

EMC Celerra FileMover Optimizes File Placement in Tiered Storage

Analyst: Michael Fisch

Management Summary

Optimized file placement is not *just* archiving. It describes the broader process of selecting data from primary storage (based on policy) and migrating it to a more cost-effective secondary tier, while preserving accessibility. Enterprises are turning to automated file placement for long-term data retention and lower cost storage in response to today's heightened focus on:

- **Regulatory compliance** – Enterprises are faced with a myriad of data-retentive regulations like Sarbanes-Oxley and HIPAA.
- **Sound business practice** – In addition to regulatory compliance, a good data retention policy has always been a solid practice for operational and legal purposes.
- **Fast data growth and growing storage budgets** – The endless growth of data and the cost to store it are reasons enough to investigate lower-cost storage platforms.
- **Slow backup and restore** – Once files are migrated from primary storage, they no longer need to be part of the backup and restore process. Therefore, backups are faster, backup windows are shorter, and restores are quicker.

In response to these challenges, EMC has developed the *Celerra FileMover API* (application programming interface). This API enables automated file movement in a tiered storage environment using Celerra NAS as the primary file storage platform. EMC offers policy migration software that leverages the FileMover API, and it is working with a number of third parties (including Enigma, Arkivio, and Veritas), who have integrated their software with it.

In addition, EMC offers a service engagement that delivers a full file movement solution based on Celerra FileMover. The policy migration software and secondary storage platform(s) are for the most part the customer's choice. They can be from a third party, though EMC does offer *Legato DiskXtender* as a policy engine and *CLARiiON* with ATA, Celerra with ATA, *Centera CAS*, and ADIC tape libraries as secondary storage platforms. In any case, EMC coordinates the overall deployment with a single schedule and statement of work.

So, if you are trying to master today's file storage challenges on a Celerra NAS platform, Celerra FileMover may be the key to a solution for you. Read on for details.

IN THIS ISSUE

> Consolidation and Archiving	2
> Celerra FileMover	2
> Conclusion	3

Consolidation and Archiving

File server consolidation is the main reason for deploying enterprise NAS solutions, like the EMC *Celerra NS* and *CNS* family of systems. These platforms solve problems inherent with many smaller, distributed file servers, or NAS appliances. Such fragmented environments are usually the unintentional result of scaling storage over time as data requirements grow. Unfortunately, it becomes cumbersome to manage activities like backup and restore, updating client share mappings, and workload balancing across file servers. Capacity utilization is also generally low because it is dedicated per server and cannot be shared among file servers.

NAS consolidation addresses these problems by providing one or few highly available system(s) for file serving. Management is easier because it is concentrated on far fewer platforms. Utilization is higher for the same reason – even more so if the NAS systems involved are gateways that share storage on a SAN with other application servers (i.e., SAN/NAS consolidation). In short, **consolidation brings efficiency and cost-effectiveness to network file sharing.**

So, how does optimized file placement fit into the picture? Layering it on top of consolidated NAS may seem counter-intuitive because it means adding tiers of storage. Nevertheless, **file placement can increase efficiency by intelligently placing data in optimal tiers of performance and price, thereby lowering the average cost of storage overall.**

For instance, **enterprises can use lower-cost storage (such as ATA disk arrays or tape systems) to store data of less utility at a fraction of the cost of high-performance storage.** Storing this data at the next more-economical tier can help enterprises meet their requirements for data retention, regulatory compliance, and speedier backups and restores. **The key to making file placement practical and feasible is intelligent software that automatically classifies and migrates data based on policy.**

EMC Celerra At A Glance

EMC's Celerra line includes the NS Series (Integrated and Gateway), and CNS. The NS Series/Integrated is mid-range NAS device with one or two Data Movers and scales to 32 TB raw. The Data Movers are specialized file servers, two of which are needed for a high-availability configuration with failover. The integrated version comes packaged with CLARiiON storage, while the NS Series/Gateway connects to either CLARiiON or Symmetrix via a SAN.

The CNS is a high-end NAS gateway that connects to both Symmetrix DMX and CLARiiON CX arrays. It has two to 14 Data Movers and scales to a massive 208 TB.

All Celerra products support software features for point-in-time copy, data movement, and anti-virus protection. The FileMover API lets software vendors tap into Celerra's ability to move files non-disruptively between tiers based on policy.

In short, file consolidation, movement, and optimization are all part of the economics of storage. Continuing to meet the increasingly heightened requirements for data management is the key to successfully driving down costs.

Celerra FileMover

For enterprises that have consolidated on an EMC Celerra NAS solution (*see box, above*) and want to optimize file placement, EMC has recently introduced its *Celerra FileMover* capabilities. The foundation is the Celerra FileMover API (application programming interface), which allows intelligent software to non-disruptively move files between the Celerra platform and other tiers based on policy. EMC's program is flexible and lets customers choose the building blocks – policy migration software and secondary storage platforms – that make-up the solution.

- **Primary storage** – This is the production data repository that makes up the primary top tier of storage in terms of performance and availability. In this case, we are talking about Celerra NAS, which is connected to either *Symmetrix DMX* or *CLARiiON CX* disk arrays.
- **Secondary storage** – This is the tier or tiers where, based on policy, designated files are stored. It has a lower cost and performance profile than primary storage. Celerra FileMover allows the customer to choose secondary storage platforms, which may be disk, tape, or optical storage from EMC or third party vendors. Of course, EMC also offers *Celerra CAS* (content-addressable storage) and *CLARiiON CX* or Celerra with ATA drives¹ for online archiving. It also resells tape libraries from ADIC for nearline and offline storage.
- **Policy Migration Software** – This is the intelligence that classifies and migrates files based on policy. The customer may also choose this software, so it can be from EMC or a third-party vendor. In this area, EMC offers *Legato DiskXtender*, a software package that features non-disruptive migration and transparent access to stored files.

Through a service engagement, EMC can program-manage the rollout of an automated file movement solution, including design, implementation, and policy establishment. It gives customers a single schedule and statement of work, even if third-party vendors are involved to implement their portions of the solution. Technical support is a coordinated effort with EMC and any third-party vendors.


The program's flexibility and scope are particularly suited to enterprises with larger, more sophisticated environments. It lets

them take a full step into file movement optimization with the assistance of EMC's expertise in storage and data management.

Conclusion

Celerra FileMover is another incremental example of EMC delivering on its strategy for Information Lifecycle Management (ILM)². ILM seeks to enable enterprises to extract the most value from information at the lowest cost by taking advantage of how the value of information changes through its lifecycle. File movement optimization does this by migrating files among different storage tiers based on policy. It lets enterprises meet business requirements for data retention and improved availability as well as lower overall storage costs.

So, if you want to leverage NAS into an ILM strategy with tiered storage, EMC's Celerra NAS platforms coupled with the Celerra FileMover capabilities offer a means to accomplish it.



¹ Note that CLARiiON CX simultaneously supports ATA and Fibre Channel drives in one array. Therefore, it could provide both primary and secondary storage in one consolidated platform. See *EMC Brings Cost Efficiency to Enterprise Storage – CLARiiON with ATA Plus SAN Copy* in **The Clipper Group Navigator** dated March 12, 2003, at <http://www.clipper.com/research/TCG2003009.pdf>.

² See *The Top 10 Things You Should Know About Information Lifecycle Management* in **The Clipper Group Explorer** dated May 11, 2004, at <http://www.clipper.com/research/TCG2004021R.pdf>.

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 eight 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. 25. (Please dial "1-25" 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.