

ADIC Hits a Home Run — Increasing Throughput for Scalar LTO Tape

Analyst: David Reine

Management Summary

It has been a hectic summer for IT. Improvements in the economy may have generated additional business, which in turn might have raised your data center’s workload. A new acquisition might have resulted in a heterogeneous integration effort by the staff to deal with all of the servers and storage units. Trying to keep all of the servers operational may have been a challenge. Trying to keep all of the disk arrays organized and operational was a nightmare, but IT succeeded. You may have had some crashes, but with a sound hierarchical storage management policy in place, IT could save space and help protect the data. The backup policy, even with new servers to manage, protected the enterprise data. Nevertheless, there was and is a problem. **With all of the new users and files, it is becoming more difficult to complete the backup in the designated window. The next acquisition, or the next expansion, would extend the backup window and restrict access to some data during operational hours.** This problem will have to wait, though, for Labor Day weekend has past, and the last fling of summer is waiting.

You have just come back from a long weekend. It was restful and the time with the family eased some guilt. No one has to reboot the “sun”, and the only “sharks” that you had to worry about were swimming in the ocean or cruising the boardwalk. Baseball was on your mind and your favorite team was losing, but had a rally going. A base hit here, a walk there, a run here, a run there. Slowly, your team started to catch up, one run at a time. **Then it happened, a home run – a grand slam – four runs at once, with a burst of players rounding the bases, all scoring.** If you view the base path as a single stream; then the bases are like a cache. The faster you fill the bases, the faster a single hit could clear them, all at the same time. **IT faces this same problem in the data center. How do you get an ever-increasing flow of traffic from point A – the disk arrays – to point B – your tape library, faster?**

Your data center budget might not support the expense of replacing the tape library at this time as a “speed-up” solution. IT cannot afford to replace the tape drives with that newer architecture. The cost of converting all of the tapes in the archives would be prohibitive. No additional drives will fit, even if there was budget space; the drive slots are at capacity already. On your return, you hear that demand might double. IT needs a back-up accelerator. However, what is it? Moreover, from where? From ADIC, perhaps. In order to find out how ADIC proposes to increase the throughput of data from your existing heterogeneous servers to your installed tape library, please read on.

IN THIS ISSUE	
➤ Data Center Backup Today.....	2
➤ A Solution from ADIC.....	2
➤ Alternatives for the Future.....	3
➤ ADIC Pathlight VX Architecture.....	3
➤ Pathlight VX Benefits	4

Data Center Backup Today

The demands being placed on IT departments across the country, to protect the electronic business assets of the enterprise, are constantly increasing. Enterprises are asking IT managers to do more with less: fewer resources, less money. Nowhere is this more evident than in the data center where the IT manager executes the applications most vital to the continued vitality of the enterprise. **One such application is backup, whether from a single server or from the Storage Area Network (SAN) supporting a network of servers, with open systems servers representing the largest portion of the data center server space today, in most enterprises.**

The execution of a backup ensures business continuity by maintaining the integrity of the various databases under IT control, despite the possibility of a human or mechanical malfunction. The ability to recover data lost as a result of this malfunction ensures continuous access from the user community, preventing the possibility of a business interruption.

Every enterprise, every data center has well-established procedures executed by mature hardware and software products. **These policies define which data set to back up, when, and to where.** In most cases, the medium is tape; the process is secure, removable and durable, providing long-term retention for historical data.

Unfortunately, the amount of data that IT needs to back up or save is constantly growing while the time window available for these processes is fixed, or shrinking. There are only 24 hours in a day, eight hours in a shift.

One of the reasons there is never enough backup time is that IT cannot drive commodity tape media at its rated speed. 30MB/sec tape drives often run at 10MB/sec, or even slower, due to the inability of open systems servers to keep them spinning at full speed and to create

a consistent stream.¹ Another reason is the need to complete backups during off-peak hours because of the degradation of performance caused by the backup environment, and the ever-decreasing off-peak hours in which to do it. In a world where enterprises operate on a 7x24x365 basis, every hour that the enterprise shuts down can cost that enterprise thousands of dollars. **Moreover, if time is of the essence for a backup, what about the recovery process?** You might be able to schedule a backup for nighttime hours, or over a weekend, but by its very nature a recovery is spontaneous and urgent. The enterprise is losing money while the recovery process is going on. Even if other applications can execute at the same time that a backup or recovery is running, those applications will degrade significantly from the I/O processes involved in the backup/recovery process.

These issues highlight the problems facing IT departments everywhere:

- How to continue to implement enterprise policies in a time-constrained world;
- How to protect the investment that has been made in network infrastructure;
- How to build in redundancy; while
- Protecting an increasing amount of data assets in a fixed amount of time?

A Solution from ADIC

There are many vendors in the market who have provided perfectly adequate solutions for storage management, in terms of both hardware, an autoloader or a tape library, and also in the area of library management software. One company that has succeeded in delivering superior solutions, is Advanced Digital Information Corporation (ADIC), recognized as the leading provider of automated tape systems to the Open Systems marketplace. ADIC'S product set includes a variety of libraries, including the *FastStor* and *Scalar* families,

¹ Because of their shared environment, the backup of open systems servers is often interrupted causing constant stop/start activity on the drives.

that connect to open systems servers. It does not include the manufacture of any tape drives. Their libraries support LTO, SDLT, and/or AIT tape formats from all of the major vendors.

Since 2001, ADIC has held the #1 position in distribution of tape automation systems used in the open systems and client-server markets, providing products for such recognizable names as Dell, HP, IBM, and Sun. ADIC has set their goal as “Make it better, but don’t add complexity”.

Alternatives for the Future

There are many alternatives in place today to increase the throughput for backup and recovery:

1. One of those involves changing the backup methodology, only saving records that change, for example. However, that does not address the time constraint for recovery and it usually adds complexity to the process.
2. A second alternative is to create more frequent replications of the data. This improves the short-term recovery time, but could be very expensive.
3. A third alternative is to create a multi-tiered storage architecture, using disks to backup and restore the enterprise’s mission-critical data. The availability of ATA disk technology² has enabled this strategy.

ATA disks provide a low cost alternative to the standard fibre-channel devices, but large volumes of ATA are:

- Still more expensive than tape media,
- Impractical for long-term data access, and
- Do not protect the investment made in the storage management infrastructure.

ADIC has viewed these alternatives and tried to incorporate the advantages of each in another path - the *Pathlight VX* to be exact, taking advantage of the best alternatives available. The Pathlight VX is the

first member of what is projected to be a series of innovative storage management products from ADIC. It is the first integrated solution to work within the boundaries of an installed storage management infrastructure. The aim of Pathlight VX is to combine the benefits of disk and tape in a single integrated backup and recovery system.

ADIC Pathlight VX Architecture

ADIC has long established itself as a leader in Storage Management Technology. That includes the areas of:

- Automated Library Expertise;
- I/O Controller Technology;
- Policy-based Data Management; and
- Software Library Virtualization.

ADIC has now broadened their scope with the introduction of the *Pathlight VX Disk to Tape Backup/Recovery Solution*: a rack-mountable, disk-based, modular storage solution that appears to applications as one or more logical tape libraries³. (See Exhibit 1, below) It integrates policy-based data management into the backup and

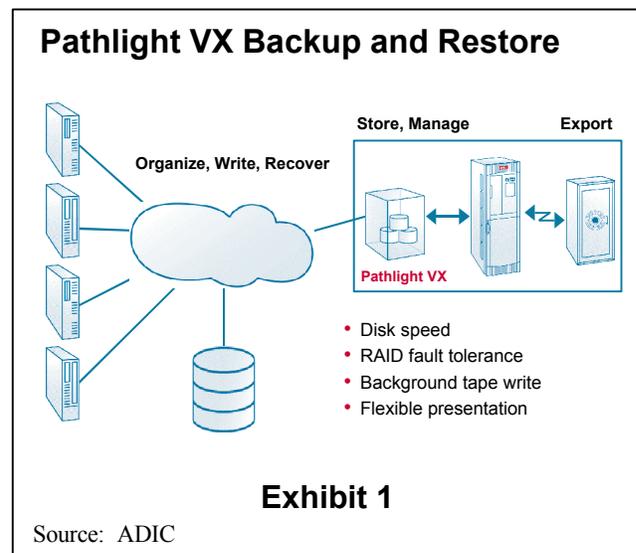


Exhibit 1

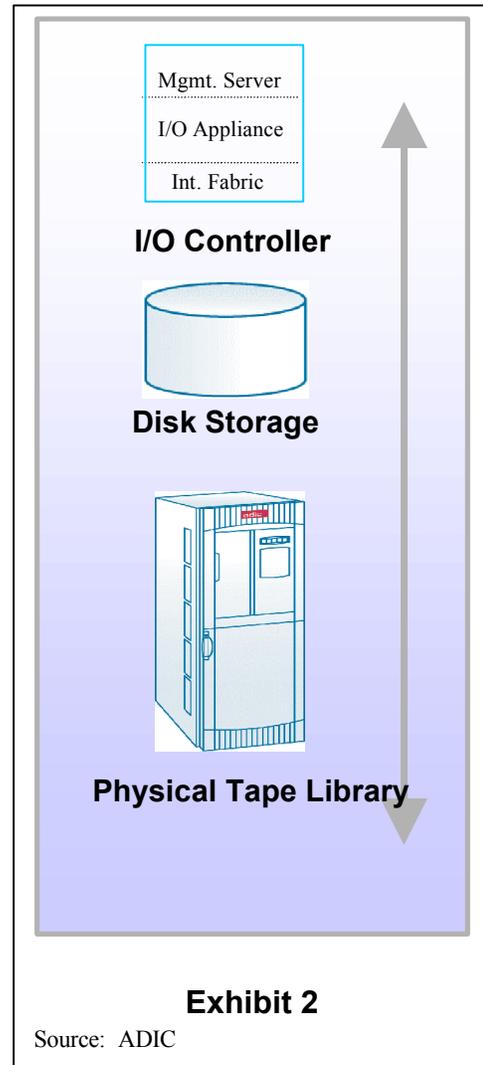
² ATA disks are a less-expensive hard drive technology than is traditionally used in the data center. However, it is good enough and fast enough for many applications.

³ Logical means that it looks and responds like a tape library, but the underlying technology (in this case disk storage) is physically different than a tape library. This is another example of virtualization.

recovery processes, using an integrated data management unit to control both the data I/O and disk/tape transparency. As a front-end to the Automated Tape Library (ATL), Pathlight provides improved read/write performance to the servers, with RAID 5 redundancy for both the backup and recovery processes, and is deployed without changing existing enterprise procedures. **The data center retains the existing servers, Storage Management software, SAN architecture, and tape library — it retains the existing automated path to tape.** The initial release of Pathlight provides support for existing or new ADIC Scalar LTO1/LTO2 libraries with the existing software, eliminating the necessity to retrain the data center personnel. The technology, however, is designed to be independent of both the media type and library brand.

With a disk capacity of up to 40 terabytes, and up to 400 logical media (virtual tape cartridges), Pathlight VX provides an intermediate storage home for the mission-critical backup and recovery data. Pathlight uses Serial ATA RAID 5 arrays to provide a secure, redundant, short-term repository for the data and to ensure the viability of the enterprise. The data center can now choose the time in which it writes the data set to the physical tape, in the background, without impacting productive work by removing that task from the backup window or interfering with a mission-critical application server. Pathlight VX can, therefore, increase throughput to and from the server(s) at a speed up to 1 TB/hour to up to 40 logical drives, using a 2 GB Fibre Channel connection.

Because of its modular construction (See Exhibit 2), incremental disk expansion can be implemented on demand without impact on the Pathlight VX front-end I/O controller that consists of a management server, I/O appliance, and internal fabric. Likewise, if higher performance is required, IT can connect additional Pathlight VX appliances to the internal fabric, providing an additional 1TB/hour of throughput to the SAN connection.



Pathlight VX Benefits

The biggest impact of the Pathlight VX introduction is the enhanced business value that it delivers to existing ADIC installations in open systems architectures, by increasing the one resource over which IT has the least control: *TIME*. By using an intermediate storage medium that is faster than tape, IT can defer the write to tape, the slowest aspect of the storage management process. This shrinks the required size of the backup/recovery window, enabling the return of system control to those applications that actually do production work (i.e. earn money) for the enterprise. Furthermore, Pathlight VX introduces a level of redundancy not previously available for this process, thus increasing the fault tolerance of the enterprise data. Pathlight also enables

the ability to stream the eventual write to tape, allowing the data center to take full advantage of the high-speed capabilities of the various tape units.

By maintaining the existing storage management architecture, the data center can eliminate the requirement for any shift in the backup/recovery paradigm. By keeping the existing storage management software and automated library, IT minimizes the cost of shrinking the backup/recovery window to the incremental cost of the acquisition of the Pathlight VX system, which is approximately \$15/20 per GB of disk cache a persistent and protected cache, preserving the data for as long as required. Furthermore, there is no need to retrain the data center staff: all processes and procedures remain the same. Quite clearly, ADIC has chosen the path of investment protection, rather than evangelizing wholesale changes to the architecture and disruption to the operation of the data center.

ADIC can introduce a Pathlight into the storage management process with an entry configuration as small as 10 TB for only \$190,000. It can grow in an on-demand manner up to 40 TB, as the requirements for interactive access to the Pathlight data increases. In addition, because the architecture remains the same, existing data on tape remains under backup software media-management policies, i.e., ensuring its long-term retention, further ensuring business continuity.

ADIC has additional releases of Pathlight planned. These will provide a path to open system tape formats other than LTO and access to open systems libraries other than those built by ADIC. **Existing ADIC customers, however, can score big with this new capability right now.** Take a closer look at Pathlight VX if the growing pains of backup are too much to handle.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➤ ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Authors

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and a MBA from Northeastern University.

➤ ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 37. (Please dial "1-37" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.