THE CLIPPER GROUP
Navigator™

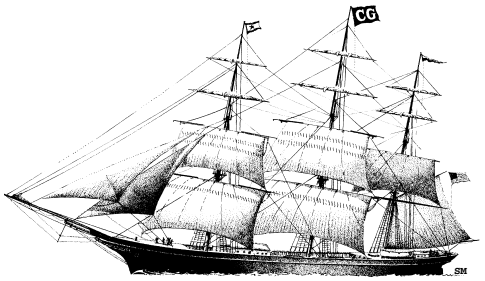

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008036R July 11, 2008

Xiotech Changes the Storage Paradigm

Analyst: David Reine

Management Summary

In the past, when shopping for a new car, we would always ask about the number of cylinders and how much horsepower the engine had, but we did not care about the car's "invisible engineering". For example, we would never ask about the technology of the engine. Did anyone care that it was an internal combustion engine designed to propel a vehicle weighing up to two tons, with another 500 pounds, or so, of fuel and anywhere from 200 to 1000 lbs. of passengers and cargo? No one cared, it simply worked. We would ask about mileage, but that was to compare relative costs between vehicles, not to consider a different engine design. Did anyone care about the amount of energy lost every time you stepped on the brakes, i.e., how much fuel was wasted? At $.39 per gallon, or even $1.39 per gallon, we were not too concerned. However, with gasoline having broken through the $4.00 per gallon ceiling, and heading upward toward a cost projected as high as $7.50 per gallon, the idea of an alternative – a hybrid engine where the vehicle can convert the energy previously wasted into usable electric power – becomes not only viable but also mandatory. The cost of fuel has changed the economic equation, changing the paradigm of automobile purchasing forever.

A similar story exists in the data center of every enterprise around the world. The availability of low-cost, open system x86 servers has led to the installation of hundreds, even thousands, of platforms in a single data center. Deployed with a single application, they utilize less than 20% of the available compute capability, resulting in server sprawl, and a terrible waste of energy, adding to the total cost of ownership (TCO). In fact, some enterprises are experiencing a problem trying to provide their data center with sufficient energy to power and cool the data center. In addition, each of these servers is connected to storage, in some cases *directly-attached (DAS)*, while in others, connected to a *storage area network (SAN)* or using *network-attached storage (NAS)*. In an attempt to get energy, maintenance, and administrative costs under control, data centers are changing their server paradigm by consolidating their server requirements onto multi-socket, multi-core platforms in order to lower their TCO. Unfortunately, most of these data centers do not recognize the opportunity available to them to do the same thing for their storage requirements. Storage arrays have become the internal combustion engine of the data center. Storage providers are doing their best to reduce the cost of disk arrays, while at the same time increasing their capacity. *But, they are doing little to improve the technology inside the array*, the basic infrastructure of drives and controllers. Yes, we have seen the introduction of iSCSI to reduce the TCO of communication access, but there has been little technological innovation to improve the reliability and performance of the storage infrastructure. Data centers implement RAID technology for *when* the disks will fail, not *if* they will fail. In fact, RAID 6, while necessary because of increased disk capacity, increases overhead and reduces performance due to lengthy rebuild time. One company, Xiotech, has changed that paradigm without changing the basic building block of the array, the disk device itself. To learn more about a disk array with a five-year warranty and no "disk events" please read on.

IN THIS ISSUE

Data Center Issues for Tomorrow

The data center is in crisis mode. Not some other data center, your data center. Server proliferation, along with the attached storage, is destroying your IT budget. You have made the decision to consolidate your servers with the latest multi-socket, multi-core servers, but that does not help the storage dilemma. With data doubling in capacity every 12 to 18 months, the IT budget for storage will soon be insufficient to meet the enterprise requirements. Capital costs are not the main issue, as the cost for storage, on a per TB basis, has been coming down. The costs associated with reliability and overhead, however, are another matter.

As storage capacity rises, the number of devices rises, proportionately. The reliability and the availability of any physical device may be the same, or even improve, but the overall expectation for any disk farm is that the data center will experience multiple failures. It is not a question of "if a disk device will fail", **it is the realization that multiple disks will fail.** Today, we are deploying disks with a capacity of up to 1TB. When a TB disk fails, the rebuilding process can take upwards of 10 to 12 hours, degrading performance on all of your business- and mission-critical applications. Moreover, we can realistically expect disk devices to continue to grow in capacity to 2TB, 4TB, and more. **Rebuild times are not going to get shorter anytime soon**. In order to protect your data, the IT staff has implemented any number of safety nets, including the implementation of RAID technology to ensure the integrity of your data. Unfortunately, the risk of a second device failure is so great now that data centers have been forced to take the extreme action of deploying a RAID-6 architecture to guard against that eventuality. With RAID-6, the IT staff must configure a second parity device, effectively *reducing* array capacity and application availability, and *increasing* the TCO at the same time.

In addition to the impact on reliability and availability is the effect of uncontrolled data growth on overhead and services. What are the costs for these today? What can you expect tomorrow? The costs for floor space, administration, maintenance, and energy continue to rise at an alarming rate. At the current rate of growth, in five years, you could be managing six-to-eight times as much data. Does the enterprise have enough floor space available in the data center to support that growth? Will the enterprise have to build a new data center at significant cost, i.e., millions of dollars? How many new administrators will you have to bring onboard to manage it? What will be the impact on the maintenance budget to support the added devices? Moreover, **will the local public utility even be able to supply the data center with enough energy to keep all of those drives spinning**, let alone the data center trying to budget for it? Depending on where your data center is located, you could be paying anywhere from $.05 per KWh to $.18 per KWh, *today*. If the projected price of gasoline is any indicator, your energy costs could triple in the years to come. Even though capital expenses are going down, operational expenses are going up even more. The data center will be forced to do even more with even less. In five years, **how much storage will you have?** What will be the cost of break-fix services? What will be the cost of the impact of the outage on the ability of the enterprise to respond to customers and partners?

Drive technology has followed Moore's Law from a capacity standpoint, but some of the storage cornerstones (RAID, SCSI, and general disk design) are beginning to show their age and are impeding progress towards improved reliability, manageability, and scalability. Processor performance increases dramatically year over year, with multiple cores and multiple threads, along with faster clock speeds, disk device capacity inches upward, up to 1TB in 2007. However, the storage architecture has remained the same for several decades, with basic disk technology the same for the past 35 years, basic RAID technology for 30 years; and SCSI for 22 years. Innovations, all, when they were introduced, however, today's data center is a different data center.

Another area of concern is scalability and the protection of the data center's storage investment. What is the scalability of your existing storage solution? If it cannot grow to meet tomorrow's needs, you must plan today to replace it with a solution that can.

The data center needs to improve the reliability, performance, and ease of use of their storage with a scalable solution that will meet the long-term requirements of the enterprise. The IT staff needs to deploy an innovative solution in the realm of storage, to the same extent that they have with server technology. One company that believes in the need for storage innovation is Xiotech.

Xiotech has recognized this need and acknowledged that they did not have the resources necessary to address these issues on their own. In order to achieve the innovation necessary to make a difference, they acquired one of Seagate's R&D groups, one that had been working on just such an innovative technology for the past five years, enabling Xiotech to change the storage paradigm. That technology is now available in Xiotech's *Intelligent Storage Element* (*ISE*), the new building block for reliable storage scalability and *ICON Manager*, providing automated integration to the data center.

Xiotech Technology

Xiotech started out with a set of goals to change the face of storage. First, and foremost, they wanted to eliminate *service events* requiring human attention and, likely, parts replacement. To accomplish this, they needed to improve the reliability of storage by a factor of 100. They realized that they could not improve the reliability of the disk device itself, think 1TB drive, but they could improve the reliability of a storage array. One byproduct of improving reliability and eliminating service events can be a longer warranty. An extended warranty would lower the storage TCO and convey confidence to the data center in the stability of the platform.

Lower costs are obviously good, but not at the expense of performance. Any game-changing design must have industry-leading performance. Moreover, it has to be scalable, capable of expanding from as little as one terabyte to a petabyte, in order to protect the investment that enterprises make in their technology. Not only reliability, performance, and scalability, the new platform has to have ease of use. Any new technology needs to take advantage of integrated automation features in order to reduce the burden on the administrative staff and enable a single administrator to manage more storage.

Xiotech designed ISE, configured with self-contained *DataPacs*, managed reliability controllers, power and cooling units, and battery modules, to meet these goals. They designed the DataPac to be self-healing, with all failures (yes, disk drives will continue to fail) resolved in place, without the requirement of a service call. Because of this innovative technology, Xiotech is able to achieve a hundred-fold improvement in reliability and offer a five-year warranty on their ISE-based storage systems, removing significant service costs, and downtime, from the data center budget. Xiotech confirmed this with 15 months of exhaustive testing, covering 208 ISE and 5,900 drives, **with no service events**. Xiotech also achieved their goal of attaining scalability to 1PB and proved, via Storage Performance Council benchmarks, the exceptional price/performance of their technology.

In order to achieve the desired levels of reliability, scalability, and performance, and to lower service costs, Xiotech set out to re-engineer the basic storage infrastructure to prevent failures and to repair those that did occur in place. They implemented improved environmentals to reduce heat and vibration and took a new approach to integrated drive and controller software, bringing cache (1GB per ISE), RAID technology, and drive management closer to the actual drives. They employ preventive reconditioning to prevent errors from occurring, and they have implemented the new ANSI T10-DIF standard for end-to-end error detection and correction. In terms of repair, ISE will automatically recreate failed data in place, at a granular level, greatly improving recovery time. Remedial reconditioning enables Xiotech's ISE to repair any surface on the fly to regain access to failed tracks. They took NASA's approach to reliability, two of everything within each ISE: two DataPacs, two managed reliability controllers to provide high performance and local RAID, cache and drive management, two battery modules for up to 96 hours of protection, and two highly efficient and reliable power and cooling units.

Each DataPac is a sealed unit with either 10 or 20 drives. Xiotech allocates 80% of each drive to the user and retains the remaining 20% for sparing. The DataPac supports drive reconditioning, granular recovery, and self-healing to achieve greater reliability. According to Xiotech, as many as 75% of the disks returned to manufacturers for repair are reported as no-trouble-found (NTF) once tested. What does this mean? First, fewer service calls mean better response time for users. Fewer unplanned, or planned, outages improve the TCO of the data center environment. Fewer replacements also mean improved security, as NTF means that repair technicians, and any number of other people, have access to your data.

Xiotech provides three different DataPacs, configurable to the specific needs of the enterprise. They configure the *Performance Tier* with either ten 146GB or twenty 73GB 15K drives for

a capacity of 1.1TB, the *Balanced Tier* with ten 300GB 15K drives with a capacity of 2.4TB, and a *Capacity Tier* containing ten 1TB 7.2K drives for a total of 8TB of active data. These DataPacs can be connected to servers through ISE via DAS or SAN, protecting the investment made in Xiotech hardware.

Xiotech's *ICON Manager* provides storage management and provisioning from a single web-based interface. It enables an IT administrator with multiple views of stored data: Virtual, Server, Storage, and Physical. In the old paradigm, storage management (i.e. provisioning) must be done from two, three, and, sometimes, four different consoles and management systems.

With ICON Manager, Xiotech provides a new storage management paradigm for automated integration. ICON Manager is a storage management system that does the following.

- **Communicates and integrates** storage with applications such as Windows and VMware;
- **Automates** and coordinates operations required between storage systems and applications for things like provisioning; and
- **Eliminates** multiple consoles, repetitive steps, expertise required, and the risk of making mistakes.

With web services, the IT staff has an easy-to-use tool to create and append management applications to ICON Manager. They can also use development platforms like Microsoft's *.NET* to create customized plug-ins for reporting, integration, and automation. Xiotech provides a full-featured suite of replication, data protection, and availability solutions, ranging from Snapshot, to CDP, to online storage services.

Xiotech has taken this innovative technology and created a pair of high-reliability, high-performance storage platforms: *Emprise 5000* for an enterprise looking for a low-cost, entry-level platform, and *Emprise 7000*, for an enterprise looking for the highest reliability and a scalable SAN solution.

Emprise Platforms

ISE and Icon Manager provide the building blocks that Xiotech uses to produce and deliver outstanding reliability in their two easy-to-use storage solutions: *Emprise 5000* and *Emprise 7000*, both coming with a five-year warranty to reinforce Xiotech's commitment to event-free storage.

Emprise 5000 is a low-cost, 3U DAS or switch-attached self-enclosed storage solution with linear scalability, from 1.1 up to 16 TB, for entry-level or special-purpose requirements, and superior performance. With a large capacity and compact size, Emprise 5000 is an ideal storage solution for a remote office or department. It is easy to deploy and use. Emprise 5000 supports either single or dual DataPacs[1] configured with either RAID-5 or RAID-10 architecture. The data center can attach up to 64 server connections to Emprise 5000, via direct- or switch-attached Fibre Channel (FC), with up to 128 volumes. Emprise 5000 has a 1GB cache with a battery-backup to hold the status of the cache for up to 96 hours. The Emprise technology supports a wide-range of operating environments, including Microsoft *Windows*, and VMware *ESX Server*, along with Red Hat and SUSE *Linux*.

Emprise 5000 offers the ISE self-healing technology with patented diagnostic and reconditioning processes to prevent issues and repair components in place, avoiding most service calls. Emprise 5000 has Snapshot data replication and supports the T10-DIF standard for end-to-end data integrity. The IT staff can hot-swap Managed Reliability Controllers, DataPacs, power supplies, cooling fans, and batteries.

Results from benchmarking submitted to the Storage Performance Council provide a clear indication that Xiotech has been able to meet their performance goals with this new technology. On a price/performance basis for the *SPC-1*[2] benchmark, Xiotech's Emprise 5000, with a RAID-10 configuration with twenty 146 GB drives distributed over two DataPacs, had the best rating of those systems submitted, with a cost ratio of $3.53/SPC-1 IOPS, just ahead of a Sun *StorEdge* system. This system has a performance rating of 5,892 SPC-1 IOPS. Xiotech also submitted testing for an Emprise 5000 system configured with forty 73GB drives with a performance rating of 8,720 SPC-1 IOPS and a cost ratio of $4.19, confirming their ability to meet high-performance requirements. Emprise 5000 also held a leading position, on a price/performance basis, on the

[1] Performance Tier, Balanced Tier, and Capacity Tier DataPacs are supported.

[2] SPC-1 is an application-level benchmark suite that generates a workload designed to emulate the typical functions of transaction-oriented database applications.

SPC-2[3] benchmark, with a cost ratio of $32.25/MB/second, edging out a Sun STK array.

Emprise 7000 is a virtualized, clustered solution that uses the same infrastructure as the Emprise 5000 system, but with management by ICON Manager, significantly more scalability, local and remote replication, and protection. The data center can create an Emprise 7000 storage environment containing up to 64 ISEs that scales from 1.1TB to 1PB of storage capacity with an architecture that virtually eliminates the vast majority of service events. With four FC interfaces standard, and four additional interfaces available, either FC or iSCSI, Emprise 7000 can support up to 248 server connections and 1,024 storage volumes. With 64 ISEs configured, Emprise 7000 supports up to 64GB of cache, with support for the same operating environments as the Emprise 5000 system.

Emprise 7000 employs the same self-healing technology as Emprise 5000, but extends data replication to a full feature set including:

- *TimeScale Copy* for local copy, mirror, and volume swap within a system;
- *TimeScale GeoRep* for async and sync remote replication between Emprise 7000 and *Magnitude 3D* systems;
- *TimeScale GeoRAID* for automatic management of two images of a volume over two locations within a storage system – with transparent failover during an outage with no host software required; and
- *TimeScale CDP* for continuous data protection with recovery to any point in time.

Emprise 7000 has significantly more storage management capabilities with ICON Manager providing point-and-click operation for ease-of-use, enabling a single administrator to manage even more capacity; ICON Services providing automated storage management in Windows environments; customized reporting; integration with other web service aware tools; and Intelligent Provisioning for automatic volume growth without the risk of over-subscription.

Conclusion

In an era when managing the operational expenses of the enterprise is more critical than managing the capital expenses, the data center staff must use every tool available to consolidate and virtualize the server and storage resources that comprise the IT infrastructure. They must also be wary of "putting too many eggs in a single basket". If you lose that "basket", you lose access to all of the mission- and business-critical data that your enterprise requires to execute day-to-day business. You must protect your data with the most reliable infrastructure available.

Every disk drive will fail; indeed, failures become a regular event in larger installations. The measuring stick of reliability for disks has always been "MTBF" – Mean Time *Between* Failures! In order to create a more reliable environment, and control your operating expenses, the data center must be prepared to change the storage paradigm, and change it radically. This is exactly what Xiotech has done with their Intelligent Storage Element and their Emprise product line. If your data center is fed up with multiple disk failures that cause significant IT outages, you need to look at a storage system that is based on the concept of no storage events and backed up by a five-year warranty. Xiotech may have the answer to your storage dilemma.

[3] SPC-2 is an application-level benchmark suite that generates multiple workloads designed to emulate the sequential data access patterns of video-on-demand.

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 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.