



NetApp's Multiple Angles on Storage Efficiency

Analyst: Michael Fisch

Management Summary

The IT industry's emphasis today on data storage efficiency and disk drive conservation could make one think there is a shortage of drives in the market. Not so – drives are abundant and plentiful and enterprises can have as many of them as they like. This is reassuring, as enterprises struggle to manage the rapid growth of data. The real challenge is the scarcity of several storage-related resources.

- **Power** – Energy costs are rising, and it is costly to power and cool racks of spinning disks year after year. In some metropolitan areas, power grids and/or local powers stations (sometimes within the building) are also constrained and cannot deliver more electricity.
- **Space** – Some data centers are literally full of equipment. To cope with growth, enterprises must expand, build another data center, or consolidate servers and storage. The last choice is almost always the most cost-effective.
- **IT personnel** – Skilled storage administrators are not as abundant as disk drives, and you need them to manage your storage.
- **IT budgets** – The end of the dot com boom and Y2K brought the end of lax and generous IT budgets. Tighter spending limits mean IT departments have to manage growth more economically.

So, data storage efficiency is really about conserving energy, space, people, and money. These resources are all limited and affect a business' profitability.

The enterprise storage vendor NetApp has addressed efficiency by building multiple layers of software technology into its storage platforms, *each of which delivers efficiency from a different angle:*

- **SATA drives** – Low-cost drives for applications that do not require top performance
- **RAID-DP** – Protects against two drive failures in a RAID group
- **Snapshot and FlexClone copies** – Space-saving point-in-time copies, readable and writable
- **Thin provisioning (i.e., FlexVol)** – Dramatically raises capacity utilization
- **Data deduplication** – Store less data by eliminating redundancy

The effect of these technologies adds or compounds, allowing NetApp storage platforms to reach a high degree of efficiency when all or several technologies are employed. These features are part of the *Data ONTAP* operating system, the common software foundation for all NetApp *FAS*, *V-Series*, and *NearStore* storage platforms. In other words, they are built in instead of bolted on. Read on for details about NetApp's storage efficiency technologies and what they can mean for your business.

NetApp Storage Efficiency Technologies

As in total quality management and other approaches for lean manufacturing, the way to achieve a major leap forward in quality – or, in this case, data storage efficiency – is through a series of incremental improvements. That is, several steps forward add up to one big leap. The following NetApp technologies each deliver a step forward in data storage efficiency.

SATA Disk Drives

SATA drives are a low-cost alternative to Fibre Channel drives. They deliver higher capacities at a significantly lower price, with a tradeoff of lower performance and robustness.

IN THIS ISSUE

➤ NetApp Storage Efficiency Technologies	1
➤ Conclusion	2

NetApp supports SATA and Fibre Channel drives in the same system, so enterprises can choose between lowest cost or best performance for each application. Backup and archive data, test and development, and select production applications that do not need top performance can be placed on less-expensive, higher-capacity SATA drives, which saves money on hardware expenditures as well as power because fewer drives are needed.

RAID-DP

RAID-DP is the key enabling technology for SATA drive support. It mitigates the risk associated with the slightly higher failure rate of SATA drives, as compared to Fibre Channel, as well as their longer rebuild times (i.e., multiple hours) due to higher drive capacities. RAID-DP or double parity protects against the possibility of a second drive failure during a rebuild by allowing for two drives failures in a RAID group, while still storing data 46% more efficiently than RAID 10 or mirroring. This also saves money on hardware.

SnapShot and FlexClone Copies

SnapShot and *FlexClone* are software features that create space-saving copies of data. Also known as differential copies, they conserve disk space by referencing the original data and copying only changed blocks. Depending on the number of copies created and the rate of data change, the savings and efficiency gains can be very large.

Data copies are useful for:

- **Data protection** – Create periodic copies for fast recovery to previous points in time
- **Non-disruptive backup** – Use a copy as source for backup to disk or tape
- **Test and development** – Use copies of production data for test and development

Snapshots are read-only, while FlexClones are readable and writable. The latter is ideal for test and development. It is worth noting that NetApp SnapShots can run with virtually no impact on system performance because of the architecture of its software and *WAFL* file system. Therefore, it becomes practical to run SnapShots as a normal course of operations.

Thin Provisioning

Thin provisioning, or *FlexVol* in NetApp parlance, addresses the common problem of allocated but unused capacity. IT administrators traditionally allocate storage capacity to applications in large chunks in anticipation of future growth. This avoids out-of-space conditions, application downtime, and the tedious process of continually adding more capacity. However, the overhead capacity remains unused, like a vacant building, and enterprises pay a high price for this purchased but effectively idle storage.

Thin provisioning is a clever way around the problem. It is a type of virtualization that allows administrators to allocate as much logical capacity as needed, but real physical capacity is drawn automatically in increments from a common pool as applications write data and fills capacity in the storage array. So, the applications think they have generous overhead capacity, but they actually receive it in small dispensations as needed. The net result is a much higher utilization of storage capacity, saving drives, power, floor space, and IT management time.

Data Deduplication

Why store the same data twice? Like wringing water from a wet towel, NetApp's deduplication technology checks incoming data for redundancy with existing data stored on the system. It removes duplicate data and inserts a pointer to the original in its place. NetApp deduplication is general purpose, meaning that it works for primary, archive, and backup data. Backup data over time typically experiences a 20:1 data reduction, while production data falls in the range of 25% to 55% reduction. In both cases, though especially for backup, it reduces the required drive capacity, power, floor space, and management time.

Conclusion

If you want to improve data storage efficiency, consider the storage vendor NetApp because it has a solid offering in this area. Its storage efficiency technologies are both comprehensive and easy to use. NetApp offers multiple software technologies, each wringing out additional efficiency and lowering the ratio of raw capacity per actual stored data, thus saving power, floor space, IT personnel and budgets. You can use them in full or in part. NetApp is also developing additional technologies (e.g., compression and power saving) planned for a future release.

Their ease of use comes from software design and integration into the Data ONTAP operating system, which is available across the NetApp *FAS*, *V-Series*, and *NearStore* storage platforms. No additional appliances, gadgets, or management tools are needed.

The time to be clever about storage efficiency is now, because brute forces approaches for handling data growth have become too expensive and untenable. Call it cost cutting or greening the data center or simplifying and consolidating your IT infrastructure. Any way you look at it, data storage efficiency is good for the business.



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.