Data Mobility Group

datamobilitygroup.com

White Paper
April 11, 2005
Dianne McAdam

Backup and Archive Storage Choices

 76 NorthEastern Blvd. Suite 29A Nashua, NH 03062 Phone: 877.463.6456 Fax: 603.563.8282

Calculating the Cost of Tape and Disk Solutions

The importance of backing up digital data records on a regular basis has long been understood. While IT administrators have used tape as the primary backup target for years, today administrators are augmenting the continued use of tape by implementing a mix of technologies, including replication, snapshots and RAID schemes, to protect their organizations' data both locally and remotely. Nonetheless, tape still remains the *de facto* standard for backup. This paper will discuss the cost implications (acquisition and environmental) for an all disk, all tape and combined disk and tape backup/archive solution over a seven-year period.

As data protection technologies evolve, backup rules are changing, and IT administrators are updating their data protection strategies to meet today's new requirements. For example, a crop of new federal and state regulations governing many industries now dictate extended retention periods for a variety of data types, including the various kinds of user created documents, email messages, voicemail messages, and digitized images.

The result is a more complicated data-protection environment—one that requires IT administrators to think out-of-the-traditional-data-protection-box. While many IT administrators mistakenly use the terms backup and archival interchangeably, the two actually describe very different processes, which can be implemented separately or in combination to ensure users have access to data when they need it and in the most time-efficient manner, without sending overall storage costs further skyward. The continuing decline in the cost of disk has some IT organizations considering disk for long term retention. However, tape has always cost less than disk and will continue to cost less as higher capacity tape drives are brought to market.

The Backup Process

The purpose of backup has not changed: to make a copy the contents of a volume, logical unit number (LUN), or file at a particular point-in-time for future use in the event the original volume, LUN, or file is corrupted or deleted. The backup copy, which can be stored on disk, tape, or optical media, is used to restore data records to the point in time of the last-known good backup.

Common Backup Facts

1. Backups are designed to protect all digital data records within an organization, including data records within the data center, on desktops, and on laptops. Backups can occur at the volume-, LUN-, or file-level; restores are done at the same level.

2. Backups are generally run on a regular basis (typically daily for most applications). Heavily modified applications and mission-critical applications may require several backups during a 24-hour period.

3. Multiple copies of backups are saved. For example, a common practice is to do a full backup once a week and then incremental backups for the rest of the week. All incremental backups are saved until the next full backup is taken. Backups of certain data record types (e.g., financial data at the end of a current fiscal quarter) are generally kept for longer periods of time (e.g., after the fiscal year has ended or longer).

4. Generally speaking, backup "data sets" are not actively deleted but rather overwritten. The backup process is typically configured in such a way that the second set of incremental backups overwrites the first set of incremental backups once the second full backup is completed. However, the number of backup sets that must be kept before any prior backup data set is overwritten can vary. If backup policies specify that four full backup sets must be maintained, then the fifth full backup overwrites the first.

5. Multiple different backup processes and software platforms can—and often do—exist within organizations. Corporations may use one backup product to protect data within the glass walls of

their data centers and a different product to back up data at remote offices or on laptops in the field. These backups can be managed centrally, from remote locations, or both centrally and remotely.

6. Two commonly used metrics to define data protection requirements are recovery time objectives (RTO) and recovery point objectives (RPO). Backups have different RTOs, or the time required to restore data. A business critical application, such as a customer service database may have an RTO of one hour or less. The RTO for the human resources database which is not essential to drive revenue may have an RTO of 8 hours.

7. Backups also have different RPOs. The RPO refers to the age of the data that is used for restores. A database application, for example, may dictate that it be restored from data that is less than six hours old from the point at which an outage or corruption event occurred. This means that the database has to be backed up at least every six hours.

The Archive Process

Simply put: an archive contains data that is no longer in use but may be required in the future. Backups consist of both active and inactive data records; archives, by comparison, contain *only* inactive data records.

Unlike backups, the purpose of an archive is *not* to restore data in the event the data is accidentally deleted or corrupted. An archive is often used to help IT administrators weed out infrequently used (or never accessed) data records (*active archives*) from databases or active-content data libraries, An archive may also be used to weed out infrequently accessed data records that are never updated, or data that needs to be kept for a long period of time (*deep archives*).

Active Archive

Databases expand over time. Consider a database consisting of entries of all customers that have bought goods and services from a particular company. Over time, some of these customers become "inactive"—that is, they stop purchasing goods and services from the company.

Keeping inactive customer data records (that is, data on customers who have not purchased goods or services, say, for five years or more) as part of an active database application not only expands host storage requirements but can also significantly degrade the performance of the application, which continues to serve active customers.

Active archiving software (such as Princeton Softech's Active Archive) allows database administrators to extract the inactive data (and associated table definitions, indexes, and relationships) to a lower-cost storage device. If the data is required in the future, the application or the administrator can access and restore it.

Again, the purpose of active-archiving is *not* to protect data but to weed out infrequently used or historical data and move it to lower-cost disk or tape storage.

Common Active Archive Facts

1. Active archiving is designed to improve the management and performance of large, expanding database applications.

2. Active archiving is not a form of data protection; it is not a replacement for backup. It is a process to move inactive or historical data records to lower-cost media. Database records still need to be backed up regularly, to protect against hardware or software malfunctions or database corruptions.

3. Unlike backup applications, active archiving does not need to be done on a daily basis, but it *should* be run on a regular basis. How often the archiving process

should be done is determined by how fast the database is growing, how quickly the data becomes inactive, and how tight the IT storage budget is. Weekly, monthly, or quarterly runs are usually sufficient.

4. Inactive data can be stored on disk, tape, or optical storage. Inactive data can be viewed from its current location or restored to the production database if required.

5. By reducing the amount of disk storage that needs to be assigned to the database, active archiving can help IT administrators save money.

6. Active-archiving products must understand the underlying structure of the database. These products not only move inactive data to secondary devices, but they also move the metadata (i.e., the tables and indexes) associated with this data. This metadata is the keeper of information about the relationship with the unmoved data.

7. The smaller the primary database, the quicker it is to back up and more importantly restore. Thus, weeding out the inactive data records improves the backup and restore performance of the data protection application used to backup the primary database.

8. Active archiving trims down databases to reasonable sizes.

Deep Archive

For years, corporations have set internal policies for how long data should be kept. For example, aerospace companies typically have policies to retain airplane designs long after the actual airplanes themselves are built, often extending until the last plane of a particular design is retired from active service.

In many cases, government regulations, such as HIPAA and SEC 17a-4, no longer leave the data retention period up to the individual organization but rather dictate the period of time. For companies that do not meet the specified retention periods, large fines will be given. Deep archival software is designed to help IT administrators store regulated (and non-regulated) data for the specified time period.

Deep archive data is *not* a backup copy or a disaster recovery copy. It is not data that is updated frequently, such as bank account transaction records. Rather, it is data that is kept for a long period of time, it is not accessed frequently, and it is never updated. This type of data is also increasingly referred to as fixed-content, unstructured, or reference data records.

Common Deep Archive Facts

1. Deep archiving is designed to store data for a specific period of time. The data may be actively deleted when the retention period is up. Not all data needs to be retained in a deep archive.

2. Archives can be stored on many different types of media (e.g., disk, tape, and optical).

3. Backups store *multiple* copies of your data; only *one* copy of any specific data record is generally maintained in an archive (see #4). The point is that you should not have to search through all the archives to determine whether or not you have the final version of the data record you are looking for.

4. Some regulations (e.g., SEC 17a-4) require organizations to keep a copy of the archived data off-site. In these instances, IT administrators will need to keep *two* copies of the data (the original and the off-site copy).

5. Some regulations (e.g., SEC 17a-4) specifically require that the archive software prevent overwriting, erasing, or altering of any data during the required retention period.

6. Backups and archives have different recovery time objectives. For example, the customer service database may have an RTO of one hour

for the backup; however, the RTO for archived data records may be much longer. If a corporation is involved in litigation and is required to produce supporting documentation as part of a legal discovery process, the RTO may be as little as one to five business days.

7. Backups and archives also have different recovery point objectives. The database application, for example, that requires it to be restored from data that is less than six hours old, has an RPO of 6 hours. The RPO for its historical, unchanging archive data records, on the other hand, is zero; there is only one final copy of archived data.

Deep Archive Software

Deep archive software (such as CommVault's QiNetix Data Archiver or KVS' Enterprise Vault) can help streamline what can be a very time-consuming and wasteful search process for data related to, for example, a particular financial transaction or for personnel records required for litigation support.

Assume that standard backup software is used to back up a Microsoft Exchange server nightly and auditors request all e-mails pertaining to a specific financial transaction. That could mean searching through weeks', if not months', worth of backup copies, depending on the history of the transaction (i.e., when the transaction actually took place and any relevant correspondence that may have preceded or followed the transaction).

Archive software greatly simplifies this task by creating indexes of stored data. E-mail archive software, for example, creates indexes of all e-mail messages, which allows users to search by "tag" fields (e.g., To, From, Subject, Date) or by text strings found within the body of the e-mail message.

Archive software also contains policies that specify how long the data must be kept, if it must be write-once read-many- (WORM) protected (considered a "best practice" and required by SEC 17a-4), and if it must be replicated to a remote location as an extra protection against loss or security breach.

Archiving Data—What's the Cost?

Archived data is retained for years, not weeks or months like traditional backup data. Therefore, organizations need to factor in the cost of storing data records over the life of the record retention period, not just the initial software/hardware costs of implementing the data archive.

Hardware Requirements

- How much data must be archived? What is the projected growth rate? How long will data be stored? Will the data be backed up to onsite/offsite tape or disk? Should extra capacity be allocated for unanticipated growth?
- Will data be stored on disk (generally Serial ATA variants), specialized appliances, or a combination of disk and tape? What is the projected average disk and tape cartridge utilization?
- Is WORM support required?
- What level of RAID protection is necessary for disk-based systems?

Purchase and Maintenance Costs

- What is the initial cost of the equipment?
- What is the warranty period and what are the maintenance costs after the warranty expires?
- What is the lifespan of the equipment? Will disk systems be replaced every three years? tape drives every five years? Will new equipment need to be purchased before the data expires and is deleted?

Software Requirements

- Is the archive software compatible with the proposed equipment? Are any other software products or options required? If so, what are the costs of these products?

Environmental Information

- What is the square footprint of the device(s)? How much does data center space cost (per square foot) in your area?
- What are the power consumption and cooling requirements for the equipment? What are current electrical costs (kilowatts per hour)?

Calculating the Costs—Tape vs. Disk

The tables below compare the costs of archiving data on tape and disk. For the sake of accuracy in comparing environmental cost factors, the cost of a Quantum tape library was compared to the cost of a readily available Serial ATA disk system.

Step1: Calculate Storage Requirements

The first step depends on the amount of data being stored and the projected growth rate. Tape and disk are never allocated at 100 percent. A tape cartridge utilization of 85 percent and disk utilization of 70 percent was assumed.

Assumptions:

- Initial Storage Requirement—50 TB
- Annual Growth Rate—10%
- Average Tape Cartridge Utilization—85%
- Average Tape Compression—2:1
- Average Disk Utilization—70%
- None of the data will expire within the first seven years.

Disk storage requirements are higher than tape storage requirements since disk utilization rates are lower than tape utilization rates.

Step 2: Determine Tape Hardware Configurations

For this particular example, two Quantum tape configurations were created. The first is a Quantum PX720 tape library with SDLT 600 tape drives. The tape library is equipped with four SDLT 600 tape drives and 190 SDLT II media tape cartridges for a total native capacity of approximately 57TB (or 114 TB of capacity at 2:1 compression).

Table 1. Calculating Storage Requirements for Disk and Tape

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7
Initial Data Stored	50 TB	55 TB	60.5 TB	66.6 TB	73.3 TB	80.6 TB	88.7 TB
New Data Added	5 TB	5.5 TB	6.1 TB	6.7 TB	7.3 TB	8.1 TB	8.9 TB
Total Data Stored	55 TB	60.5 TB	66.6 TB	73.3 TB	80.6 TB	88.7 TB	97.6 TB
Disk Storage Requirements	78.6 TB	86.4 TB	95.1 TB	104.7 TB	115.1 TB	126.7 TB	139.4 TB
Tape Storage Requirements	64.7 TB	71.2 TB	78.4 TB	86.2 TB	94.8 TB	104.4 TB	114.8 TB

Table 2A: Purchase Price for PX720/SDLT Tape Solution

Hardware	Cost
1 PX720 Library[1]	$104,000
4 SDLT 600 drives ($12,495 per drive)	$61,000
190 cartridges ($130 per cartridge)	$28,237
	Total Cost: $193,237[2] *Note: List pricing is used for all hardware configurations.*

Figure 1. Quantum PX720 tape library

The Quantum PX720 tape library also supports LTO tape technology and an equivalent LTO configuration is outlined below. Here, the PX720 tape library is equipped with four LTO-2 tape drives and 285 LTO-2 tape cartridges for a total native capacity of approximately 57 TB (or about 114 TB of compressed capacity). One LTO-2 tape cartridge can store up to 400 GB of compressed data which is two-thirds the capacity of single SDLT cartridge. Since the LTO-2 cartridge is smaller in capacity, an additional 95 cartridges (and 95 library slots) is needed to store the same amount of data. Above, the PX720 library was upgraded to the intermediate model to store the additional cartridges.

Table 2B: Purchase Price for PX720/LTO Tape Solution

Hardware	Cost
1 PX720 Library[3]	$75,140
4 LTO-2 drives ($12,495 per drive)	$49,980
285 cartridges ($75 per cartridge)	$21,375
	Total Cost: $146,495[4] *Note: List pricing is used for all hardware configurations.*

Step 3: Calculate Tape Environmental Costs

The cost of floor space varies widely across geographic regions, as do electricity costs. For this example, the floor space cost was estimated to be $20 per square foot (sq ft)/ per month and the electricity cost to be $0.10 per kilowatt/ hour (kwh).[5]

Also, since 10 disk systems with their expansion units require four racks, for the purpose of these calculations, the disks are stacked on a 39.4” by 23.8” rack.

Assumptions:

Computer room cost per square foot—$20 sq ft
Cost of electricity—$0.10 kwh

[1] MSRP for PX720 base frame with entry slots, 1 drive cluster and a communication module. This also includes Web-based remote management and monitoring software, redundant power supplies and fans installation and a one year 5x9xnext day on-site warranty.
[2] MSRP pricing is used for all hardware configurations. All pricing is applicable as of this paper’s publication date.
[3] MSRP for PX720 base frame with entry slots, 1 drive cluster and a communication module. This also includes Web-based remote management and monitoring software, redundant power supplies and fans installation and a one year 5x9xnext day on-site warranty.
[4] MSRP pricing is used for all hardware configurations. All pricing is applicable as of this paper’s publication date.
[5] The United States Government surveys current electrical costs for each state. This report is available at *http://www.eia.doe.gov/cneaf/electricity/epm/table5*.

Calculations:

Electricity is required to both power and cool the hardware. To calculate the yearly electrical cost, add the power and cooling requirements for each device.

The yearly electrical cost=[[Power requirements (kwh) + cooling requirements (kwh] x the number of units] x 8760 (number of hours in a year) x $0.10 (the cost per kwh)

Total environmental cost=Cost of footprint x 12 (number of months in a year) + yearly electrical costs

Table 3A: Environmental Costs for PX720/SDLT Tape Solution

	Footprint	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
PX720 Library	10.4 sq ft	$2,496	1.6 kwh	1.6 kwh	$2,803	$5,299
4 SDLT[6] Drive	0	$0	.018 kwh	.018 kwh	$126	$126
					Total Yearly Electrical Cost: $2,929	**Total Yearly Equipment Cost:** $5,425

Table 3B: Environmental Costs for PX720/LTO Tape Solution

	Footprint	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
PX720 Library	10.4 sq ft	$2,496	1.6 kwh	1.6 kwh	$2,803	$5,299
4 LTO-2[7] Drive	0	$0	.017 kwh	.017 kwh	$119	$119
					Total Yearly Electrical Cost: $2,922	**Total Yearly Equipment Cost:** $5,418

[6] The SDLT 600 SCSI drive uses 30 watts when writing data, and 14 watts when idle. Eighteen watts was used as an average which reflects writing 25 percent of the time and idle the rest of the time.

[7] The LTO-2 SCSI drive uses 29 watts when writing data and 13 watts when idle. Seventeen watts was used as an average which reflects writing 25% of the time and idle the rest of the time.

Table 4. Purchase Price for Disk Solution

Hardware	Unit Cost	Total Cost
10 Disk Systems	$22,385	$223,850
30 Expansion Units	$22,185	$665,550
Rack[9]	$4,500	$18,000
		Total Cost: $907,400

Step 4: Determine Disk Hardware Configurations

For the disk configuration, a commonly available SATA disk system was chosen that consists of 10 controllers, each equipped with three expansion units. Each disk system (equipped with three expansion units) has a capacity of 14TB, making the overall total capacity of the configuration 140TB.[8]

Step 5: Calculate Disk Environmental Costs

Step 6: Calculate the Seven-Year Total

The seven-year total cost for the tape and disk solutions is calculated. The first year costs include the purchase price and the environmental costs for that year.

Table 5. Environmental Costs for an All-disk Solution

	Footprint	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
Disk Systems	26 sq ft	$6,251	.33 kwh	.39 kwh	$6,307	$12,558
Expansion Units	0	$0	.33 kwh	.39 kwh	$18,921	$18,921
					Total Yearly Electrical Cost: $25,228	**Total Yearly Equipment Cost:** $31,479

Table 6. Seven-Year Cost for Disk and Tape Solution

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Total
SDLT Tape	$142,245	$5,425	$5,425	$5,425	$5,425	$5,425	$5,425	$174,795
LTO Tape	$151,913	$5,418	$5,418	$5,418	$5,418	$5,418	$5,418	$184,421
All Disk	$938,879	$31,479	$31,479	$31,479	$31,479	$31,479	$31,479	$1,127,753

[8] The IBM DS4100 SATA disk system with the DS4000 EXP100 expansion unit was used in this example to provide accurate pricing and environmental information.

[9] The dimensions for the IBM 7014 rack were used. Four racks are required to support this configuration

The PX720 library when configured with SDLT or LTO drives costs only 16 percent of the equivalent disk configuration over a seven year period!

Future Disk and Tape Pricing

There has been a great amount of press coverage regarding the erosion in the cost of disk over the last few years fueling speculation that disk will soon cost less than tape.

Surveys that track end user pricing note that the price of disk subsystems has been declining at about 30 percent annually. However, those same surveys reveal that the cost of tape subsystems has also declined at a similar rate.

Fred Moore, president of Horison Information Strategies, publishes an annual paper that reports on current and future trends in the storage industry.
According to his latest report (titled "Storage Navigator"):

- Enterprise disk (FC, FICON, SCSI) costs about $40-$70/GB
- Midrange disk (SCSI, FC) costs about $15-$35/GB
- Economy disk (SATA, JBOD) costs about $3-$15/GB
- Automated tape costs about $.50-$3.00/GB
- Disk will increase from its current 300-400 GB capacities to over 4-5 TB by 2013
- Tape cartridges will increase from its current 300+ GB capacities to over 6-8 TB by 2013.

The predicted increases in tape cartridge capacities will allow tape to maintain its cost advantages over disk.

Some Caveats

The cost calculations above represent a simple view of the total cost for a seven-year solution; however, the purpose of the exercise is to highlight the effects of environmental costs on long-term retention solutions.

A more comprehensive calculation would also include:

- The cost of maintenance after the warranty period expired
- The cost to replace disk and tape after several years (if that was the typical buying pattern of the IT organization)
- The cost of migrating to new disk or tape technology. How will data be migrated from an old device to a new device? Will the migration be completed in-house or will professional services be required?
- The cost of the software. It was assumed that the same software would be used for all solution examples.
- The cost of phasing in equipment purchases. It is not necessary to purchase or lease all equipment the first year.
- The amount of data that would expire every year.
- The cost associated with backing up data to local or remote tape or disk.
- The additional cost for mirroring disk (RAID-1).

Which Media to Choose?

The above calculations show that the all tape solution costs only *16 percent* that of the all disk solution, factoring in the initial purchase price and environmental costs. Therefore, if low cost is an organization's highest priority, tape is the logical choice. For data that needs be archived for long periods of time but does not require very quick retrieval access times, tape is the most economical choice. If quick access is more important than the overall cost of the solution, disk may be the better choice. However, not all data requires the same speedy access.

Another increasingly popular option is to implement a hybrid disk/tape solution. In a backup environment, disk-to-disk backup systems allow disk to initially be the target for the backup image. Older copies of the backup can later be migrated to tape for longer term, less expensive storage.

Table 7: Disk and Tape Storage Capacities

	Raw	Useable
Disk Capacity	28 TB	19.6 TB at 70% utilization
Tape Capacity	92 TB	78.4 TB at 85% utilization
Total Capacity	120 TB	98 TB

Table 8: Cost of DX100

Hardware	Requirements	Cost
DX100	28TB	$294,550

Tiered storage solutions can also be implemented when archiving data. Data is kept on disk for a short period of time, and then migrated to lower cost tape for the duration of its retention period. To calculate the costs for this type of configuration, the Quantum DX100 Disk-based Backup System was used as the target for 20 percent of the organization's data (critical, quick access data). The DX100 is a SATA disk-based backup device that emulates a tape library and easily integrates into a customer's existing environment. The rest of the data would be kept for longer periods of time on the PX720 tape library.

Calculate the DX100 Costs

Since the DX100 emulates tape drives, customers have reported disk utilization greater than 70 percent. However, the disk utilization was kept at a conservative 70 percent.

The base DX100 consists of 2 arrays (8 TB) of storage. Five arrays (4 TB) were added to configure a 28 TB system that can be mounted in one standard 24" by 36" rack.

Table 9: Environmental Costs for DX100 Solution

	Footprint (8 Racks)	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
DX100	6 sq ft	$1,440	2.9 kwh	6.3 kwh	$ 8059	$ 9499

Figure 2.
Quantum DX100

Table 10A. Cost of PX720/SDLT Tape Library

Hardware	Cost
1 PX720 Library	$ 62,140
2 SDLT drives	$ 24,990
153 cartridges	$ 19,890
	Total Cost: $107,020

Table 10B. Cost of PX720/LTOTape Library

Hardware	Cost
1 PX720 Tape Library	$ 75,140
2 LTO-2 drives	$ 24,990
230 cartridges	$ 17,250
	Total Cost: $117,380

Calculate the Tape Costs

With the hybrid configuration shown on the previous page, 20 percent of the data records are stored on disk and older records are migrated to tape. This disk/tape solution requires fewer tape drives and cartridges than the all tape solution. Additional slots are available in either the PX720/SDLT 600 or PX720/LTO-2 configuration to accommodate future growth.

Table 11A: Environmental Costs for PX720/SDLT Tape Solution

	Footprint	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
PX720 Library	10.4 sq ft	$2,496	1.6 kwh	1.6 kwh	$2,803	$5,299
2 SDLT 600 drives	0	$0	.018 kwh	.018 kwh	$63	$63
					Total Yearly Electrical Cost: $2,866	**Total Yearly Equipment Cost:** $5,362

Table 11B: Environmental Costs for PX720/LTO Tape Solution

	Footprint	Footprint Cost Per Year	Power Requirements Per Unit	Cooling Requirements Per Unit	Electrical Costs Per Year	Total Environmental Cost
PX720 Library	10.4 sq ft	$2,496	1.6 kwh	1.6 kwh	$2,803	$5,299
2 LTO drives	0	$0	.017 kwh	.017 kwh	$59	$59
					Total Yearly Electrical Cost: $2,862	**Total Yearly Equipment Cost:** $5,358

Calculate the Seven-Year Total

The combination of DX100 and PX720 library (with either LTO-2 or SDLT drives) is only 45 percent the cost of the all disk solution in Table 6 yet provides even greater capacity.

In this example, the entry level PX720 tape library with SDLT drives can store an additional 22 TB without upgrading the library to the intermediate version. A simple software upgrade to the intermediate level provides storage for another 115 TB of data without increasing the footprint of the library. In the LTO-2 example, the PX720 is already configured at the intermediate model since LTO-2 cartridges are smaller in capacity than the SDLT cartridges. This PX720/LTO-2 configuration can store an additional 80 TB without requiring any changes.

Again, the same caveats explained above apply to these calculations. This example is designed to give users a general idea about the long-term costs of backup or archiving data using three different configurations.

The Bottom Line

Which configuration is the right one for your organization? It really boils down to how many different data record classifications your organization requires, what your access requirements are for each classification (i.e., how quickly you want to be able to retrieve data records in searches, etc.), the frequency of access required, and the projected amount of data you will need to retrieve annually; all balanced with how much you are willing to spend.

Tape backups are the logical choice for corporations that require the lowest cost option. Tape is also the most economic choice for data records that need to be archived for long periods of time but does not require very quick retrieval access times. An all tape solution initially costs 84 percent less than an all disk solution, and continues to cost less over a seven year period. Future improvement in tape capacities that match or exceed improvements in disk capacities allows tape to maintain its cost advantages. Another advantage is the backward read and write compatibility of tape. IT administrators need to consider tape automation solutions that can satisfy existing backup requirements as well as new archiving requirements, without sending storage costs through the roof.

Quantum's PX720 tape library has a flexible architecture allowing customers to support numerous connectivity interfaces and drive types within one library frame. For example, customers can choose LTO technology to support their backup applications and higher capacity SDLT technology for their archiving needs today. When

Table 12. 7-Year Cost for Hybrid Disk and Tape Solution

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Total
SDLT Tape	$112,382	$5,362	$5,362	$5,362	$5,362	$5,362	$5,362	$144,554
LTO Tape	$122,738	$5,358	$5,358	$5,358	$5,358	$5,358	$5,358	$154,886
DX100	$304,049	$9,499	$9,499	$9,499	$9,499	$9,499	$9,499	$361,043
Total for SDLT and DX100	$416,431	$14,861	$14,861	$14,861	$14,861	$14,861	$14,861	$505,597
Total for LTO and DX100	$426,787	$14,857	$14,857	$14,857	$14,857	$14,857	$14,857	$515,929

new drives and connectivity options become available, customers can upgrade to the new technology without purchasing a new library, therefore, extending the life of the library. New LTO-3 technology doubles the capacity of LTO cartridges, without doubling the cost. Customers can convert from LTO-2 to LTO-3 technology and double the capacity of the PX720 without requiring any changes or extra library frames. Future generations of high capacity SDLT drives and high performance LTO drives allow customers to create a tiered tape storage solution, with high capacity tape drives for archiving applications and high performance tape drives for backup applications, within the same tape library. Accommodating both tape technologies in order to optimize the users infrastructure is enabled by the PX720's mixed media feature. A feature of the PX720, library partitioning, allows customers to consolidate backup and archive applications supporting multiple, heterogeneous platforms within one library—simplifying management.

Corporations governed by compliance regulations are required to retain data for longer and longer periods of time. However, long term data retention is not restricted to regulated companies. Small to medium-sized private companies are also being counseled by their legal advisors to retain data for extended periods of time. The data growth rate of 10 percent used in this paper is a very conservative number. Many customers are experiencing growth rates greater than 50 percent per year and expect that trend to continue.

The scalable architecture of the PX720 library can accommodate this rapid growth. If a customer exceeds the total capacity of a single frame, multiple PX720 libraries (up to five) can be connected with the CrossLink Mechanism. This equates to 965 TB of native SDLT 600 storage, 725 TB of native LTO-2 storage and over one petabyte native LTO-3 storage capacity.
For mission-critical applications that require quick restore/retrieval times, disk is the better choice. Not all applications are mission-critical and have speedy retrieval access requirements. IT administrators should evaluate integrating disk with their current tape-based backup infrastructure to provide quick restores for those applications that demand it. Since Quantum's DX100 emulates tape drives, storage administrators can implement disk-based backups without requiring changes to their existing backup processes.

The best option for corporations to implement is a tiered storage architecture, leveraging a hybrid disk/tape solution. The DX100/PX720 with SDLT 600 tape drives hybrid solution costs 50 percent less than the seven year all disk solution, yet provides more capacity. Using the above example, this solution could store more than ten years of data while costing half as much. Quantum's DX100, when coupled with the PX720 offers the best of both worlds: fast retrieval from disk for mission-critical applications and cost-effective storage on tape for all other applications, and long-term archive.