IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Aligning Data Protection Strategies with Business Value

Sponsored by: StorageTek

Richard L. Villars Robert Amatruda

October 2004

Executive Summary

The loss of valuable enterprise data can be catastrophic to an organization. An investment in data protection for mission-critical data sets is justified because a data outage would immediately affect business value metrics such as revenue and profitability.

Not all data is equally valuable, however, and not all data protection architectures and methods are equally robust. The important tasks for IT planners are to understand the strengths, weaknesses, and costs of data protection solutions and to match those solutions to different classes and types of enterprise data.

While it is prudent to invest in technology that will keep critical systems available in the face of many different threats, it is equally prudent to use simpler, less expensive techniques for data that is less critical to the enterprise. The keys are to leverage a common process for determining which data sets should be protected with what method and to have a common set of tools to manage all of the solutions in a coordinated way.

IDC identifies three basic data protection architectures: disk to disk, disk to tape, and disk to disk to tape. Each architecture addresses different data protection risks. Mirroring data to other disks, for example, provides rapid recovery should one of the disk systems suffer an outage. Other techniques, such as storing point-in-time replicas of data, address other threats, such as data corruption and viruses.

IDC recommends a three-step process to address data protection issues. First, prepare an inventory of enterprise data sets with an eye to the current replication techniques and the relationship of these data sets to business value metrics. IDC suggests using the Storage Network Industry Association's (SNIA's) five-step data classification system. Second, prioritize the data sets and map them to the three data protection architectures. Third, prepare a data protection plan that invests in technology in proportion to business value risk.

StorageTek provides a suite of data protection products that supports the different architectures and methods described herein. StorageTek products are available at different price points, capacities, and performance levels so that proportional investments can be made for different classes of enterprise data.

IDC believes that enterprises have an opportunity to improve their business continuity plans with a thoughtful review of today's data protection options. Measuring investment in data protection against business value is a prudent strategy.

Tying Data Protection to Business Value

Investments in data protection should be guided by their contributions to business value, that is, to the fundamental measures of business success for the enterprise. Three heuristics are important to keep in mind:

- **Understand data diversity.** In today's data-rich business environments, information varies on many dimensions. The life cycle for data such as enterprise financial transactions may be measured in years, while the life cycle for pricing data in a financial market may be measured in minutes or hours. During that life cycle, the need to retrieve data also varies. When financial transactions are aggregated, individual entries are rarely needed again. Pricing data, on the other hand, may be accessed thousands of times during its short life cycle. The need for privacy reveals another dimension. Some data sets are intended for public display, such as financial market prices, whereas other enterprise data must be protected from unwanted access.

- **Investigate business continuity.** In evaluating different data sets, consider what the business impact would be if the data sets became unavailable. In a retail business, for example, loss of point-of-sale data would halt all transactions and the company would effectively be out of business until the data is restored. Loss of inventory data might not have as immediate an impact with shelves becoming increasingly bare in a few days' time. Loss of accounts receivable data could disrupt cash flow if the lost data is not restored within a week.

- **Estimate the cost of unavailable data.** Understanding the differential costs associated with the loss of different sets of data is the key to data protection planning. And costs vary significantly for the same kind of data at different enterprises. In a multinational professional services firm, for example, the cost of losing access to email can be catastrophic as teams are paralyzed and unable to do their work. In a manufacturing company or a retail company, however, loss of email data may have a much smaller negative impact.

The primary goal for data protection is to mitigate overall risk to the enterprise by ensuring that the most valuable data is well protected from loss and that data sets that are critical to the business can be restored quickly and accurately. Less valuable data should also be protected, but only in proportion to its value to the enterprise.

DATA PROTECTION

The fundamental mechanism for data protection is replication. Replicas of data sets are prepared regularly and stored where threats to one copy of the data are unlikely to threaten other copies. Replicating data to a second network attached storage array protects the data from the threat of an array failure, for example, but fails to protect the data from a datacenter power outage or a network fault. Replicating data to a second datacenter provides greater assurance that sitewide outages will not disrupt business continuity, as shown in Figure 1. Replicating data to tape and storing that tape in a secure archive also improves data protection.

FIGURE 1

Isolating Data Replicas

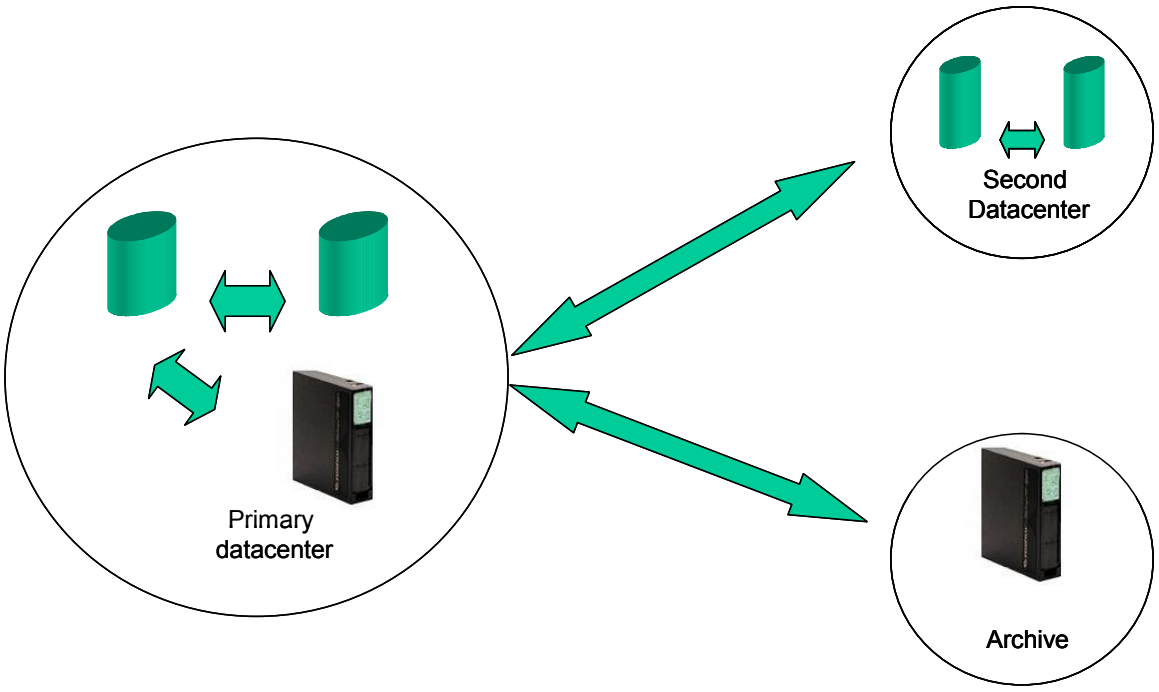

Source: IDC, 2004

The storage architecture determines the speed with which replicas can be put into service when the primary data becomes unavailable. Disk-to-disk replication supports the fastest recovery times, but has historically been very expensive Disk-to-tape solutions provide greater economy and flexibility because tapes are removable media. Hybrid disk-to-disk-to-tape solutions provide ways to optimize access and flexibility as well as to relieve primary storage systems from the task of writing data to tape. The advent of lower-cost, high-capacity disk systems based on technologies such as Serial ATA are improving the cost metrics for both disk-to-disk options, but tape remains a critical and cost-effective part of solutions for long-term archiving.

Weighing the plusses and minuses of each solution, the data protection designer chooses among these architectures. This includes decisions about how often recovery points are needed, how long the enterprise can afford to wait for recovery, and how long replicas of data must be retained.

Data Protection Architectures and Methods

Three data protection architectures, shown in Figure 2, provide the foundation for data protection planning. These architectures vary in cost, offer different benefits, and mitigate different risks. All are likely to find a place in a comprehensive enterprise data protection plan.

FIGURE 2

Data Protection Architectures

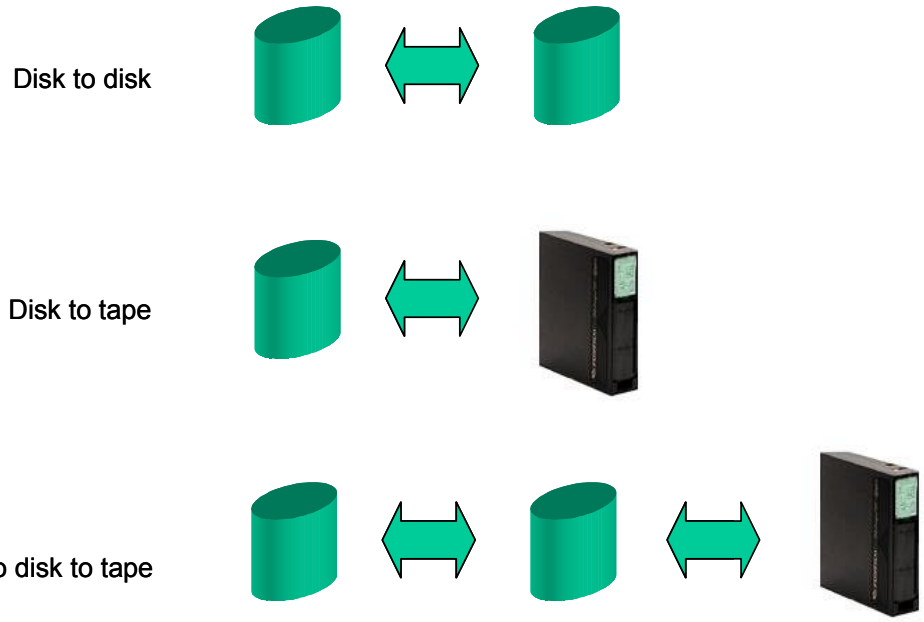

Source: IDC, 2004

Disk-to-Disk Replication

Disk-to-disk replication uses a second disk system to store a replica of data stored on the primary storage system. Information stored on a second disk system mitigates the risk of a primary storage system failure. When the second disk system is at another datacenter, then the risk of sitewide losses and outages is reduced.

Two techniques are commonly used in a disk-to-disk replication architecture: mirroring and point-in-time replicas.

- Mirroring, which is the process of writing data to two disks simultaneously, is an efficient way to maintain two copies of a data set. Should a disk fail, then its mirror is immediately ready to support the application. However, mirroring does not protect data from threats such as an application error or a malicious attack such as a virus. Should the source of data be corrupted, then all mirrors would be corrupted automatically, leaving no trusted replica as backup.

- Point-in-time replicas are snapshots of data stored on a regular basis to form recovery points. At a recovery point, a copy of the data set is stored on a secondary storage device. The data set may be a complete replica or, for efficiency, an incremental update to be stored that is used in conjunction with a complete replica taken at an earlier time. Should the primary storage system become unavailable, then recovery is accomplished by reverting to the most recent snapshot.

Costs and Benefits of Disk-to-Disk Replication

Because the backup data resides on disk, it can be accessed or restored quickly, and systems that depend on the data can return to service. Thus, disk-to-disk replication is a preferred architecture for mission-critical systems such as enterprise Web commerce sites.

Disk-to-disk systems provide the fastest recovery times; therefore, this architecture is best suited to critical business systems that the enterprise cannot afford to have out of service. While more costly than disk-to-tape systems, disk storage solutions that leverage low-cost, high-capacity disk technologies are increasingly competitive with tape storage solutions.

Companies that want to gain greater isolation among data replicas might choose to disperse disk-to-disk systems to secondary datacenters over hundreds or even thousands of miles. These distant datacenters are unlikely to share threats such as severe weather, earthquakes, or other calamities. This added resiliency comes at a cost, however, which is the bandwidth necessary to move data among the primary and secondary datacenters.

Disk-to-Tape Replication

In a disk-to-tape architecture, replicas of data primarily stored on disk are written to tape on a regular basis. In addition to mitigating the risk of a disk outage or accidental deletion, tape systems provide the opportunity to move data sets to secure offsite locations and to read data sets at alternative datacenters.

Costs and Benefits of Disk-to-Tape Replication

Writing to tape is a much slower process than replicating data from disk to disk, not because of limitations in tape drive speeds, but because, applications must typically be brought offline while being backed up. Mirroring is also not possible and, for large data sets, recovery points will be limited by the time available to prepare backup tapes. Writing to tape is even slower when multiple copies are made and tapes are verified to make certain that data sets have been stored successfully.

Recovering data from tape is also much slower than retrieving it from a secondary disk. The actual data transfer from tape is quite fast; however, finding and mounting the appropriate tapes can lead to recovery delays. Fast robotic tape libraries can help overcome some of these limitations. IT organizations ordinarily keep tapes with recent replicas close at hand, either mounted in multidrive tape systems or stored in robotic tape silos; however, to recover data sets on tapes sent offsite for archiving may take hours or even days.

Disk-to-tape replication has the unique advantage of creating a portable copy of a data set. Portability provides greater flexibility because tape cartridges can remain close at hand when the need to access to information is greatest and then be moved to archival storage for greater safety and economy. Due to portability, tape backup could form the basis for recovery at a third-party datacenter.

Tape remains the lowest-cost, highest-density storage alternative. For large amounts of data that need to be protected over months or years, tape remains the medium of choice.

Disk-to-Disk-to-Tape Replication

Disk-to-disk-to-tape systems are hybrid systems that use secondary disks for short-term storage and rapid recovery while writing to tape for long-term storage and archiving. While more complex than the previous architectures and more expensive than disk-to-tape, these systems address weaknesses found in disk-to-disk and disk-to-tape architectures. Two techniques are common when using the disk-to-disk-to-tape architecture: point-in-time replicas and virtual tape.

- Point-in-time replicas, mentioned as a disk-to-disk technique, are snapshots of data moved from primary disk storage to secondary disk storage to form recovery points. In a disk-to-disk-to-tape environment, these snapshots may also be moved to tape to form longer-lived recovery points. This staged backup to tape has the added advantage of virtually eliminating backup windows since the backup to tape is done via a snapshot, not the active volume.

- Virtual tape systems utilize the secondary disk as if it were a tape environment and include the ability to provide more tape drive images than the number of physical tape devices available. Due to the higher performance of disk systems, replicas can be made much faster and more conveniently. These replicas can remain on disk if access is expected over the short term. Eventually, and often more efficiently, virtual tape system software migrates the data sets to tape. This is a simpler implementation of disk-to-disk-to-tape replication, as some of the complexity of traditional disk-to-disk-to-tape solutions is managed internally by the virtual tape system.

Costs and Benefits of Disk-to-Disk-to-Tape Replication

Flexibility is the hallmark of disk-to-disk-to-tape systems. Data sets can be mirrored or stored at multiple recovery points quickly. Replicas stored on the secondary disk system can be moved to tape without disrupting the primary system, thus creating a larger window for backup operations. With this architecture, the IT organization can plan for a comprehensive data life cycle that uses faster disk storage for short-term data protection and less expensive, removable tape storage for long-term archiving.

While many disk-to-disk systems require the use of expensive primary storage on both the primary and secondary systems, disk-to-disk-to-tape systems use higher-performance disk storage systems for primary storage and lower-cost, higher-capacity disks (e.g., SATA disks) for secondary storage. While SATA-type disks offer lower performance, they are well suited for streaming data, which is what is needed in a secondary storage system when writing data sets to tape.

Another key differentiator of an attractive disk-to-disk-to-tape system is the degree of automation provided. Manual systems will be less expensive to buy, but higher operating expenses can consume much, or even all, of the cost savings. Systems capable of following data protection policies automatically may appear more expensive to purchase (due to added software functions) but are less expensive to

operate in both the short and long run. As always in IT investments, acquisition cost is only part of the total cost of ownership.

Point-in-time replicas offer other benefits in addition to data protection. In many organizations, data collected by one organization is useful to another. Customer care data may be collected as part of a product support function, for example. That same data may be useful to the engineering department to investigate product quality. Similarly, a replica may be useful to the IT development group for application testing.

At the same time, disk-to-disk-to-tape systems, especially those that go beyond virtual tape, can be among the most complex and most expensive. Their purpose should be to provide both a high level of data protection for the enterprise data that is most valuable and a faster time to recovery for the applications that need that data. Table 1 compares and contrasts the three data protection architectures.

TABLE 1

Three Replication Architectures

	Disk to Disk with Mirroring	Disk to Disk with Recovery Points	Disk to Disk to Tape	Disk to Tape
Recovery time	Very fast	Fast	Variable	Slow
Recovery points	1 per mirror	1 per mirror	1 per mirror and 1 per backup cycle	1 per backup cycle
Typical retention period	Hours	Days	Days to years	Years
Relative cost	High; when remote, very high	Medium	Highest	Lowest

Source: IDC, 2004

Continuous Data Protection

Continuous data protection (CDP) is an emerging disk-to-disk replication methodology that utilizes continual logging of all data changes at the block level. Each changed element is saved to a secondary disk with a time stamp. Point-in-time virtual mirrors can be created as needed with this time stamp. The major benefit of CDP is the ability to create multiple mirrors while using far less storage capacity. The number of recovery points is limited only by the logging period, which is typically measured in seconds.

When used in conjunction with a disk-to-disk-to-tape architecture, CDP supports disk-to-tape backups for any time of the day. For example, in the early hours of the morning, storage management software can assemble a virtual mirror for the state of a data set at the close of business the previous day and dispatch that snapshot to tape.

IDC believes that CDP products will play an important role in overall data protection strategies in the coming years, but they must be well-integrated with existing data protection systems.

A Note on Data Protection Methodology

These basic data protection architectures must be augmented with a thoughtful methodology to ensure streamlined operations. For many organizations, for example, organizing data sets into consistency groups is a sensible move. A consistency group is a collection of interdependent data sets. For example, an ERP application may require access to several data sets. To successfully return this ERP application to operations after a data outage, organizations need to restore all of these data sets. By identifying and storing collections of interdependent data in consistent groups, companies can make the overall process of data protection more efficient and reliable.

Three Steps to Improved Data Protection

IT organizations interested in improving data protection should consider three initial steps, as follows:

1. Prepare an inventory of data sets, their associated applications, their expected impact on business value, and current replication technique, if any. Consider the cost of downtime and classify the data. SNIA offers the following classification system:

   - Class 1 — Data that is not important to operations. Store with 90% data availability.

- Class 2 — Data that is important to productivity. Store with 99% data availability.
- Class 3 — Data that is important business information. Store with 99.9% availability.
- Class 4 — Data that is vital business information. Store with 99.99% availability.
- Class 5 — Data that is mission-critical. Store with 99.999% availability.

2. Prioritize the inventory of data sets organized by classes of data. Figure 3 shows how the three data architectures and continuous data protection fit with the classes of data.

FIGURE 3

Data Protection for Classes of Data

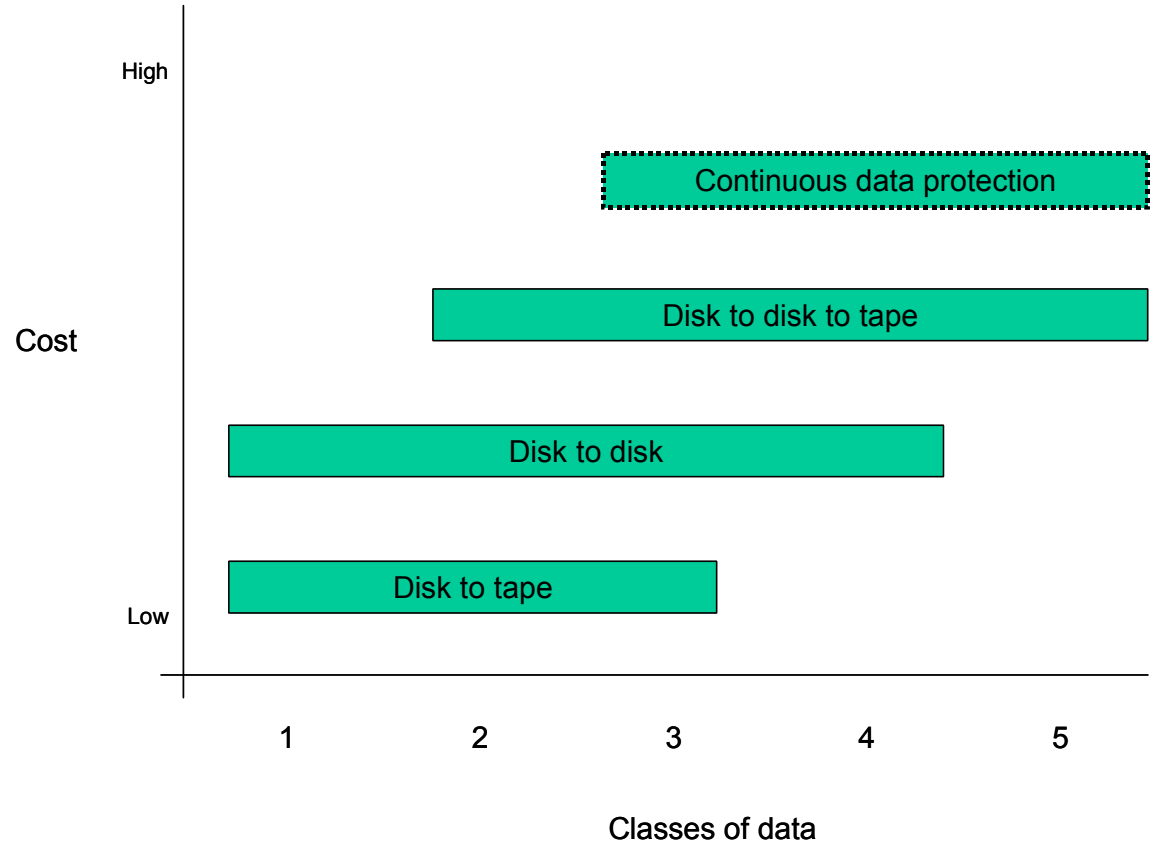

Source: Adapted from SNIA, 2004

3. Build an enterprise data protection plan that protects high-value data sets with more robust data protection technology and low-value data sets with more economical data protection technology.

   - ❑ Data not crucial to daily operations, not easily recreated, or not involved in compliance regulations should be protected with the most economical technology available, which is often tape backup. It is important that IT organizations save resources when dealing with less important data, since there will be no return on a high investment when an outage occurs.
   - ❑ Data that is increasingly important deserves improved protection technology. This means moving to hybrid environments that use both disk and tape replication.
   - ❑ For highly valuable data, more than one protection method may be needed. Multiple mirrors with frequent point-in-time replication may be needed to ensure that real-time commerce systems are highly available. Those commerce systems may also require tape backup to meet longer-term archival requirements.

DATA PROTECTION SOLUTIONS FROM STORAGETEK

StorageTek provides a suite of data protection products that meet the diverse needs of IT departments. Table 2 shows how these products align with the architectures discussed in this white paper and highlights techniques that they support.

TABLE 2

StorageTek Data Protection Products

	Disk to Disk with Mirroring	Disk to Disk with Recovery Points	Disk to Disk to Tape with Recovery Points	Disk to Disk to Tape with Virtual Tape	Disk to Tape
Disk products	• FlexLine™ 200 and FlexLine™ 300 series disk • FlexLine™ SVA • MirrorStore data replication appliance	• FlexLine™ 200 and FlexLine™ 300 series disk • Flexline™ SVA • FlexLine™ 600 series disk	• FlexLine™ 600 series disk • Streamline modular library systems	• Virtual Storage Manager® • Virtual Storage Manager® for Open Systems	• FlexLine™ SVA • FlexLine™ 200 and FlexLine™ 300 series disk

TABLE 2

StorageTek Data Protection Products

	Disk to Disk with Mirroring	Disk to Disk with Recovery Points	Disk to Disk to Tape with Recovery Points	Disk to Disk to Tape with Virtual Tape	Disk to Tape
Tape products			• L-series tape libraries • T-series tape drives • LTO Ultrium tape drives • SDLT tape drives	• SL8500 and SL500 modular library system	• L-series tape libraries • SL8500 and SL500 modular library system • T-series tape drives • LTO Ultrium tape drives • SDLT tape drives

Source: StorageTek, 2004

StorageTek is well positioned to provide products and services that help its customers build and execute a data protection plan. A well-considered plan will minimize investment for less valuable data sets while making sure critical data remains available to support business continuity.

Directly reflecting StorageTek's information life-cycle management strategy, the company's data protection offerings support the full range of data protection options for an enterprise storage environment. Disk and tape technologies are both important when designing a data protection strategy. Moreover, a thoughtful blend of disk and tape storage can enhance the efficiency of tape storage. Using a secondary disk storage system to stage tape backups allows for greater utilization of tape capacity and better organization of data sets on tape.

Most of StorageTek's products support more than one data protection option, adding greater flexibility in implementation and simplified management through standardization.

Challenges For StorageTek

StorageTek must find ways to demonstrate the value of data protection in both business and technical terms. In spite of evidence that data can become unavailable and trigger catastrophic business problems, many enterprises are reluctant to invest in robust backup systems. Even within IT organizations, one can find a "write once

read never" mentality among technicians who have come to trust system reliability beyond its limits.

Risk mitigation is a fundamentally difficult task. While some parameters in the equation are known with great precision (e.g., the mean time between failure for storage components), other parameters are difficult to assess (e.g., the unexpected business effects of losing access to a data set).

Finally, StorageTek must ensure that its customers can effectively coordinate their use of the companies growing range of data protection solutions. As noted above, IT managers must apply different data protection policies to different data sets. Without effective management tools to coordinate these activities and monitor their status, IT managers face rising administrative costs and, more important, won't be able to guarantee data integrity.

FINAL THOUGHTS

Data protection requires solutions that are designed and deployed after a business-value analysis of enterprise data sets and a thoughtful design that matches more costly technology to more valuable data sets. To make the business case for investments in data protection, IT organizations will need to develop a portfolio of data sets matched with appropriate replication architectures.

When selecting data replication solutions as part of this architecture, they also must deploy a set of common administration tools that will enable them to consistently manage and monitor all data protection processes based on a common set of policies.