

Virtualization of Tape Storage

Abstract: Realizing the benefits of Disk to Disk to Tape backup strategies incorporating disk-based virtual tape devices.

Table of Contents

Introduction_____	1
Benefits of D2D2T_____	1
Speed_____	1
Reliability_____	2
Benefits of Virtualization_____	3
Centralized, Shared Storage_____	3
Lower TCO_____	4
Scalability_____	5
Security_____	5
Maximum Throughput_____	5
Seamless Integration_____	6
Conclusions_____	7

Introduction

First introduced by Cybernetics in the early 1990s, the High Speed Tape Cache (HSTC™) disk-to-disk-to-tape backup solution has long been considered the ultimate data protection strategy. More recently, advances in disk performance and capacity at declining prices, coupled with the advent of iSCSI technology have made Cybernetics' newest generation HSTC™ not only affordable, but an utterly compelling choice for every IT site. The Cybernetics HSTC™ empowers IT to develop a strategy that maximizes the benefits of both disk and tape mediums. This newest HSTC™ solution overcomes every obstacle to disk-based backup: the cost of distributed disk capacity, the need for offsite storage and point in time restores, and the risk of catastrophic data loss in the event of disk drive failure.

Benefits of D2D2T

Backup to disk is the simplest, most efficient method for quickly saving and restoring data. Tape has the lowest cost per gigabyte, as well as the integrity for repeated loads and unloads, the impact resistance for transportability, and the shelf life required by most retention policies. Ideally, a backup solution employs both tape and disk devices, without doubling or tripling the cost of data protection.

Speed

Despite significant advances in tape technology for improving the speed of writing to tape, disk speed has historically outpaced tape speed. For sites with a large investment in systems and an infrastructure designed to achieve maximum performance, the transfer rate of a tape backup drive can be a major obstacle to completing backup within the allotted window. When staging the backup to disk, the only speed limitation is the rate at which the host can stream the data, resulting in a backup completed within the absolute minimum window.

For sites maintaining legacy systems, the host and network infrastructure can be a serious impediment to sustaining reasonable rates of transfer to tape. Streaming tape drives simply cannot perform with reasonable speed or reliability under these conditions. If data is not delivered to tape efficiently enough to keep the drive streaming, the tape drive must stop and reposition to avoid wasting space on the tape. This stopping and repositioning, also known as "back-hitching", causes transfer rates to plummet and puts undue wear and tear on the tape drive. The end result is an increase in service problems and a reduced life expectancy for the tape drive.

The HSTC™ can accept backup at any rate with absolutely no negative impact on the disk hardware. Once the virtual tapes are cached on disk, they are offloaded directly to tape at the tape drive's maximum streaming rate to achieve removable media backup within the smallest window with the least wear possible on the tape drive.

With the cost of downtime estimated at values ranging from thousands to millions of dollars per hour, the ability to recover quickly from a catastrophic data loss is absolutely crucial. A virtual tape can be stored on the HSTC™ disk for any length of time, regardless of whether or not it is archived to physical tape, to restore data at maximum speed directly from disk cache. Furthermore, users can maximize the benefits of the disaster recovery features in backup software packages such as ARCserve or Backup Exec, or the IPL feature of OS/400, to create bootable tapes stored on the HSTC™. In the event of a catastrophic drive failure, the bootable virtual tape can be used to recover to a brand new disk drive without the lengthy process of reloading the operating system and applications, for complete recovery in a fraction of the time it would normally take.

Reliability



With a low cost per GB and the security of removable media, direct-attached SCSI tape drives have long been the medium of choice in data protection. However, tape drives, tape libraries and tape media incorporate a relatively high number of moving parts, making a tape solution more prone to mechanical failure. A D2D2T backup strategy employs the strengths of disk storage and the power of redundancy to eliminate tape as a single point of failure. RAID disk solutions incorporate parity for fault tolerance - redundant disk drives, power supplies and fans can be hot-swapped with no down time. Backup and restore operations can continue to and from disk, even in the event of a tape drive failure. A tape drive can be removed, repaired, and re-installed without interruption to normal backups. With Cybernetics' HSTC™ Offload feature, offline tape archiving operations are seamlessly resumed to secure all backup copies to removable media for offsite storage and/or point in time restores. Any number of backup copies can be archived to removable media tapes with no impact on the host systems. All data that is offloaded to physical tape can be restored directly from tape using any tape drive - the drive attached to the HSTC™ or any other iSCSI network or direct-attached tape drive.

Cybernetics' iSAN™ Ethernet SAN technology deploys iSCSI disk and tape storage to take data protection to the next level, and completely eliminates any single point of backup failure. Users across the enterprise have direct access to network-attached tape drives and virtual tape drives and can easily re-route any backup job directly to any Cybernetics' iTape™ tape device or HSTC™ virtual tape device to take full advantage of the simplicity and flexibility of shared storage networking.

Benefits of Virtualization

Centralized, Shared Storage

The HSTC™ is an ideal solution for IT professionals™ confronted with the task of implementing and managing centralized, shared backup and archival capabilities. Through the use of onboard Ethernet connections and the iSCSI protocol, it is now possible to install an HSTC™ device at one location to serve as the backup target for multiple heterogeneous hosts spread across a building or multiple sites. Independent HSTC™ units installed at remote locations can even serve as off-site backup for each other.



The HSTC™ offers the flexibility to define multiple virtual tape drives and libraries within one central backup device, making them visible to hosts across the enterprise for direct access to dedicated, independent virtual tape systems. For example, it is possible to define multiple virtual tape libraries, each with its own virtual tape drive, and simultaneously backup to these virtual devices from multiple independent hosts. With the capability of Cybernetics' iClient™ to make practically any host iSCSI compliant, it is possible to share the HSTC™ among many heterogeneous hosts. The need to maintain multiple physical tape drives of varying formats is completely eradicated.

Lower TCO

Cybernetics HSTC™ implements a core strategy to avert the inordinately high potential costs of disaster, but it is designed with incredible efficiency to reduce the real, tangible costs of every day backup operations:

- The cost for add-on virtual tape drives is a mere fraction of the cost of adding physical tape drives, yet results in even greater benefits in throughput and reduced backup/restore time.
- The HSTC™ can produce tremendous savings in tape media costs. Having the security of virtual tapes stored in disk cache, IT managers can selectively choose which backup files to archive to tape. For example, a daily backup of 100-200GB can cost more than \$100 per day in tape media. If backup copies are archived to removable media just once a week, the resulting savings can be \$400 or more per week for just 100GB.
- The management of older, inefficient distributed backup strategies can be extremely labor-intensive. Trimming the workload for IT professionals by only one or two hours a day results in significant cost-savings over a relatively short period of time.
- The HSTC™ creates a central backup solution with one common media for all platforms across the network, eliminating the need to purchase, maintain and support an inventory of varying types of tape drives and tape media.
- Tape virtualization eliminates the need for costly new software products and there are no additional steps required to accomplish redundant backup. The HSTC™ uses any COTS (Commercial Off The Shelf) software designed for writing to SCSI tape drives. There is no learning curve for IT associated with implementing this device.
- Future investments in backup capacity no longer involve adding servers, backup software licenses, and support staff. The HSTC™ is expandable to a total of 36 RAID modules for hundreds of terabytes of capacity. Continuing advances in the cost-efficiency of RAID hardware sustain expansion with real long-term savings.
- Time and money-saving incremental backup strategies are far more feasible with Cybernetics' HSTC™. There is no concern about storing and retrieving all of the media for the intricate task of reconstructing data from incremental backup files. The HSTC has all incremental backup tapes stored on disk and ready for instant access in a restore, allowing IT managers to choose incremental backups throughout the week, postponing full backup and tape archives for weekly operations.
- The HSTC™ defers the cost of upgrading to higher capacity tape drives or libraries by extending the life expectancy of current tape devices which can now be dedicated to the sole task of archiving data at maximum streaming rates in an offline mode.

- Cybernetics' HSTC™ is based on iSCSI technology incorporating all the benefits of Fibre Channel Storage Area Network (SAN) in an Ethernet network environment. The expensive Fibre Channel infrastructure (switches, HBAs, cabling and special software) is no longer necessary to implement block level device sharing.

Scalability

As information technology advances, data storage capacity requirements grow at exponential rates. It is increasingly difficult to accurately predict capacity needs for the future. The HSTC™ can scale from gigabytes to terabytes, making planning for the future much easier and minimizing the risk of wasted investments in equipment. It eliminates the common dilemma over whether to petition management for spending approval on an expensive solution that may be overkill, or to settle for an affordable quick fix that is sure to be obsolete in the too near future. The HSTC™ is designed for expansion. Larger disk drives can easily be added to existing RAID modules, or additional RAID modules and disk drives can be attached as the storage capacity requirements grow. The HSTC™ is easily upgraded to support new virtual tapes, virtual drives and virtual and physical libraries. The connectivity features and configuration flexibility allow the HSTC™ to quickly and easily adapt to the changing requirements of the growing network or enterprise.

Security

Unlike a standard Network Attached Storage (NAS) disk array, the HSTC™ does not present itself as a disk to any host on the network. The virtualized device is identical to a tape drive or tape library, so there is no volume for a virus to attack. The HSTC™ enjoys a level of security unattainable with NAS, which is disk storage and therefore vulnerable to the common malicious threats to this medium. In addition, users do not access HSTC™ backup tapes as data files, and therefore backup files cannot be accidentally deleted. Data that is backed up to the HSTC™ is completely secure, both before and after it is archived to removable tape media.

Maximum Throughput

The HSTC™ is designed with maximum throughput in mind. Backup is completed at disk speed, with the only limitation being the speed at which the host can stream the data to the HSTC™. Furthermore, virtualization adds greatly to the efficiency of the backup device because virtual tapes do not have to rewind, virtual libraries do not have to pick tapes, and virtual drives do not have to load and unload tapes. These time-consuming operations necessary in physical tape drives and libraries are bypassed with the HSTC™. Unlike NAS disk staging, which carries significant file system overhead, the HSTC™ uses iSCSI technology to facilitate the same type of block-level transfers implemented in the more expensive Fibre Channel SAN technology. Gigabit Ethernet enables the HSTC™ to achieve SAN speeds using ubiquitous, affordable Ethernet infrastructures.

Once the network backup is completed to the HSTC™, data is archived to physical tape in an offline mode over a high-speed local LVD160 SCSI bus at block level with no overhead. Therefore, the data streams to physical tape at its maximum transfer rate for the fastest possible backup to removable media.

Creating large virtual libraries with many tape drives allows for independent and simultaneous backup of many servers. It is no longer necessary to run the network backup sequentially through servers of variable performance. Multiple independent backup streams can be created to truly utilize the available bandwidth.

Seamless Integration

Integrating the HSTC™ in an existing environment is a simple, straightforward task, which does not require any changes to the existing backup software packages. The onboard SCSI (LVD and/or HVD) and Gigabit Ethernet connections ensure physical connectivity to any host in the data center. iSCSI technology allows any server on the network to access the HSTC™ as a locally-attached SCSI tape drive. If a legacy host does not natively support iSCSI, a Cybernetics iClient™ can be connected to the host to instantly enable that host to send iSCSI data at Gigabit speeds.

Since the HSTC™ appears to the system as a tape device, IT is not required to make any adjustments to backup jobs and policies. The HSTC™ simply attaches to the network or to a SCSI bus, and existing backup jobs route data to the new target. HSTC™ is compatible with virtually any backup software and supports most industry standard tape formats allowing IT to maximize use of current hardware and software licenses. Hosts connect to the HSTC™, and then interact with the virtual tape systems on the device exactly as though directly-attached to a local SCSI card. The same backup software used for writing to direct-attached SCSI tape drives will continue to operate just as it has in the past. The software is not even aware that backup jobs are targeting a remote disk device - all operations continue just as though backing up to a local SCSI tape system.

Management of the HSTC™ is surprisingly simple. The HSTC™ operator connects to the device through a web browser to control all aspects of HSTC™ operations. From this management console, the operator manages physical and virtual tape systems and selects the strategy for offloading data to removable tape.

Conclusions

First released in the early 1990s, Cybernetics' HSTC™ has evolved into the ideal data protection solution that our customers have been seeking for years. Combining the new flexibility of our iSAN™ Ethernet connectivity with the powerful tape virtualization features of the HSTC™, this innovative new generation of the HSTC™ incorporates all the best features of disk and tape for data security with the incredible benefits of Ethernet SAN shared storage networking. The ability of the HSTC™ to adapt and scale with the growing network gives IT incredible new power over responsibilities that have become increasingly difficult to manage under the constraints of the direct-attached storage model. This newest HSTC™ melds performance, security and cost-benefits to make the previously illusive goal of centralizing and streamlining disaster recovery an affordable reality.

Copyright © 2004, Cybernetics. All rights reserved.

The information in this publication is subject to change at any time. This publication is intended for use by authorized recipients only and may not be recreated in any way without the consent of Cybernetics, in accordance with U.S. copyright law.