

CYBERNETICS

Virtualization Of Tape Storage

January 2008

Abstract

Disk-based virtual tape backup devices, when combined with removable tape for archiving, are a particularly effective data protection strategy, delivering the very tangible benefits of both disk and tape for securing data efficiently and cost-effectively at all points in its lifecycle. This paper discusses the advantages of a D2D2T strategy in reference to Cybernetics' miSAN V Series, iSAN V Series, and HSTC models.

Introduction

Cybernetics introduced the High Speed Tape Cache (HSTC) virtual tape appliance in the early 1990s, making it one of the first D2D2T solutions on the market. By creating numerous virtual tape drives and tapes out of an array of high performance, low cost disks, and by writing in simultaneous, parallel data streams, the HSTC allowed backup to disk – or virtual tape - in the least amount of time while it delivered near instantaneous restores. Virtual tapes remained cached until they were no longer needed. The HSTC's offload feature could be employed to backup selected virtual tapes to removable, archival tape - completely offline, and at maximum speeds. Supporting multiple host initiators with SCSI connectivity, the HSTC was a breakthrough in high performance data protection. Today, HSTC technology is the powerful software engine behind Cybernetics' latest virtual tape solutions: the miSAN V Series, the iSAN V Series, and the HSTC.

Benefits Of D2D2T

Backup to disk is the simplest, most efficient method of quickly saving and restoring data, while tape offers the lowest cost per gigabyte, the integrity for repeated loads and unloads, the impact resistance for transportability, and the shelf life required by most retention policies. But a backup solution that employs both disk and tape, without doubling or tripling the cost, has become standard for data protection because it delivers speed, capacity, and reliability with respect to the data's overall importance to the organization at various points in time.

Speed and Efficiency

Despite significant advances in tape technology for improving the speed of writing to tape, RAID disk speed has historically outpaced tape speed. Sites with a large investment in systems and an infrastructure designed to achieve maximum performance may find the transfer rate of a tape drive to be a major obstacle to completing backup within the allotted window. For sites maintaining legacy systems, the host and network infrastructure can also be a serious impediment to sustaining reasonable rates of transfer to tape. When data is not delivered fast enough to keep a tape drive streaming, the drive must stop and reposition to avoid wasting space on the tape. This stopping and repositioning, 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.

A D2D2T solution minimizes or eliminates such problems. When staging a backup to virtual tape, backup can be performed at RAID disk speed; the only limitation is the speed at which the host can stream the data. Furthermore, virtualization adds greatly to the efficiency of the overall backup operation because the data is laid down in a highly efficient sequential tape format without file structure overhead. In virtual backup there is no need to pick tapes, load and unload tapes, or rewind. The time-consuming operations necessary with physical tape drives and libraries are bypassed altogether. Unlike NAS disk staging,

which carries significant file system overhead, Cybernetics VTLs use iSCSI technology to facilitate the same type of block-level transfers implemented in the more expensive Fibre Channel SAN solutions, while 1000/10000 Gigabit Ethernet enables Cybernetics VTLs to achieve SAN speeds using ubiquitous, affordable Ethernet infrastructures.

It is also worth noting that 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 on the disk. 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.

Once the virtual tapes are cached on disk, they can be offloaded directly to removable tape. With the miSAN and iSAN V Series, archival backup is a completely serverless operation, so there is no drain on the network. Because this backup is written at the tape drive's maximum streaming rate, backup is accomplished within the smallest window and with the least possible wear on the tape drive.

Capacity

Data capacity grows exponentially, and budgetary and space requirements will necessarily limit the amount of disk that can be acquired and installed. A great backup solution incorporates scalability for increasing backup capacity. Cybernetics VTL solutions make the most of current disk capacity and provide for modular expansion to be added as necessary, while squeezing the highest capacity possible out of existing disk space. Advances in Cybernetics' VTL functionality, notably the introduction of hardware data compression to the miSAN and HSTC in May 2007, extends the virtual tape capacity significantly, making virtual tape backup a more cost-effective solution. Virtually unthinkable with tape cartridge sets, incremental backup technology is highly effective and reliable with virtual tape.

Hardware data compression, standard on the iSAN V Series and HSTC, and optional on the miSAN V Series, can increase the capacity of virtual backup to disk and to archival tape by a factor of two. This doubling of capacity has a profound and positive impact on the overall price-performance of the VTL in that it delivers the additional capacity in the same amount of rack space, while dramatically shrinking the backup window. And, clearly, the ability to keep twice the amount of data on virtual tape for instantaneous restore has positive disaster recovery implications. In reference to disaster recovery, it is important to note that archival tapes written with Cybernetics' hardware data compression are non-proprietary and universally compatible, and can be restored anywhere, with or without the VTL.

Virtual tape on disk revolutionizes incremental backup. Historically, full system backup has been the default method in tape-only environments simply because it is too cumbersome and too risky to try to recover from a tape set of the incremental or changed data backup. If just one tape fails, the entire restore will fail. But daily full system backups are wasteful of resources and taxing on physical tape drives and tapes. With Cybernetics' VTLs, incremental backups are completely risk-free, because all virtual tapes remain on disk where they are instantly accessible at random disk speed, with the solid reliability of RAID 5+ Hot Spare. Data compression and support for incremental backups work together to ensure that large quantities of the most important data are stored on disk, and instantly accessible, for the longest possible time. With the cost of downtime ranging from thousands to millions of dollars per hour, the ability to recover quickly from catastrophic data loss by restoring data directly from cache is absolutely essential.

Reliability


Tape drives, tape libraries and tape media incorporate a relatively high number of moving parts, making a tape-only 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, and redundant disk drives, power supplies, and fans can all be hot-swapped with no down time. Virtual picker arms do not drop or jam tapes and virtual tapes do not break or stick in the drive. The physical tape drive can be removed, repaired, and re-installed without interruption to virtual backups to disk. When the tape drive is re-installed, offline tape archiving operations can resume normally. The fact that archival backup to removable tape takes place offline at maximum speeds means less repositioning, and therefore minimal wear and tear on the tape drive or library as well as the tape media.

Benefits of Virtualization

Centralized, Shared Storage

A Cybernetics VTL 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 a virtual tape device at one location to serve as the backup target for multiple heterogeneous hosts spread across a building or at multiple sites. Independent Cybernetics VTLs installed at remote locations can even serve as off-site backup for each other.

Traditional Direct Attached Storage Network



HSTC™ Centralized Storage Network



Cybernetics VTLs offer 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 virtual library among many heterogeneous hosts. The need to maintain multiple physical tape drives of varying formats is completely eradicated.

Scalability

It is increasingly difficult to accurately predict capacity needs for the future, and system administrators must often choose between appealing to management for an expensive solution that might be overkill versus settling for an affordable quick fix that will be obsolete too soon. Cybernetics VTL solutions are expandable through additional SCSI or iSCSI RAID modules, for capacity into the hundreds of terabytes. 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 miSAN V Series, iSAN V Series, and HSTC models are easily upgraded to support new virtual tapes, virtual drives, and virtual libraries, so they can evolve with the changing requirements of the growing network or enterprise.

Security

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

Also, Cybernetics VTL models are available with integral 256-bit AES hardware data encryption for securing virtual tapes and/or physical tapes from unauthorized access. Cybernetics encryption key management is flexible, intuitive, and easy to implement.

Seamless Integration and Ease of Operation

Integrating a Cybernetics VTL into an existing environment is a simple, straightforward task, and does not require any changes to the existing backup software packages. The onboard SCSI and Ethernet connections ensure physical connectivity to any host in the data center. iSCSI technology allows any server on the network to access a Cybernetics VTL as a locally-attached SCSI tape drive. If a legacy host does not natively support iSCSI, SCSI ports support direct host connectivity, or a Cybernetics iClient™ can be connected to the host to enable that host to send iSCSI data at 1000/10000 Ethernet speeds.

Because Cybernetics VTLs appear to the system as tape devices, system administrators are not required to make any adjustments to backup jobs and policies. The VTL simply attaches to the network or to a SCSI bus, and existing backup jobs route data to the new target. Cybernetics VTLs are compatible with virtually any backup software and support most industry standard tape formats, allowing IT to maximize use of current hardware and software licenses. Network hosts connect to the VTL 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.

Management of a Cybernetics VTL is surprisingly simple. The operator connects to the device through a web browser to control all aspects of backup operations. From this management console, the operator can manage virtual and physical tape systems and select the strategy for offloading data to removable tape. The operator can view the status of backup jobs, and can suspend, reschedule, delete rescheduled

jobs, and create new jobs individually - without any impact on other operations. The HSTC software engine maintains a database of virtual and physical tapes, which include information including the date of last offload, next scheduled offload, as well as detailed information on tape usage.

Lower Total Cost of Ownership

Cybernetics' virtual tape libraries implement a core strategy to avert the inordinately high potential costs of disaster, but they are 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.
- Virtual tape libraries 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.
- Cybernetics virtual tape libraries create 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.
- Legacy tape drives can be SCSI attached to the Cybernetics VTL for purposes of easy and affordable legacy media conversion. A legacy tape such as DLT or SDLT can be read in to a virtual tape, and can then be archived in the new tape format such as LTO or AIT.
- Tape virtualization eliminates the need for costly new software products and there are no additional steps required to accomplish redundant backup. miSAN V Series, iSAN V Series, and HSTC solutions can utilize any commercial, off-the-shelf software package designed for writing to SCSI tape drives. Therefore, there is no learning curve for IT associated with implementation of the Cybernetics VTL.
- Future investments in backup capacity no longer involve adding servers, backup software licenses, and support staff. The iSAN V Series and HSTC are expandable to 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 very feasible with Cybernetics' VTLs. There is no concern about storing and retrieving all of the media for the intricate task of reconstructing data from incremental backup files because the VTL keeps all incremental changed data backup tapes stored on disk and ready for instant restore. This allows IT managers to perform incremental backups on a daily basis and to reduce full backup and tape archiving to a weekly operation, with no threat of lost data.

Conclusion

Cybernetics was at the forefront of D2D2T backup in the early 1990's when it introduced the HSTC virtual tape appliance. More recently, Cybernetics has been a pioneer in iSCSI technology. Now, HSTC technology drives the performance of the company's latest virtual tape solutions, the miSAN V Series, the iSAN V Series, and the HSTC. Combining the flexibility of Ethernet connectivity with powerful tape virtualization features, these innovative solutions incorporate all the best features of disk and tape into an intelligent data protection platform with all the benefits of shared storage networking. The ability of these Cybernetics tape libraries to adapt and scale with growing networks gives IT incredible new power over responsibilities that have become increasingly difficult to manage under the constraints of the direct-attached storage model. They combine performance, security and compelling cost-benefits to make the previously illusive goal of centralizing and streamlining backup and disaster recovery an affordable reality.