



EMC Intros New Disk Library Models, Plus Deduplication and Spin-Down for Efficiency

Analyst: Michael Fisch

Management Summary

EMC recently announced two new models for its *Disk Library* family, the *DL3D 1500* and *DL3D 3000*, which provide disk backup over a LAN. EMC also introduced new Disk Library features for improving storage efficiency and lowering costs for its *DL4000* series – data deduplication, drive spin-down, and low-power drives. Read on for details.

Disk Backup Is Faster

Disk is significantly faster than tape for backups and recoveries. Arising from the universal trend of continuous data growth, enterprises are stressed to perform backups within the allotted windows and to recover data in a timeframe acceptable to the business. Since disk backup speeds up the process both ways, enterprises are turning to it en masse, either as a replacement for tape or as an initial staging point before moving data to tape for long-term archiving.

Virtual tape libraries (VTLs), also called *disk libraries*, are a popular disk backup solution because they slide right into existing backup environments. They are disk platforms that emulate tape, so the backup software thinks it is still moving data directly to tape, though at a much faster rate. Storage management tasks, such as volume provisioning, RAID layouts, spin-down, and data deduplication, become transparent to the backup operations. Nothing has to change in the environment or backup processes, which probably took significant work and time to develop. Disk libraries make it easy to switch to disk backup.

At the same time, disks consume more energy and space than tapes. They have magnetic platters that continuously spin, while tapes only spool when being read or written to. To mitigate this difference in resource consumption and lower costs, EMC has introduced three new features into its Disk Library family: deduplication, spin-down and lower-power drives. These allow Disk Libraries to store data more efficiently by reducing the number of disks required and consuming less power.

EMC DL3D 1500 and DL3D 3000

EMC announced two new Disk Libraries, the *DL3D 1500* and *DL3D 3000*, that provide disk backup over IP LANs. As you can see in *Exhibit I* below, these platforms offer 6 and 8 Ethernet ports at 1 Gbit/s, respectively, for NFS and CIFS connectivity over an IP LAN. The maximum rate of backup is 720 GB/hour and 1.44 TB/hour, respectively. These platforms primarily target enterprises without Fibre Channel SANs or whose backup application is already configured over an IP LAN. However, they also can connect to a Fibre Channel SAN.

IN THIS ISSUE

➤ Disk Backup Is Faster	1
➤ EMC DL3D 1500 and DL3D 3000.....	1
➤ EMC DL3D 4000	2
➤ Conclusion	2

Exhibit 1: Specifications for New EMC Disk Library Models

	EMC DL3D 1500	EMC DL3D 3000
Ethernet ports at 1 Gbit/s	6	8
Fibre Channel ports at 4 Gbit/s	2	4
Max backup rate	720 GB/hour	1.44 GB/hour
Storage platform	CLARiiON CX3-10	CLARiiON CX3-40
Usable capacity	4 to 36 TB	8 to 148 TB
Supported drive	1 TB SATA, 7200 RPM	1 TB SATA, 7200 RPM
Disk Protection	RAID 6	RAID 6

Source: EMC

The Disk Libraries are based on *Clariion CX3* storage arrays, for which EMC reports 99.999% availability, due to their exceptionally robust architecture. The *DL3D 1500* scales to 36 TB of usable capacity, while the *DL3D 3000* scales to 148 TBs. They use 1 TB SATA drives in a RAID 6 configuration, which can sustain two simultaneous drive failures without data loss.

Other advanced features include deduplication and replication. By using sub-file, variable-length deduplication, the Disk Libraries typically can reduce the amount of backup data stored over time by a factor of 20:1 or more. This feature employs a sophisticated algorithm to find redundancies in data, remove them, and use an index to point to the original data stored. Additional compression algorithms can further increase efficiencies. Deduplication may be performed in real time to maximize storage efficiency or post-processing to maximize backup performance, based on user policy per virtual tape library or NAS share. The benefit of deduplication is a massive reduction in the amount of storage required for backup, which saves money and/or permits a more aggressive backup schedule.

Disk Libraries can also replicate data asynchronously over an IP link for remote archiving and disaster recovery. Since it replicates the deduplicated data, it can operate on a much smaller and lower-cost WAN link.

EMC DL3D 4000

EMC will release in July 2008 the *Disk Library 3D 4000*, a deduplication option for new or existing *DL 4000* series platforms. EMC will also provide spin-down as a no-charge software upgrade for DL 4000 plus support for a low-power 1 TB SATA II drive. The DL 4000 is a larger, SAN-based platform that scales to 674 usable TBs with a maximum backup rate of 8 TBs per hour.

The deduplication capability is the same one offered in the DL3D 1500 and DL3D 3000. As

target-based deduplication, it does not require altering the backup software and greatly reduces the amount of data that needs to be stored or sent over a WAN connection, if replication is employed.

Spin-down is the storage equivalent of turning down the lights when no one is home. When the drives in a particular LUN are not accessed for a user-defined period of time, the Disk Library will place the drives to sleep¹ to save energy and reduce heat dissipation. If there is a read request at a later point, the system will turn them back on. This is a great feature for disk backup because backup data is stored but accessed infrequently, if at all, so there is a high likelihood the spin-down feature would be activated. Unlike some competitive systems, the drives can all be spinning at the same time for fast access to all data.

Finally, DL 4000 now supports a large, 1 TB drive that rotates at 5400 RPMs, instead of 7200 RPM. Like a car engine running at a lower RPM, this drive consumes less energy. EMC claims that switching from the 7200 to 5400 RPM drive and employing spin-down can reduce energy consumption (and thus electricity bills) by 47%.

Conclusion

Storage efficiency has become a major emphasis because it saves disk-related costs and reduces the consumption of power, cooling, and data center floor space. At the same time, many enterprises are moving to disk backup. The efficiency features that EMC has introduced in its Disk Libraries – deduplication, spin-down, and low-power drives – can help in both areas. The new 3D 1500 and 3D 3000 models also provide a lower entry level for LAN backup to disk.



¹ Sleep means the drive platter stops spinning and the electronics are on standby.

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

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over 12 years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" 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, The Clipper Group Voyager, Clipper Notes, 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.