THE CLIPPER GROUP
Voyager™

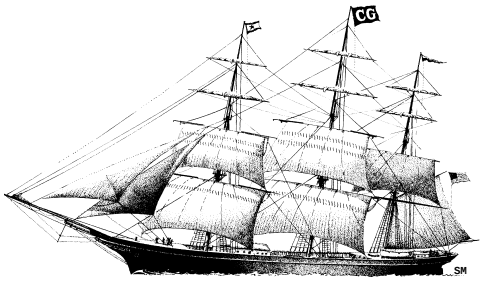

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2005046 July 28, 2005

Hitachi Data Systems Fine-Tunes Data Vastness For the Smaller Enterprise

Analyst: Anne MacFarland

Management Summary

Even smaller enterprises are faced with many dimensions of vastness in their data environments. Virtualization can control the too-many-ness of data bulk but it, alone, cannot address all the ramifications of *vast*. These days, cost-effective storage comes as *very-large capacity disks* that, like huge boxes of laundry detergent, can be unwieldy. The problems of data availability and seek times could be addressed by using smaller disks, but the premium paid in capital, environmental, and management costs would be extreme. The option of doing all the necessary data access optimization in server cache or on solid-state disk is even more expensive. The answer is to embrace very large disks, to virtualize pervasively, and then to address the problem of latency, as reflected in data availability and access times, by *tuning.*

With very-large disks, the traditional unit of failure (the disk drive) is the gateway to much more data, and rebuild time with traditional forms of RAID is longer – usually too long for business process tolerance. **A new kind of RAID, double-parity RAID, addresses this problem of bigness.** In addition, with more GB per array controller, **effective use of cache on the controller becomes an important key to assuring prompt responses to application requests for data.** When array caches were constrained to megabytes, cache concerns were focused on loading and flushing algorithms. Now that array caches have grown gigabyte-large, it becomes necessary to give applications a focused area in which to look for data – a *cache partition.*

Consider how you are using your compute power. *Do you need to do more modeling of scenarios and real-time analysis of operations? Are you streaming large data files and supporting the communication that ties an organization together? Does your business live or die by its databases?* Many enterprises would answer *all of the above*. Their assets must work double-time to support their enterprise. So, with today's big disk capacities, there is a need to double-tune production storage.

Hitachi Data Systems' new mid-range storage offerings give many enterprises the advantage of a virtualized architecture and two tools that once would have been limited only to high-end enterprise arrays. *Double-parity RAID* fosters better data availability. *Controller cache architecture and partitioning* give optimized data access for multiple applications, allowing an administrator to create and enforce different qualities of service to meet the needs of different applications. If you are tired of squandering your budget on arrays of inadequate functionality, read on.

IN THIS ISSUE

Tuning for Availability - Double Parity RAID

RAID is a basic tool for ensuring resilience against read errors and disk failure in data storage[1]. It protects the data by writing it redundantly, which takes a little longer. The simplest form of RAID is mirroring, which doubles the storage capacities required for a given amount of data, but adds little performance penalty to a particular write, as the write instruction is simply sent simultaneously to different disks. In parity RAID, the data is stored once, usually across many drives. Then information about it, called a parity bit, is stored separately in a parity stripe of data on a different disk. The parity calculation makes writes take a little longer, but recovery from a read error is almost instantaneous. During recovery from a disk drive failure, the parity data allows information from a failed disk to be rebuilt from the information remaining on the other drives.

When drives were smaller and the pace of business was slower, single-parity RAID did the job fine. However, as drive capacities have increased, the time to rebuild a drive within a single-parity RAID scheme has grown to an unacceptable length, particularly when it involves slower SATA drives. What the production environment needs is a more resilient, faster-rebuilding form of RAID Therefore, many vendors, Hitachi Data Systems among them, have turned to forms of double-parity RAID.

With double-parity RAID, you have two sets of parity information, calculated by two separate algorithms. Deriving the information to solve for two unknowns takes more I/Os, but with the clock speeds of modern processors, this is a reasonable strategy. **Double-parity RAID allows data to be rebuilt if two drives fail. If only one drive has failed, the rebuild can be done in half the time.** Thus, the use of double parity covers both the more usual (one failed drive) and worst-case scenario (two failed drives).

The two parity bits of double parity RAID form a belt-and-suspenders kind of solution. Moreover, this is not the end of this trend. RAID gurus are already developing higher-parity forms of RAID, known collectively as *RAID N+K*, to further enhance how data can be protected and rebuilt.

Cache Partitioning - Tuning to the Tempo of the Application

The vast data stored on disks of large capacity demands an appropriately large controller cache to enhance access to information. How you tune your storage array with cache partitioning depends on the way data is used by your applications.

The databases that underlie operations of almost all organizations work by many small reads and small writes. Quick access is imperative to support the complex transactions that support e-business.

How Hitachi Data Systems' Cache Partitions Work

Cache management is critical for providing quick access to data held in large databases. Large cache would seem to accelerate data access, but that is not necessarily the case. It works this way. Data in controller cache is processed in standard-size 16 KB chunks, no matter what its inherent size. If your data is small, say 4 KB or 8 KB (as may be true with atomic database values), dividing the cache makes the size of the isolated partition more congruent with the data size. This can increase hit rates, like catching fish in a bucket as opposed to a pond.

Exhibit 1 -
Cache Management Strategies

Workload	*Cache Strategy*
Real-time Analysis and Modeling (a.k.a. High Performance Computing)	Use large partitions to match large data stripes. Turn off cache mirror.
Media Streaming	Turn off cache mirror to improve throughput
Database	Align cache structure to data structure

[1] RAID is the most basic form of data protection. It is different from and complemented by data replication, and the consequent ability to failover to a known good copy of data and rebuild from logs (a more involved process that is getting more timely with continuous data protection)and the still less timely traditional backup, that is great for disaster recovery where an entire application environment must be recovered.

In the controller cache on Hitachi Data Systems' new arrays (described shortly), there is a 16 KB master partition that houses the cache management capabilities. Other partitions can range from 4 KB to 512 KB, depending on the nature of the application that is using them. With the separation of the control partition from the data partitions, the controller can manage more storage and the requests of more applications (i.e., scale better) than if the partitions were general-purpose. With cache partitioning management on the controller, the cache in the controller can be partitioned by application, ensuring that applications get what they need. A lower-priority application can be prevented from monopolizing all the cache that it would like to use. For each partition, the cache mirror can be turned on and off. Moreover, the partition sizes of all the partitions, except the master partition, are flexible.

Thus, **administrators can match the cache partition to a particular stripe of data. A large capacity stripe would be best served by a large capacity cache partition.** Training in the niceties of cache partition management is covered in HDS Academy training, which is available to both customers and to Hitachi Data Systems' distribution channel.

The Argument for Higher-Performing, More-Unified Storage

Most enterprises today depend on the information stored in their databases and files. The competitiveness of their business rests on timely access to information, to accelerate the response to opportunities, demands and threats, and the removal of waste and its costs from the business. There are multiple balance points between performance and parsimony. These balance points form the basis for the practice of Data or Information Lifecycle Management (DLM or ILM)[2].

Over time, usage of any particular piece of data in the database usually diminishes. It will still be used for analysis and research, but the rapidity of response may be less urgent. DLM and ILM respond to the aging of data by moving the data to less costly, less slower-performing storage. That said, as enterprise focus moves from quick-time response to real-time response, and analytics become a part of workflows, the need for ASAP access to certain kinds of data has grown. Instead of waiting days to have the data extracted, transformed, and loaded (ETL) into a data warehouse for analysis, enterprises are accessing the data where it lies, transforming it on the fly, and providing analytical results to their sales folks to optimize customer-facing opportunities. A unified storage environment with good management of cache partitions becomes important.

The use of file data has also been affected by the need for real-time response to customers. Database-accessing sales opportunities are often enhanced, these days, by presentations and other materials that are accessed as streamed files, and by modeling, which is often a form of high-performance computing. In such situations, the throughput is the key – something that can be greatly enhanced by the ability to turn off cache mirroring.

The New Hitachi Data Systems Storage Products

It is important, therefore, to consider tuning capabilities as well as price and capacity when considering mid-range storage. Hitachi Data Systems' new *Network Storage Controller (NSC55)* is a smaller, rackmounted version of the *TagmaStore Universal Storage Platform* product family, featuring TagmaStore's crossbar switch architecture[3], Fibre Channel internal storage and the back end ports to attach and manage other storage arrays as well (up to 16 PB). With the ability to manage external storage, this product can act as the unifier of disparate legacy equipment behind a common front end. By the end of 2005, the NSC will have up to 72 TB of internal storage, 64 GB cache, and FICON mainframe connectivity (16 ports), as well. Of more importance to tuning, the NSC55 features 73, 146, and 300 GB drives. The latter give the most internal capacity.

[2] For more on Hitachi Data Systems' approach to tiered storage, see **The Clipper Group Navigator** dated April 12, 2005, entitled *Hitachi's Tiered Storage Manager Optimizes Storage Provisioning*, and available at http://www.clipper.com/research/TCG2005019.pdf.

[3] The massively parallel crossbar switch, with its massive I/O throughput, was first introduced in the Lightning 9900 in 2000. At that time, Hitachi revamped its storage management capabilities to include the time-stamping, disk-based journaling, transparent data migration, port virtualization, and virtualization of third party storage that enabled an array that could meet the demands of the most imperious enterprise. In TagmaStore, these capabilities were re-launched in an array that could front-end third-party storage, sweeping an entire environment behind a single, robust, point of control.

Its new *Adaptable Modular Storage* product comes in two models, the *AMS 500* (with a raw capacity of 88.5 TB) and the *AMS 200* (to 40.5 TB raw), which can be upgraded to the 500. The AMS can support drawers of either FC or SATA, with one drawer of FC needed for cache flushing. The AMS200 offers from 1 to 4 GB of cache, while the AMS500 offers or 2 to 8 GB of cache. The AMS 500 has 4 Gb/s connectivity; the AMS 200 has only 2 Gb/s. AMS supports the same Fibre Channel drive capacities as the NSC (listed above), and SATA drive capacities of 250 and 400 GB. It does not manage externally attached storage, but its disk drive ambidexterity allows it to support data lifecycle management in a single product.

The new *Workgroup Modular Storage product (WMS100)* features both 250 or 400 GB SATA-1 drives. RAID 0 is not recommended for this SATA-only storage environment. The cache is 512MB to 2 GB. This product is targeted at branches and remote offices. This junior end of the Hitachi Data Systems mid-range product line can participate fully as part of a larger, remotely managed Hitachi Data Systems storage environment.

As well as storage virtualization on the controller, all products will feature embedded NAS, native disk-based WORM, and the usual assortment of traditional RAID (0,1+0,1,5). iSCSI will be supported on all by early 2006. The projected street price for these products extends from $20,000 for the WMS to $125,000 for the NSC55 with 5 TB of internal storage.

Conclusion

Why are these capabilities important to smaller enterprises? It all comes down to the variety of uses an enterprise now makes of its data, described on page one, and the need, in the face of competition from behemoths and others in low-cost business environments, to do every single part of "it all" well. It all comes down to a matter of money and time, with time being even more implacable than financial constraints.

Virtualization allows management to be done on a vaster scale. With the availability given by double parity RAID, using large disks saves floor space, and using SATA saves up-front costs and environmental costs as well. Using large, partitionable controller cache can expedite business processes. In conjunction with the virtualization, NAS and WORM capabilities, Hitachi Data Systems' fine tuning offers mid-range enterprises the capabilities they need to better meet the too-many, too-urgent demands of their enterprise.

If more data capacity has brought more headaches, consider the relief that more capabilities, particularly those tuned to the demands of large disk drives, can bring.

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

Anne MacFarland **is Director of Infrastructure Architectures and Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" 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**, 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.