



MARCOM
ENGINEERING

Affordable Tape Automation

*The Key to Removing Business
Risk Through Reliable Disaster
Recovery*

Marcom Engineering
6956 South Garrison Street
Suite 300
Littleton, CO 80128
(303) 399-9940
www.marcomengineering.com

Project: Exabyte VXA Autoloader
White Paper

Client: Exabyte

Contact: Chris Cottle

Version: 2

Date: August 9, 2004



Exabyte Corporation
1685 38th Street
Boulder, CO 80301

www.exabyte.com

An Exabyte White Paper

Affordable Tape Automation

The Key to Removing Business Risk Through Reliable Disaster Recovery

Introduction

Without question, carefully planned and executed backup procedures are critical to every business. Data loss can mean the loss of accounts receivable, pending sales, and customer loyalty. As evidence of data loss devastation, 90% of all companies that experience data loss because of disaster are out of business within two years, and nearly 50% never reopen their doors at all¹.

In addition to the importance of data protection to ensure business preservation, recent corporate misbehavior has led the government to require the protection and retention of detailed business records. Over 10,000 state and federal regulations deal with the retention of records², across almost every industry.

Yet, despite the consequences of a poor backup strategy, surveys reveal disturbing statistics about data restoration events. The Gartner Group reports that 40 to 50% of all backups are not recoverable in full, and that 60% of all backups fail in general³. Even in large enterprise data centers, nearly one quarter of respondents report that 20% or more of their tape-based recoveries fail⁴.

Numerous surveys reveal that second to misconfiguration of the backup software, the leading cause of backup failure is human error⁵. As a monotonous, thankless job, the shuttling of tapes in and out of a backup tape drive is frequently delegated to employees without the benefit of professional training, who perform the task as their lowest priority on the way out the door. Minimizing the dependence on human intervention by consolidating the tape loading effort reduces the risk of human error by 80%, and proportionately increases the overall reliability of the backup process.

Eliminating the daily insertion and removal of backup tapes by implementing robotic tape automation systems is the key to eliminating the risk of human error from the data backup process. Although standard in large scale IT operations and data centers, tape automation systems are uncommon in small to medium-sized businesses. A 2004 Exabyte survey reveals that 70% of small to medium-sized businesses recognize the need for a tape automation system but forgo implementation because it is cost prohibitive. This paper explores new technologies as a solution for deploying affordable tape automation.

¹ University of Texas, Center for Research on Information Systems, 1994 Survey

² Compliance: The effect on information management and the storage industry, May 2003, Peter A. Gerr, Brian Babineau and Patrick C. Gordon, The Enterprise Storage Group.

³ Gartner Group, January 2002, Adam Couture

⁴ Enterprise Storage Group, "The Evolution of Enterprise Data Protection", January 2004

⁵ DM Direct Newsletter, DM Review, January 14, 2000, Drew Rob



The Manual Backup Process

Considering the number of times that human intervention is required in a single tape drive backup process, it is not surprising that backup reliability is negatively impacted. Figure 1 illustrates that a backup operator handles tapes 20 times in ten individual sessions over a two-week period. The number increases if a restore operation is required, with a greater potential for error considering that the restore operator is likely different than the employee tasked with inserting the nightly backup tape.

Each time human intervention is required, one of the following may occur:

- 1) The proper tape is inserted, and the ejected tape is correctly filed.
- 2) The proper tape is inserted, and the ejected tape is misfiled, which will impact future backup or restore operations.
- 3) The wrong tape is inserted, preventing the current day's backup from occurring. In the worst case, the backup software is configured to write over existing data on the tape.
- 4) The employee is absent or neglects to change the tape, preventing the current day's backup from occurring.
- 5) The tape drive detects a mechanical anomaly while mounting the newly inserted tape, and ejects. A stand-alone tape drive unit cannot reinsert and retry the tape without operator intervention.

The backup process is generally reliable in its first month or two. Habits are formed, and a few restores are performed to verify the process. However, as weeks turn into months, the task of swapping tapes loses its aura of importance, particularly after a few months pass without the need to restore data. Once the backup process begins to erode, months or years may pass before a problem is revealed, and often at the worst possible moment.

Labor costs of the manual backup process seem deceptively minimal yet accumulate significantly over time. In the best case, adding a tape interchange task to an employee's workload results in a disruption of at least fifteen minutes a day, the equivalent of over one work week per year. Any media or tape loading error requires additional hours per event to recover, often requiring attendance outside normal business hours.

Finally, the manual backup process becomes impractical once the volume of data exceeds the capacity of a single tape. As such, a backup operator is required to mount subsequent tapes within an already tight event window hours after backup initiation. At this point, backup operations suffer from either neglect or sporadic performance, or the investment in the inadequate backup tape drive is abandoned for a new, larger device.



Figure 1 – Illustration of 20 tape movements over a two-week backup cycle.

Introduction to Backup Automation

The inadequacy of the manual backup process in an on-demand world of explosive data growth led to the introduction of the autoloader. Autoloaders specifically address the human intervention challenge of successfully moving data to and from tape by automating the backup and retrieval processes. An autoloader is a compact storage device that contains a tape drive, tape cartridge slots, and a robotic mechanism that moves tapes between the slots and the tape drive. Unlike libraries, which contain several tape drives and many cartridge storage locations, an autoloader typically contains only a single tape drive and up to ten cartridge slots.




Figure 2 – Illustration of two sets of tape movements over a two-week backup cycle.

Controlled by backup software, autoloaders enable specific backup routines to run automatically while simplifying data restoration. The autoloader's robotic mechanism operates as a random-access "changer" device with software that controls libraries, or as a sequential-access "stacker" device with tape drive control software. The random-access feature distinguishes the autoloader from a simple stacker device, which only processes cartridges in sequential order.

With ten cartridge slots, a single tape loading operation, occurring weekly or even bi-weekly, addresses as many as ten backup events. By reducing the frequency of human involvement from daily to weekly, the autoloader eliminates the opportunity for human error or neglect by 80%. Filling the autoloader with a new batch of tapes at extended intervals is intrinsically more reliable than depending on daily tape swapping, and even accounts for employee absence and turnover.

Conveniently, a weekly tape loading routine also organizes the tape sets for off-site rotation. A single new tape loading session produces a collection of the most recent backup copies prepared for transport to a disaster-proof location. Furthermore, handling tapes as complete sets greatly reduces the probability of tape misfiling errors.

The autoloader also simplifies the majority of restoration operations caused by the loss or deletion of recent files. Because the initial full backup and several daily incremental or full backups remain in the autoloader after their completion, in most cases the tape containing the file to be restored is still loaded and available online. Furthermore, when the restore is completed, the backups will continue on schedule without operator intervention.

Invariably, tapes rapidly deteriorate with frequent human handling due to greater exposure to contamination and mishandling. In contrast, autoloaders, based on tests performing

one million insertion/extraction cycles, affect no significant attrition to the tape cartridge or loading mechanism. Ultimately, autoloaders maximize tape cartridge life and reliability by ensuring consistent and precise cartridge insertion and extraction.

The autoloader is invaluable when the size of the backup data set expands beyond the capacity of a single tape. The backup application automatically directs the autoloader to immediately load a new tape upon filling the first tape. Without an autoloader, the backup operator must be present at the completion of the first tape, to load the next tape, hours after backup initiation.

The daily manual exchange of backup tapes imposes a minimum fifteen-minute interruption of an employee's primary duties, amounting to 65 hours of lost productivity per year. By reducing this workload to only fifteen minutes per week, or 13 hours per year, tape automation annually adds 52 hours of employee time for other tasks. Assuming a total employee compensation cost of \$60,000, deploying tape automation realizes \$1,500 in annual increased productivity. As such, a solution under \$3,000 will pay for itself in less than two years through reduced labor costs. These calculations do not include the labor savings an autoloader deployment achieves by precluding the requirement for after-hours recovery by immediately addressing media and tape loading errors.

The additional reduction of business risk by employing an autoloader presents the greatest return on investment, as the financial ramifications of catastrophic data loss exponentially exceed the cost of the autoloader. Because the autoloader multiplies the reliability of the backup process, any financial justification for an investment in the backup process multiplies proportionately, again increasing the return on investment.

Automating the backup and retrieval processes facilitates increased backup reliability and enhanced operational efficiency. As such, autoloaders ensure optimal data protection with the added benefit of cost efficiency and investment protection.

Advanced Features in Automated Backup Solutions

Over the past twenty years, autoloaders have evolved from scaled-down versions of their sibling large-scale tape library systems to simplified and streamlined designs specifically intended for small and medium-sized businesses. Such purpose-driven designs have lowered the manufacturer's suggested retail price for feature-rich autoloaders such as the Exabyte VXA-2 PacketLoader, driven by VXA Packet Technology, to below \$3,000.00. This dramatically low price is due to a combination of a new tape technology and elegantly designed advanced robotics, both providing greater reliability at a lower cost.

Rather than reading and writing data in a continuous track, which is the technique utilized by older tape technologies such as DDS, DLT and AIT, Exabyte's VXA Packet Drive reads and writes small, individually addressed packets, similar to the method used to transmit data over the Internet. Tape drives that read and write in a constant stream require very tight mechanical tolerances to maintain perfect alignment of the tape with the read-write heads.

When reading the tape, four tape heads simultaneously read each packet, each slightly offset from the other. By overscanning with four heads, tape alignment and distortion errors are eliminated. Should an error be detected, the drive applies a four-level Reed-Solomon Error Correction Code (ECC) that corrects errors in most circumstances, allowing the drive to achieve a bit-error rate of 10^{-13} . The combination of overscanning and four-level ECC correction makes the VXA Packet Drive 100 times more likely to recover damaged data than other tape technologies.

As each packet is written, a trailing head re-reads the packet to detect write errors. Upon detecting a write error, the drive re-writes the packet without stopping or rewinding. Because the packets are individually numbered, they are automatically rearranged into the correct order. Further, Packet Drive technology allows the drive to adjust the tape speed to match the host data transfer rate, thus eliminating the need to rewind the tape and restart recording. Hence, Packet Drive technology minimizes job delays and excessive wear on the media and drive mechanism⁶.

Because the VXA Packet Drive reads and writes individual packets rather than a continuous stream, the manufacturing tolerances are 270 times larger than competing tape technologies. This larger tolerance directly translates to a lower cost of manufacture, whereby pricing the VXA Packet Drive at about one-third the cost of comparable tape technologies.

Most autoloaders targeted at small to medium-sized businesses are based on a magazine mechanism that selects from a stack of tapes, which the magazine contains. Exabyte's VXA-2 PacketLoader is based on a far simpler carousel mechanism that stores tapes in a circle on a platter. Where the magazine mechanism loading apparatus must precisely move the magazine up or down from one tape to another, the carousel mechanism merely rotates a platter to position the selected tape in front of the drive, eliminating most moving parts. This streamlined robotics and creative engineering again reduces the cost of manufacture while increasing the device's reliability.

The Exabyte carousel accommodates more tapes than a similarly sized magazine mechanism – ten versus eight or fewer. With the capacity for ten tapes, the VXA-2 PacketLoader

⁶ Exabyte Corporation, VXA: Re-Inventing Tape Storage, April 2002

supports two weeks of single-tape weekday backups. Reducing human involvement to bi-weekly rather than weekly again reduces the probability of error by half.

The simplicity of the carousel design facilitates optimal tape drive placement at the rear of the autoloader, behind the carousel. The rear, fixed position enables easy field-replacement of the tape drive for maintenance or upgrades. Upgrading the tape drive to future VXA formats doubles the drive's capacity and throughput, extending the useable life of the autoloader.

The carousel mechanism simplifies the implementation of a bar-code reader, establishing an automatic tape inventory as the tapes are loaded. Bar-coded tapes facilitate on-screen tape load verification, and allow backup software to identify and select specific tapes for backup and restore. With bar-coded tapes, an operator has the ability to instantly verify the carousel's contents, an operation that otherwise requires ejecting and reloading of all tapes or cartridges.

Where past autoloaders were relatively simplistic devices merely able to respond to commands from backup software with little interaction, the Exabyte VXA Autoloader is a complete stand-alone system incorporating an Ethernet interface, and Web and email hosts. By incorporating a complete Web host, the Exabyte VXA Autoloader enables remote management and configuration via any Web browser across the Ethernet network. Installation requires the simple input of device IP addresses via the front panel, or if a DHCP server is available, the autoloader will configure its own IP addresses. From that point, all configuration and management is performed via a local or remote Web browser.

The most commonly utilized feature of remote management is the carousel's tape inventory display to confirm that the correct tapes are loaded. Further, an advanced feature of remote management allows an IT administrator or system integrator to remotely access diagnostics and event logs for reactive or proactive maintenance, without dispatching a technician on-site. Should the tape drive begin to report excessive tape read retries, the customer can easily replace the drive on-site. Even the software utilized to operate and manage the autoloader, called firmware, can be replaced remotely, which uploads new device operations features, supports new devices and accessories, or upgrades the Web based management software application.

The integration of an email server enables the VXA-2 PacketLoader to automatically notify the appropriate personnel of unusual events and device failures. For example, configuring the autoloader to email notifications of excessive tape or loading mechanism retries proactively addresses failing components before they reach the point of unrecoverable failure.

With built-in intelligence, the VXA-2 PacketLoader is capable of dedicating slots to a spare tape and a cleaning tape. In the event of an unrecoverable tape error, the spare tape is

automatically deployed and the backup restarted, ensuring that the backups continue on schedule without human intervention. In addition, the spare tape is deployed to contain overflow should a backup unexpectedly exceed the capacity of a single tape. If a cleaning tape is present, the autoloader will automatically insert and play the tape periodically.

Until the introduction of the VXA-2 PacketLoader, features such as bar code labeling and integrated intelligence appeared only on data center class autoloaders costing an order of magnitude more, and thus targeted at large IT operations. The inclusion of this extensive suite of features in a device targeted at small to medium-sized businesses brings greater data protection to those organizations at an easily justified price point.

Conclusion

Eliminating potential sources of error from the backup process increases backup effectiveness. With the availability of autoloaders, excessive operator intervention in the backup process is an avoidable pitfall. By reducing human involvement by 85-90%, an autoloader greatly increases reliability while eliminating the daily mundane, error prone, and often neglected task of tape swapping. The autoloader's integrated intelligence for emailing notifications of failing components or tapes, and ability to utilize spare tapes for replacement or overload, further ensures that backup reliability remains constant over time.

All organizations, regardless of size, must value data backup as mission critical. An effective backup strategy must begin with the expectation that the need for data restoration, whether enterprise-wide or isolated, is inevitable. The risk of financial loss due to unrecoverable data is such a compelling reality that optimizing backup reliability as a top priority is necessary to secure business continuity and investment preservation. The key to comprehensive data restoration is consistent and reliable data backup. VXA Packet Drive and PacketLoader Technologies offers any sized business an economically attainable solution that reduces the probability of human error, and improves overall backup operations efficiency and reliability.