



Migrating from Tape to Disk Backups



DATA PROTECTION WHITE PAPER

Title: Migrating from Tape to Disk Backups
Product: Revinetix Volo™ and Sentio™ Series
Target Audience: IS / IT Managers, Systems Engineers

ABSTRACT

As backup software breaks away from its historically tight integration with tape, IT administrators are implementing disk-based backup products that are optimized to address new priorities. The new disk-based backup products geared to SMBs are being enhanced with enterprise-class product features and come with prices that are getting less and less expensive, making it feasible to back up from disk to removable disks and do away with tape backups altogether. In fact, when you factor in the initial investment in the tape drive along with the tapes themselves and the lifespan of each type of media, as well as administrative overhead, you will find that disk to disk backups are more cost effective than tape.



1011 Murray-Holladay Road | Salt Lake City | Utah | 84117 | 888-264-5116 | www.revinetix.com

TABLE OF CONTENTS

Abstract	2
Table of Contents.....	3
Introduction.....	4
Evolution of Storage and Data Protection	4
Disk versus Tape Technology	6
Implementing Disk-based Backup Solutions	99
The Revinetix Data Protection Solution.....	111
Summary	13
Related information	14

888.264.5116

www.revinetix.com
info@revinetix.com

INTRODUCTION

New trends in storage make it easier than ever for companies to expand their data capacity. However, scaling up your data backup requires that you also think in terms of scalability of performance, reliability, security, and usability issues.

The basics of backup management haven't changed much over the past year, but new products and technologies have made it easier to address these peripheral issues, thereby making extreme scalability in disk-based data backup much more feasible and desirable than the traditional tape backup.

As storage capacity increases, a larger backup capability is required. Now is a good time for IT administrators to seriously consider migrating to new disk-based data protection technology to address the challenges of creating a fast, user-friendly, reliable, secure and scalable backup strategy.

Unfortunately, small and medium sized businesses have often been underserved by storage products. Either the product is too expensive to match the budget of the business or the solution is designed for a small or home office and is incongruent, unreliable and difficult to use. Small and medium sized businesses need products with features designed for the large enterprises but in a footprint of an appliance that is compatible with their budgets.

EVOLUTION OF STORAGE AND DATA PROTECTION

To date, tape-based backup solutions have been the default medium for offsite storage. Tape has been a central fixture in corporate data centers for the last 50 years, serving as the primary data backup-and-restore media. But increasing data storage needs, shrinking backup windows, the need to recover critical applications quickly, the emergence of Serial ATA technology, and declining disk costs have combined to push tape out of the spotlight in favor of disk.

Disk Backups - As a primary target for backup, disk has quickly become accepted as the optimal media. Backing up directly to disk will speed up the backup process and also reduce restore time for individual files.

In addition, it is also widely accepted that data should be stored offline in some form of removable media. For most companies, their archiving strategy consists of moving data from disk to tape at some point. This backup strategy is commonly known as disk to disk to tape or Disk2Disk2Tape.

However, there is strong market demand for a different solution that has no dependence on tapes. This backup strategy avoids the inherent problems of tape archiving by using removable disks in what is known as Disk2Disk2Disk products.

Comparing the Costs - Disk-based backups don't suffer from the same incremental restore penalties experienced by tape drives. For example, if an enterprise keeps 12 weeks of backup of 100GB of file data that changes at a rate of 5% per day. With tape, a reasonable backup schedule would mean at least 1 full backup per week, plus nightly incremental backups. This adds up to $12 * 100 + 12 * 6 * (100 * 0.05) = 1560\text{GB}$, a factor of 15.6 times the original dataset size, assuming all the tape capacity is used perfectly.

With disk-based solutions, the enterprise only needs to perform one full backup every 4 weeks, with nightly incremental backups in between. This adds up to $3 * 100 + (9 * 7 + 3 * 6) * (100 * 0.05) = 705\text{GB}$, less than half the storage capacity used by the tape. If tape capacity costs \$3/GB, the effective cost to compare with a disk-based backup solution would be $3 * (1560 / 705) = \$6.64/\text{GB}$. A difference of that magnitude would easily bridge the gap between tape and disk costs on a simple \$/GB basis. Take into account all the other benefits that disk-based data protection solutions have to offer and you can see that disk-based backup is a compelling proposition.

Licensing and other Costs - Some costs are not as easily calculated. Important factors to consider include the software features, licensing fees, installation and configuration and the ability to archive backup data to off-site storage. The licensing costs can vary dramatically across the different backup solutions. The complex licensing structure charge based on the

number of clients, the types of plug-ins, or the applications being backed-up. A simplified approach is to have a single site license that covers unlimited clients and applications.

Building a Foundation with Data Protection – The foundation of a data protection strategy requires expandability, so that as the organization and data increases, you can expand storage space to meet your needs without interfering with the ability of users to access previously stored data.

Government regulations and court precedents are forcing companies to retain more data for longer periods of time. But there's more to it than storage space. A data protection solution requires not just the ability to store more and more data, but also the ability to organize, retrieve and secure that data; otherwise the stored data is not very useful.

The security of stored data is a big concern — especially in regulated industries where HIPAA, Sarbanes-Oxley, the GLB Act and other legislation require that the privacy of certain information be protected. In addition, reliability is of utmost importance. Loss of important data can result in irreparable harm to your organization. In fact, a report showed that 93% of all small to medium sized businesses (SMBs) who lose their data backups go out of business in three years¹.

However, while disk has received a great deal of attention lately as the new backup-and-restore media, some are still questioning whether disks are an all-out replacement for tape in an existing backup process.

DISK VERSUS TAPE TECHNOLOGY

There is no better computer security defense than having a known, good, safe data backup, right? Then why is tape backup software so complex and buggy? Why do backup drives frequently fail? Why do backup tapes fail so often?

¹ University of Texas Center for Research of Information Systems, Datamation June 14, 1994

If you are still clinging to tape backup solutions, you better make sure your tape drive is under warranty, as they don't seem to last very long anymore. That expensive tape drive seems to fail as frequently as inkjet printers. It's almost a consumable now, much like the tapes and backup cartridges themselves. You can buy the most expensive backup cartridges available, and you're lucky to get a few hundred hours of backup time before they start spitting errors and lost CRC sectors. There are lots of good reasons why backup cartridges fail so often, but it still doesn't solve the problem.

Understanding the Facts - The truth is that most IT managers can't wait to replace all their tape drives and cartridges within their data centers with hard disk drives. (A survey in 2006 conducted by Peripheral Concepts reported 58% of IT managers wish they could get rid of tape backups). Holding them back from doing this is the common belief that disks can't match the attributes of tapes:

- **First**, tape can be removed. Data should be removed from the system to help prevent intentional or unintentional corruption from a virus or system error.
- **Second**, tape can be transported. It can be taken to an offsite location. This provides a second level of disaster recovery. If the main data center suffers an outage, data can be restored from offsite tapes.
- **Third**, tape media quality is thought to have improved since the appearance of round tape reels in the 1950s.
- **Fourth**, tape is believed to be a cost effective storage media.

These attributes are indeed important but the facts are that tape is not the only backup media with these attributes.

The Truth About Tape Backups - A recent survey of IT executives on tape backup solutions reported these findings:

- **75%** of respondents indicated that their companies suffered unrecoverable loss of corporate data they thought was successfully backed up to tape due to unreadable, lost or stolen media.

- **63%** said they encountered unreadable tapes when they tried to retrieve data with 76% of those cases reporting a direct impact to their business from loss of productivity to punishments for regulatory compliance infractions.

Moving to Disk-Based Backups – Early opinions indicate that the disk-based backup market grew 50% in 2006. The high adoption of disk-based data protection solutions is a result of increasing data capacities and its unique ability to meet backup windows that cannot be met with tape. The needs of businesses small and large have outpaced the capabilities of tape backup technologies in recent years.

Major advantages of disk technology over tape technology include the following:

- **Faster restores**
The random access nature of disk drives enables the instant merging of incremental backups. Contrast this with tape, where multiple incremental backups are often spread across multiple tapes and likely to be far from the start of the tape. The RAID configuration of the disks in a backup appliance allows for redundancy and increased throughput by backing-up and restoring multiple clients simultaneously.
- **Shorter backup windows**
Disk-based backups reduce the backup windows in two ways. First, random access and higher reliability of the disk media means you can leverage more incremental backups per full backup. Full backups can be scheduled much less frequently, or selected to coincide with an expected lull (i.e., over the weekend) in operations. Since incremental backups typically need to save a small fraction of the total data, most backup windows become much shorter. Second, multiple clients can be backed up simultaneously with disk-based backup which results in potentially higher throughput.
- **Economies of incremental backup**
Leveraging more incremental backups not only helps with performance, but also with cost. Due to the additional time penalties of incremental tape restores and because tape restores frequently fail, almost all IT shops keep many copies of full (Level 0) backups in their tape libraries. Most tape backup schedules use only a small to none at all number of incremental backups for each full backup set. That means there are often many redundant copies of the same file on tape. This significantly increases the \$/GB of the tape media by wasting capacity on unnecessary duplicated file storage.
- **Accelerated backup and archiving**
Another benefit of disk-based backups is performance. Disk-based solutions leverage RAID performance and reliability to quickly backup and restore data.

- **Greater reliability**
Research has found that as many as 70% of tape recovery attempts fail². Disk-based solutions benefit from very reliable commodity disks, made even more reliable by using RAID to protect against disk failures. There are no tapes to misplace, or robotics to jam and recalibrate. Disk drives are self-contained and less sensitive to environmental conditions, so they can be reliably used for a longer period of time in more diverse environments. For most SATA drives, the manufacture's warranty is 5 years.
- **Easier management**
Disk-based solutions are easier to manage overall. The interfaces and concepts of file system, disk, and network attached storage (NAS) are familiar even to entry-level IT staff, meaning less training and fewer mistakes from the start.

IMPLEMENTING DISK-BASED BACKUP SOLUTIONS

New technology increases the power of disk to support backup, data retention, and disaster recovery planning, while delivering the performance that has long been a key differentiator of disk-based solutions.

The challenge is finding a solution that gives you the speed advantage of disk and the long term retention qualities of tape – and you don't want to deal with the hassles of integrating multiple components yourself.

Choosing the most appropriate disk-based backup solution is a critical task for the simple reason that there are numerous, varying options available. As a guide, users should ask vendors the following;

- How is the product managed?
- Is there a remote management facility?
- Is the management software included in the price of the product?
- Are the components (hardware, software) of the solution fully integrated?
- What is the licensing model for the product?
- Which operating systems are supported?
- Is data backed up on a volume or file level, or both?

² "Disk to Disk Backup: Pass the Data Please", Computer Technology Review, June 2003

- What email databases are supported, if any? Are specific agents available for applications such as Microsoft Exchange? How much do they cost?
- What is the minimum and maximum capacity of the system? Can you easily upgrade to larger capacities? What RAID levels are supported?
- How long does it take to install the product?
- What is the throughput of the system?
- How many backups, and how many versions of backups, can the configuration support?
- When restore is required how granular is the restore process? Can data be restored to a specific date-in-time or is a full restore required?
- Is compression supported to reduce capacity requirements?
- Can the disk drives be removed for offsite storage?
- Can the solution perform a bare metal restore?

Migrate from Tape to Disk in 3 Easy Steps

For small and medium size businesses, migrating from tape technology to disk-based solutions is easier than you may think.

1. Start with Discovery

- a. Understand your areas of frustration and how your ideal solution would behave.
- b. Create a design plan showing your current network infrastructure with the existing backup solution.
- c. Modify the plan to show how a disk-based backup solution would be implemented

2. Do your Homework

- a. Research the market place to find the product with the features you need while also discovering new features previously unknown to you. Use webinars and online documentation to understand how the product will be used.
- b. Evaluate the products that best fit your situation.

- c. Create cost comparisons, keeping in mind today's needs and future growth requirements.
 - d. Find a company willing to review your plans and assist in optimizing your migration to a tape-free network.
3. Execute on your plan
- a. Read the technical documentation and ask questions when ever possible.
 - b. Document the migration process.
 - c. Document an archiving process for offsite storage.
 - d. Make the switch and enjoy a good night's rest knowing you have the best data protection solution available.

THE REVINETIX DATA PROTECTION SOLUTION

The core competencies of Revinetix' data protection solutions are tailored for the small and medium business (SMB) market. Revinetix' Disk2Disk2Disk backup solution integrates seamlessly into existing networks and contains the features that SMBs are looking in a comprehensive data protection solution.

The benefits of Revinetix products start immediately by removing price and performance barriers, thereby enabling IT managers to deploy a tiered storage approach for optimum data protection. The combination of backup and recovery into a single appliance gives organizations of all sizes a highly effective Disk2Disk2Disk strategy.

Faster Backups - Revinetix' Disk2Disk2Disk products provide the best of two worlds, the speed of disk and the safety of offsite archiving. Revinetix gives you the benefit of a super-fast disk backup, minimizing backup windows, and incredibly quick restores, restoring from randomly accessible disk. Users get their data and systems back almost immediately.

Disk-Based Archiving – Disk2Disk2Disk appliances come with removable drives that enable mission critical data to be archived in a safe, offsite location. Offsite archiving is made possible through hot swappable removable disks that are then placed in protective containers for transportation and storage. At anytime, archived data can be assessed by inserting the disk back into the appliance.

Comprehensive Scheduling - IT departments are now able to perform a backup to disk quickly, in a way that can be catalogued off-site, while maintaining a disk image of the backup on-site. The advanced scheduling software found in Revinetix products let you plan days, weeks, months and even years ahead. An advanced scheduler also makes on and off-site backup truly an automated and unattended process. This provides SMB departments with the tools needed to make preventing data loss easy.

Disaster Recovery - With Revinetix, recovering from a complete failure can be as quick and easy as booting up your system. With every appliance, Revinetix includes their proprietary Bare Metal Restore (BMR) software. This replicating software is an IT manager's best friend. By creating a complete disk image to be stored on the appliance, the BMR software can quickly add new users, or get users back up and productive after a fatal system crash.

Disaster Recovery Assistance Program - Under this program, a customer may call the Revinetix support line and inform the support staff of the emergency. Within two business days of notification, Revinetix will ship a temporary appliance to help the customer restore their data from their archive media.

Advanced Archiving Technology - By eliminating redundant data, Revinetix appliances allow users in small to medium businesses to retain 10 to 50 times more backup data on fast recovery disks and cost-effectively store data for months or even years.

In addition to data de-duplication, Revinetix' integrated software layer includes a high performance embedded operation system (RevOS), high speed data compression, asynchronous replication, interface flexibility, and built-in monitoring, alerting and diagnostic tools. This integrated software layer offers an extensible foundation for future intelligent

backup, recovery and archive solutions that will improve data protection for a broad range of customer environments from small remote offices to large enterprise data centers.

Revinetix designed the Volo and Sentio Series to be easily installed and managed in a wide range of IT environments, and the company's long-standing expertise in backup, recovery and archive solutions makes it uniquely positioned to meet customers' disk-based backup needs within the broader storage context.

SUMMARY

Revinetix' Disk2Disk2Disk data protection appliances offer a complete backup solution for protecting all your critical information in an enterprise network environment. Revinetix is a low cost solution addressing the storage needs for large and small. Revinetix appliances all come with intelligent storage management providing superior performance, unparalleled scalability and reliability.

Standard features on every Revinetix product include:

- Turn-key solution with fully integrated components
- Unlimited site license for servers/desktops
- Plug-in support for Microsoft Exchange
- Fully automated backups
- Offsite archiving with removable "hot swappable" disk drives that are stored in protective enclosures
- Incremental and differential backups
- Rapid file restoration with point-and-click interface
- Advanced replication capabilities that allow for LAN and WAN based disaster recovery strategies. The appliances can be configured to replicate data by client to other Revinetix appliances located remotely, making for a seamless and completely automated disaster recovery plan.
- Bare Metal Restore software that allows IT administrators to quickly add new users, or get users back up and productive after a fatal system crash.

Conventional disk storage cannot provide enough cost-effective capacity to let users retain more than a few days of backup data, which limits the benefits disk provides. Additionally,

the amount of backup data in distributed sites and cost of bandwidth have prevented disk-based remote replication from being a viable solution, leaving data at risk in the event of a site loss or other localized threat. Revinetix addresses these issues with a scalable appliance-based backup and replication solution that further expands the company's disk-based backup portfolio.

RELATED INFORMATION

To understand the many aspects involved with serial ATA drives, visit the Serial ATA Working Group (www.serialata.org).



© 2008 Revinetix Inc. All rights reserved.

The information contained in this document represents the current view of Revinetix Inc. on the issues discussed as of the date of publication. Because Revinetix must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Revinetix, and Revinetix cannot guarantee the accuracy of any information presented after the date of publication. This document is for informational purposes only. REVINETIX MAKES NO WARRANTIES, EXPRESSED OR IMPLIED, IN THIS DOCUMENT. Revinetix and the Revinetix logo are trademarks of Revinetix Inc. Other product or company names mentioned herein may be the trademarks of their respective owner(s).