



Making a Business Case for Tape

by
Jon William Toigo
Independent Consultant and Author

SDLT
Super DLTtape™ Technology

Table of Contents

Introduction.....	3
Setting the Stage.....	5
The Business Case for Tape in the Distributed Computing Environments.....	8
<i>Risk Reduction</i>	8
<i>Cost-Savings</i>	11
<i>Business Enablement</i>	12
The Business Case for Tape in the Networked Storage World	13
Conclusion	15

INTRODUCTION

In the wake of the tragic events of September 11, 2001, interest in disaster preparedness and business continuity planning has risen to new heights. Organizations, both public and private, are re-examining their IT infrastructure to assess its resiliency to unplanned interruptions that can bring mission critical business processes to a sudden halt.

Experience drawn from numerous case studies has demonstrated the importance of recovery and continuity planning. In the wake of the World Trade Center bombing in 1993 and the disruption of operations caused by that incident for the 450 businesses occupying the WTC, 147 companies failed to recover. They simply ceased to operate within one year – in large part because they had no disaster recovery plan.

Current projections hold that the same percentage of firms will disappear in the 2002 timeframe as a direct result of the most recent incident. Simply put, those who prepare for the possibility of an interruption are more likely to survive than companies that take no proactive measures to avoid preventable disasters and to cope with the outcome of disasters that cannot be avoided.

Renewed attention to the importance of disaster recovery planning has thrust technologies such as magnetic tape back into the forefront of business thinking. Tape backup strategies are being seized upon once again as a cost-effective means to protect a organization's most irreplaceable asset: data.

However, history teaches that the perception of the need for disaster preparedness tends to diminish over time. Without additional "high profile" disasters or the emergence of compelling legal mandates to fuel continued vigilance, business managers tend to de-prioritize risk mitigation projects and to allocate scarce resources to other business initiatives. Thus, to assume that the business value of tape technology can be expressed solely in terms of its risk reduction value flies in the face of facts.

The Harvard Business School suggests that a comprehensive assessment of the business value of any technology acquisition requires the consideration of three value categories: cost-savings, risk reduction, and business enablement. In other words, to make a solid business case for the acquisition and deployment of any technology, including tape, IT managers need to demonstrate value to the company in as many of these categories as possible.

This paper will review tape technology from the standpoint of this business value model. It will provide useful background for the decision maker who is seeking to build a value-based business case for tape technology – Quantum Corporation's Super DLTape, in particular – for presentation to financial and business managers within the organization.

The challenge of such an endeavor is to advance the business value case in the face of contrarian views advanced both by business competitors within the tape technology

market and by competitors outside the tape space, in the realm of optical disc and magnetic disk technology. This paper will not address every argument currently being advanced against tape. Rather, general themes will be covered where relevant to provide the decision maker with a better basis for evaluating options.

The good news is that the business value case for tape, as the following pages suggest, is a robust one. Far from being what one uninformed advocate of storage area networks (SANs) has recently termed “an aging technology from a previous generation of computing,” contemporary tape technology features capacities and performance that is keeping pace with disk technology, while sustaining a 4-to-1 price advantage over disk. Ironically, given the view of the aforementioned SAN vendor, sharing a tape library remains among the top five reasons that companies provide to justify their deployment of current generation SANs.

More than anything, the development of a business case for tape should be founded on common sense. Pundit prognostications and marketing hype have little relevance to the exercise, except as they influence perceptions and need to be addressed and dismissed.

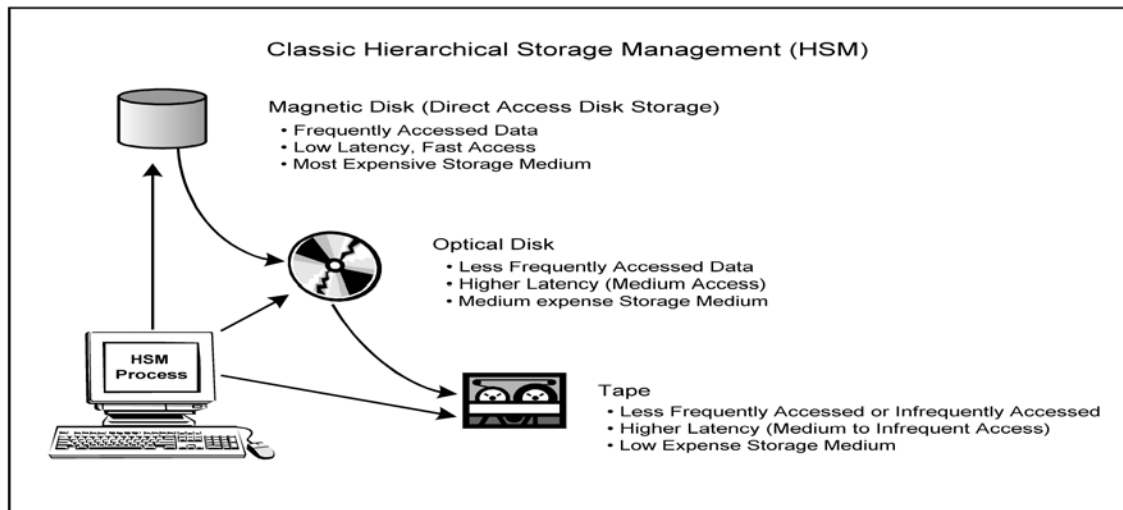
SETTING THE STAGE

Tape technology has a lengthy pedigree in the data storage marketplace. Tape was among the original media used as “primary storage” for early computers. Over time, rotational media – magnetic disk – took over this role and tape was relegated to the role of “secondary storage” (backup) and “tertiary storage” (data archive).

The reasons for the rise of disk are well understood. Rotational media (disk) offered faster access than streaming linear media (tape) to stored data. As CPU speeds increased, and cost per MB for disk-based storage decreased, magnetic disk became the preferred “primary storage” medium. However, the use of magnetic disk for primary storage has not diminished the utility of tape. In fact, quite the opposite has occurred.

Over time, tape technology has proven to be a robust technology that can adapt to the changing architecture of modern business computing. From its earliest days as a primary storage medium, tape evolved to become a secondary storage platform for production data within a hierarchical managed storage paradigm (HSM) advanced by IBM and other mainframe vendors throughout the 1960s and most of the 1970s.

In an HSM role, tape was used to store less frequently accessed production data that had been migrated from more expensive direct access storage devices (DASD). This role for tape may return into vogue as storage itself moves into a networked architecture (see below).



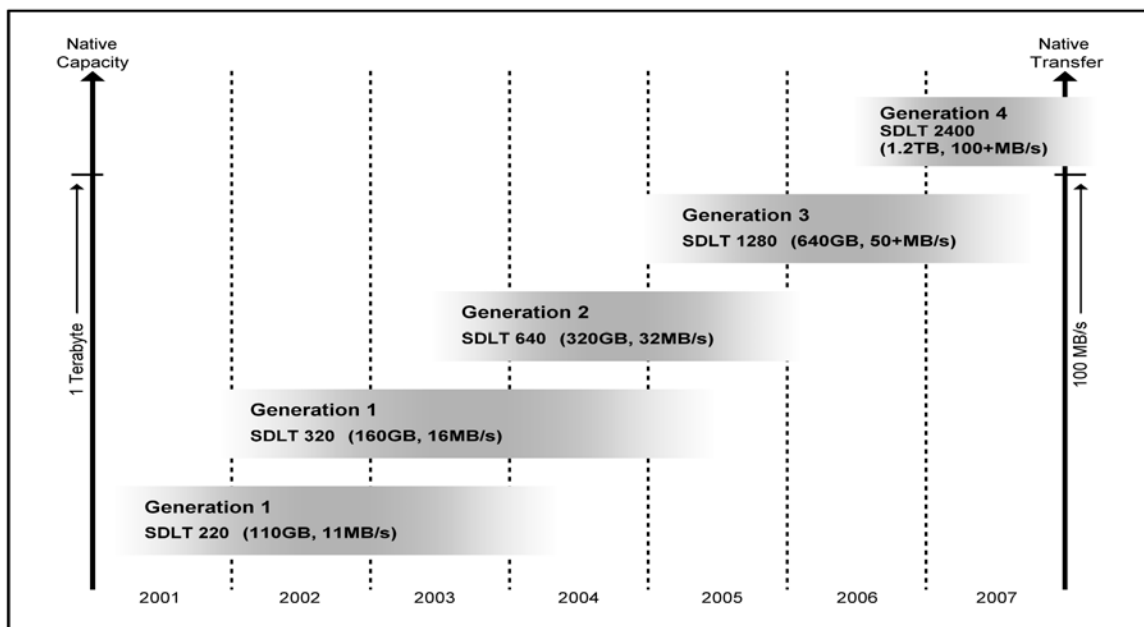
With the rise of distributed “open” systems and networked client-server applications in the 1980s, tape technology accompanied servers out of the glass house of the corporate data center and into the workgroup environment. It was in the open systems world that

tape came into full flower as a flexible medium suited to the performance, capacity, and resiliency requirements of both server-based and PC-based computing.

Revenues for tape in the open systems market have soared, while revenues for older tape technologies aimed at mainframe data center markets have flattened substantially. Improvements in tape technologies aimed at this open systems segment have mapped closely to the improvements in server disk drive and storage interconnect technologies.

As evidenced in Quantum's Super DLT product, tape has mirrored many of the same technologies implemented by disk drive manufacturers, including advances in read/write head technologies and media coatings that enable more data to be written and read more reliably to a fixed amount of recording surface.

- Super DLTape implements PRML technology, first implemented on high-end server disk drives, to improve the read/write accuracy and efficiency of tape drives, while supporting ever-increasing data densities.
- Super DLTape's native capacity, which can be effectively doubled using compression algorithms, is growing from 160 GB per tape (320 GB compressed) in current generation drives to more than a Terabyte in Generation 4 drives anticipated by 2007. In doing so, Super DLTape plans to double cartridge capacity every 18 months, thus keeping pace with end users' exploding storage requirements. It is worth mentioning that Quantum is alone among tape manufacturers to provide such an aggressive capacity improvement roadmap.



- Super DLTtape technology also implements advanced interfaces such as Ultra2 SCSI, LVD and HVD, first employed on disk drives. The short term roadmap for the technology goes further to implement Ultra160 and Ultra320 SCSI and Fibre Channel to facilitate the high speed transfer of data from tape to disk and vice versa. Current transfer rates of uncompressed data stand at 16 Megabytes per second (MB/s), with Super DLTtape roadmaps projecting 32 MB/s in Generation 2 drives and 100+ MB/s in Generation 4 SuperDLT products.

In short, tape technology—led in the open systems space by Quantum Corporation's DLT and Super DLTtape products—has been keeping pace with advances in magnetic disk technology for nearly two decades. Whether as part of a hierarchical storage management scheme within the carefully controlled environment of a corporate data center, or as a highly available, highly flexible data backup medium in the distributed computing world, tape has proven to be a versatile and dependable medium for data storage.

Moreover, tape will continue to play a role as current generation storage topologies continue to evolve and mature toward a network-based architecture. Before turning to the future role of tape in networked storage topologies, it is worthwhile to consider the business case for tape in current generation distributed computing environments.

THE BUSINESS CASE FOR TAPE IN THE DISTRIBUTED COMPUTING ENVIRONMENT

The business case for any technology derives from three characteristics: risk reduction, cost-savings and business enablement. The business case for tape technology in the current context of distributed computing is a strong one, and is outlined below.

Risk Reduction

Risk reduction means several things, but we primarily think of the risk reduction value of tape in terms of the protection of data assets. Tape is, of course, ideally suited to the role of data protection and for reasons that are straightforward and readily understood.

With most IT assets, IT planners have the option of protecting the asset through either a strategy of replacement or one of redundancy. A server, for example, can be replaced with another similarly equipped device following a disaster, or an identical server can be pre-staged at a remote recovery facility. The difference is one of time and expense.

With data, which is an irreplaceable asset, redundancy is the only effective option for protection. From the standpoint of data, redundancy translates to data copy. A copy of the data must be made and moved from the environment where the original data is stored to keep it safe from disaster events that might consume the original. Since data is changing on a constant basis, copy-and-remove is, by necessity, an on-going process.

Tape provides part of its risk reduction value by providing an effective means for storing data copies and a convenient and resilient package for transporting the data copy to a remote recovery facility out of harm's way. This is well understood and accounts for why local tape—tape drives, autoloaders or libraries installed in the distributed computing setting—has become the primary mechanism for data protection in the client-server world today. Less well perceived is why tape is viewed as superior to other strategies for copy-and-remove processes, such as centralized data vaulting or remote disk-to-disk mirroring.

Part of the reason has to do with the limited bandwidth available in the Local Area Networks (LANs) of most businesses. The bandwidth of the LANs in most organizations is carefully matched to the demands of client-server application messaging. Conducting large-scale data movements, such as those associated with tape backup operations, across the LAN would impose too great a burden on these nets. Even in networks with adequate bandwidth, cross-LAN disk copies can introduce unacceptable latency into the operation of distributed applications.

This, coincidentally, is a chief reason why centralized backup/recovery architectures (and LAN-based hierarchical storage management) have not fared as well in the distributed computing world as local tape. And it also explains why tape, with its

capability to be deployed locally as well as centrally, has come into great use as the copy-and-remove strategy of choice.

Tape has also tended to be preferred to remote disk mirroring in distributed computing environments as a copy-and-remove strategy. While remote disk-to-disk mirroring has the apparent advantage over tape of eliminating the time and effort required to recover data to a useable form following a disaster, the cost to provision a Wide Area Network for remote disk is prohibitive.

The issue is sometimes raised by advocates of remote disk mirroring that their strategy is superior to tape from the standpoint of “time-to-data” —that is, the amount of downtime a company will need to endure while tapes are transported to the remote recovery facility, prepared for use, and their contents are transferred to disk platforms at the facility. Critics claim that burgeoning data growth and the advent of Terabyte-sized databases invalidate tape as a recovery medium. The facts, however, stand in stark contrast to this claim.


- Many tape critics use the data transfer rate of a single tape drive to calculate the time required to move data from tape to disk. In other words, they leave consumers with the impression that tape-to-disk data transfers are limited to the transfer speed of an individual drive. This is fallacious and ignores the capabilities of medium to high end tape libraries to operate multiple drives concurrently and the ability of medium to high end disk storage arrays to receive and process multiple tape streams concurrently. These capabilities are a multiplier of transfer rates and can effectively reduce the time required to re-host even large datasets to recovery platforms.

Number of Drives	SDLT 320 (Native)	
	Capacity (GB)	Performance (GB/hr)
1	160	59.25
2	320	118.5
3	480	177.75
4	640	237
5	800	296.25
6	960	355.5
7	1120	414.75
8	1280	474
9	1440	533.25
10	1600	592.5
11	1760	651.75
12	1920	711
13	2080	770.25
14	2240	829.5
15	2400	888.75
16	2560	948
17	2720	1007.25
18	2880	1066.5
19	3040	1125.75
20	3200	1185

- Disaster recovery planners typically conduct a thorough process of business function classification to identify which systems and data are used to support mission critical business operations. Only this data needs to be recovered in an extremely short timeframe—typically a subset of the overall data storage maintained by a company.
- The modern database, according to industry experts, commonly contains a mixture of static and changing datasets. Static datasets, which may comprise upwards of 80 percent of an active database, are accessed infrequently and are rarely subject to change. Thus, the volume of data that can be pre-staged to a recovery center and pre-installed to storage platforms while DR personnel are in route may by itself render this argument moot. DR personnel may only need to transfer a subset of data to the storage platform in order to get up and running.

While remote disk mirroring is entirely possible, its solution cost generally is much greater than tape and its application is generally limited to data copies between homogeneous arrays and only in the case of the most time sensitive applications. The costs of remote mirroring typically include the cost to operate network facilities between remote sites and the cost for three, rather than two, disk arrays.

Three arrays are required so that applications are not impacted by the latency created when writing data across a geographical distance. Typically, two arrays are configured in a local, highly synchronous mirror, while a separate process is used to move data copies between array number two and remote array number 3. This configuration usually produces a remote mirror that is asymmetrical or asynchronous. In other words the data on the remote array is usually “older” than the data on the production array or its local mirror: how much older is a function of the bandwidth of the network interconnect.



If arrays are mirrored asymmetrically and the time difference between them exceeds a few minutes, it is likely that tape will still play a role in the solution. DR personnel will need to use tape to copy the data that is “out-of-sync” between the arrays, and must carry those backups to the recovery site to bring datasets to a current state.

Thus, tape continues to play a key role in risk reduction, interpreted as data protection, in distributed computing. Risk reduction, however, also has other meanings. In

addition to the risk of downtime and disaster, another risk confronting businesses is that of technology obsolescence. IT managers seek to make technology investments that will yield service for a predetermined period of time—4 to 6 years in the case of most tape technology acquisitions—and that will provide a smooth, evolutionary path forward after that time.

Super DLTtape from Quantum, like its DLTtape predecessors, provides just such investment protection. The Super DLTtape 220 drive is fully backward read compatible to previous DLTtape generations including DLTtape 8000, 7000, 4000 and DLT1. With respect to interchangeability the Super DLTtape 220 is fully interchangeable with Super DLTtape and DLTtape IV media. Thus, Quantum's DLTtape and Super DLTtape products provide full investment protection and enable users to restore business critical data recorded several product generations back.

A collateral advantage of tape from an investment risk reduction perspective is that tape is an enabler of other storage acquisitions. For example, migrating from a storage array offered by vendor A to one offered by vendor B need not be hampered by a lack of utility software for transitioning data from the older array to the newer one. Tape provides data portability and enables the expeditious re-hosting of data to new platforms as platform types change. This is also a strength of tape from a cost-efficiency perspective as discussed in the following section.

Cost-Savings

From a cost-savings perspective, tape technology has a strong value argument to make. Tape remains the least expensive medium for storing backup copies of data, despite some claims about the decreasing cost for disk-to-disk data copying. (Comparisons of disk-to-tape and disk-to-disk are fallacious: the technologies are increasingly complimentary and each, by itself, may offer a specific fit for a given application.) The cost-savings that can be derived from tape in operation are several:

- Tape can be used to store archival data and data that is rarely accessed by the organization, providing an excellent hedge against the need to expand expensive disk resources rapidly in response to data growth.
- Tape provides significant cost savings over other copy-and-remove strategies, such as the remote mirroring configuration described in the preceding section. In the case of most applications, tape is entirely suitable – and, in some cases, absolutely required – as a medium for data recovery. As previously stated, the cost to mirror data is a function of the number of storage platforms involved (three are typically required for a synchronous/asynchronous deployment) and the cost for a wide area network communications facility of sufficient bandwidth to minimize latency. The cost for such a facility currently exceeds \$1 Million per month, precluding the use of such a strategy in all but the most time-sensitive applications.

- Cost savings may also result from the use of tape to migrate data between un-like platforms during storage platform migrations. While most array vendors provide utility software for moving data between two platforms from their own product family, moving to a new vendor's platform usually constrains the value of this software. Tape can facilitate the transfer, saving cost in terms of effort, and providing the capability to select whatever disk platform is considered best suited to applications regardless of brand.
- Certain tape technologies, including Super DLTtape from Quantum, offer performance and capacity enhancing features, including PRML technology, MR read-write heads, compression algorithms, etc., that can dramatically reduce the costly overhead associated with other tape technologies, while maximizing media utilization and reducing tape media costs.
- Super DLTtape also provides assured compatibility of media between drives saving the costs of testing and verification of tape readability between drives of the same type and between current technology drives and media recorded on prior generation drives. This is also an advantage of Super DLTtape technology over its nearest rivals, such as LTO, where less tightly defined drive specifications have the potential to limit the interchangeability of media between drives from different vendors.
- Super DLTtape supports a broad range of operating systems and server platforms, enabling the standardization of tape technology throughout a heterogeneous infrastructure and saving the cost for supporting multiple tape technologies at the recovery center.

Business Enablement

The business enablement value of a technology is defined as the capability afforded by the technology to improve the efficiency of existing business pursuits, or to pursue new business objectives—either as a direct result of the technology or as an indirect result. With respect to contemporary IT infrastructure, the acquisition of a best-of-breed tape technology does not directly enable new business or improvements in the efficiency of existing business pursuits. Its benefits are more indirect:

- Increased business opportunity: Deploying best-of-breed tape technology to support disaster preparedness may provide the indirect benefit of aiding the business in qualifying for certain lines of business such as government contracting, supply provisioning, or outsourcing. In these arrangements, increased attention is paid to the preparedness of service partners or suppliers as a prerequisite for inking deals. The proper application of tape technology may, therefore, indirectly contribute to the capability of the business to obtain new business contracts.

- Enhanced business responsiveness to market change: A best-of-breed tape technology can support, in ways already described, smooth platform transitions within the storage infrastructure of the company. Indirectly, this may facilitate the company's responsiveness to market shifts by enabling greater flexibility in overall IT support for business processes and initiatives.
- Greater cost control and efficiency in business processes: A best-of-breed tape technology can facilitate the migration of older data from disk media, forestalling the need to add expensive disk-based storage to the IT infrastructure and controlling costs. Indirectly, this may lower the cost of infrastructure allocated to specific business processes, thereby increasing cost efficiency.

The preceding are only a few of the indirect business enablement benefits that may be derived from tape.

THE BUSINESS CASE FOR TAPE IN A NETWORKED STORAGE WORLD

Of course, the world of business computing is not standing still. Today, many companies are seeking to re-centralize servers and to consolidate and centralize key storage platforms back inside the "glass house" of the corporate data center. This does not reflect the abandonment of client-server architecture as much as a recognition that, in many cases, a degree of management discipline was lost when systems and storage were originally decentralized. Re-centralization, it is believed, will enable more IT assets to be managed by existing IT staff resources, containing labor costs.

Against this backdrop of change, storage technology is also changing. Recognizing the downtime associated with traditional server-attached storage, the storage industry and IT management are working to develop and deploy networked storage solutions. Early implementations of this concept are storage area networks (SANs) and network-attached storage appliances (NAS).

Conceptually, a storage area network establishes a "back-end network" to which servers and storage platforms are connected. As stated earlier, a key motivator of early adopters of SAN technology has been to share a high-end tape automation platform or library among numerous open system servers and storage platforms. By delivering managed tape services to a broad range of platforms, IT planners hope to derive

- Greater cost-savings – tape libraries will be more fully utilized and will have a longer useful life in the company;

- Enhanced risk reduction – centralized and managed tape backup will assure that backups are performed and will simplify the creation of organized tape sets for delivery to offsite storage or directly to recovery centers; and
- New business enablement – storage, including tape storage, will be able to be delivered as a service subject to service level agreements and charged back to business units. This, in turn, will enable the more precise modeling of enabling IT infrastructure costs as part of the total cost of doing business and provide decision-makers with greater insights into operational costs on a line of business basis.

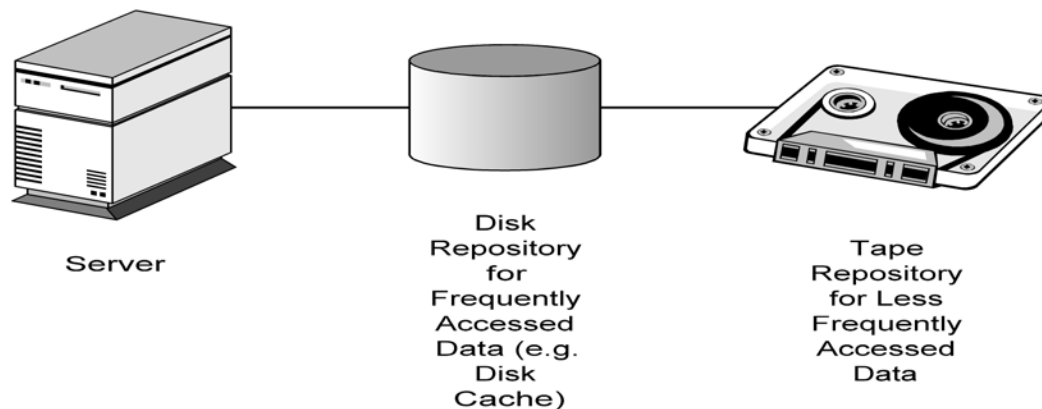
Back-end SANs also provide a medium for LAN-free and, potentially, server-free data backups and other large scale data movements. This unleashes an entirely new opportunity for Hierarchical Storage Management (HSM) in the open systems space, and potentially an even stronger business value case for tape beyond traditional roles of backup and archive.

In HSM strategies once employed within the data center, older and less frequently accessed data was migrated to less expensive tape media, where it provided near on-line availability. This was critical given the expense and bulk of disk based storage, or DASD.

Today, with some organizations confronting data growth rates of 70 percent or more per year, the need for better disk capacity management is again an issue. While storage platforms themselves have gotten smaller physically, the paucity of effective storage management software tools – particularly in heterogeneous storage platform environments—typically mandates that more IT staffers be hired as new storage acquisitions are made. Put another way, labor costs are increasing concurrently with storage capacity.

The availability of a high bandwidth, back-end network for storage enables the use of tape in an HSM context once again. Resilient and dependable tape technologies such as Super DLTape may well see service as “near on-line” repositories for some data currently stored to disk, providing the IT organization with the ability to forestall some disk storage platform acquisitions and to make the best use of the platforms they already have.

Near On-line Configuration



Interesting combinations of tape and disk are already on the drawing boards of many major storage vendors in which a disk array serves as a “cache” for a resilient tape platform. In some implementations, this is done to enable a tape cartridge to be virtualized – that is, to enable an image to be created of all the data to be written to a cartridge so that media is completely utilized. Once the virtual image is complete, the tape is written. This approach attacks the cost and volume of tape media that are used by the company.

An alternative approach is to use tape as a repository for data files supporting a Web site. A disk array front end contains the information most often requested by Web based users, while the tape library contains less frequently accessed documents. This Web server design can be extremely cost effective in applications such as on-line product literature libraries or streaming media sites where file requests are not extremely time-sensitive. Provided an extremely robust and high performance tape technology, like Super DLTape, is used to provide the back-end Web repository, the performance of the solution should be consistent with current Web response times.

These are examples of risk-reducing, cost-saving and business-enabling applications of tape in a networked storage world. Others will likely emerge over time as storage networking technology matures.

CONCLUSION

Rather than provide hypothetical quantifications of business value, this paper has suggested practical examples of business value benefits that may derive from the acquisition of a best-of-breed tape solution like Super DLTape from Quantum. Supported by a growing number of case studies taken from the shops where more than two million DLT and Super DLTape drives and more than 80 million DLT and Super DLTape media cartridges are in use, the business value case of Quantum tape technology can be made persuasively.

Quantum tape technology continues to offer organizations a complete business value proposition based on risk reduction, cost-savings and business enablement. With its new SuperDLTtape products, and enhanced backup solutions leveraging the latest tape automation techniques and disk-based caching topologies, the business case for Quantum's tape products continues to be compelling.

For companies requiring a more quantitative assessment of business value, Quantum would be happy to support this requirement with knowledgeable and qualified account representatives.

For more information, visit www.SuperDLTtape.com.