

# Commentary

April 28, 2006

## EMC's CLARiiON Disk Library Raises the Bar for Virtual Tape Libraries

*The virtual tape library (VTL) industry has faced an industry challenge that has — up to now — prevented IT organizations from extracting full value from their investment: How to optimize the VTL-enabled backup process and get the most out of their data restoration process. Only a full logical integration between disk and tape libraries can meet the challenge. The latest edition of the CLARiiON Disk Library delivers integration to answer the challenge.*

### Doing Great, But ...

For a great many IT organizations, the backup/restore process is under a great deal of stress and strain, such as the “running out of night” problem where backups cannot be completed during the available backup window. A virtual tape library (VTL) can reinvigorate the backup/restore process by improving process manageability and reliability. For example, eliminating media errors improves reliability. In addition, a VTL inserts a faster backup/restore process on disk in front of the traditional tape library.

As a result, the VTL market has been a very robust one. And EMC has had a strong leadership position in the VTL market since the inception of the CLARiiON Disk Library (CDL) two years ago, as evidenced by the fact that EMC has installed more than 35 petabytes (PB), across 600 customers, of CDL disks by the end of 2005.

Yet the industry has faced a challenge that has hobbled the ability of IT organizations to get all of the benefits from a VTL that they expected when they purchased it. Existing solutions force IT to make the unpleasant choice between either not relieving the stress on the current backup environment or not getting all the restore capabilities that had been expected. That failure comes from the inability to tightly integrate a VTL and a tape library into the backup/restore process.

As part of EMC's recent announcement of the 4<sup>th</sup> generation architecture for CDL, EMC provides a solid solution to this industry challenge. In addition, EMC has increased the breadth of the CDL product family, increased platform support, and introduced new functionality/features.

### What is the Challenge?

A VTL is not a replacement for a tape library; rather, it is complementary, and therein lies both the promise and the challenge.

# Commentary

## *Standing Alone Is Not the Answer*

VTL disk and a tape library can operate either standalone or integrated. Standalone means that the VTL and the tape library run independently; that is, any communication between them is over the SAN or LAN and not through a direct connection. Thus, in a typical situation, the backup server that controls writing to media adds to its task of writing to the tape library the job of writing to the VTL's disk. The backup server writes a disk copy to the VTL to speed up backups with improved reliability, and then goes off to the tape library to create a copy to go possibly offsite for disaster recovery protection.

The backup server uses the VTL copy to create the clone on the tape library. The benefit is that the physical tape is exposed to the full backup functionality of the backup software. The backup software can 1) change barcodes so that it knows what copy is being used to restore data, 2) concatenate multiple small logical tapes on disk into one physical tape (also known as tape stacking or virtual tape), and 3) provide health management and fault notification services.

All that functionality comes at a cost. The backup software has to write a clone copy of the VTL-created copy to the tape library during the day. While that clone process does not bother application server performance, the process adds to SAN or LAN network traffic during the busiest part of the day. That results in higher networking costs (hosts and network equipment) to accommodate the extra traffic. And that is unacceptable to many organizations.

## *Integration Is the Answer, but ...*

The answer is to integrate the VTL and a tape library. In this architecture, the backup server writes to the VTL and the VTL writes to the tape library whenever it chooses. This avoids overburdening the backup server with double duty, and allows the VTL to fine-tune dynamically what goes on disk and what on tape. IT can now take advantage of the disk solution (at the front end) as planned and transfer the job of copy-to-tape to a non-blocking, on-demand solution that is not concerned with or worried about network traffic.

The complicating factor is in a "vanilla" integrated solution, the backup server (a.k.a. master server) did not create the tape copy and is not aware of its creation, so therefore it does not have the information on the tape copy in its catalog. The master server needs that information to determine which media must be mounted and read in order to perform a restore. Moreover, the master server needs to be able to distinguish which of two identical copies, both having the same "barcode", is the valid one.

But because of the concerns over the ability to restore data properly when needed, many enterprises have been forced to use the less desirable standalone solution.

## **CDL Meets the Challenge**

The solution is to couple the backup software and the VTL software in a fully integrated solution (through Storage Node software), which is the approach that EMC has taken with CDL (Figure 1).

The goal is to preserve the use of the full functionality of the backup software while at the same time directly integrating the VTL and the tape library to

eliminate any unnecessary network traffic.

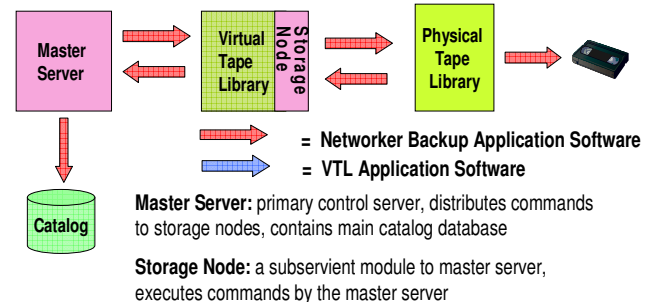
The basic principle in VTL-tape library integration is that the backup application — through the master server — needs to have full control and visibility over the entire backup environment. The backup application can do this by embedding a storage node or media server, which is software subservient to commands issued by the master server, on the VTL server. The storage module “sees” all the virtual components of the VTL. Thus the storage node software can be requested to write to tape and inform the master server of what is happening so that the master server can maintain the proper information in its catalog.

This solution makes the VTL and tape library finally complementary and truly integrated for the first time. No workarounds are required. This integrated solution provides both backup-server burden and backup/restore-process stress relief compared to a “vanilla” standalone solution, and allows the data restoration flexibility that a partially-integrated solution can not deliver.

### *Strength to Meet the Challenge*

Why was EMC able to solve this problem? By and large, the vendors who offer the leading VTL products and those that offer the leading backup software solutions are not one and the same. EMC is a clear exception. EMC owns one of the most popular backup software solutions — Networker — and has great influence over its chosen VTL software solution. EMC therefore could coordinate the development process effectively.

**Figure 1: The CDL Integrated VTL Solution**



Source: EMC April 2006

However, now that EMC has demonstrated the viability of this approach with Networker, it has announced that it is in the process of working with other leading backup/restore products to include them in CDL.

### **At a Distance**

Typically, the integrated approach involves both disk and tape backup. However, the solution can also be disk to disk, i.e. one CDL can replicate to another CDL. That is really important for disaster recovery where a local CDL can communicate to a remote CDL. The disk copy at the remote site can now be used to restore in conjunction with a remote copy of the master server and its catalog.

CDL to CDL replication over IP can now be performed by the embedded storage node. This provides greater functionality and control over distance solutions. Before when the VTL managed the replication, it did not have the capabilities that the storage node has as part of the backup software process.

# Commentary

## Not a Shred of Evidence

Of course, backup application integration and CDL replication at a distance are not the only new capabilities that the CLARiiON Disk Library family offers. Tape shredding (logical data, not physical tape, destruction, since the shredding is carried out on disk that continues to be used) is a security feature that enables data erasure compliant with Department of Defense 5220.22-M standards. Virtual tape shredding allows IT to maintain the best practices that are set by the firm's security czar as when data is properly deleted from tape the same data can be deleted from tape.

## Little Brother Added to CDL Family

The CLARiiON DL210 was recently introduced as the entry-level disk library in the CDL to complement the high-end DL740 as well as the middle-of-the-road DL710 and DL720. The DL210 scales from 4 TB to 24 TB usable capacity with the ability to support 500 GB SATA II drives, but offers the same software functionality (such as remote replication and compression) and performance (up to 380 MB per second) as its kin. Those capabilities should be attractive to midrange data centers as well as remote/branch offices that still need to maintain backup data locally.

## Dotting the i's

EMC states that the CDL is one of the first VTL to support IBM's Series i platform (formerly known as the iSeries, and before that as the AS/400). All CDL models support the Series i platform. CDL supports both the native backup software and BRMS backup applications to which Series i users are accustomed, as well as supporting IBM tape drives and tape libraries.

Even IT organizations that do not use Series i platforms should be pleased with this new capability, as it shows EMC's commitment to the breadth and depth of platform coverage.

## Conclusion

The latest edition of EMC's CDL ends the need to choose between a network-burdening standalone or non-guaranteed restore partial integration VTL solution. Instead, the CDL offers a full logical integration between a CDL and a tape library.

Enterprises that are considering a VTL for the first time or are thinking of taking further advantage of the CDL that they already have may want to give the CLARiiON Disk Library family of products greater attention.

---

**David Hill**

---

Analyst Name: David Hill  
Topic Area: Data Protection

Mesabi Group LLC  
26 Country Lane  
Westwood, MA 02090  
[www.mesabigroup.com](http://www.mesabigroup.com)

*Mesabi Group LLC is an affiliate of Valley View Ventures that aims to provide thought leadership and sound advice to both vendors and users of information technology. .*

Phone: (781) 326-0038  
email the author: [davidhill@mesabigroup.com](mailto:davidhill@mesabigroup.com)

---

The information contained in this publication has been obtained from sources Mesabi Group LLC believes to be reliable, but is not warranted by Mesabi Group LLC. Commentary opinions reflect the analyst's judgment at the time and are subject to change without notice. Unless otherwise noted, the entire contents of this publication are copyrighted by Mesabi Group LLC, and may not be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written consent by Mesabi Group LLC