

White Paper

FUJITSU Storage ETERNUS CS8000 and IBM® TSM (Tivoli® Storage Manager)

ETERNUS CS8000 in TSM environments – an ideal solution for storage consolidation, data backup and disaster resilience.



Content	
1 Introduction	2
2 Overview	3
2.1 Optimize the overall backup and recovery process – including real tape	3
2.2 What our customers say	4
3 Benefits	5
4 How is ETERNUS CS8000 used optimally in a TSM environment?	6
4.1 Minimizing primary disk pools	6
4.2 Reducing competition for tape drive resources	6
4.3 Effective use of TSM collocation on tape	7
4.4 Increasing TSM LAN-free connectivity	7
4.5 Faster mounts and instant positioning	8
4.6 Deduplication for tape data	8
4.7 Compression on disk	8
4.8 Industry leading data availability	8
4.9 Move physical tape workload to ETERNUS CS8000 and free TSM server resources	9
4.9.1 Reduction in the TSM Reclamation process	9
4.9.2 Optimization of migration processes	10
4.9.3 Automated media refreshing and technology changes	10
4.9.4 Relieving the load from the TSM server	10
4.10 Omission of the copy pool and TSM offsite tape exports	11
4.11 Reduce the number physical tape drives	12
4.12 Prevent tape errors by decoupling the physical tape drives from the TSM server	12
4.13 Save physical tape cartridges and library slots by ETERNUS CS8000's high utilization of tape media.	12
4.14 ETERNUS CS8000 backup and archive solution including tape	12
5 Certified for IBM TSM	13

1 Introduction

ETERNUS CS8000 radically simplifies data protection in heterogeneous IT environments allowing consolidating backup systems on a large scale and offers an easy way to introduce disaster resilience for backup and archiving. As an integrated appliance it centrally manages backup and restore of data from mainframes, UNIX and x86 systems combining the best of disks and tape in a uniform way.

Tivoli Storage Manager (TSM) is IBM's company-wide storage management and backup application that supports backups to local tape storage media, local disk storage, or NAS storage.

TSM is very powerful and feature rich technology. Compared to similar solutions it has a special design. Two of the major distinctions are, that TSM basically introduces an 'incremental forever' backup strategy and is able to intelligently improve own operations by integrating disk storage as a first backup target.

TSM offers many options in setting up and running a backup and recovery environment, therefore implementation details and usage scenarios mostly differ. Because no TSM installation is alike, also different opinions appear when the question arises if 'tape virtualization' can offer significant benefits for operating TSM.

A common architected virtual tape library (VTL) uses disk storage and emulates a tape robot with its associated tape drives and tape media. This has many potential advantages that are described later in this paper. In practice a disk based VTL's capacity is limited and in most cases real tape is still needed to be used by TSM.

The ETERNUS CS8000 data protection appliance is much more than a plain VTL.

It offers all the benefits of emulating a tape environment on a disk based device too, but in addition it has real tape technology directly attached and automatically integrates low cost tape storage into the overall virtualization concept. With this, the whole backup process, down to finally storing data on physical tape, can be optimized by ETERNUS CS8000.

In heterogeneous environments, ETERNUS CS8000 is the only data protection appliance that combines the presenting of a VTL interface with integrated usage of real tape storage.

Together, Fujitsu ETERNUS CS8000 and IBM TSM supply an ideal solution for storage consolidation, data exchange, data backup and disaster resilience. This paper likes to point out some of the potential advantages Fujitsu's - tape aware - ETERNUS CS8000 data protection appliance offers.

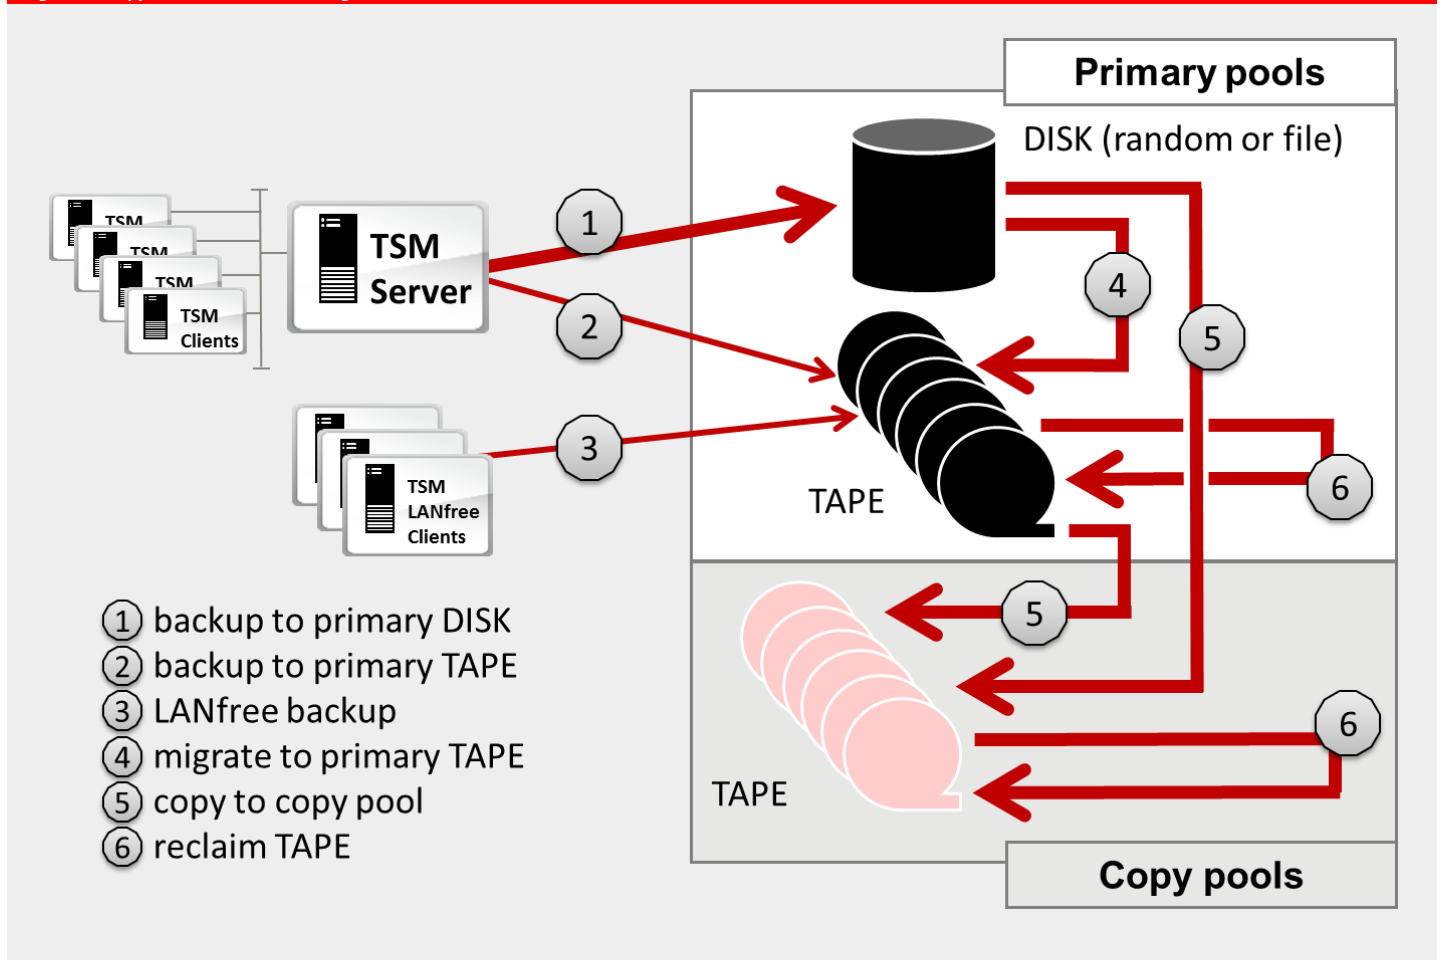
2 Overview

2.1 Optimize the overall backup and recovery process – including real tape

To get a first glance overview the following illustrations show the classical TSM processing (Figure 1) compared to TSM processing optimized by ETERNUS CS8000 (Figure 2). The most important possible improvements are discussed separately in more detail.

The typical TSM setup uses disk- and tape- technology to build the primary backup pools. To ensure data availability within TSM, the data is again copied to copy pools. Usually all the backup storage tiers are dedicated to TSM, the TSM Server has full control over them. Besides the actual backup processing to the primary pools, TSM migrates data between tiers and optimizes storage utilization. All data movement that results from copying, migrating and housekeeping uses significant resources and takes time.

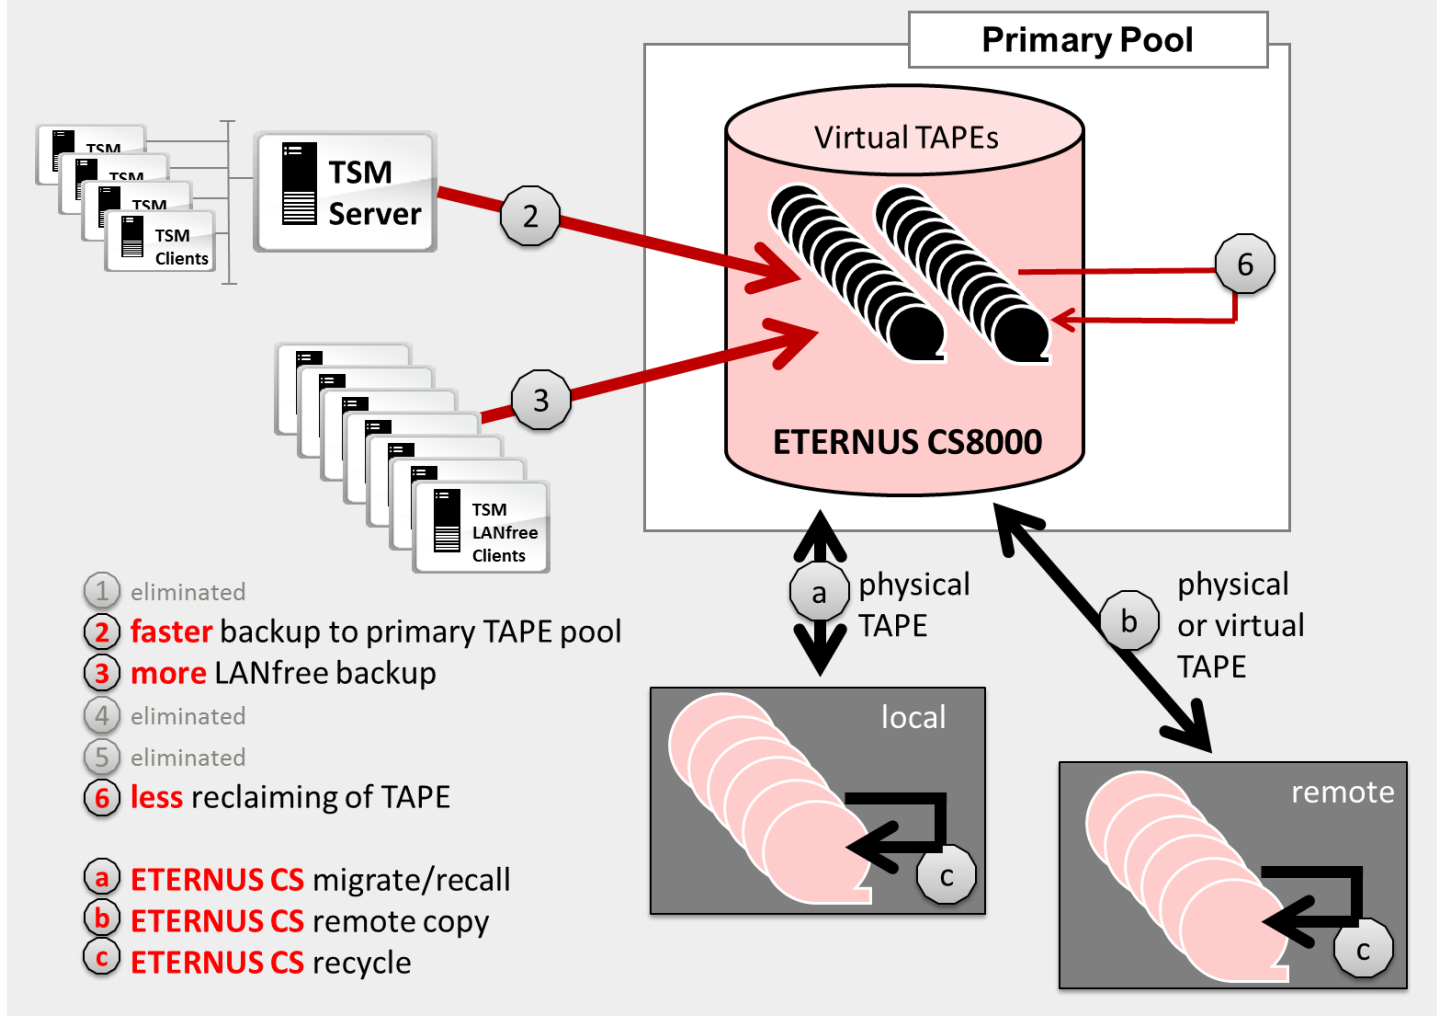
Figure 1: Typical TSM Processing without a VTL



In general the implementation of a virtual tape solution can significantly optimize TSM processing. A VTL allows eliminating or massively reducing the primary disk pool, it provides a higher parallelism, faster performance and increases availability.

Beyond the VTL functions the ETERNUS 8000's direct real tape integration allows the workload of the internal TSM data management processes (migrations, creating copies, major part of reclamation) to be shifted to the ETERNUS CS8000 and thus frees TSM resources. Multiple physical tape copies on different media, or even in different sites, can autonomously be created by ETERNUS CS8000.

Figure 2: Optimized TSM Processing by using the - tape aware - ETERNUS CS8000



2.2 What our customers say

Our current ETERNUS CS8000 TSM customers observe cost savings by reducing the expensive primary disk pool to a minimum lower the number of TSM instances by offloading work from TSM dramatically decrease the number of real tape drives and highly simplifies administration and monitoring

'Various tape technology changes could be achieved fully transparent for the backup/restore processes for various platforms (Open Systems / Mainframe) thanks to ETERNUS CS8000. This has minimized the migration effort and the risk of data loss.' (leading German bank)

3 Benefits

The combined disk- and tape- architecture of an ETERNUS CS8000 virtual tape system delivers TSM serious superiorities over using a classic tape infrastructure or a "disk based only" VTL.

Overview: Advantages with ETERNUS CS8000

General VTL - Advantages when using ETERNUS CS8000

- Minimizing primary disk pools
- Reducing competition for tape drive resources
- Effective use of collocation on tape
- Increasing TSM LAN-free connectivity
- Faster mounts and instant positioning
- Deduplication for tape data and compression on disk.

Additional Advantages when using ETERNUS CS8000 with real tape

- Provide an industry leading data availability by taking advantage of
 - synchronous mirroring between sites
 - replicating over long distances
- Move physical tape workload to ETERNUS CS8000 and free TSM server resources by
 - reducing TSM reclaim, migration and copy
 - refreshing real tape technology independently
- Omission of the copy pool and TSM offsite tape exports
- Reduce the number physical tape drives by being able to utilize modern tape drive performance.
- Prevent tape errors by decoupling the physical tape drives from the TSM server
- Save physical tape cartridges and library slots by ETERNUS CS8000's high utilization of tape media.
- ETERNUS CS8000 backup and archive solution including tape

4 How is ETERNUS CS8000 used optimally in a TSM environment?

4.1 Minimizing primary disk pools

One reason TSM uses disk as the first storage tier is that disk can handle multiple IOs from many sources concurrently. Tape drives can only process one single data stream at a time. Furthermore, if a single data stream is not fast enough to keep a tape drive busy, the tape performance suffers extensively. In this respect, disk is absolutely superior, because many slow backup streams can be directed to a single disk device without decreasing the overall performance.

For TSM a disk buffer is important, because at normal processing TSM does not multiplex backup streams when writing to tape. Using Disk allows TSM to increase the number of concurrent backup sessions to temporary disk data storage until a high-performance and space saving backup to tape is possible.

Recovery times, especially created by an incremental backup strategy, can also be substantially shortened as long as the backup data is still in the TSM disk pool.

Since ETERNUS CS8000 emulates tape technology on disk storage, all advantages of using disk are present for tape access also. There is no urgent need to write to a separate disk buffer as a first "data landing platform". The virtual tapes provide the same performance and parallelism as native disks.

Because of this, TSM can directly backup to virtual tape and the expensive primary disks buffer can be reduced to a minimum.

4.2 Reducing competition for tape drive resources


Limiting the tape drive infrastructure can be the bottleneck in a TSM infrastructure. This impact is not caused by the performance of a single drive (modern tape drives can read and write with more than 300 MB/s natively), but is rather due to the number of available drives, the number of TSM "mount-points".

Various processes, like backup/restore, copy, migrate and reclaim, compete for the available tape drives.

When library sharing is used, other TSM servers or storage agents may also compete for the same drives. This competition subjects the planning of a TSM data backup infrastructure to great challenges and demands a high administrative outlay.

Up to 1280 virtual drives are offered by a single ETERNUS CS8000 providing sufficient tape drive resources to be made available for all potential situations. The number of available physical tape drives therefore does not represent any limitation for a TSM infrastructure. More Jobs can run in parallel.

Figure 3: Tape integration with ETERNUS CS8000



4.3 Effective use of TSM collocation on tape

To backup the TSM clients fast and space efficient, TSM uses the so called “progressive backup methodology” (= basically an “incremental forever”), that only backs up data changed or new since the last backup.

When doing this, the backup data of a single client will be saved together with data of other clients. Over time, the clients’ data will be dispersed over many tapes. In case a full recovery of a single client is needed, the data to be recovered has to be recalled from many tapes and a high number of time consuming mounts and searches on tape occur.

Besides the usage of active data pools (that are also perfectly suited to be stored on virtual tape instead of expensive disk), this behavior can also be minimized by a TSM function that collocates all backup data of single clients (or a group of client or a filespace) on separate tapes.

If a client needs to be fully recovered, TSM can retrieve the data from one or only a few tapes and thus only a single or low number of mounts must be executed.

To use the collocation function on tape, a dedicated tape is needed for each TSM client or client group processed.

ETERNUS CS8000 solves this conflict very effectively:

- Many virtual tape drives can satisfy many data streams in parallel when creating collocated tapes.
- Also slower backup streams do not decrease performance, because the data is written to the ETERNUS CS8000’s disk cache.

In real world collocation on tape also has a disadvantage:

It can waste storage capacity by leaving large parts of a tape empty because data of a single client is just not enough to fill a whole tape or the last tape of a set. ETERNUS CS8000 also solves this because a logical volume (tape media seen by TSM) only occupies as much data on all ETERNUS CS8000’s internal storage tiers (disk and physical tape) as it is actual in size.

Additionally, the collocated backup data can remain in the ETERNUS CS8000 cache until the next backup or even longer. For such a purpose, ETERNUS CS8000 offers policies such as ‘SAVE_DELAY’. The logical Volumes in cache are available for fast access at all times, like those which were stored in a TSM disk pool. Independent of the duration in the cache, ETERNUS CS8000 autonomously and transparently copies the data to physical tape media. Again, the policies ‘SAVE_DELAY’ is beneficial: By being set to e.g. ‘Max backup interval + 12 hours’ it may minimize the usage of physical tapes and dive mounts.

4.4 Increasing TSM LAN-free connectivity


In TSM environments the TSM clients send the backup data via an IP network to a TSM server. The TSM server then sends the data to the physical disk- or tape backup storage.

Since the availability of 10 GbE IP networks data transmission from client to server is not a real bottleneck anymore. Even though it still makes a lot of sense to let large and fast clients directly save their data without going through the backup server in the middle. The backup server still initiates and manages the backup process but the client uses a direct SAN path to the backup device and no backup server resources are used to transfer mass data.

If disk is used as the LAN-free backup target a shared filesystem must be introduced to be able to share the disk storage between the clients. Implementing a shared disk pool can be difficult to setup, this is why most LAN-free implementations are using (virtual) tape as the backup target.

ETERNUS CS8000 permits the user to easily back up the large data streams of a TSM client via a TSM storage agent. These clients have their own FC connection enabling them to back up to their own, dedicated and assigned virtual tape drives. The provision of dedicated ETERNUS CS8000 drives for each TSM storage agent allows to run more LAN-free tasks in parallel and avoids the common challenges of drive sharing.

Figure 4: Tape integration with ETERNUS CS8000 (LAN-free)



4.5 Faster mounts and instant positioning

When operating large real tape robots in conjunction with the usage of high capacity tapes, significant time is spent to mount, load, rewind, and unload the physical tape media. If the tape is not accessed from the beginning also positioning and searching the media is needed. Depending on the technology used this can add up to 2 or more minutes for each tape access and become a critical factor when processing urgent backup- or recovery actions.

As long as a logical volume resides in the ETERNUS CS8000's integrated disk cache, the system reduces these mount- and data access times to an absolute minimum (< 1 second). Eliminating 2-3 minutes for every tape access can speed up the overall backup or restore jobs dramatically.

If the data is no longer available in the cache, ETERNUS CS8000 can prefetch logical volumes beforehand to eliminate additional time for retrieving data from physical volumes. So, all additional volumes that belong to the backup are simultaneously read into cache in parallel and are then available for an instant access.

When TSM expires or frees a tape volume it is best practice to scratch the associated logical volume also in the VTL media inventory. Doing this enables the ETERNUS CS8000 to next time mount this logical volume right away in cache, also in case it was not cache resident anymore and the former version is still stored on physical tape. No recall from physical tape is necessary, since a new empty version of it is instantly created and mounted.

4.6 Deduplication for tape data

ETERNUS CS8000 supports data deduplication as another powerful function. Deduplication helps to have much more data stored on fast disk cache. The more data resides on the disk cache, the higher the hit rate of TSM for specific tape requests get. The higher the hit rate the faster mount requests are served, resulting in faster backups and restores.

However, deduplication is not a one-size-fits-all approach for all backup challenges. Some data (e.g. database log files) is just not suited for deduplication which not only may lead to low or moderate savings of disk space but especially to longer lasting deduplication - housekeeping times and a potentially increased time for restoring data.

In contrast to other "dedup-only" VTL solutions, ETERNUS CS8000 gives the choice to select whether data to be stored deduplicated or conventionally. Combining this selection choice with a very large dedup repository and the option to also copy the deduped data to tape, ETERNUS CS8000 makes dedup very efficient and much more powerful for TSM.

4.7 Compression on disk


In contrast to tape, standard disk systems usually are not able to compress data. By storing the data in tape format, ETERNUS CS8000 can use common tape compression algorithms when writing to the disk buffer. Typically 2-3 times the data can be stored on the disk cache compared to the use conventional TSM disk buffers.

4.8 Industry leading data availability

ETERNUS CS8000 can be installed across separate locations, and a TSM server is able to use virtual tape drives from both locations at any time. Even in case of a downtime in one location the TSM server is able to continue operation and meet all backup and restore requirements.


The synchronous cache mirroring option of the ETERNUS CS8000 and the feasibility to "dual- or triple copy" tapes in more than one library makes all data available in both locations without any restriction and without TSM writing the data twice.

Figure 5: Replication with ETERNUS CS8000 (synchronous)



It is also possible to cascade several ETERNUS CS systems. This means that an ETERNUS CS8000 system does not write its data through the backend into a physical tape library but into the frontend of a second ETERNUS CS8000 system. In other words, the virtual tape drives of a remote unit are used as the tape backend of the primary. The logical volumes written on the second unit are written in native TSM format so that a remote TSM server can access those easily. Since it is an asynchronous replication, the feature enables the data to be replicated over very long distances. Replication can be performed in both directions at the same time.

Figure 6: Replication with ETERNUS CS8000 (asynchronous)



The administrative outlay for both integrated DR functions is notably lower than with conventional methods, such as volume cloning by the TSM server, because the mirroring and replicating mechanisms are tightly anchored in the ETERNUS CS8000 and do not have to be additionally maintained.

4.9 Move physical tape workload to ETERNUS CS8000 and free TSM server resources


4.9.1 Reduction in the TSM Reclamation process

Data on tape regularly expires. TSM optimizes the utilization of the tape media by consolidating the backup data that is still active to new media. This action, also known as the reclamation process, can get very time consuming in large environments. It depends on the size of the tape volume and the number of available tape drives.

When using ETERNUS CS8000 with TSM, the need for this resource intensive process can be minimized. Through the use of significantly smaller volume sizes within ETERNUS CS8000 (e.g. 2GB), it is much more often that all data of a single virtual tape expires at almost the same time. Those tapes are just become scratch again and no reclamation is needed. Customers observe up to 90% reduction in TSM reclaiming workload.

Irrespective of this, ETERNUS CS8000 autonomously reorganizes the physical media in a continuous rule-based background process in order to eliminate expired virtual volumes or scratch volumes. This ensures maximum utilization of the physical media.


Figure 7: Reclamation with ETERNUS CS8000



4.9.2 Optimization of migration processes

More tape resources, like a large number of virtual drives provided by ETERNUS CS8000, also optimize, ease and accelerate TSM migration processes. In TSM the process of data movement between a hierarchy of storage pools, including storage pools on disk and tape, is called migration. This migration process is object based, optimization is possible through parallelization of multiple migration processes. It requires the existence of an appropriately large number of drive resources. Furthermore, the coordination must be intensively planned with other processes that are competing for the same tape resources (e.g. reclamation, creation of copy pools, backup and restore procedures, internal data base backups. When using an ETERNUS CS8000 Data Protection Appliance, TSM migration can be completely eliminated.

Figure 8: Migration with ETERNUS CS8000




4.9.3 Automated media refreshing and technology changes

Automatic media refreshing also takes place with ETERNUS CS8000, in which ETERNUS CS8000 can periodically copy inactive tapes to new tapes in order to ensure data availability at all times. ETERNUS CS8000 will also automatically replace tapes that exceed a maximum of permitted read- and write- errors. In contrast, TSM resets all tape error statistics when a tape gets back into scratch status and maintaining media is more difficult and work intensive.

Thanks to the use of ETERNUS CS8000 the physical drive technology used becomes invisible for the TSM server infrastructure. If additional drives or a technology refresh for drives or robotics are needed, ETERNUS CS8000 assures that the migration can be done on the fly, without impairing the running TSM production.

Figure 9: Media-refresh with ETERNUS CS8000



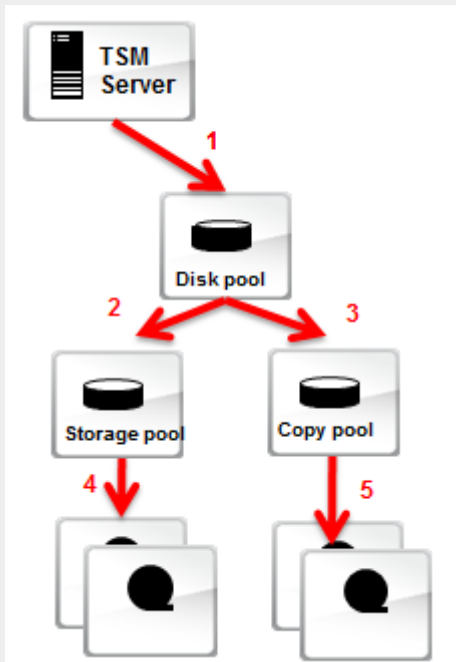
4.9.4 Relieving the load from the TSM server

Due to the offloading of IO-intensive TSM processes to the ETERNUS CS8000 the TSM servers are notably relieved of their load and can then devote themselves to their actual functions, the management and administration of data backup and recovery. Less investment in hardware resources to operate TSM instances are needed.

4.10 Omission of the copy pool and TSM offsite tape exports


TSM creates a 2nd or 3rd copy of data in the primary storage pool on one or several "copy pools", mostly comprising of tape media. One of the benefits of the "copy pools" is the synchronously or asynchronously possible creation of redundant copies of tapes to improve data availability, for tape export or for disaster recovery purposes.

Figure 10: Challenge "copy pools"



ETERNUS CS8000 permits the automatic creation of multiple copies of a backup volume without the necessity of a "copy pool" and without making demands of the TSM server. These copies can be located in a single or multiple libraries in different locations. Individual copies can be outsourced. The ETERNUS CS8000 copies thus fulfill the same purposes as the copies created by TSM, but without making demands of the resources of the TSM server. Creation of the copies is completely transparent for the TSM infrastructure. Furthermore, the data can be provided as highly available, i.e. across location boundaries. Like TSM, ETERNUS CS8000 is also able to restore data from onsite tapes, from exported tapes or, in the event of a disaster recovery, from offsite.

Figure 11: Automatic creation of copies with ETERNUS CS8000



As a side effect the elimination of the copy pool reduces also the size of the TSM database and therefore all TSM processes that querying the database get optimized indirectly.

4.11 Reduce the number physical tape drives

Modern tape drives, like LTO, T10000 or TS11x0, provide a very high throughput. They are able to read/write uncompressed data with 160-250 MB/s. If the drive can compress the data, the performance can even be multiplied. Faster technologies, which are able to operate with over 500 MB/s are expected soon.

In many cases the media servers are not able to generate data streams that are that fast. But if backup streams are not fast enough to utilize the drive's performance, more tape drives are needed in parallel to achieve the backup needs.

The dedicated ETERNUS CS8000 tape backend processors are fast enough to make full use of the high tape drive performance and therefore need considerably less drives to process the data. Actual ETERNUS CS8000 users were able to reduce their number of physical tape drives by up to 50%, which resulted in significant cost savings for them.

4.12 Prevent tape errors by decoupling the physical tape drives from the TSM server

While reading or writing to a physical tape drive critical media errors can occur. Sometimes the tape drive itself can fail also. A quite harmless consequence is, that the backup process is disrupted and needs to be restarted or fully repeated. In the event of no additional data copy is present, the way more badly case of losing data is also possible.

The virtual tape drives of the ETERNUS CS8000, that are seen by the TSM Server, cannot physically brake due to the full redundant architecture of the unit. If a physical tape drive or a physical media fails, ETERNUS CS can, transparently to the application, overcome this situation by automatically use a different drive or access an existing copy of the data. Through this decoupling TSM does not see any physical tape errors anymore, availability gets increased and service level can be met.

4.13 Save physical tape cartridges and library slots by ETERNUS CS8000's high utilization of tape media.

ETERNUS CS8000 stacks logical volumes very effectively to physical tapes. A logical volume only occupies the space on real media that it needs to. Physical tape capacity is not wasted, by always storing the whole logical volume if this was only partly used.

Also with a wide range of choices to control and steer physical tape reclaiming optimal processing and high tape media effectiveness is achieved. The better the media can be utilized the less media and library resources are needed. The effect will become even more significant as the capacity of single media increases in future.


ETERNUS CS8000 customers experience significant media savings already today

4.14 ETERNUS CS8000 backup and archive solution including tape

Besides providing the described VTL functions for backup and restore, ETERNUS CS8000 also offers NAS file services. Having both access types in place, the ETERNUS CS8000 is an optimal platform for consolidation. The file storage provides fast Ethernet connections (1 or 10 Gb/s,) and is accessed via CIFS or NFS protocol. Applications, e.g. databases, with integrated native backup- or archive capabilities can use ETERNUS CS8000 NAS files services easily and directly saving the fees for additional TSM storage agents.

Similar to the ETERNUS CS8000 VTL services, the data is stored on file systems in the disk cache (alternatively also in the mirrored cache with the HA solution). The data can automatically be saved to physical tapes. Various policies can be set to e.g. move data permanently to tape and create an online archive that appears infinite. For security reasons, it is also possible to immediately displace the data to tape, but still keep a version on disk for fast access

Figure 12: Unified backup and archive platform ETERNUS CS8000



5 Certified for IBM TSM

ETERNUS CS8000 is "Ready for Cloud and Smarter Infrastructure" and a TSM solution released by IBM.

More information is available on the IBM Tivoli pages:

<http://www.ibm.com/software/ismlibrary?NavCode=1TW10SM4G>.