place" data recoveries on a deduplicated disk partition might be untenable due to the latency that is introduced into the process. Blocks of data would need to be continuously un-deduplicated and re-deduplicated. The random IO would kill performance and the deduplication engine itself.

In fact, performance can be so poor that some deduplication appliance vendors recommend installing stand-alone disk in front of their systems just for the purposes of instant recovery. This would obviously introduce additional cost and management complexity into the backup environment.

On the other hand, if backup images are immediately available in their native form in a front end cache or landing zone, then virtual administrators can simply boot the virtual machine directly off of the disk-based backup system. This enhances the value of the investment in the disk backup storage platform as it can now be used as a way to further improve RTO's for critical virtualized application infrastructure.

Conclusion

Organizations that are interested in implementing data deduplication to improve efficiencies, may want to consider how deduplication can adversely impact their backup and recovery windows. If elongated backup windows and application RTO's are a concern, then deploying a disk staging landing area without deduplication could prove beneficial. As importantly, however, the system should also include a separate disk zone for storing deduplicated backup images so that organizations can attain the benefits of efficiently maintaining data on disk for extended retention. It also fosters the efficient replication of backup data to a secondary location for DR purposes.

Furthermore, the challenges of effectively and efficiently managing unrelenting data growth and extended backup windows may require a break from traditional backup storage architecture design. The widespread adoption of scale-out storage architectures in primary storage environments is clear evidence that a new paradigm is needed to help organizations scale storage capacity and performance more linearly. Likewise, backup storage infrastructure needs to evolve to meet these same challenges.

Data compute intensive processes like data deduplication arguably make the need for GRID based scale-out designs even more critical for effectively scaling backup infrastructure. GRID-based Scale-Out backup architectures, like those from ExaGrid, give mid-sized to small enterprise data centers the best of both worlds - fast recovery on a classic disk staging area with a separate data duplication storage zone.

This hybrid approach to backup provides organizations with a way to reduce backup windows, speed up data recoveries from disk and garner all the disk retention and WAN replication bandwidth efficiencies available through data deduplication. This helps organizations flexibly grow their backup infrastructure without resorting to time consuming and costly forklift upgrades.