TANEJA GROUP
TECHNOLOGY ANALYSTS

S O L U T I O N   P R O F I L E

Using Tape Virtualization to Improve Backup Performance

August 2008

THE NET NET

If you're like most companies, you've probably got a data protection infrastructure that relies heavily on tape. And you're probably struggling with backup window, recovery point objective (RPO), recovery time objective (RTO), and recovery reliability issues. Disk-based backup approaches promise to address these issues, but introduce some issues of their own. On the one hand, the use of disk in backup infrastructures allows the introduction of other technologies important in data protection, such as storage capacity optimization technologies (like data de-duplication) and replication, and for that reason it can pay handsome dividends. On the other hand, adding disk adds cost, may change processes somewhat, and needs to be done in a way that complements your existing tape-based infrastructure. Tape continues to be a viable resource for backup environments, and when supplemented by disk in the right way offers improved performance and reliability for backup related tasks.

Virtual tape library (VTL) technology is arguably the easiest way to introduce disk into a tape-based backup infrastructure. A VTL is a set of disks that emulate tape; you see a disk array, your backup software sees a tape drive or a tape library. The use of tape virtualization will improve backup performance and reliability while allowing you to preserve investment in your existing tape infrastructure. But choosing the right VTL technology depends upon what particular backup issues are most important in your environment. This paper considers these issues, then provides a quick review of Sun Microsystems' tape virtualization solutions, identifying how their offerings can best be leveraged to produce an optimized backup infrastructure.

Understanding Tape's Strengths and Weaknesses

Tape has been the mainstay of backup infrastructures for more than thirty years. It offers very low costs on a $/GB basis for data storage, and can support effectively unlimited capacity. Today's tape subsystems can also support very high data transfer rates if the tape drives can be kept streaming. For full system restores, a streaming tape drive can actually provide better restore performance than that available from disk. Tape is a sequential access media though, and because of that has inherent weaknesses when it comes to certain backup tasks that favor random access.

When we refer to backup, that includes both backup and restore tasks. Generally, backups are performed every day, and if you're like most backup administrators, you probably spend more time backing up your data than you do restoring it. Daily backup jobs are comprised of backing up one or more data sets from one or more servers on a regular basis – random access tasks that do not lend themselves very well to the type of high speed, sequential writes at which today's high end tape drives excel. There is a lot of "start" and "stop" as backups run, and this leads to an occurrence called "shoe shining" on tape drives. Shoe shining refers to the start/stop activity regularly seen on tape drives during backup jobs: the tape drive begins writing data in short bursts, but each time, before it begins writing anew, it must reposition the tape back to the point where the writes stopped. The mechanical nature of the tape drive leads it to slightly overshoot the stopping point each time, driving the need to reverse the tape to reposition back to where the previous write stopped. If the drive did not reposition each time, a significant amount of storage space would be wasted on the tape.

Shoe shining puts a lot of mechanical strain on a tape drive. A tape drive that can be kept streaming will perform much more reliably than one which repeatedly shoe shines over time, decreasing the useful life of the tape drive as well as the tape media itself. Tape media reliability issues would be significantly lower if a tape drive could be kept from shoe shining. Through the use of multiplexing, an approach where multiple files, possibly from multiple backup clients, are backed up concurrently, backup administrators can keep a tape drive streaming and thereby minimize the amount of shoe shining a tape drive must perform. But approaches like these add significant complexity, not only in planning backups but also in performing restores.

On the restore side, most restore requests are for one or more discrete objects, such as a deleted or corrupted file. While backup administrators must be prepared to perform full system restores, the need to actually perform them is quite rare. Because of tape's sequential access method, you cannot go directly to a file for a restore from tape. You will need to load the tape and run through it sequentially until you reach the location of the file that was requested. This adds time to the restore process as it puts a high load on both the tape drive and the tape media relative to the amount of data actually accessed and transferred. When full system restores are required, a large amount of data must be sequentially sent to a particular server. Tapes are very good at this type of request since they can be kept streaming, putting relatively little load on the tape drive and media for the amount of data transferred.

Understanding Disk's Strengths and Weaknesses

Disk, because of its random access nature, is much better suited to both initial backups and object-level restores. Disk's ability to accommodate the short, bursty nature of most backup jobs makes it a perfect "landing pad" for most initial backups, while its ability to randomly access specific files makes it a faster, more reliable medium to handle object-level restore requests, most of which are taken from recent backups. If all backups are performed directly to disk, and then the latest backup is always retained on disk, several advantages accrue. First, backups tend to complete noticeably faster, shortening the backup window. Second, because they can be completed more quickly, backups can be performed more often, leading to improved RPO. Third, if we assume that most restores are for one or more objects as opposed to full system restores, disk's random access capability leads to much better RTO. And finally, the use of disk for most restores should lead to improved restore reliability relative to an all-tape infrastructure.

Where disk falls down, however, is in two areas: cost and provisioning. Raw disk capacity for high performance, Fibre Channel (FC)-based primary storage can be as much as two orders of magnitude more expensive than tape on a $/GB basis, while SATA-based secondary storage is generally at least one order of magnitude more expensive. And to be used as a backup target, raw disk must be provisioned effectively, a time consuming administrative task not required when tape is used as the backup target.

But the use of disk has two other huge advantages which recommend its appropriate use in backup infrastructures. First, disk-based data can be reduced through the use of storage capacity optimization (SCO) technologies which can achieve realistic data reduction ratios of 20:1 or more over time against backup data. The level of data reduction achievable through SCO has a significant impact on the overall $/GB; for a SATA-based configuration which prices out at $10/GB for raw capacity, the use of SCO technology can result in a cost of $.50/GB for usable capacity. This makes a disk-based backup tier possibly only 3-5x more expensive than a tape tier (assuming 20:1 data reduction for disk and 2:1 data reduction for tape). Still not nearly as inexpensive as tape, but possibly inexpensive enough to consider the use of disk to store some percentage of backup data where disk's random access capabilities provide relevant advantages.

Second, the use of disk provides access to replication technologies that make data distribution between sites faster and more secure. Data can be scheduled for replication immediately upon completion of backup jobs, offering significantly better RPOs from remote sites and less transport risk than that achievable with more conventional approaches that ship boxes of physical tapes to remote facilities.

Blending Disk and Tape for Optimal Backup Performance

While disk offers some significant advantages in backup environments, it is not the best medium to meet all backup requirements. Ironically, while the same can be said of tape, it was the single medium used for *all* backups for many years, a fact that has contributed to many current complaints about tape-based backup. If disk and tape are blended, however, with each performing the backup tasks to which it is best suited, a high performance, cost-effective backup infrastructure can be created and maintained over time.

Disk excels in its random access performance, so it makes sense to use it to handle the random access nature of initial backups and frequent object-level restores. But disk is higher cost, so it makes sense to understand the data access patterns driven by your backup and restore requirements, and deploy disk only where its performance characteristics are required. Tape, on the other hand, offers very low cost storage capacity, as well as good restore performance for the types of restores that are often required in DR situations (full system restores). Once backups have been completed to disk, they can be streamed at high speeds to tape drives for longer term storage without impacting backup performance at all. In this way, data can be retained over long periods much more cost-effectively than on disk and yet still be available for restores if required. Tape drives used in this manner will perform much more reliably than tape drives that are used as initial landing pads for backups since shoe shining will be minimal.

Most organizations are dealing with more data than ever before, and it is growing faster than ever before. These are the types of organizations where blending disk and tape can have a real benefit. For these organizations, the most cost-effective way to incorporate disk into an existing tape-based backup infrastructure is to take a "disk to disk to tape" approach. Initial backups are performed to disk, retained for some period of time, and then eventually migrated to tape for longer term storage.

Disk can be introduced into your backup infrastructure either as "disk-as-disk" targets or "disk-as-tape" targets. Both offer better performance and reliability than physical tape for initial backups and frequent, object-level restores, but impose very different management overhead. Backup software is natively built to recognize, interoperate with, and manage tape devices, and must often be complemented with additional cost options to be able to interact directly with disk drives as disk targets. Even with these enhancements, the addition of this disk tier is very visible to the backup software, and may require changes to your existing backup processes and procedures. Finally, disk-as-disk targets also impose provisioning tasks not required with tape targets, requiring additional administrative time and effort to effectively integrate disk.

Integrating disk-as-tape through the use of VTL technology, on the other hand, affords all the benefits of disk technology, including access to both SCO and replication technologies, while still retaining tape's management semantics. When deployed as a "landing pad" for initial backups, VTLs afford the performance characteristics of disk, with their associated benefits in backup window, RPO, RTO, and recovery reliability, while you continue to use your existing backup software and its catalog to manage backups and restores. In our opinion, VTLs offer a compelling and arguably the easiest way to introduce the advantages of disk into a tape-based backup infrastructure, since they require little to no modification of your existing backup schedules and processes.

Storage Capacity Optimization (SCO)

SCO technologies reduce the amount of raw storage required to store a given amount of data, primarily by finding and replacing redundant data with references. Backups are a particularly relevant area in which to use SCO because there are generally very high levels of redundancy between successive daily backups. If data is only changing on average 2-3% per day, it is not unreasonable that more than 95% of the data across 30 days of daily backups is redundant. SCO technologies find this redundancy, generally replacing redundant pieces of data with references while reliably retaining a copy of the original data. “Single instancing” finds and replaces redundancy at the file level, while “data de-duplication” works at an even more granular level, breaking files down into smaller chunks and running comparisons at this level to achieve even higher data reduction ratios. Data de-duplication achieves higher ratios because it finds a significant amount of redundancy between two files where only a title or date may have changed, whereas single instancing would see those as two different files.

There are two general SCO architectures available today: in-line and post-processing. In-line architectures add latency as the data is captured, analyzed, and converted on the fly prior to being written to disk for the first time; post-processing architectures pick the data up after it has been written to disk to process it, so while they do not add latency prior to the initial disk write, they do require additional time and storage capacity before the data is in SCO form. The major advantage to in-line SCO is that can minimize the overall disk storage capacity required with the risk of possibly slowing down the initial backup somewhat. The major advantage to post-processing SCO is that it will not impact backup performance in any way, but it may require more storage capacity and overhead as data must be read from disk a second time to complete SCO operations. Unlike in-line SCO, this places less dependency on the hard disk drive technology (FC vs SATA) and may lengthen the period of time until the data is available in capacity optimized form. Typically on reads, both architectures take an in-line approach when converting data back to its original form, although this generally only adds several microseconds to the read latencies for any particular piece of data.

There are variations to these basic operations. For example some SCO solutions allow the customer to set policies to determine when de-duplication is performed. An example of such a policy might be one where the last backup is retained on disk in non-de-duplicated form, but all older backups are de-duplicated. Since most restores are done from the most recent backup, this would allow them to incur without the latency required to return de-duplicated data to its original form.

The economic advantages of SCO technology are obvious. Why buy 10TB of raw disk capacity to store 10TB of data when you can store it in 500GB using SCO technology? There are several issues of which to be aware when considering SCO technology though:

SCO has a performance impact: today’s highest performance SCO VTL solutions can process data at no more than 400MB/sec. Note that the industry’s highest performance VTLs that do not use SCO technology may support data rates up to 2400MB/sec. If you require high backup performance, then SCO VTLs may not be an option.

Data must be reliably recoverable, ensuring not only that conversion back to its “full” form always results in the exact same sequence of 1s and 0s that were originally converted, but that the data will continue to be available in the face of inevitable disk failures; look for technologies that verify data on the fly to address the first issue, and the use of RAID and/or logical redundancy schemes to address the second.

SCO technologies are not free, so its important to understand just when the amount of data you want to retain on disk can justify its use. These calculations will vary based on a number of variables, such as the cost of the SCO solution you deploy, the $/GB cost of disk against which you want to deploy SCO technology, the data reduction ratios achievable in your environment, and understanding any additional savings that may accrue as a result of its use (such as any cost savings associated with replicating smaller data sets between sites, etc.).

Replication

Replication technologies distribute data, either in real time or via established policies, between sites across wide area networks (WANs). For long distance distribution, like that which is generally required between remote facilities, asynchronous replication over IP is recommended. VTL products generally offer replication as an option, enabling asynchronous replication over IP between VTLs.

Replication makes data distribution between sites faster and more secure. When data can be electronically transmitted to a central site for remote vaulting, it can generally arrive there at least 1-2 days faster than a physically transported tape. This can have a very positive impact on the RPOs associated with disaster recovery (DR) plans. The use of SCO and replication together offers the additional benefit of significantly reducing the amount of data that has to be sent across WANs; transfers between VTLs at different sites have already been capacity optimized, making the most of available bandwidth. And because the data can be encrypted without impacting the initial backup, it can be transmitted very securely without negative operational implications.

Deploying Tape Virtualization for Optimal Backup Performance

When considering how best to implement tape virtualization technology into your existing backup infrastructure, it is important to understand whether backup performance or data replication is your single most critical backup issue. This makes it easy to determine where to implement VTL technology, which offers up to 2400MB/sec of performance, and where to implement SCO VTL technology, which can significantly reduce WAN bandwidth requirements but only offers up to 400MB/sec of performance.

Higher backup performance tends to be more critical in environments with very large data sets that are growing rapidly. Data centers and other large server environments are good examples of situations where higher backup performance is critical. Distributed environments where data is being transferred between sites, and minimizing the amount of data sent across WANs is a strong consideration, are environments where the use of SCO technology is a good fit.

Distributed environments tend to deal with less data per backup client, and SCO VTLs lower performance may be more than sufficient to handle initial backups.

Evaluating Sun's Offerings for a Blended Backup Infrastructure

Sun's pedigree in tape subsystems includes a long history of dominance in the tape industry (through their StorageTek acquisition in 2005), and their tape subsystems lead the industry today in terms of installed base. Even so, Sun understands the key role that disk will play in backup moving forward, and offers the disk-based backup technologies necessary to meet today's evolving backup requirements. These technologies include VTL, SCO, and replication.

Sun's VTL product line includes two different families: VTL Prime and VTL Plus. Both families support easy expandability from the smallest to the largest unit, significant tape management and emulation capabilities, replication, a wide range of backup software, and simple co-existence and investment protection throughout the family. Two key features differentiate the two product families though: backup performance and support for SCO technology. The VTL Prime family includes SCO technology and supports backup performance up to 400MB/sec, while the VTL Plus family does not include SCO technology (although it does include hardware compression) but supports backup performance up to 2400MB/sec.

When the requirement for replication is key, VTL Prime offers strong advantages. It can be deployed in remote office branch office (ROBO) environments where data can be capacity optimized locally prior to replication, minimizing bandwidth requirements to transfer backups to a central site. VTL Prime supports tape-format aware data de-duplication that operates in a post-processing manner, providing realistic data reduction ratios of 20:1 or more in most backup environments. This leads to a usable effective capacity of 25TB to 580TB per VTL.

When the requirement for backup performance is paramount, VTL Plus fits the bill. VTL Plus is Sun's highest performance VTL offering, supporting up to 6 times the backup performance of VTL Prime. Hardware compression provides some level of data reduction along with the high performance, but the data reduction ratios will be on the order of 2:1, leading to generally higher bandwidth requirements for remote replication between VTLs. Additionally, VTL Plus supports the following capabilities:

- Support for up to 64,000 virtual tape cartridges per node, making it Sun's highest capacity VTL offering
- NDMP 4.0 (i.e. Direct to Tape) support which allows NDMP 4.0-compliant backup applications to track data throughout the infrastructure as it is migrated, per policies controlled by the backup software, from VTLs to physical tape. It also allows VTLs to support backup software-based features, such as tape stacking, without having to modify native VTL capabilities, bringing new

features to NDMP 4.0-compliant VTLs more rapidly.

- Tape caching, which allows the copy of of virtual tape volumes to physical tape volumes to be controlled and managed by VTL Plus independent of the back-up application. VTL would essentially be placed in the role of a disk-based "landing pad" to physical tape (when NDMP 4.0 cannot be used for this)
- Active-active configurations for high availability
- Support for legacy environments like ACSLS and IBM AS/400 (iSeries, System i)

Sun's two VTL solutions can also be deployed together. In a data center where both high backup performance and the need to keep a considerable amount of recent backup data locally on disk before migrating it to physical tape is a requirement, a VTL Plus may "front end" a VTL Prime. In this scenario, all backups would be initially written to the VTL Plus, but after some amount of time would be migrated to the VTL Prime where they could be retained at much lower $/GB costs. This offers high performance for initial backup ingestion, but then also allows a considerable amount of disk-based backup data to be retained locally at very low cost, supporting better RTOs than would be achievable if data was moved to an offsite location and dumped to tape.

Taneja Group Opinion

The opportunity to add disk to tape-based backup infrastructures offers backup administrators a powerful tool for improving backup performance and reliability while optimizing the overall spend on backup storage for low cost. Using tape virtualization, such as that provided by VTL solutions, is in our opinion the best way to do that. VTL technology provides the advantages of disk in addressing backup window, RPO, RTO, and recovery reliability issues with a minimum of disruption, allowing all data to continue to be managed by your existing backup software as if it was still on tape. And it allows SCO and replication technologies to be used to further reduce the secondary storage and data transport costs associated with a backup infrastructure. Tape virtualization offers compelling functional and economic benefits, and should at least be considered by any IT organization that is backing up several terabytes or more of rapidly growing data on a regular basis.

While it is necessary to protect a company's data by having a backup infrastructure in place, backup does not drive revenue and therefore should be managed as a cost center. Backup performance, secure data distribution, and secondary storage costs are all key issues to be managed. Sun's range of tape virtualization solutions gives their customers the tools to manage these issues, optimizing their environment to meet their customers' paramount concerns in the most cost-effective manner.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com