



ADIC Parlays Partnership with EMC – VTL with Improved Performance, Reduced Cost

Analyst: David Reine

Management Summary

The never-ending thirst for the acquisition of more and more data continues to bog down even the most prepared data centers. Requirements for backup and restore, disaster recovery, and archiving continue to plague enterprises of all sizes. It is not only the continued growth of the enterprise with additional customers, additional products, and additional processes, but also the transition of the Information Technology (IT) department from one based simply on “best practices” for that industry to one attempting to adhere to the standards established for the implementation of regulations such as Sarbanes-Oxley and HIPAA.

These government regulations for publicly-held enterprises, along with industry standards, have contributed to a doubling of storage requirements, and in some cases even more, in order to ensure the availability (and recoverability) of mission-critical and “non-mission-critical data”. In addition to the amount of storage, the length of time, or window, required to protect the enterprise from a loss of this data is also expanding, geometrically. *Replication* of an enterprise’s entire data set, of which mirroring is a good example, on the other hand, can be accomplished quickly through the mirroring of all of its files and databases. This RAID1 approach to information protection, however, can be very expensive, as it doubles total disk requirements, and furthermore, does not protect against human error, data corruption, or failure of the array, as the IT staff would simply be reproducing the same error from the primary data set. Snapshots, on the other hand, do not require anywhere near the same amount of disk capacity as mirroring because of their use of metadata and the preservation of a “before” image. However, snapshots do not protect you if the primary data set is corrupted or destroyed. In addition, snapshots create the additional problem of managing the audit trail as multiple updates occur against a shared database.

The IT staff has to change the data protection paradigm of the data center. They need to devise new processes for the backup and recovery of one of the most vital assets of the enterprise. To this end, in 2003, Advanced Digital Information Corp (ADIC) introduced us to the Pathlight VX, a Virtual Tape Library (VTL) with a disk storage array and integrated tape library. In fact, Pathlight VX was the first integrated solution to provide existing backup environments with the benefits of an integrated solution. ADIC has now introduced the second version of their VTL, Pathlight VX 2.0, with improved capacity and throughput at reduced cost. To see if Pathlight can show your data center the way, please read on.

IN THIS ISSUE

➤ Data Center Replication Requirements	2
➤ Pathlight VX Revisited	3
➤ Pathlight VX 2.0	4
➤ Conclusion	4

Data Center Replication Requirements

Sometimes information analysts are prone to overstate the obvious, such as: *Automobile accidents are the leading cause of deaths on the highway.* (Not too many people are killed on the highways by airplanes.) Likewise, consider: *The leading cause of disk storage proliferation is the replication of existing files and databases, driven by any number of application needs, from high availability, to consolidation to data protection, to archiving.*

Data replication, in fact, is a leading tool for solving business problems¹ and will be used frequently in any data center involved in Information Lifecycle Management (ILM). One of those problems deals with business continuance, ensuring access to enterprise data no matter what.

One aspect of business continuance is high availability, with failover scripts in place if any server or disk array fails. Here, mirroring a database can help to ensure that the enterprise has a no-single-point-of-failure architecture, with the same data being replicated, dynamically, multiple times across the cluster. Another aspect of business continuance involves completing a data backup within the assigned window, while at the same time, assuring that the data can be recovered expeditiously in order to protect a major operational asset of the enterprise. By replicating files and databases in active storage, the data center can ensure both a rapid backup and restore. Unfortunately, multiple replications, such as point-in-time, can quickly fill up all available disk storage capacity, an expensive resource. A disk replication, therefore, must be only viewed as an initial step in the data protection process, with export to removable media another key component.

When dealing with the backup/restore process, it is important to keep two goals in mind. The first deals with recovery time

objectives (RTO) while the second involves recovery point objectives (RPO). When the data center creates a mirrored environment, RTO is essentially zero; you have a duplicate volume on the SAN available to you. In this case, the recovery point is essentially the current status of the array, with all of its potential warts exposed. **Mirroring does not protect you from data corruption.** Also, there is no audit trail of change history as would be available if the data center were taking periodic snapshots, with the before condition being retained for a specified length of time. In addition, mirroring doubles the amount of storage required for each mirror.

Snapshots contain a stream of metadata plus a copy of the old block before it was overwritten. **This provides an excellent history of change, but it does not protect against the loss of the primary data or contribute to a rapid recovery.** What it does do is save space, with only a fraction of the original file being reproduced.

If time is the key consideration in the process, then the data center needs to consider the physical structure of the solution: local or remote. With a local solution, direct connect or SAN, **the data center has immediate access to all copies of the data.** If there is a remote site involved, then bandwidth has to be a consideration, especially if the remote environment includes tape. This is especially true if the data set spans multiple tapes or has been written to a multiplexed array of tapes. What may be satisfactory in a disaster recovery mode may not suffice in business continuance, or recovery, mode. Recovery from remote tape over an asynchronous network can be slow. Obviously it is important to have defined the differences between your business continuance objectives (very quick, file-based, multiple iterations) versus your disaster recovery objectives (complete, but okay if it takes hours). **It all comes down to resolving three issues which plague every enterprise: the effectiveness of protection (security), the system performance, and cost.**

¹ See **The Clipper Group Explorer** dated December 16, 2004, entitled *Data Copying – A Toolbox of Business Solutions* and available at <http://www.clipper.com/research/TCG2004101.pdf>.

The availability of a new generation of ATA drives has introduced a new factor into the enterprise storage equation, with the availability of high-capacity, low-cost disk storage. The cost of Fibre Channel devices has restricted their use to the primary storage tier in an ILM architecture, with the secondary tier reserved for tape, LTO, SDLT or other low-cost media, with a tape cartridge significantly less expensive. With the advent of new ATA drives, a low cost, but higher capacity alternative to Fibre Channel, the allure of disk as a secondary medium has strengthened. While ATA may have a reduced I/O level and lower reliability than Fibre Channel, thus limiting their use in high-performance, mission-critical applications, ATA has superior performance characteristics for data recovery when compared to tape, even if it is still slightly more expensive. While this may extend the purview of disk into some secondary storage solutions, the cost to keep any disk spinning will ensure the place of tape in long-term data protection solutions. **ATA does make an ideal staging platform for the transition from online storage to offline tape where high-capacity, low-cost, and long-term retention are the keys.**

A data center needs to be able to address data security, performance and cost in a single package. ADIC, with its claimed #1 ranking in worldwide library market share, did just that with their introduction of Pathlight VX in 2003.²

Pathlight VX Revisited

Pathlight VX (1.0) was introduced as a disk-based storage appliance that looked like a tape library to commonplace backup software from vendors such as Veritas and Legato. It provided disk speeds to write and read activities normally associated with tape. This includes supplying RAID 5 protection to otherwise unprotected data. Pathlight VX was designed as an integrated solution specifically for the disk-to-tape backup

² See **The Clipper Group Navigator** dated October 2, 2003, entitled *ADIC Hits a Home Run – Increasing Throughput for Scalar LTO Tape* at <http://www.clipper.com/research/TCG2003051.pdf>.

Exhibit 1 – Pathlight VX's Benefits to IT

- Faster performance – especially with regard to the restore process;
- Disk storage enables the delayed creation of backup tapes – out of the backup window;
- Includes RAID – to enhance the overall system reliability;
- Eases management – by maintaining existing processes.

environment to resolve several problems within the IT environment:

- Data grows, but time does not – Even though the amount of data needing protection was doubling and tripling, the size of the backup window was constant;
- When a file or database was lost, it took too long to retrieve and restore;
- It was difficult to build redundancy into the backup/recovery process; and
- Both backups and restores degraded the processing of mission-critical applications.

With a storage capacity of up to 40 TB and addressability for up to 400 logical tape cartridges, Pathlight VX could drive information at the rate of 1 TB/hour over 40 Fibre Channel disk drives, providing major benefits to the IT department. (See Exhibit 1, above.) Not only that, but **Pathlight VX is integrated into the existing backup process**, working with the currently available data protection applications, maintaining a familiar management process. By integrating disk as a VTL, ADIC is able to retain an automated path to the internal tape library (from ADIC or STK), through the integrated disk array. **Pathlight VX provided the data center with a complete disk-to-tape integrated backup solution.**

Pathlight VX 2.0

Pathlight VX 2.0 is a direct result of the partnership between ADIC and EMC, and continues the process initiated with Pathlight VX 1.1, with the inclusion of a *CLARiiON* ATA RAID array as the primary storage resource of the integrated appliance. With potential capacity expanded to over 2800 TBs, the data center can be assured of non-disrupted growth for years to come with upgradeable libraries from either ADIC, with its Scalar Family, or StorageTek, with the *L180* or *L700*.

Pathlight's newest version also automates the life cycle management process for existing backup processes by using embedded policy-based management and creates application-readable media for export to a remote site. With a throughput of up to 2.0 TB/hour, ADIC has doubled the backup/restore performance of Pathlight from its inception in both backup and restore modes. This integrated solution provides configuration flexibility and also continues the tradition of a single point of management established by ADIC in release 1.0, maintaining a low total-cost-of-ownership. As a disk only ATA appliance, Pathlight VX fits into a standard 19" rack and occupies 12U of vertical space, and costs only \$15.80 per GB for a 45 TB system. If the data center decides to combine a 3.8 TB disk capability into a 45 TB solution, then the cost of the integrated system is lowered to \$7.50 per GB. If the data center implements a 300 TB solution with 23.2 TB of disk, then the average cost is reduced to \$2.00/GB, **providing performance, economy, and security in a single integrated appliance.**

Conclusion

ADIC continues to lead the way in the area of integrated, open systems solutions for backup and recovery. With version 2.0 of Pathlight VX, ADIC introduces a flexible, scalable solution that increases the performance while lowering the cost in an integrated solution composed of off-the-shelf components from ADIC, EMC and STK. With a two-year warranty, 7x24 on-

site service, Pathlight VX enables the data center to balance its recovery point and recovery time objectives to ensure the security of the data within the permitted backup/restore window of opportunity, ensuring business continuity.

ADIC has developed a system which:

- **Takes advantage of the high performance of the CLARiiON disk array to double the potential backup/restore throughput of the data center while maintaining the security and integrity of the data through RAID reliability;**
- **Scales to 5.6 Petabytes of compressed capacity to protect the enterprise's investment in Pathlight;**
- **Reduces the required initial investment by 50%;**
- **Improves data protection by providing management for the complete data life cycle; and,**
- **Most importantly, saves the enterprise significant dollars by working within the framework of existing processes and applications, eliminating the need to seek out replacement packages and the required retraining of the workforce.**

If you have not already checked out Pathlight VX as the data center recovery tool for your enterprise, you should. Perhaps this is one appliance that performs as advertised for you and can save your enterprise where it really counts – on the bottom line.



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 Author

David Reine is Director, Enterprise Systems 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 an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 23. (Please dial “1-23” 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.