



## Hitachi Reinvigorates Enterprise Storage with the HDS *Virtual Storage Platform*

Analyst: Anne MacFarland

### Management Summary

The Internet has expanded the horizons of all businesses. There are more sources from which to procure, and more markets that can be addressed. This hugeness can be overwhelming – but it is also a grand opportunity. For large businesses, a similarly large opportunity – to respond to market forces and opportunities sooner and better – lies in the new richness and variety of relevant-to-corporate data, the democratization of its use, and the popularity of analytics.

The challenge of the variety and use of data is different from previous data challenges – those met by deftly-timed batch jobs and workflows, where faster was usually the answer. Now, cost-effective delivery of precisely the appropriate information to the appropriate users (applications and human audiences), while enforcing security and privacy, is a high priority. *Faster* is still a part of the solution, but data services have to become more capable and efficient to support a solution that meets the whole challenge.

Compare the situation to that of a professional athlete – as opposed to a weekend player. The weekend player strives for sufficiency and occasional éclat. A professional athlete – say, a tennis player – lives or dies by the precision of stroke and ball placement that he or she can achieve, and succeeds by achieving that fine precision over many games, sets, and tennis tournaments. To achieve this scale of success, it is important to be efficient – to have a full range of capabilities without expending unnecessary effort or adopting habits that could cause injuries or bad outcomes.

In the past few years, particularly for large organizations, the growth of stored data has reached a point at which storage-array-based capabilities, and their tight integration with applications and virtualized server environments, has become key to using information in a timely, strategic way. These capabilities, in addition to large capacity, have become a basic characteristic of the product category known as *Enterprise Storage*. Like the skills of a professional athlete, enterprise storage capabilities help an organization turn its operational and external data into new forms of actionable assets, such as alerts that support the organizational logistics that keep a business successful. Such ambitions demand sophistication in data protection, replication, migration, placement, and movement. The whole of this information must be managed. **Today, enterprise storage is much more than just a peripheral.**

Many enterprise arrays strive to provide this enterprise-strength functionality, but they differ in how they do so. Hitachi Data Systems has provided back-end support for multiple heterogeneous applications and servers, and front-end control for multiple heterogeneous storage arrays and media types for years. With its new *Virtual Storage Platform (VSP)*, Hitachi has now added and enhanced both greater capacity and capability. It supports more precision and efficiency in storage operations, while also reducing operational costs. For more particulars, please read on.

### IN THIS ISSUE

➤ Enhancing Many Kinds of Scale .....	2
➤ Supporting Greater Precision in Storage Operations .....	3
➤ The Hitachi Mandate for Cost Efficiency .....	4
➤ Conclusion .....	4

## Enhancing Many Kinds of Scale

### *Enhanced Scale*

With VSP, HDS has built out a network-switched architecture that includes both greater capacities and new capabilities. Higher-density 2.5” drives give more capacity (40% more TBs per square foot), more IOPS per watt (162%), and less power draw (48%) than its predecessor, the USP V. The connectivity supported has doubled, and how that connectivity can be used to support a business with volatile operations has grown more sophisticated. (More on this, later.) It supports larger volume sizes (15 times those of USP V), which, as with the predecessor, can be striped and opportunistically moved across the entirety of the pooled, virtualized storage to give the I/O characteristics that are needed, when they are needed.

A switch and cache matrix of front- and back-end directors, cache adapters, and virtual storage directors (controllers) all leverage a global distributed cache to support the broad-scale yet precise data access optimization that enterprises require. It scales to over a TB of cache and up to 2048 HDDs. The increased density of this behemoth comes in 19” racks with front-to-back cooling, supporting a hot aisle/cold aisle deployment.

Hitachi VSP offers controller-based encryption (in hardware) for data at rest on any kind of media. No external key management is needed. It also supports not just LDAP and RADIUS access controls, but role-based access control differentiation among security administrators, storage administrators, and maintenance folks.

### *Scale Up from Lower*

Hitachi’s VSP can be bought as a diskless system of redundant, potent controllers, to be used with existing storage media. This provides a low point of entry that lets Hitachi’s enterprise array give moderately-sized companies the storage effectiveness that large enterprises find so valuable. It can be delivered as a service. It can be delivered with traditional internal disk drives. Moreover, as with previous generations of Hitachi enterprise storage, it can front a variety of networked storage assets.

In all scenarios, the controller virtualizes, migrates, and manages the data, using a global distributed cache. This global cache lets the system optimize data placement more comprehensively, to meet the demands of surges in data use. It can load balance across more physical devices, to meet day-to-day operational variations.

### *Scale Independently in Three Dimensions*

Because of both the matrix architecture and Hitachi’s comprehensive storage virtualization, storage performance, capacity, and connectivity can scale up *as needed* (i.e., independently) – something that most appliance-based data services cannot do. In addition, the data and control information are separated, so that I/O is not impaired by data service operations.

### *Managing Many Dimensions of Scale Out*

The Hitachi VSP supports many kinds of scale-out architectures. Four strategies give enterprise customers the ability to leverage a full array of scale out tactics, in pursuit of a competitive edge.

- **Scale out with tight coupling** is supported by the Virtual Storage Directors that share the global cache. This option supports more sophisticated forms of resilience. In this scenario, host servers can access storage dynamically from a virtualized pool.
- **Secure scale out** is achieved by the use of dynamic partitions to create the secure isolation needed in multi-tenancy environments. Dynamic partitions can preclude the escalation of privileges that might otherwise occur in virtualized servers on a single processor with different security requirements. Now, Hitachi *Resource Groups*, an evolution in Hitachi’s *Virtual Partition Manager*, support the more granular and flexible partitions demanded in cloud operations.
- **Scale out for performance** via parallelization of process is another mode, used for workloads like analytics and for processes such as Hadoop’s map-reduce paradigm.
- **Scale out for low cost**, i.e., leveraging of commodity arrays, gains sophistication by HDS’ *Scale Deep* strategy.

### *Scale Deep*

*Scale Deep* comes in both hardware and software. In hardware, it is the ability to add external storage to a virtual pool, prolonging the life of existing assets. In software, *Scale Deep* is a matter of greater precision in storage management, enabled by metadata, and by automation that allows standard data services (data placement, striping, and thin provisioning) to be achieved with greater precision than is possible with manual approaches.

This is a significant addition to the sophistication of the *Command Suite* software that aggregates the management of many tasks into a simple, more integrated whole. It enhances sev-

eral storage management strategies. Deployment can be faster with less risk, as automation means there are fewer steps to manage. Greater agility and utilization are achieved, which paves the way for enhanced precision of storage operations.

## Supporting Precision in Storage Operations

### *Dynamic Tiering – Taking Wide Striping to the Next Level*

Wide striping across disk drives has been used for some time to amplify the I/O to particular volumes of data. It mitigates the problem of hot spots. It also can reduce the time of a disk rebuild by sharing the work across many drives. With Hitachi's new *Dynamic Tiering*, sub-volume tiering is supported, and allocation is by pages, and only of those being used. Dynamic Tiering volumes are based on thin provisioning, either in the same stripe or in the virtual pool of resources – including external resources.

This sets the stage for Hitachi's three-tier optimization strategy that manages data to meet the demands of business more precisely.

- The first tier consists of *virtualized pools based on drive type or media*. The highest performance media is SSDs. Other media include SAS for working data and SATA for less active data. Array automation handles the lifecycle management – a key value for customers. Previously, a significant amount of analysis would be required to optimize what data should be moved between tiers.
- The second tier of optimization analyzes data center operations over the previous 24 hours, both for a “heat index” of data use and to construct topography of I/O rates. These analyses, a configuration option, are the basis of page movement between all three tiers<sup>1</sup>, to *optimize for operational reality*.
- Then, a third tier of analysis looks at *page movement patterns* that show, at greater detail, the patterns of operational reality over time. The analysis cycle can be from an hour to a day.

These three tiers give a precise storage-eye view of how a business uses its information. They can be used for self-diagnosing heat management of

heavily used data. They can support anticipatable events better – every organization, no matter how chaotic it seems to the participant, has patterns. Optimizing to support these patterns reduces the exceptional situations that stress the system and cost money.

### *The Hitachi Command Suite v. 7*

*Hitachi Command Suite v. 7* uses an agentless approach for host discovery and other basic management tasks (element management, performance monitoring) across both storage and server environments. A host agent is available for more advanced management options, such as performance management.

Command Suite addresses a broad range of servers – including the IBM mainframe – and provides key management capabilities for a wide variety of networked storage assets, including tape. This wide scope has huge value for large companies who have aggregated disparate assets – and for new companies who want to use what best fits their business needs.

*Hitachi Device Manager* is the base of the Hitachi Command Suite and an integral component of the *Hitachi Basic Operating System* for the VSP. The integrated *Task Manager* lets customers define common tasks, to be run immediately or scheduled, that work across a broad range of physical assets and network connectivity.

Hitachi has evolved the Command Suite to reduce the pain and costs and risks of storage management. The challenge is to keep the sophistication while supporting ease of use. It offers both a newly enhanced GUI and a streamlined command line interface (CLI) for those who prefer that mode. Common tasks have been streamlined. Enhanced workflows incorporate best practices and are presented via wizard-based interfaces. Command Suite also has streamlined the basic templates for provisioning and virtualizing, bringing inherent discipline to these tasks, which can be executed across a full breadth of physical assets. It also supports selection of specific objects for inclusion in pools as well as defining the storage that supports the pools. These make the *Tiered Storage Manager* easier to use without sacrificing its sophistication or potency. Specificity and choice are both well supported.

With VSP, Hitachi's Basic Operating System (BOS) is priced by frame, not by capacity. Once VSP and BOS have been purchased, upgrades to BOS are free under the maintenance contract.

---

<sup>1</sup> There are two granularities of data migration supported by the Hitachi array. At an infrastructure level, the mobility is one of blocks. VSP also has the capability to move data at the object or file level.

## The Hitachi Mandate for Cost Efficiency

The big news in the VSP release is the scale and precision of data service operations that it supports. However, there are other kinds of bigness, and more forms of efficiency, particularly when you take the long view over time.

While acquisition costs are important at the time of procurement, it is the operational costs that cause chronic distress and cumulative pain. Hitachi is one of those who have made increasing efficiency – particularly energy efficiency – a mandatory part of product development. These efforts are clear in the energy and capacity efficiencies mentioned earlier in this bulletin. In hard times, there are many seemingly good reasons to postpone efficiency improvements. Hitachi's corporate mandate assures that increasing cost efficiency will be a regular cadence that will continue in the future.

## Conclusion

Once upon a time, storage was a peripheral. It now is essentially the heart of the business. How well it meets the dynamic needs of the business is critical. For the most part, this is a matter of management *capabilities* as well as management *practices*. The challenge of supporting a seemingly permanent state of business volatility – and its consequently high rate of non-disruptive change while preserving performance of the core transactional operations – can only be met by excellent data services, automated by management that is comprehensive in scope, precise in achieving the desired outcomes, and parsimonious in how these results are achieved. If you crave deft control in chaotic times, consider Hitachi's Virtual Storage Platform.



### ***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](http://www.clipper.com).***

### ***About the Author***

**Anne MacFarland** is a Senior Contributing Analyst 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 in 2001 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](mailto: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, 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.

### ***Disclosures***

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

### ***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.