



A New Year Dawns for I.T. — IBM Adds New Paradigm to its Storage Architecture

Analyst: David Reine

Management Summary

In the past, we made most of our major purchasing decisions based upon traditional thinking, using the same criteria for our *next* purchase that we used for the *last* one. We purchased automobiles, for example, based upon engine capacity, trunk space, and braking distance. All are reasonable criteria and perfectly valid for our *next* purchase. However, the cost of energy has spiked, skewing all of these criteria behind economy, i.e., miles per gallon, enhancing the position of an entire new breed of car – the hybrid. Our decision making process has evolved, adding new criteria that complement the old ones. Any auto manufacturer who wants to continue in business will need to respond to this changing climate by positioning their product set to include these products. In fact, due to the nature of the automobile business, they will need to be working on them years before they reach the consumer.

In the data center, 2008 sees a similar evolution in the enterprise infrastructure, from a network of under-utilized servers and storage devices to an architecture based upon consolidation and virtualization, with operational efficiencies implemented to optimize the TCO of the data center. Server virtualizers – such as *VMware*, *Xen*, and *Virtuozzo* – have enabled the IT staff to reduce the number of servers that have proliferated across the enterprise landscape and increase *server utilization*, thus reducing administrative overhead and energy consumption and improving enterprise profitability. The unrestrained growth of data, however, continues unabated, with independent silos of information standing very tall, and under-utilized, throughout the network. The I.T. staff faces a continuing challenge from this next-generation of data to create a shared and virtualized storage environment to accommodate a new storage paradigm based upon the rapid growth of unstructured data. The capability to collect and store information is no longer the overriding challenge. **The objective, today, is to find a better, simpler, and less costly way to utilize and manage that information in the next-generation data center.** As we can see on page 2, this is exactly what IBM is doing with their recent acquisition of XIV.

Today's data centers are replete with a variety of unique storage devices whose only common trait is an inability to communicate with each other. Often based upon a proprietary architecture, these expensive arrays have limited scalability and complex provisioning requirements. Usually configured with a single communications protocol, such as Fibre Channel, SAS, or iSCSI, they require constant performance tuning, increasing the TCO to the enterprise, and they continually expose the data center to the risks resulting from device failures. What the I.T. staff requires is a flexible infrastructure, designed upon standardized, low-cost storage platforms capable of scaling out to petabytes of storage. This new-generation of data storage must be high-performance, simple to manage, capable of being fully utilized, and accessed from multiple servers using different protocols. As with any next-generation architecture, these new storage arrays must be self-healing, with autonomic tuning capabilities, and the ability to process evolving storage workloads.

Many mission-critical applications will continue to process data in 2008 exactly the same way that they did in 2007, for mainstream activities such as file services, high-performance computing (HPC), and On Line Transaction Processing (OLTP). New business- and mission-critical applications, however,

now process digital archiving, clustered computing, video streaming, and IPTV workloads, as well as medical images and digital video surveillance. These new applications, known loosely as *Web 2.0 applications*, access “blobs” of variable length data (video, voice, and image) via a non-traditional (often collaborative or self-sourcing) processing model over the Internet. The number and frequency of inputs and accesses is unpredictable. The data center must be prepared, however, to support multiple requests for the same data *in parallel*. In addition, the next-generation of web communities and hosted services, such as social networking sites, target collaboration, and data sharing among users. Storage providers must satisfy this evolving and expanding requirement; and this is where IBM appears to have set its sights with a continuation of its aggressive strategy of acquisition, this time targeting XIV, an Israeli storage vendor, and its innovative, new technology, reestablishing its significant position as a storage provider for the years to come.

Evolving the Data Center

IBM has stated that its goal is to deliver XIV's intellectual property, as an IBM-branded solution, integrated into IBM's storage portfolio, in 2008. (In the meantime, it will continue to be sold in its current incarnation, called *Nextra*¹. It is a block-based system of clustered arrays built with open system components in a scale-out architecture, with x86 processors, Ethernet switches, and SATA drives, to address emerging workloads. The Nextra solution provides virtually unlimited snapshots, I/O load balancing, and automatic configuration in a virtualized system designed for high-performance throughput, high reliability, and outstanding dynamic scalability for disk capacity, I/O interfaces, cache, CPU power, and I/O bandwidth. Nextra distributes all the data in 1MB chunks across all of the devices in the disk infrastructure, improving performance, ensuring that one or two disk devices do not have to bear the brunt of the new workload. This also increases system reliability by minimizing hot spots and enabling the seamless replacement of any component. By distributing these 1MB chunks from a virtual pool across all devices, Nextra can eliminate the overhead associated with the storing of variable length records, and the problems associated with dynamic scalability. Nextra can heal itself in the event of a failure and self-tune for best performance. As a direct result of these optimiz-

¹ Nextra now becomes one of many storage architectures supported by IBM.

ing efficiencies, Nextra reduces administrative overhead and energy consumption, reducing TCO and increasing enterprise profitability. **The data center can now access previously-allocated but unused “capacity” rather than having to order more “drives”, with the pool of storage virtualized for any application to access, simplifying both expansion and rebuild**, with the time to recreate a 750GB drive reduced to 20 minutes.

Tomorrow's data center will consist of a mix of workloads, some old and some new. The enterprise storage environment must be able to support both – file services and video imaging, HPC and digital surveillance, OLTP and Web 2.0. Being “good enough” is never good enough. IBM has had good success in recent years with their storage products. However, like the automobile, every successful product must evolve in order to stay competitive.

Conclusion

The acquisition of XIV provides IBM with a proven, innovative, next-generation storage architecture, along with an experienced development team to integrate their solution into the IBM storage portfolio. We cannot ignore the fact that as an integral part of the acquisition, IBM also acquired one of the leading individual innovators in storage system architecture in Moshe Yanai. This should provide IBM with thought leadership and an outstanding competitive advantage, as who would know IBM's main competitor in the storage business better than Mr. Yanai.

Some analysts seem to have been asking the same question recently: “Is IBM in the storage business? Seriously?” It appears that IBM has answered that question, again in the affirmative. With their existing portfolio enhanced by XIV, they are *in position* to be a significant player in enterprise storage for years to come. **We must now wait to see how IBM will deploy this new architecture within its storage product lines. How will IBM use this new architecture to unify their various storage platforms to simplify customer solutions?**

Hockey great Wayne Gretzky said that a *good* hockey player plays where the puck is; but a *great* hockey player plays where the puck is going to be. We know where IBM is playing in storage today, at or near the top of the heap in many categories. Time will tell if they can be *the great one*.



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 Authors

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- *Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" 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.