THE CLIPPER GROUP
Navigator™

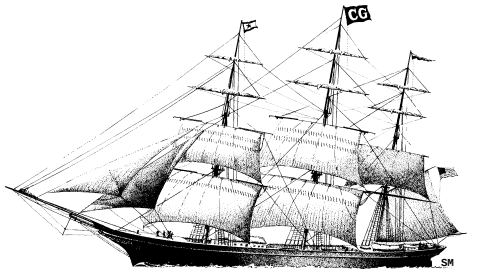

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2011028 August 5, 2011

XIV Gen 3 — IBM Lowers TCO and Raises Performance and Functionality

Analyst: David Reine

Management Summary

Lowering the cost of everything that we buy is an ambitious goal, but one well worth investigating every time that we go shopping. Whether we are talking about an everyday purchase, such as an appliance, or a major investment, such as an automobile, we are always looking for a way to reduce the cost. One way we do that is to go to a "big box" store that sells in volume and can afford to offer items at lower prices. They minimize their acquisition costs and they can also minimize their transportation costs by consolidating distribution by region, minimizing the freight costs, getting goods to market faster, wherever possible. It is also important to look at the total cost of ownership (TCO) whenever we make a purchase. How much does it cost to run the item? Energy consumption will vary from one product to another, whether we are talking about the electricity required to run a TV or the gasoline necessary to drive your car. In some cases, it may be less expensive to replace an older item that is still functional, if the cost to operate it over time is more than the replacement cost, not to mention the additional functionality that we usually get with the newer product.

Trying to lower the TCO is nothing new for the CIO of any enterprise data center. Lowering the TCO, and making the IT infrastructure more efficient, has been an on-going exercise. Virtualizing the data center servers has been but one step in improving the efficiency of the IT environment. Virtualization of the server architecture, however, has also helped to drive data storage requirements through the roof. Storage requirements in enterprise data centers have been expanding, are expanding, and will continue to expand into the foreseeable future – and there is nothing that the IT staff can do to stop it. New applications are generating more and more unstructured mission- and business-critical information that needs to be mined and acted upon by the enterprise seeking to become smarter in their IT operations. With billions of new devices such as smart phones and tablets being put into operation every year sending millions of new images and data across the Internet, many data centers must find a way to improve performance while ingesting additional petabytes of additional structured and unstructured data that needs to find a home in any of the myriad data centers that populate the Internet. These are but two examples of the manner in which enterprise applications and web use is rapidly expanding the requirements for storage. The enterprise data center has to find a way to increase storage capacity and improve performance and functionality, while at the same time, lowering the TCO for that storage and functionality.

One company that has recognized that need is IBM. IBM has just announced an upgrade to Generation 3 of their *XIV Storage System*, increasing data center efficiency and performance by increasing backplane performance by 20 times with *InfiniBand* and by refreshing the internal processing capability with the latest Westmere CPUs from Intel. To learn more about how IBM is responding to the growing storage and application needs of the enterprise data center, please read on.

IN THIS ISSUE

Virtualizing Storage in the Data Center

Today's enterprise data center is faced with a myriad of problems, as are all data centers. Trying to do more with less is a never-ending challenge. The virtualization of enterprise servers in the data center has alleviated one set of problems. The IT staff has reduced the number of servers proliferating throughout the enterprise and they have improved the efficiency of each server by deploying upwards of ten mission- and business-critical applications on each new, multi-core, multi-threaded virtualized server. Highly virtualized environments provide the enterprise with the benefits of maximum consolidation, reducing the data center footprint, increasing resiliency, along with automatic and dynamic load balancing. In fact, the IT staff has succeeded so well in consolidating servers that they have reduced the number of administrators required to manage the server network. Have no fear, however, as the virtualization of the server has increased the workload on the shared storage environment, and the storage management challenges, as well.

Storing enterprise data is not the problem. *Finding efficient and economical storage is!* Disk, and array, capacity goes up every couple of years with 1TB devices being replaced by 2TB disks, and now we see the arrival of 3TB disks making their way into the enterprise infrastructure. Unfortunately, as the total expandability of any disk array goes up, so does the amount of unused and inaccessible capacity. The data center needs to find a disk array that will scale to the highest amount of usable capacity for both structured and unstructured data, while keeping the TCO of the storage infrastructure as low as possible. The data center needs to find the highest performing storage system possible, for all types of data, in a virtualized environment. By virtualizing storage, the data center can achieve the same benefits that were achieved by virtualizing the server environment.

To satisfy the increased demand that multiple workloads place on storage, today's enterprise data center needs to improve array performance, storage efficiency, and platform reliability in order to achieve maximum application efficiency in a workload-optimized environment[1]. Therefore, the IT staff must improve storage virtualization and data protection. They must also strive to reduce storage complexity in order to help lower the TCO, while, at the same time, improving the SLA level for a wide variety of mission- and business-critical workloads. In addition, the IT staff needs a solution that is easy to use with as much automation as possible. By improving the performance of the workload-optimized storage, the data center can get more real work accomplished, and get it done faster. With a simplified management, the enterprise can get an improved time to value.

High-end enterprise storage platforms (which tend to be proprietary and cost the most) have always provided the scalability and the reliability that the enterprise data center required to maintain business continuity. These high-end platforms, such as EMC's *Symmetrix VMAX*, Hitachi's *VSP* (*Virtual Storage Platform*), and IBM's DS8000 have all of the functionality that the enterprise data center could possibly need. Unfortunately, they also carry a price tag and a TCO that the enterprise cannot afford to extend to all business- and mission-critical applications and uses. Today's IT staff is looking for a solution that takes advantage of the advances in open systems technology, commodity components and can scale-out in a modular fashion. All of these features can be found in the *IBM XIV Storage System*.

IBM XIV Generation 2

Has it really been only three-and-one-half years since IBM introduced us to the XIV architecture?[2] In September 2008, IBM acquired XIV and their innovative, high-performance storage design. This announcement provided a boost for the IT infrastructure, sharing and virtualizing the storage environment with commodity components with a low cost per TB, in order to counter the challenge of the rapid growth of unstructured data and improve storage efficiency. This challenge was not so much for the collection and storage of data as it was finding a more efficient and less costly way to utilize that data, and an easier way to manage information in the next-generation data center.

XIV was designed with a highly-parallel architecture in order to provide superior perfor-

[1] Workload optimization represents a continuing investment in systems integrated and optimized across chips, hardware, and software, targeted for a specific range of workloads, each representing a different set of requirements.

[2] See *The Clipper Group Captain's Log* dated January 15, 2008, entitled *A New Year Dawns for I.T. – IBM Adds New Paradigm to its Storage Architecture* and available at http://www.clipper.com/research/TCG2008003.pdf.

mance and a hot-spot-free environment, utilizing all resources, including disks, RAM, and CPU, without manual tuning. The second generation of XIV[3] continued to take advantage of an active-active design with N+1 redundancy of all system components to deliver the highest level of resiliency. With automatic load balancing, capacity utilization, and an intuitive graphical user interface (GUI), XIV Gen 2 helped to reduce administrative overhead and lower the TCO.

Each XIV module is an independent computer with multi-core processing, cache RAM, and disk drives. The innovative caching technique used by XIV Gen 2 enabled concurrent service by all cache units. A high cache-to-disk bandwidth enabled aggressive pre-fetching and cache management. Modules communicated with each other over an internal network. With linear scalability from 72 drives to 180 for a fully configured XIV system, it had ample capacity for your data and scaled nicely. What the IT staff required, and what XIV also provided, was a flexible infrastructure using standardized, low-cost storage building blocks, capable of scaling out to petabytes of capacity while delivering high performance and ease of management. XIV was delivered with a truly innovative and fully autonomic GUI to reduce complexity and simplify storage management. The design enabled XIV storage arrays to be fully utilized, and accessible from multiple servers using different protocols. XIV Gen 2 was introduced with autonomic tuning capabilities and was self-healing. In fact, all of the XIV software features, from snapshots, mirroring, and thin provisioning with space reclamation, to advanced reporting and virtualization plug-ins, are included with XIV at no additional charge.

XIV was designed to support the growth of large collections of block data, such as databases and special-purpose servers for Microsoft *Exchange* and *SQL Server*, as well as unstructured data that had become standard in a broad range of new workloads for business- and mission-critical applications in a variety of operating environments, including *AIX*, *Windows*, *Linux*, *HP-UX*, and *Solaris*. These included medical images, video streaming, digital surveillance, and virtualization environments such as *VMware*, *vSphere*, *Hyper-V*, and *PowerVM*, as well as Exchange and *Oracle* application workloads, and more. It is a block-based system of clustered arrays (modules), within a grid architecture, enabling seamless integration into the data center or cloud to improve performance at a reasonable price. XIV Gen 2 was built with open systems components in a scale-out architecture, with Intel *Xeon L5410* CPUs (*Harpertown*) running at 2.33 GHz, Ethernet switches, and SATA drives, built to address these emerging workloads. All XIV nodes scale in parallel. Thus, when the data center added additional XIV nodes, in addition to capacity, they also added additional processing power and cache, as well as I/O pipes.

Exhibit 1 — XIV Storage System Features

- **Parallel Architecture** – to deliver consistently high performance with automated balanced data placement across all key system resources, eliminating hot spots without manual tuning;
- **Resiliency** – providing data protection and availability through active-active N+1 redundancy of system components with rapid self-healing and fast rebuilds;
- **Ease of Use** – enables a simplified management via automated tasks and an intuitive GUI to eliminate the complexity of managing enterprise storage; and
- **Reduced TCO** – enabled by high-density, low-power disks with optimal utilization.

Source: IBM

XIV provides support for a virtually unlimited number of differential and writeable snapshots, replication and disaster recovery, I/O load balancing, and automatic configuration in a virtualized environment. It was designed for high-performance throughput, high reliability, and dynamic scaling. With 1MB chunks automatically distributed across all of the devices, XIV was able to self-tune and self-heal. (See Exhibit 1, above, for advanced XIV features.)

With advanced thin provisioning, XIV improved system utilization without affecting performance. In fact, XIV Gen 2 improved performance and reliability, as compared to the first

[3] The first generation was what XIV offered prior to being acquired by IBM. The second generation was the first offered by IBM, although IBM never referenced it as *Gen 2*, but we have chosen to do so for clarity. The third generation is what has just been announced – its next generation – what IBM calls *XIV Gen 3*.

generation, ensuring that one or two devices did not have to bear the brunt of the workload, eliminating hot spots in the infrastructure. Thus, the IT staff has been able to access previously allocated but unused capacity, rather than having to order additional drives.

XIV Gen 2 can be configured with up to 180 1TB or 2TB SATA drives for a maximum capacity of 161TB via up to (24) 4Gbps Fibre Channel (FC) ports with network connectivity via up to 6 1Gbps iSCSI ports. In addition, the data center can deploy up to 240GB of XIV memory using 16GB memory modules with 240Gbps of cache-to-disk bandwidth.

Even with this performance, capacity, throughput, and resource utilization, more is needed by many data centers. So, IBM has raised the ceiling, again, in order to reduce the TCO for the enterprise data center. They have now evolved to the next step in data center efficiency with *XIV Gen 3*.

IBM XIV Generation 3

While XIV has been characterized by some as a Web 2.0 scale-out storage platform, IBM has been investing a sizable portion of their $3B annual investment in R&D for hardware design and system software, to ensure that XIV Gen 3 is a Tier-1, full-service enterprise-class disk storage system that can match the delivery requirements that traditionally required a more expensive high-end array – in terms of performance, scalability, utilization, and energy use. It uses the same patented grid architecture as the previous models, but it has been optimized for use with mission-critical high-transaction workloads running on VMware, *Linux*, *UNIX*, and *Windows* servers in order to deliver both high-capacity and high-performance.[4] XIV Gen 3 has all of the features of the previous generation, but has been upgraded with the latest hardware technology to provide more performance and a lower TCO per terabyte.

Based upon commodity components and industry-standard interfaces, XIV Gen 3 has been upgraded by adding InfiniBand connectivity between modules[5], 8Gbps Fibre Channel ports rather than 4Gbps, a 2.4GHz, quad-core Intel *Xeon E5620* CPU (*Westmere*), expanded memory, and 2TB native SAS disks. In fact, Gen 3 systems have up to 24 FC ports (as well as 22 iSCSI ports), 60 multi-threaded Westmere cores, up to 360GB of cache[6], with 24GB of memory per module, and twice the cache bandwidth of Gen 2 at 480Gbps. At the same time, XIV Gen 3 consumes 25% less energy on a Watts per thread basis[7]. In fact, Gen 3 consumes 20% less power and generates 20% less heat than Gen 2.

With Intel's Westmere driving Gen 3, IBM gets a two-to-four times boost in performance for workloads dealing with transaction processing, sequential read and writes, and file/print services, over the previous generation. Part of this improvement can be attributed to the increased number of threads available in the E5620 (twice as many as in the L5410), part to the improved I/O capability and expanded memory, and the doubling of the cache-to-disk bandwidth, as well. In addition, similar gains can be expected in mission- and business-critical applications such as Microsoft *Exchange* and *Hyper-V*, *Oracle Data Warehouse*, and *SAS Analytics Reports*.

In addition to the standard functionality provided with all XIV systems, XIV Gen 3 has additional functionality to provide multi-system monitoring with enhanced, centralized administration, and improved replication. It also provides a plug-in for *VMware vCenter Server* to make it easier to manage XIV using VMware software. IBM has also improved the XIV interface, which already was a first-class GUI. In fact, IBM has adopted the XIV GUI for other storage platforms in their arsenal.

In the not so distant future, by 1H2012, IBM expects to have an optional feature for 500GB of SSD capacity per module, for a total of 7.5TBs in a fully-configured XIV. IBM will use the SSDs as a transparent DRAM cache to accelerate module performance, and not for an additional tier of storage to manage. IBM has also revealed that there are 15 spare controller sockets and 720 GB of spare memory space per XIV Gen3 rack, for future IBM software enhancements.

[4] IBM's DS8000 remains the preferred solution for mainframe connectivity.

[5] InfiniBand interconnect provides a 20x increase in module-to-module communication.

[6] This is a 50% increase over the previous generation.

[7] Harpertown consumes 50 watts per CPU with a maximum of 21 CPUs per XIV for a total of 1050 Watts. Westmere consumes 80 Watts per CPU with a maximum of 15 CPUs for a total of 1200 Watts. However, Westmere has double the threads resulting in a consumption of 12.5 Watts per thread for XIV and 10 Watts per thread for Gen 3.

XIV Gen 3 represents up to a three times performance increase on read/write hits (in excess of 500,000 read-hit IOPS), up to a four times increase on sequential loads (in excess of 10 GBps), and up to a three times reduction in latency (to less than 170 µsec for a single I/O). The average list price for a fully configured XIV Gen 3 is around $2M, as compared to the average list price for a Gen 2 model, which is now $1.5M to $1.6M, about a 12% reduction at list from earlier pricing. **Gen 3 delivers two to four times the performance for just 25% more, using the Gen 2 revised pricing, at list**. How does this compare to the competition? IBM has positioned XIV to be more cost-efficient than the high-end platforms mentioned earlier, in order to be more competitive in the open systems market.

XIV Gen 2 will continue to be offered in the IBM product set for at least one more year in support of customers who do not need the expanded capacity and improved performance of Gen 3.

Conclusion

IBM continues to invest in enterprise-class IT infrastructure – and invest heavily. They have proven to be committed to advances in system design, with an emphasis on smarter computing relative to the movement and management of data. They have implemented overall system improvements for higher performance and bandwidth and lower latency in support of your high-volume, transaction processing applications. They have enabled the enterprise data center to reduce backup time and boost performance for Business Intelligence by up to four times. They have made an improvement of up to two times in snapshot and thin provisioning density.

With 4500 XIV systems now deployed worldwide, IBM has established this platform as a viable force in the data storage community, allowing IBM to capture new-name storage accounts. Now, they have optimized it to be the logical upgrade for any IBM client currently running on an existing XIV platform. In addition, with about 1,100 *new* storage customers running XIV, Gen 3 has proven itself to be a cost-effective alternative to Symmetrix VMAX from EMC, VSP from Hitachi, and the HP *P9000* series disk platform, among others. As indicated above, IBM has priced XIV Gen 3 to be very competitive against the EMC VMAX.

With additional performance, more cache, more ports and faster ports, InfiniBand connectivity, and improved capacity with SAS drives, IBM's XIV Gen 3 platform ideally is suited to provide the enterprise data center with the storage performance that it needs today and tomorrow. In addition, with improved storage efficiency, reduced space and energy requirements, along with improved data resiliency, XIV Gen 3 lowers the TCO for IT infrastructure for the enterprise data center. Moreover, with unused CPU and DRAM slots available, XIV is designed for long-range enterprise planning, and not for planned obsolescence. If you are looking to upgrade your existing storage environment, XIV Gen 3 may be the right horse to ride in this race. If you haven't looked at XIV before, now is the time to take a close look at XIV Gen 3.

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

David Reine **is a Senior Contributing Analyst for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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.

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.